

INSTRUCTIONAL DRAFT

Commitment Admissibility and Integrity in Disturbance-Coupled Exploration Systems

**Irreversible thresholds, ignorance dominance, and the
governance of frontier infrastructure**

INSTRUCTIONAL DRAFT

Abstract

Exploration and infrastructure systems increasingly operate in environments where the target condition cannot be fully known before action is required. Lunar volatile exploration, geothermal development, carbon storage, critical minerals, seabed systems, orbital infrastructure, grid interconnection, autonomous systems, and geological disposal all exhibit this structure. In each case, evidence is indirect, interpretation remains non-unique, and verification often requires disturbance, deployment, access, emplacement, or institutional commitment. The act that produces knowledge can also alter the system, create exposure, concentrate capital, establish access pathways, generate dependency, or set precedent.

Existing decision frameworks address risk, readiness, value of information, real options, adaptive management, resource classification, assurance, and optimization. These frameworks are necessary, but incomplete where learning and irreversible commitment are structurally coupled. They do not fully govern whether a commitment should be allowed to form before uncertainty has ceased to be decision-dominant. They also do not determine whether an initially defensible commitment remains within its admissible basis as evidence changes, assumptions expire, dependencies accumulate, or authority conditions degrade.

This paper introduces two linked governance conditions: commitment admissibility and commitment integrity. Commitment admissibility governs entry. It asks whether a proposed irreversible commitment may be considered under the evidence, uncertainty, and authority conditions present at the threshold. Commitment integrity governs persistence. It asks whether an existing or emerging commitment remains defensible as the conditions that supported entry change over time.

The framework defines disturbance-coupled exploration systems, irreversible thresholds, admissibility flips, ignorance dominance, dependence topology, refusal credibility, authority coherence, precedent externality, and five regimes of irreversibility: subsurface, placement, topology, authority, and precedent. Lunar south polar in-situ resource utilization is used as the forcing case because volatile evidence can justify bounded exploration without justifying infrastructure commitment. Geological disposal provides the limiting case for long-duration commitment integrity. Critical minerals demonstrate prospectivity-to-commitment conversion. Together, these cases show that frontier systems require governance before commitment hardens and integrity review after commitment begins.

INSTRUCTIONAL DRAFT

Table of Contents

<i>Commitment Admissibility and Integrity in Disturbance-Coupled Exploration Systems</i>	1
Irreversible thresholds, ignorance dominance, and the governance of frontier infrastructure	1
Abstract	2
1. Introduction: When Learning Becomes Commitment	16
1.1 The failure mode	17
1.2 Entry and persistence	17
1.3 Why disturbance matters	18
1.4 Why outcomes are insufficient	19
1.5 Contribution of the paper.....	20
1.6 Structure of the paper.....	20
2. Disturbance-Coupled Exploration Systems	21
2.1 Defining properties	22
2.1.1 The target condition is not directly observable at commitment scale	22
2.1.2 Available evidence admits multiple plausible interpretations.....	23
2.1.3 Verification requires intervention, access, deployment, or disturbance	23
2.1.4 The intervention changes the system or creates exposure	24
2.1.5 Future decisions become path-dependent.....	25
2.2 Disturbance and exposure are not the same	25
2.3 Evidence classes in disturbance-coupled systems	26
2.4 Ambiguity persistence and decision dominance	27
2.5 The exposure trajectory	28
2.6 Disturbance-coupled systems and institutional behavior	29
2.7 Implications for governance	30
2.8 Section conclusion	30
3. Analytical Approach	31
3.1 Plausible-state reasoning	31
3.2 Irreversible-threshold analysis	32
3.3 Dependence topology	33
3.4 Authority coherence analysis	34
3.5 Disturbance-coupled evidence analysis.....	35

INSTRUCTIONAL DRAFT

3.6 Cross-domain transfer testing	36
3.7 Case selection.....	37
3.7.1 Lunar south polar ISRU as the forcing case	37
3.7.2 Geological disposal as the limiting case for integrity	37
3.7.3 Critical minerals as the transfer case	38
3.8 Analytical outputs	38
3.9 Boundary of the approach	39
3.10 Section conclusion.....	39
4. Existing Frameworks and the Missing Category.....	40
4.1 Risk management.....	40
4.2 Readiness frameworks	41
4.3 Real options.....	42
4.4 Value of information	43
4.5 Adaptive management	44
4.6 Resource classification	44
4.7 Planetary protection, mission assurance, and safety frameworks.....	45
4.8 Project finance and investment diligence	46
4.9 Systems engineering and architecture practice.....	47
4.10 The missing category: threshold governance	48
4.11 Comparative summary	49
4.12 Section conclusion.....	49
5. Formal Primitives and Structural Logic.....	50
5.1 Commitment.....	50
5.2 Decision Boundary	51
5.3 Plausible State Space	51
5.4 Exposure	52
5.5 Reversal Burden and Reversal Capacity	53
5.6 Refusal Credibility	54
5.7 Authority Coherence.....	55
5.8 Admissibility Flip	56
5.9 Ignorance Dominance	56

INSTRUCTIONAL DRAFT

5.10 Dependence Topology.....	58
5.11 Commitment Admissibility	59
5.12 Governance Posture.....	60
5.13 Admissible Basis.....	61
5.14 Commitment Integrity	61
5.15 Precedent Externality.....	62
5.16 Relationship Among Primitives	63
5.17 Section Conclusion.....	64
6. <i>Commitment Admissibility: The Entry Problem</i>.....	65
6.1 Why entry requires governance	65
6.2 Admissibility is distinct from feasibility.....	66
6.3 Admissibility is distinct from risk tolerance	66
6.4 Admissibility is distinct from evidence strength.....	67
6.5 The admissibility set	67
6.6 The admissibility flip condition	68
6.7 Ignorance dominance.....	69
6.8 Minimum evidence before admissibility	69
6.9 Refusal authority as an admissibility condition	70
6.10 Value of information and the admissibility boundary	70
6.11 Admissibility outcomes	71
6.12 Admissible to consider is not approved to proceed	72
6.13 Admissibility under coupled irreversibility	72
6.14 Admissibility and communication	73
6.15 Boundary conditions and expiration	73
6.16 Section conclusion.....	74
7. <i>Governance Posture: Between Admissibility and Action</i>	74
7.1 Why admissibility needs a posture layer.....	75
7.2 Governance posture is not decision authority.....	76
7.3 The three governance postures.....	77
7.3.1 Proceed-Compatible.....	77

INSTRUCTIONAL DRAFT

7.3.2 Deferral-Indicated	78
7.3.3 Refusal-Required.....	79
7.4 Posture defensibility conditions.....	80
7.4.1 Evidence authority.....	80
7.4.2 Uncertainty dominance	81
7.4.3 Exposure containment.....	81
7.4.4 Refusal credibility.....	81
7.4.5 Dependency bounds	82
7.4.6 Authority coherence	82
7.4.7 Timing and urgency	82
7.4.8 Communication effects.....	83
7.4.9 Precedent containment	83
7.5 Posture classification across irreversibility regimes	84
7.6 Examples of posture classification	84
7.6.1 Lunar ISRU	84
7.6.2 Grid interconnection	85
7.6.3 Critical minerals.....	85
7.6.4 Geological disposal.....	85
7.7 Posture and institutional responsibility	86
7.8 Posture expiration and reclassification.....	86
7.9 The posture layer and commitment integrity	87
7.10 Section conclusion.....	88
8. Regimes of Irreversibility.....	88
8.1 Why irreversibility requires regimes.....	89
8.2 Subsurface irreversibility	90
8.2.1 Admissibility implication	90
8.3 Placement irreversibility	91
8.3.1 Admissibility implication	92
8.4 Topological irreversibility.....	92
8.4.1 Admissibility implication	93
8.5 Authority irreversibility	94

INSTRUCTIONAL DRAFT

8.5.1 Admissibility implication	94
8.6 Precedent irreversibility	95
8.6.1 Admissibility implication	96
8.7 Coupled irreversibility	96
8.7.1 Admissibility implication	97
8.8 Regime-specific failure patterns	98
Subsurface failure	98
Placement failure	98
Topological failure.....	98
Authority failure	98
Precedent failure	98
Coupled failure	98
8.9 Regime-specific governance questions.....	99
8.10 Relationship to commitment admissibility	99
8.11 Relationship to governance posture.....	100
8.12 Relationship to commitment integrity	100
8.13 Section conclusion.....	101
9. <i>Commitment Integrity: The Persistence Problem</i>.....	101
9.1 Why admissibility is not enough.....	102
9.2 The admissible basis	103
9.3 Integrity is distinct from performance	104
9.4 Integrity degradation pathways	104
9.4.1 Evidence drift.....	105
9.4.2 Assumption expiry	105
9.4.3 Dependency accumulation	106
9.4.4 Authority degradation	107
9.4.5 Precedent expansion.....	108
9.5 Integrity states	108
9.5.1 Maintain.....	109
9.5.2 Constrain.....	109
9.5.3 Re-evaluate	110
9.5.4 Terminate	111
9.6 Integrity and the five regimes of irreversibility	111

INSTRUCTIONAL DRAFT

9.6.1 Subsurface integrity	111
9.6.2 Placement integrity.....	112
9.6.3 Topological integrity	112
9.6.4 Authority integrity	112
9.6.5 Precedent integrity	112
9.6.6 Coupled integrity	112
9.7 Integrity and time	113
9.8 Integrity and evidence freeze	114
9.9 Integrity and precedent control.....	114
9.10 Integrity and termination	115
9.11 Integrity and institutional memory	115
9.12 Relationship to admissibility and posture	116
9.13 Integrity failure modes	117
1. Basis drift.....	117
2. Scope creep	117
3. Dependency creep	117
4. Authority erosion	117
5. Communication ratchet	117
6. Precedent escape.....	117
7. Trigger breach	117
8. Stale reliance	117
9. Coupled escalation	118
9.14 Section conclusion.....	118
10. Forcing Case: Lunar South Polar ISRU	118
10.1 Why lunar ISRU is the forcing case.....	119
10.2 Detectability limits	120
10.3 Ambiguity persistence	121
10.4 Environmental response	122
10.5 Architecture under constraint	124
10.6 Surface commitment classes.....	125
10.7 Signal-to-lock-in chain	126
10.8 Applying the five regimes of irreversibility.....	127
10.8.1 Subsurface irreversibility	127
10.8.2 Placement irreversibility	127
10.8.3 Topological irreversibility	128
10.8.4 Authority irreversibility	128

INSTRUCTIONAL DRAFT

10.8.5 Precedent irreversibility.....	128
10.8.6 Coupled lunar irreversibility.....	128
10.9 Minimum evidence before lunar commitment	128
10.9.1 Multi-modal convergence.....	129
10.9.2 Scale match	129
10.9.3 Ambiguity reduction	129
10.9.4 Disturbance justification	129
10.9.5 Refusal preservation	129
10.9.6 Dependency containment	129
10.9.7 Precedent containment.....	129
10.10 Applied admissibility logic	130
10.11 Governance posture in the lunar case	130
Bounded remote analysis	130
Bounded reconnaissance.....	131
Intrusive verification	131
Fixed power or support placement	131
ISRU-dependent architecture.....	131
10.12 Commitment integrity in the lunar case	132
Evidence drift	132
Assumption expiry	132
Dependency accumulation	132
Authority degradation	132
Precedent expansion	132
10.13 Precedent in lunar ISRU.....	133
10.14 Lunar implications.....	134
10.15 Section conclusion	135
<i>11. Limiting Case: Geological Disposal and Nuclear Waste.....</i>	<i>135</i>
11.1 Why geological disposal is the limiting case.....	136
11.2 Entry admissibility in geological disposal.....	137
11.3 The admissible basis in disposal systems	137
11.4 Subsurface integrity	138
11.5 Placement integrity	139
11.6 Topological integrity.....	140
11.7 Authority integrity	141
11.8 Precedent integrity.....	141
11.9 Integrity degradation pathways in geological disposal.....	142

INSTRUCTIONAL DRAFT

11.9.1 Evidence drift	142
11.9.2 Assumption expiry	142
11.9.3 Dependency accumulation	143
11.9.4 Authority degradation.....	143
11.9.5 Precedent expansion	143
11.10 Integrity states in geological disposal	143
Maintain	143
Constrain	144
Re-evaluate	144
Terminate	144
11.11 Why geological disposal extends the theory	145
11.12 Comparison with lunar ISRU	145
11.13 Section conclusion	146
12. Transfer Case: Critical Minerals.....	146
12.1 Why critical minerals are a transfer case	147
12.2 Prospectivity is not commitment authority.....	148
12.3 Ambiguity persistence in critical minerals	149
12.4 Disturbance and evidence generation	150
12.5 Land, access, and placement irreversibility	151
12.6 Topological irreversibility in critical minerals	152
12.7 Authority irreversibility and strategic urgency.....	153
12.8 Precedent irreversibility in critical minerals	154
12.9 Critical minerals and admissibility	154
Exploration-stage admissibility	155
Land-control admissibility	155
Access admissibility	155
Processing admissibility	155
Infrastructure admissibility	155
Capital admissibility	155
Extraction-pathway admissibility.....	155
12.10 Governance posture in critical minerals.....	155
Bounded exploration	156
Expanded drilling.....	156
Land consolidation	156
Processing assumptions	156
Offtake or strategic supply commitments	156
Infrastructure commitment	156

INSTRUCTIONAL DRAFT

12.11 Commitment integrity in critical minerals.....	156
Evidence drift	156
Assumption expiry	157
Dependency accumulation	157
Authority degradation	157
Precedent expansion	157
12.12 The critical minerals failure mode.....	157
12.13 What the critical minerals case adds to the theory.....	158
12.14 Section conclusion	158
<i>13. Cross-Domain Transfer Matrix.....</i>	<i>159</i>
13.1 Transfer logic	160
<i>13.2 Geothermal Systems</i>	<i>160</i>
13.2.1 Signal and evidence	160
13.2.2 Ambiguity persistence.....	161
13.2.3 Disturbance / exposure action	161
13.2.4 Irreversibility regimes	162
13.2.5 Admissibility question.....	162
13.2.6 Integrity question.....	162
<i>13.3 Carbon Storage</i>	<i>162</i>
13.3.1 Signal and evidence	163
13.3.2 Ambiguity persistence.....	163
13.3.3 Disturbance / exposure action	163
13.3.4 Irreversibility regimes	164
13.3.5 Admissibility question.....	164
13.3.6 Integrity question.....	164
<i>13.4 Seabed and Marine Infrastructure.....</i>	<i>164</i>
13.4.1 Signal and evidence	165
13.4.2 Ambiguity persistence.....	165
13.4.3 Disturbance / exposure action	165
13.4.4 Irreversibility regimes	166
13.4.5 Admissibility question.....	166

INSTRUCTIONAL DRAFT

13.4.6 Integrity question.....	166
13.5 Grid Interconnection.....	166
13.5.1 Signal and evidence.....	166
13.5.2 Ambiguity persistence.....	167
13.5.3 Exposure action.....	167
13.5.4 Irreversibility regimes.....	168
13.5.5 Admissibility question.....	168
13.5.6 Integrity question.....	168
13.6 Orbital Systems.....	169
13.6.1 Signal and evidence.....	169
13.6.2 Ambiguity persistence.....	169
13.6.3 Exposure action.....	169
13.6.4 Irreversibility regimes.....	170
13.6.5 Admissibility question.....	170
13.6.6 Integrity question.....	170
13.7 Autonomy Near Irreversible Thresholds.....	170
13.7.1 Signal and evidence.....	170
13.7.2 Ambiguity persistence.....	171
13.7.3 Exposure action.....	171
13.7.4 Irreversibility regimes.....	171
13.7.5 Admissibility question.....	171
13.7.6 Integrity question.....	172
13.8 Cross-domain transfer matrix.....	172
13.9 What the transfer matrix shows.....	174
13.10 Transfer limits.....	174
13.11 Section conclusion.....	175
14. Institutional Implications.....	175
14.1 Evidence is not authorization.....	176
14.2 Refusal and deferral are valid outcomes.....	177
14.3 Decision records must be reconstructible.....	177

INSTRUCTIONAL DRAFT

14.4 Evidence freeze and change control	178
14.5 Precedent must be governed.....	179
14.6 Authority must remain aligned with consequence	179
14.7 Decision exposure should be identified early	180
14.8 Minimum evidence should be tied to the next threshold.....	181
14.9 Governance posture should be separated from executive decision.....	182
14.10 Commitment integrity requires active preservation.....	182
14.11 Institutional design implications	183
14.11.1 Decision threshold registers	183
14.11.2 Assumption registries	183
14.11.3 Commitment trigger registers	183
14.11.4 Evidence authority labels.....	183
14.11.5 Refusal conditions	183
14.11.6 Precedent controls.....	183
14.11.7 Decision change logs.....	184
14.12 Implications by institutional type	184
14.12.1 Mission agencies	184
14.12.2 Infrastructure developers	184
14.12.3 Investors and investment committees	184
14.12.4 Regulators and permitting bodies	184
14.12.5 Boards and executive sponsors	184
14.12.6 Scientific and technical teams	184
14.12.7 Autonomous systems governance bodies.....	185
14.13 What the framework does not imply.....	185
14.14 Section conclusion	185
15. <i>Why Outcomes Cannot Validate Threshold Decisions</i>	186
15.1 The outcome fallacy	186
15.2 Success does not prove admissibility	187
15.3 Failure does not prove inadmissibility	188
15.4 Outcomes as integrity evidence, not retroactive justification	189
15.5 Threshold provenance	190
15.6 The four outcome-decision combinations	191
15.7 Why irreversible systems amplify outcome bias.....	191
15.7.1 Evidence baseline alteration	191
15.7.2 Sunk-cost reinterpretation	192
15.7.3 Narrative hardening	192

INSTRUCTIONAL DRAFT

15.7.4 Memory loss	192
15.7.5 Precedent simplification	192
15.8 Decision quality is evidence-bound and time-bound	192
15.9 Implications for scientific and institutional claims	193
15.10 Outcomes and the Precedent Card	193
15.11 Section conclusion	194
16. Discussion.....	195
16.1 What the framework adds	195
16.1.1 Feasibility is not admissibility	195
16.1.2 Evidence is not authorization	196
16.1.3 Uncertainty is not ignorance dominance	196
16.1.4 Exploration is not commitment	197
16.1.5 Entry is not persistence	197
16.1.6 Outcomes do not validate thresholds.....	198
16.2 Relationship to existing domains of practice.....	198
16.3 Generality and limits.....	199
16.4 The danger of over-formalization	199
16.5 The role of judgment.....	200
16.6 Implications for science	201
16.7 Implications for infrastructure and capital	201
16.8 Implications for frontier governance	202
16.9 Limitations	202
16.10 Future work	203
16.10.1 Formalization of admissibility sets	203
16.10.2 Metrics for reversal burden and reversal capacity	203
16.10.3 Domain-specific threshold taxonomies	203
16.10.4 Commitment integrity diagnostics	203
16.10.5 Empirical case reconstruction.....	204
16.10.6 Governance record design	204
16.10.7 Relationship to assurance and regulation	204
16.10.8 Organizational adoption.....	204
16.11 Theoretical contribution	204
16.12 Practical contribution.....	205
16.13 Section conclusion	205
17. Conclusion: The Cost of Error Is Entrenchment	206

INSTRUCTIONAL DRAFT

17.1 The central claim	206
17.2 What the forcing cases show	207
17.3 What the framework changes.....	208
17.4 The role of refusal.....	209
17.5 The final contribution	210
17.6 Closing statement	210

INSTRUCTIONAL DRAFT

1. Introduction: When Learning Becomes Commitment

Exploration often begins with a signal.

A hydrogen indication near the lunar south pole. A geophysical anomaly beneath a geothermal prospect. A resource model suggesting mineralization. A storage formation that appears suitable for carbon injection. A seabed route that looks feasible. A point of interconnection that appears commercially attractive. A repository site that seems geologically favorable. An orbital deployment path that appears technically achievable.

Signals organize attention. They make some futures more plausible than others. They give institutions a basis for further investigation. They also create pressure. Once a signal appears promising, the next institutional impulse is usually to refine, test, secure, sequence, finance, or advance the opportunity.

In many systems, that impulse is reasonable. More information improves the decision. Additional work reduces uncertainty. Early action preserves optionality.

In disturbance-coupled exploration systems, the sequence is more dangerous.

The actions required to improve knowledge may also begin creating commitment. Drilling can reveal subsurface structure while anchoring capital and site exposure. Excavation can verify a volatile signal while disturbing the system being evaluated. Injection can test storage behavior while initiating pressure, migration, liability, and permanence. A rover traverse can produce knowledge while establishing a corridor. Queue entry can reveal grid constraints while creating site, capital, and interconnection exposure. A pilot deployment can test feasibility while creating precedent. Repository characterization can refine geological understanding while increasing institutional, political, and infrastructural commitment.

The problem is not uncertainty alone. All exploration contains uncertainty.

The problem is that uncertainty reduction can become inseparable from exposure creation.

A project may remain exploratory in stated intent while becoming commitment-bearing in practical effect. A “test” can begin site fixation. A “preliminary route” can structure future access. A “non-binding” partner discussion can create expectation. A “small” capital step can create continuation pressure. A “promising” resource signal can begin governing infrastructure before the resource has become decision-grade.

Many institutions treat uncertainty as a parameter to be managed after commitment. They assume that once a project is sufficiently attractive, capable teams can manage residual uncertainty through design, engineering, monitoring, finance, permitting, or adaptive execution.

INSTRUCTIONAL DRAFT

That assumption fails when the decision to learn, deploy, disturb, or sequence also changes the future decision space.

In such systems, the central question is not only whether the project can succeed.

The central question is whether the commitment should be allowed to form before uncertainty loses decision-changing power.

1.1 The failure mode

Conventional institutional processes often move from evidence to planning, from planning to execution, and from execution to risk management. This progression works when the project remains reversible during the learning period. It fails when the learning period itself creates irreversible exposure.

The failure mode has a repeated structure.

First, indirect evidence creates a plausible opportunity.

Second, institutional attention concentrates around the opportunity.

Third, a preferred site, route, resource interpretation, or architecture begins to form.

Fourth, capital, partners, access, support systems, or public language align around that path.

Fifth, future refusal becomes more costly, less credible, or institutionally embarrassing.

Sixth, the original uncertainty remains unresolved, but the organization is already inside a commitment pathway.

At that point, optimization and execution quality cannot recover the optionality that has already been lost.

A technically excellent team can optimize a premature commitment.

A strong capital partner can finance a structurally fragile pathway.

A capable mission architecture can harden around an unearned resource assumption.

A rigorous engineering process can refine a project whose threshold decision was never admissible.

The error is not necessarily poor execution. The error is commitment formation before the evidence, authority, and dependency structure could support it.

1.2 Entry and persistence

This paper develops two linked governance concepts to address that failure mode.

The first is **commitment admissibility**.

INSTRUCTIONAL DRAFT

Commitment admissibility governs entry. It asks whether a proposed irreversible commitment may be considered under the evidence, uncertainty, authority, and reversibility conditions present at the threshold.

A commitment is admissible only when unresolved uncertainty no longer has the power to materially change whether the action should proceed, defer, or be refused within the defined boundary. A commitment can be technically feasible, economically attractive, strategically important, and still inadmissible if plausible unresolved states would imply materially different commitment decisions.

The second is **commitment integrity**.

Commitment integrity governs persistence. It asks whether an existing or emerging commitment remains within its admissible basis as evidence changes, assumptions expire, dependencies accumulate, support systems harden, precedent expands, or refusal authority weakens.

An initially admissible commitment may later become unstable. The evidence that supported entry may degrade. Assumptions may expire. Dependencies may exceed their original bounds. Capital or public language may create continuation pressure. The authority to refuse, revoke, re-site, or re-sequence may become non-credible. A bounded action may begin functioning as precedent beyond its original scope.

Commitment integrity therefore asks a different question from admissibility. It does not ask whether entry was defensible at the time. It asks whether continuation remains defensible as the system hardens.

Together, admissibility and integrity cover the commitment lifecycle:

- before entry, whether the commitment may be considered;
- at the threshold, which governance posture remains defensible;
- after entry, whether the commitment remains within the basis that made it defensible.

1.3 Why disturbance matters

Disturbance is central because it links knowledge to exposure.

In many exploration systems, evidence cannot become more precise without physical intervention. Lunar volatile conditions may require probing, trenching, drilling, excavation, heating, or repeated traverse. Geothermal reservoirs may require drilling. Storage formations may require injection. Seabed routes may require survey and installation preparation. Critical mineral prospects may require drilling, access, and sampling. Orbital systems may require deployment to reveal traffic and coordination behavior. Autonomous systems may require operational exposure to reveal edge-case behavior.

INSTRUCTIONAL DRAFT

These actions can improve knowledge. They can also alter the system, create sunk costs, establish access paths, privilege specific locations, generate public or institutional expectation, or normalize future action.

Verification is therefore not neutral. It can produce evidence while changing the evidentiary baseline. It can reduce ambiguity while increasing path dependence. It can reveal the system while committing the institution to the pathway through which the system was revealed.

This is why disturbance-coupled systems require a governance layer before commitment hardens.

The issue is not whether learning should occur. The issue is whether the form of learning preserves refusal authority or silently consumes it.

1.4 Why outcomes are insufficient

One reason these systems are difficult to govern is that outcomes can mislead.

If a project succeeds, the original commitment may appear justified. If it fails, the original commitment may appear mistaken. But in irreversible systems, outcome evaluation is not sufficient to judge threshold decision quality.

A successful lunar infrastructure project may still have been inadmissible if it hardened around unresolved volatile assumptions that could have supported materially different decisions. A profitable mineral project may still normalize premature commitment standards. A carbon storage project may operate without early failure while still relying on containment assumptions whose long-term admissible basis has changed. A repository may appear stable during early monitoring while institutional memory, authority continuity, or public legitimacy degrade over time.

The quality of the threshold decision must be reconstructed from the evidence state at the time of commitment.

That reconstruction must include:

- the plausible states still consistent with available evidence;
- the irreversible threshold being crossed;
- the exposure created by action;
- the dependency topology created by sequencing;
- the credibility of refusal or reversal;
- the authority condition at the threshold;
- and the precedent implications of the decision.

This is the reason commitment governance cannot be reduced to prediction or performance evaluation. It is a provenance discipline.

INSTRUCTIONAL DRAFT

1.5 Contribution of the paper

This paper makes four contributions.

First, it defines **disturbance-coupled exploration systems** as a class of systems in which learning, intervention, exposure, and commitment formation are structurally linked.

Second, it introduces **commitment admissibility** as an entry condition for irreversible commitments under persistent uncertainty. Admissibility is not approval, optimization, feasibility, or prediction. It is the condition under which a proposed commitment may enter governed consideration.

Third, it introduces **commitment integrity** as a persistence condition. Integrity determines whether an existing or emerging commitment remains within its admissible basis as evidence, assumptions, dependencies, authority, and precedent evolve.

Fourth, it decomposes irreversibility into five regimes: subsurface, placement, topology, authority, and precedent. This decomposition shows why different systems can share the same decision problem even when their evidence environments differ.

The paper develops the framework conceptually, then applies it to three case classes. Lunar south polar ISRU serves as the forcing case for entry admissibility. Geological disposal and nuclear waste serve as the limiting case for long-duration commitment integrity. Critical minerals serve as the transfer case for prospectivity-to-commitment conversion. Additional domains, including geothermal, carbon storage, seabed systems, grid interconnection, orbital infrastructure, and autonomy, are used to test structural transfer.

1.6 Structure of the paper

Section 2 defines disturbance-coupled exploration systems.

Section 3 describes the analytical approach, including plausible-state reasoning, irreversible-threshold analysis, dependence topology, authority coherence, disturbance-coupled evidence generation, and cross-domain transfer testing.

Section 4 positions the framework relative to risk management, readiness, real options, value of information, adaptive management, resource classification, assurance, project finance, and systems engineering.

Section 5 defines the formal primitives: commitment, irreversible threshold, plausible state space, exposure, reversal burden, reversal capacity, refusal credibility, authority coherence, admissibility flip, ignorance dominance, dependence topology, commitment admissibility, commitment integrity, and precedent externality.

Section 6 develops commitment admissibility as the entry problem.

INSTRUCTIONAL DRAFT

Section 7 develops governance posture as the layer between admissibility and action.

Section 8 defines the five regimes of irreversibility.

Section 9 develops commitment integrity as the persistence problem.

Section 10 applies the framework to lunar south polar ISRU.

Section 11 uses geological disposal and nuclear waste as the limiting case for commitment integrity.

Section 12 examines critical minerals as a transfer case.

Section 13 extends the framework through a cross-domain transfer matrix.

Section 14 discusses institutional implications.

Section 15 explains why outcomes cannot validate threshold decisions.

Section 16 discusses contributions, limitations, and future work.

Section 17 concludes by arguing that in frontier systems, the cost of error is not failed execution alone. It is the entrenchment of commitments made before the system was sufficiently understood, and the persistence of commitments after their admissible basis has expired.

next

Thought for 4s

2. Disturbance-Coupled Exploration Systems

A disturbance-coupled exploration system is a system in which the target condition cannot be fully known before action is required, and the action required to reduce uncertainty also creates physical, institutional, capital, operational, or precedent exposure.

This definition is deliberately narrower than “uncertain system.”

All exploration involves uncertainty. Not all exploration systems couple uncertainty reduction to irreversible commitment. In some settings, additional information can be gathered without materially changing the system, narrowing future options, degrading refusal authority, or creating dependency. In disturbance-coupled systems, the act of learning is itself structurally implicated in commitment formation.

The defining problem is that knowledge and exposure do not accumulate at the same rate.

INSTRUCTIONAL DRAFT

Evidence may improve gradually.

Commitment may harden discretely.

Authority may degrade silently.

Dependency may form before it is formally recognized.

Precedent may be created by actions described as exploratory.

A system becomes disturbance-coupled when the path to better evidence passes through actions that can make later refusal, reversal, re-siting, re-sequencing, or abandonment materially more difficult.

2.1 Defining properties

A disturbance-coupled exploration system has five defining properties.

2.1.1 The target condition is not directly observable at commitment scale

The condition being evaluated cannot be observed directly at the scale required for the decision.

This may be because the target is buried, remote, distributed, dynamic, hazardous, inaccessible, politically constrained, or observable only through proxy signals.

Examples include:

- volatile distribution beneath or within lunar regolith
- geothermal reservoir structure
- permeability and pressure behavior in a storage formation
- critical mineral continuity
- seabed geotechnical conditions
- grid constraint topology
- orbital traffic behavior under deployment density
- repository performance over deep time
- autonomous system behavior under rare edge cases

The issue is not merely that observations are incomplete. The issue is that the decision often requires commitment at a scale different from the scale of observation.

A lunar volatile indication may be regionally meaningful but not site-commitment grade.

A geophysical anomaly may be prospect-relevant but not drilling-commitment grade.

A queue opportunity may be commercially interesting but not interconnection-commitment grade.

A repository site may be geologically promising but not yet integrity-grade across institutional time.

The gap between observation scale and commitment scale is where admissibility becomes necessary.

INSTRUCTIONAL DRAFT

2.1.2 Available evidence admits multiple plausible interpretations

In disturbance-coupled systems, the same evidence may remain consistent with multiple materially different realities.

This is ambiguity persistence.

The evidence narrows the search space but does not eliminate all decision-relevant alternatives. In many cases, the remaining alternatives do not merely imply different technical estimates. They imply different commitments.

A hydrogen signal may imply accessible ice, patchy volatile-bearing regolith, diffuse hydrogen, hydrated material, or mechanically difficult material. These states imply different site, power, excavation, processing, and infrastructure decisions.

A geothermal anomaly may imply a productive reservoir, a thermally attractive but hydraulically poor system, a fault-controlled resource with high drilling risk, or a zone whose productivity cannot be sustained.

A grid interconnection opportunity may imply bounded upgrade exposure, congestion-dominated economics, restudy risk, cost allocation exposure, or site-grid coupling failure.

A mineral prospect may imply an economically coherent deposit, a geologically interesting but discontinuous body, a processing-constrained resource, or a deposit whose infrastructure burden exceeds its strategic value.

The point is not that the evidence is weak. Evidence can be real, useful, and scientifically valid while remaining insufficient for commitment.

Ambiguity becomes governance-relevant when plausible interpretations would lead to different actions at the threshold.

2.1.3 Verification requires intervention, access, deployment, or disturbance

In disturbance-coupled systems, uncertainty cannot be reduced sufficiently through passive observation alone.

The system must be entered, disturbed, drilled, sampled, traversed, injected, deployed into, instrumented, or otherwise physically engaged.

This is why the governance problem is not solved by asking for “more data.” The relevant question is:

What kind of action is required to obtain the data, and what exposure does that action create?

INSTRUCTIONAL DRAFT

A drill hole may produce essential subsurface information, but it also creates site commitment, capital escalation, permitting posture, and continuation pressure.

An excavation may reveal volatile state, but it can alter thermal conditions, disturb regolith, mobilize volatiles, and contaminate later measurements.

An injection test may reveal pressure behavior, but it can also initiate plume evolution, monitoring obligation, liability exposure, and permanence.

A grid queue entry may reveal upgrade exposure and deliverability constraints, but it may also create deposits, study timelines, site coupling, gen-tie assumptions, and withdrawal costs.

A pilot deployment may reveal operational behavior, but it may also normalize a mode of use, create counterparty reliance, and set precedent.

Verification is therefore not neutral. It has its own commitment topology.

2.1.4 The intervention changes the system or creates exposure

The intervention required for learning can change the physical system, institutional environment, capital sequence, or future decision space.

Physical changes may include:

- regolith disturbance
- volatile mobilization
- thermal alteration
- pressure perturbation
- contamination
- seabed disturbance
- drilling pathways
- injection-induced changes
- infrastructure emplacement

Institutional changes may include:

- board alignment
- partner reliance
- regulatory posture
- public expectation
- capital sequencing
- offtake assumptions
- delegated authority
- loss of refusal credibility

Topological changes may include:

INSTRUCTIONAL DRAFT

- corridor formation
- access repetition
- power placement
- logistics dependency
- grid coupling
- shared-capacity exposure
- processing dependencies
- support infrastructure

Precedent changes may include:

- normalization of first use
- treatment of a pilot as proof of permissibility
- public framing of a site as primary
- reuse of one decision as a broader rule
- expectation that future actors may proceed similarly

This is why disturbance-coupled systems cannot be governed solely by technical risk reduction. The act of reducing technical uncertainty may increase governance exposure.

2.1.5 Future decisions become path-dependent

Once action begins, later decisions inherit the consequences of earlier choices.

This path dependence may be physical, financial, institutional, operational, or narrative.

A site that receives repeated access becomes easier to continue than abandon.

A route that receives planning effort becomes the default corridor.

A queue entry begins coupling the project to study outcomes and upgrade exposure.

A drill program begins concentrating capital and expectation around the target.

A support node privileges one operating geometry.

A pilot project becomes evidence that a type of action is acceptable.

A public statement makes future refusal more costly.

Path dependence is especially dangerous because it can form through actions that appear individually reversible.

A single action may be defensible.

A sequence of individually defensible actions may become indefensible.

The system-level commitment can emerge before any actor declares that commitment has occurred.

This is the reason dependence topology is central to the framework.

2.2 Disturbance and exposure are not the same

INSTRUCTIONAL DRAFT

Disturbance is one pathway into exposure, but it is not the only one.

A system can become commitment-bearing without physical disturbance. Grid interconnection, capital sequencing, site control, public positioning, partner reliance, and regulatory posture can all create irreversible exposure without changing the physical environment.

Conversely, not all physical disturbance is equally commitment-bearing. Some disturbance may be bounded, reversible in practical terms, or governed under conditions that preserve refusal authority.

The relevant question is not simply:

Does this action disturb the system?

The relevant question is:

Does this action create exposure that the current evidence, authority, and dependency structure cannot yet justify?

This distinction matters because it prevents the framework from becoming anti-action or anti-exploration. The framework does not prohibit disturbance in general. It asks whether the disturbance is being used as bounded evidence acquisition or whether it is beginning to authorize commitment.

Disturbance can be admissible when:

- the action is bounded;
- the evidence gain is material;
- alternatives remain credible;
- refusal authority is preserved;
- the site, route, or architecture is not prematurely privileged;
- termination conditions are explicit;
- the action does not create unbounded precedent.

Disturbance becomes commitment-bearing when:

- it privileges one site or geometry;
- it creates sunk costs;
- it alters the evidence baseline;
- it produces continuation pressure;
- it requires support systems that harden into dependency;
- it makes re-siting or refusal non-credible;
- it is treated as justification for infrastructure commitment.

2.3 Evidence classes in disturbance-coupled systems

INSTRUCTIONAL DRAFT

A recurring failure in these systems is evidence over-authorization.

Evidence that is valid for one purpose is used to justify a stronger purpose.

A signal that supports prospecting is treated as site-selection evidence.

A site-comparison model is treated as infrastructure evidence.

A bounded test is treated as architecture validation.

A resource indication is treated as dependency permission.

A queue position is treated as project progress.

A pilot is treated as precedent.

To prevent this, evidence must be classified by what it can legitimately support.

A useful evidence authority ladder is:

1. **Interest-generating evidence**
Evidence that justifies attention or further study.
2. **Prospect-ranking evidence**
Evidence that helps compare candidates or prioritize investigation.
3. **Site-comparison evidence**
Evidence that supports relative preference among sites, routes, or architectures.
4. **Bounded verification evidence**
Evidence sufficient to justify limited learning actions under explicit constraints.
5. **Design-supporting evidence**
Evidence sufficient to inform technical design assumptions within defined bounds.
6. **Commitment-authorizing evidence**
Evidence sufficient to support irreversible or hard-to-reverse commitment within a defined boundary.

Disturbance-coupled systems often fail when evidence from levels 1 through 4 is allowed to function as level 6.

The framework does not demand certainty. It demands that evidence not be asked to govern commitments beyond its authority.

2.4 Ambiguity persistence and decision dominance

Ambiguity persistence is the condition in which multiple plausible states remain consistent with available evidence.

Ambiguity becomes decision-dominant when those plausible states imply materially different commitment decisions.

This distinction is critical.

INSTRUCTIONAL DRAFT

Some ambiguity is tolerable. If all plausible states support the same bounded action, then ambiguity may not prevent admissibility. If plausible states differ in detail but not in commitment consequence, action may remain defensible.

Ambiguity becomes dominant when one plausible state supports proceeding, another supports deferral, and another requires refusal.

In that case, the issue is not the amount of uncertainty. It is the role of uncertainty in determining admissibility.

This is the difference between ordinary uncertainty and ignorance dominance.

Uncertainty is decision-relevant when it affects estimates.

Ignorance becomes decision-dominant when unresolved plausible states can still change whether a commitment is admissible.

A lunar volatile signal may remain ambiguous, but if all plausible states justify only bounded reconnaissance, then the admissible posture is stable. If one plausible state would justify infrastructure placement and another would make that placement premature or damaging, then the ambiguity is decision-dominant.

A grid project may face uncertain upgrade costs. If all plausible outcomes remain within bounded exposure, the commitment may be admissible. If some plausible outcomes create non-recoverable site-grid lock-in or withdrawal non-credibility, the uncertainty may be decision-dominant.

A repository may have residual geological uncertainty. If that uncertainty is bounded within containment and stewardship assumptions, integrity may be maintained. If new plausible states undermine the basis for long-duration containment or institutional responsibility, integrity may fail.

2.5 The exposure trajectory

Disturbance-coupled systems often move through a recognizable exposure trajectory.

Signal → Preference → Access → Support → Disturbance → Dependency → Commitment

The exact sequence varies by domain.

In lunar ISRU:

Volatile signal → site preference → access repetition → power/support placement → disturbance → corridor formation → excavation planning → ISRU dependency → infrastructure lock-in

In geothermal:

INSTRUCTIONAL DRAFT

Resource indication → lease/site focus → access and pad planning → drilling → reservoir commitment → offtake/capital dependence → infrastructure lock-in

In grid interconnection:

Site interest → POI assumption → queue entry → study exposure → upgrade liability → gen-tie corridor → capital escalation → site-grid lock-in

In critical minerals:

Prospectivity → land control → access planning → drilling → resource interpretation → processing assumptions → offtake/capital sequencing → extraction commitment

In geological disposal:

Candidate formation → site focus → characterization → access infrastructure → licensing posture → repository design → emplacement pathway → long-duration stewardship commitment

These chains matter because the commitment does not appear all at once. It accumulates.

The governance task is to identify where the chain becomes commitment-bearing before the institution treats the chain as inevitable.

2.6 Disturbance-coupled systems and institutional behavior

Disturbance-coupled systems are difficult to govern because institutional behavior is biased toward continuation.

Once evidence supports a plausible opportunity, institutions tend to seek confirmation, momentum, alignment, and funding. Each step may be rational locally. But the sequence can create commitment before admissibility is established.

Several institutional mechanisms accelerate this process:

- **confidence signaling:** teams must show progress to maintain support;
- **capital timing:** funding windows encourage entry before uncertainty is resolved;
- **stakeholder alignment:** partners and investors prefer clarity over ambiguity;
- **technical momentum:** engineering teams begin solving the assumed problem;
- **public language:** exploratory interest becomes public commitment;
- **sunk effort:** prior work makes alternatives appear wasteful;
- **authority diffusion:** no single actor retains clear refusal responsibility;
- **precedent pressure:** earlier pilots or decisions make future action appear normal.

The result is a common institutional inversion:

INSTRUCTIONAL DRAFT

The organization becomes increasingly committed as evidence remains incomplete.

This inversion is the target of commitment governance.

2.7 Implications for governance

Disturbance-coupled systems require decision governance before commitment hardens.

That governance must do five things.

First, it must distinguish evidence from authorization. A signal can justify exploration without authorizing commitment.

Second, it must identify the irreversible threshold. The relevant boundary is not always construction or execution. It may be site control, queue entry, corridor selection, disturbance, capital alignment, public language, or dependency formation.

Third, it must classify whether unresolved ambiguity remains decision-dominant. The key question is whether plausible remaining states still change the admissibility of the commitment.

Fourth, it must preserve refusal credibility. A project that cannot credibly stop, defer, re-site, or re-sequence is already inside a commitment pathway.

Fifth, it must track commitment integrity after entry. Even if entry is defensible, continuation may become indefensible as evidence, dependencies, authority, and precedent evolve.

These requirements distinguish commitment governance from ordinary analysis. The goal is not to generate more information in the abstract. The goal is to prevent information, action, and authority from becoming coupled in ways that create irreversible commitment before the system is sufficiently understood.

2.8 Section conclusion

Disturbance-coupled exploration systems are not defined by uncertainty alone. They are defined by the structural coupling of uncertainty reduction and exposure creation.

They are systems where the act of learning can disturb the target, privilege a location, concentrate capital, form dependency, degrade authority, or create precedent.

In such systems, the central governance problem is not whether exploration should occur. It is whether exploration can proceed without prematurely authorizing commitment.

That problem leads to the next question: how should such systems be analyzed?

Thought for a few seconds

INSTRUCTIONAL DRAFT

3. Analytical Approach

This paper develops a structural decision framework for systems in which evidence, disturbance, exposure, authority, and commitment formation are coupled. The approach is conceptual rather than predictive. It does not estimate the probability of project success, rank alternatives, optimize portfolios, or prescribe engineering design. It identifies the conditions under which a commitment may become inadmissible before action, or lose integrity after action begins.

The framework proceeds from a simple observation: in disturbance-coupled systems, the important decision is often not whether the project can ultimately succeed. The important decision is whether the next hard-to-reverse step is defensible under the evidence and authority conditions present at the threshold.

To analyze that condition, the paper uses six linked forms of reasoning:

1. plausible-state reasoning
2. irreversible-threshold analysis
3. dependence topology
4. authority coherence analysis
5. disturbance-coupled evidence analysis
6. cross-domain transfer testing

Together, these form the analytical basis for commitment admissibility and commitment integrity.

3.1 Plausible-state reasoning

Plausible-state reasoning begins from the premise that available evidence often does not identify a single underlying reality.

A signal may be informative without being unique. A model may constrain possibilities without resolving the state that matters for commitment. A measurement may reduce uncertainty locally without supporting system-scale action. A successful test may prove capability without proving that the broader commitment is defensible.

The framework therefore asks:

What plausible states remain consistent with the evidence, and would those states imply different commitment decisions?

This is different from asking only for a best estimate.

INSTRUCTIONAL DRAFT

A best estimate may be useful for planning, but commitment admissibility depends on whether remaining plausible states can still change the decision. If the same proposed commitment is defensible under one plausible state and indefensible under another, then the uncertainty has not lost decision-changing power.

This is the logic of an admissibility flip.

A commitment exhibits an admissibility flip when:

- the proposed action is the same,
- the evidence boundary is the same,
- multiple plausible states remain,
- and those plausible states imply different admissibility outcomes.

For example, a lunar volatile signal may support several plausible states: an architecture-grade deposit, a patchy operational resource, a diffuse or inaccessible signal, or mechanically difficult material. The same infrastructure commitment may be defensible in the first state and premature in the others.

A grid interconnection opportunity may similarly support several plausible network futures: bounded upgrade exposure, congestion-dominated economics, restudy-driven delay, or site-grid coupling failure. The same queue-entry decision may be tolerable in one state and commitment-bearing in another.

Plausible-state reasoning therefore shifts the question from:

What is most likely?

to:

Which unresolved states would change the legitimacy of commitment?

That shift is central to the framework.

3.2 Irreversible-threshold analysis

An irreversible threshold is the point at which reversal becomes non-credible under plausible states.

The threshold is not always construction, deployment, or execution. It may occur earlier. It may occur when a site becomes the default, when capital begins sequencing around an assumption, when a corridor is treated as the route, when a queue position couples the project to a network pathway, when public language creates expectation, or when intrusive verification changes the system being evaluated.

INSTRUCTIONAL DRAFT

Irreversible-threshold analysis therefore asks:

Where does ordinary planning become commitment-bearing?

This requires distinguishing between nominal reversibility and practical reversibility.

A project may be reversible in theory but not in practice. It may be possible to withdraw, re-site, redesign, or defer, but doing so may require disproportionate capital loss, reputational damage, regulatory retreat, counterparty disruption, political embarrassment, or abandonment of accumulated infrastructure.

The framework evaluates reversibility through three linked concepts:

- **reversal burden:** what must be absorbed to reverse;
- **reversal capacity:** what the institution can actually absorb;
- **refusal credibility:** whether the institution can still say no without unacceptable loss.

A threshold is crossed when reversal burden exceeds reversal capacity under relevant plausible states.

This is why early actions matter. A non-binding step can still be commitment-bearing if it changes the practical refusal environment. A small test can still create exposure if it anchors site preference, creates partner reliance, or alters the evidence baseline. A preparatory action can still harden into commitment if it makes the next step easier to justify and the alternative harder to defend.

Irreversible-threshold analysis identifies the point where the decision must be governed, not merely analyzed.

3.3 Dependence topology

Commitment often forms through sequences rather than single acts.

A project rarely becomes irreversible all at once. Instead, earlier actions create dependencies that later decisions inherit. Each step may appear bounded, reversible, or rational in isolation. The system-level commitment emerges through the topology of dependence among those steps.

Dependence topology maps:

- what an action depends on;
- what begins to depend on that action;
- which assumptions become embedded;
- which alternatives become less credible;
- which actors gain leverage;

INSTRUCTIONAL DRAFT

- which pathways become easier to continue than abandon.

This is especially important in systems where evidence and infrastructure co-evolve.

In lunar ISRU, volatile interpretation may create site preference. Site preference encourages repeated access. Repeated access creates pressure for power and support. Support enables disturbance. Disturbance concentrates future evidence and activity. Concentrated activity supports corridor formation. Corridor formation enables excavation planning. Excavation planning creates ISRU dependency. ISRU dependency hardens infrastructure.

No single link in the chain necessarily carries the whole commitment. But the chain as a whole can convert exploration into lock-in.

In grid interconnection, site interest may lead to a point-of-interconnection assumption. That assumption leads to queue entry. Queue entry reveals constraints while creating deposits, study exposure, and timing pressure. Study exposure influences gen-tie routing and site control. Capital then begins aligning around a pathway that may not be withdrawable without loss.

Dependence topology therefore asks:

What future decisions will inherit this action?

This matters because admissibility must be evaluated not only at the moment of action, but across the sequence that action initiates.

3.4 Authority coherence analysis

A commitment is not governed only by evidence. It is governed by authority.

Authority coherence exists when formal decision rights, practical control, downside responsibility, revocation capacity, and refusal credibility remain aligned.

In many institutional settings, these elements diverge.

A board may hold formal authority while developers create practical momentum.

A mission office may retain decision rights while partners structure expectations.

An investor may bear downside while operators control escalation.

A public agency may retain formal authority while public language makes reversal politically non-credible.

A project team may claim reversibility while capital sequencing makes exit unrealistic.

Authority coherence analysis asks:

Can the institution still exercise the authority it claims to possess?

INSTRUCTIONAL DRAFT

This is not a legal question alone. It is a structural question.

The framework evaluates whether the decision authority can still:

- refuse;
- defer;
- revoke;
- re-site;
- re-sequence;
- constrain scope;
- terminate;
- communicate uncertainty honestly;
- absorb the cost of reversal;
- avoid being trapped by counterparties or prior representations.

Authority can degrade before execution. It can degrade through public commitments, partner reliance, capital deployment, regulatory posture, internal incentives, or institutional embarrassment. It can also degrade through diffusion, where no actor retains clear responsibility for stopping a pathway once it begins.

This matters because a commitment may be technically reversible while institutionally irreversible.

A commitment is not admissible if the authority required to refuse, constrain, or terminate it cannot survive the act of entry.

3.5 Disturbance-coupled evidence analysis

Evidence is not evaluated only by its content. It is evaluated by the action required to obtain it and by the exposure created through that action.

Disturbance-coupled evidence analysis asks four questions:

1. What evidence is needed to reduce the dominant uncertainty?
2. What action is required to obtain that evidence?
3. Does that action disturb the system or create exposure?
4. Does the evidence gained justify the exposure created?

This is where the framework differs from ordinary evidence sufficiency analysis.

In many systems, the evidence required to make a better decision cannot be obtained from outside the commitment pathway. The project may need to drill, excavate, inject, deploy, traverse, disturb, or enter a queue to learn what matters. But those same actions may alter the target, narrow options, or create continuation pressure.

INSTRUCTIONAL DRAFT

The question is not whether more evidence would be useful. It is whether the evidence can be obtained without crossing the same threshold the evidence is meant to justify.

If not, the system may be under ignorance dominance.

This analysis also distinguishes between evidence levels:

- evidence that justifies interest;
- evidence that justifies prospect ranking;
- evidence that justifies site comparison;
- evidence that justifies bounded verification;
- evidence that supports design;
- evidence that authorizes commitment.

Many failures arise when lower-level evidence is used at a higher authority level.

A hydrogen signal becomes treated as infrastructure permission.

A resource model becomes treated as capital authorization.

A local measurement becomes treated as site-scale confidence.

A queue position becomes treated as project progress.

A pilot success becomes treated as precedent.

Disturbance-coupled evidence analysis prevents evidence from carrying a burden it has not earned.

3.6 Cross-domain transfer testing

The framework is not built by analogy alone.

It transfers across domains only when the same structural conditions are present:

1. inference is indirect;
2. interpretation remains non-unique;
3. verification requires disturbance, deployment, access, or intervention;
4. intervention creates irreversible exposure;
5. commitment creates path dependence before uncertainty collapses.

This transfer rule prevents overextension.

The framework is not equally relevant to all decisions. It is strongest where the next action changes the future decision space. It is weaker where learning is reversible, alternatives remain open, authority remains intact, and exposure does not materially harden.

INSTRUCTIONAL DRAFT

Cross-domain transfer testing is used to evaluate whether the same governance problem appears across lunar ISRU, geological disposal, critical minerals, geothermal, carbon storage, seabed systems, grid interconnection, orbital infrastructure, and autonomy.

The point is not that these domains are substantively identical. They are not.

The point is that each may share the same decision structure:

evidence remains incomplete, learning requires action, action creates exposure, and exposure hardens before uncertainty loses decision-changing power.

The framework transfers by decision structure, not by domain.

3.7 Case selection

The paper uses three major case roles.

3.7.1 Lunar south polar ISRU as the forcing case

Lunar south polar in-situ resource utilization compresses the full entry problem.

The evidence is indirect.

The target state is ambiguous.

Verification may require disturbance.

Support systems create dependence.

Infrastructure pressure is strong.

Precedent implications are significant.

A volatile signal may justify prospecting and bounded verification. It may not justify site fixation, power placement, excavation dependency, corridor formation, or ISRU infrastructure commitment.

This makes lunar ISRU the forcing case for commitment admissibility.

It tests whether evidence that supports exploration has earned the right to govern infrastructure.

3.7.2 Geological disposal as the limiting case for integrity

Geological disposal and nuclear waste systems are the limiting case for commitment integrity.

Here, the issue is not only whether entry is admissible. The issue is whether the commitment can remain defensible across long time horizons, institutional change, monitoring limits, public legitimacy, containment assumptions, and intergenerational responsibility.

INSTRUCTIONAL DRAFT

A repository commitment cannot be evaluated only at entry. It must be governed through persistence.

This makes geological disposal the canonical case for commitment integrity.

It tests whether the original admissible basis survives as evidence, assumptions, authority, and stewardship conditions evolve.

3.7.3 Critical minerals as the transfer case

Critical minerals demonstrate prospectivity-to-commitment conversion.

A promising resource indication can harden into land control, access roads, permitting posture, offtake expectations, processing assumptions, water and power dependencies, strategic-resource designation, and capital exposure.

The critical minerals case tests whether the framework transfers from frontier exploration to strategically urgent terrestrial resource development.

It asks whether prospectivity has earned the right to govern irreversible extraction pathways.

3.8 Analytical outputs

The framework produces two conceptual outputs and three institutional outputs.

The conceptual outputs are:

1. **commitment admissibility**
whether a proposed commitment may enter governed consideration;
2. **commitment integrity**
whether an existing or emerging commitment remains within its admissible basis.

The institutional outputs are:

1. **exposure classification**
where exposure begins;
2. **governance posture classification**
which posture remains defensible after admissibility clears;
3. **integrity classification**
whether commitment should be maintained, constrained, re-evaluated, or terminated.

The institutional outputs are not the paper's primary object, but they show how the theory can be operationalized without becoming an optimization or execution framework.

INSTRUCTIONAL DRAFT

3.9 Boundary of the approach

The framework does not provide a universal decision rule.

It does not determine the best project.

It does not estimate expected value.

It does not replace engineering, law, finance, risk assessment, resource classification, mission design, or regulatory review.

It does not predict outcomes.

It does not claim that irreversible action is always wrong.

Its scope is narrower.

It governs whether a commitment should be allowed to form under unresolved uncertainty, and whether a commitment that has formed remains within its admissible basis.

The framework is most useful where:

- uncertainty remains decision-dominant;
- learning requires exposure;
- refusal credibility is at risk;
- dependency formation is underway;
- authority may degrade;
- precedent may be created;
- and outcome success cannot repair a defective threshold decision.

It is least useful where:

- the action remains fully reversible;
- the evidence is direct and sufficient;
- alternatives remain intact;
- authority remains coherent;
- no material exposure is created;
- and the decision does not create precedent or dependency.

This boundary is important. Without it, the framework could become overbroad. Its value lies in disciplined application to commitment-bearing thresholds.

3.10 Section conclusion

The analytical approach developed here is structural.

INSTRUCTIONAL DRAFT

It begins with plausible states rather than point estimates.
It identifies irreversible thresholds rather than assuming reversibility.
It maps dependence topology rather than treating actions as isolated.
It evaluates authority coherence rather than assuming formal authority is sufficient.
It tests whether evidence generation itself creates exposure.
It transfers across domains only where the same commitment structure is present.

This approach supports the two central contributions of the paper.

Commitment admissibility governs whether entry may be considered.

Commitment integrity governs whether continuation remains defensible.

The next section positions these contributions against existing decision frameworks and identifies the missing category they address.

Thought for 6s

4. Existing Frameworks and the Missing Category

The framework developed in this paper is not proposed as a replacement for established decision methods. Risk management, readiness assessment, real options, value of information, adaptive management, resource classification, assurance practice, project finance, and systems engineering all address important parts of high-consequence decision-making.

The problem is that none of them fully governs the threshold where evidence, authority, disturbance, dependency, and irreversible commitment begin to couple. That gap is the missing category this paper addresses.

The central claim is narrow:

Existing frameworks are necessary but incomplete where the act of learning or advancing a project can itself create irreversible exposure.

In those settings, the relevant question is not only whether a project is feasible, valuable, ready, fundable, or manageable. The prior question is whether the commitment should be allowed to form under the evidence and authority conditions present at the threshold.

4.1 Risk management

INSTRUCTIONAL DRAFT

Risk management is usually applied after a project, mission, asset, or operating pathway has been defined. It identifies hazards, estimates likelihood and consequence, assigns controls, and determines whether residual risk is acceptable.

This is essential work. It can improve safety, reduce loss, allocate responsibility, and make execution more robust.

But risk management often assumes that the project boundary is already legitimate. It asks how to manage the risks of a defined action. It does not always ask whether the action should have been allowed to become the governing commitment in the first place.

In disturbance-coupled systems, this distinction matters.

A project can have a mature risk register and still be inadmissible if the project definition rests on unresolved uncertainty that could still change the commitment decision. A lunar ISRU architecture can identify excavation, power, thermal, logistics, and dust risks while still over-authorizing the volatile evidence that justified the architecture. A geothermal project can manage drilling risk while still depending on a reservoir interpretation that has not earned drilling commitment. A grid project can model cost and schedule risk while still having entered a queue before network constraint ambiguity is bounded.

Risk management asks:

Can the hazards of this pathway be managed?

Commitment admissibility asks:

Should this pathway be allowed to harden before the uncertainty it depends on has lost decision-changing power?

These are different questions.

4.2 Readiness frameworks

Readiness frameworks evaluate whether a capability, technology, system, resource, mission, or organization is mature enough for the next stage.

Technology readiness, construction readiness, operational readiness, mission readiness, and resource readiness all provide valuable discipline. They help institutions avoid deploying immature systems or relying on unproven capabilities.

But readiness is not the same as admissibility.

INSTRUCTIONAL DRAFT

A system can be ready to act while the commitment remains premature. A drill rig can be ready. A rover can be ready. A lander can be ready. A processing technology can be ready. A queue application can be ready. A repository design process can be ready. None of those readiness conditions by themselves determines whether the commitment should form.

Readiness evaluates capability maturity.

Commitment admissibility evaluates whether the evidence and authority structure can support the irreversible burden of action.

This is especially important where technology maturity creates pressure to act. A capable system can pull the institution toward commitment because it reduces execution uncertainty. But execution readiness may obscure unresolved site, subsurface, network, authority, or precedent uncertainty.

The framework therefore separates:

Can the system act?

from:

Should the commitment be allowed to form?

That separation is central.

4.3 Real options

Real options analysis is one of the closest existing frameworks because it explicitly values flexibility, staged commitment, and the option to wait.

It recognizes that uncertainty can make delay valuable, that investment can be staged, and that optionality has economic significance. This is highly relevant to irreversible systems.

However, disturbance-coupled systems introduce a complication.

The action that appears to preserve an option may also consume it.

A preliminary site step may be described as option-preserving while privileging one location. A queue entry may be treated as buying optionality while coupling the project to network studies, deposits, and withdrawal costs. A pilot project may be treated as a low-cost option while creating public precedent and partner reliance. A lunar surface demonstration may be framed as exploratory while beginning to organize access, support, and architecture around one site.

Real options asks:

INSTRUCTIONAL DRAFT

What is the value of waiting, staging, or preserving flexibility?

Commitment governance asks:

Does the option-preserving action itself create exposure that makes later refusal less credible?

The difference is subtle but important. In disturbance-coupled systems, optionality can decay not only through delay, but through the very actions taken to maintain momentum.

The framework therefore treats option preservation as a governance condition, not only an economic one.

4.4 Value of information

Value of information evaluates whether additional knowledge is worth obtaining before making a decision.

This is essential in uncertain systems. It helps determine whether further investigation, sensing, modeling, testing, drilling, surveying, or monitoring is justified.

But in disturbance-coupled systems, information has an exposure cost that may not be captured by a conventional information-value framing.

The key issue is not only whether additional information would improve the decision. It is whether obtaining that information requires crossing or approaching the irreversible threshold the information is meant to justify.

A lunar excavation may reveal volatile form while disturbing the volatile environment. A geothermal well may reveal reservoir behavior while anchoring capital and site exposure. A carbon storage injection test may generate containment evidence while initiating plume and pressure evolution. A seabed survey or pilot intervention may produce route evidence while disturbing the baseline and influencing permitting posture. A queue entry may produce information about grid constraints while creating withdrawal and site-coupling exposure.

Value of information asks:

Is the information worth obtaining?

Commitment admissibility asks:

Can the information be obtained without creating exposure that the current evidence cannot yet justify?

INSTRUCTIONAL DRAFT

This does not reject value-of-information logic. It bounds it.

In disturbance-coupled systems, the value of information must be evaluated against the irreversibility cost of obtaining it.

4.5 Adaptive management

Adaptive management is designed for systems where uncertainty can be reduced over time through monitoring, learning, and adjustment. It accepts uncertainty at entry and relies on iterative correction.

This is valuable where future adjustment remains possible.

But adaptive management becomes weaker when adjustment is non-credible after entry. Some systems do not allow later correction without disproportionate cost, institutional failure, environmental damage, regulatory retreat, or authority collapse.

A repository cannot rely only on future adaptation if institutional memory, monitoring capacity, or public legitimacy cannot be guaranteed across the relevant horizon. A lunar site cannot fully adapt after support systems, corridors, and infrastructure have hardened around an incorrect volatile assumption. A grid project may not easily adapt after queue entry, gen-tie routing, site control, and capital sequencing create coupled exposure. A seabed project may not adapt cleanly after ecological baseline disturbance and installation geometry form.

Adaptive management asks:

Can the system learn and adjust over time?

Commitment integrity asks:

Does the commitment remain within its admissible basis as learning, dependencies, and authority conditions change?

This is the persistence problem.

The framework does not deny that adaptation is valuable. It asks whether adaptation remains credible after the threshold is crossed.

4.6 Resource classification

INSTRUCTIONAL DRAFT

Resource classification systems evaluate confidence in resource presence, quantity, quality, continuity, accessibility, economic viability, or recoverability. In mining, petroleum, geothermal, water, and planetary resource contexts, classification disciplines are essential.

But resource classification does not automatically determine commitment authority.

A resource can be indicated, inferred, prospective, strategically important, or scientifically compelling without being sufficient to govern infrastructure, capital, extraction, or mission architecture.

In lunar ISRU, hydrogen indication does not automatically imply extractable water. Permanently shadowed regions may support volatile retention without resolving form, concentration, accessibility, continuity, or mechanical behavior. In critical minerals, prospectivity may support further exploration without justifying land control, access roads, processing assumptions, offtake expectations, and capital sequencing. In geothermal, resource indicators may support drilling design without proving that drilling commitment is admissible under the current evidence boundary.

Resource classification asks:

How confident are we in the resource?

Commitment admissibility asks:

Has the resource evidence earned the right to govern irreversible commitments?

This is the evidence-to-authorization gap.

The framework therefore does not replace resource classification. It determines whether the classified evidence is being used within its proper authority.

4.7 Planetary protection, mission assurance, and safety frameworks

Planetary protection, mission assurance, and safety frameworks govern contamination, mission risk, reliability, failure tolerance, and operational constraints.

They are essential for space systems and other high-consequence domains. They help ensure that missions do not contaminate environments, fail prematurely, violate scientific constraints, or produce unmanaged operational risks.

But they do not fully address commitment admissibility.

INSTRUCTIONAL DRAFT

A mission can comply with contamination constraints and still prematurely harden infrastructure around an unresolved resource assumption. A system can pass assurance requirements and still create precedent or dependency that exceeds the evidence basis. A payload can be reliable while its success is over-authorized into infrastructure commitment. A surface operation can be safe while still generating site, corridor, support, and precedent exposure.

Planetary protection and assurance ask:

Can the mission proceed without violating protection, safety, reliability, or assurance requirements?

Commitment governance asks:

Should the evidence, architecture, and authority conditions be allowed to harden into irreversible commitment?

This question is adjacent to assurance, but not identical to it.

4.8 Project finance and investment diligence

Project finance and investment diligence assess whether a project is bankable, investable, creditworthy, resilient, and aligned with capital requirements.

They evaluate cost, schedule, revenue, market, counterparty, regulatory, legal, technical, and execution risks. They are indispensable for infrastructure and resource development.

But they often operate after a project pathway has already begun to form.

The diligence question may become:

Is this project financeable?

when the prior question should be:

Should this project pathway have become the financing object at all?

Capital can accelerate commitment. It can convert uncertainty into momentum, compress timelines, create public expectation, and make refusal appear costly. A project may look increasingly real because capital has aligned around it, even if the evidence basis remains structurally weak.

In grid interconnection, capital may align around a site and POI before network constraint topology is sufficiently bounded. In critical minerals, capital may align around strategic-resource urgency before prospectivity earns extraction commitment. In lunar infrastructure, capital or

INSTRUCTIONAL DRAFT

agency partnership may align around a volatile interpretation before the resource has become architecture-grade.

Project finance asks:

Can this commitment be financed and structured?

Commitment admissibility asks:

Should capital be allowed to harden around this commitment under current uncertainty?

Commitment integrity then asks:

Does the capital sequence remain within the admissible basis as evidence and dependencies evolve?

4.9 Systems engineering and architecture practice

Systems engineering evaluates requirements, interfaces, architecture, dependencies, trade spaces, verification, validation, and lifecycle behavior.

It is central to complex projects.

But systems engineering can also become a mechanism through which premature assumptions harden. Once an architecture begins to form around a resource, site, corridor, power geometry, communications network, or processing pathway, the act of engineering can convert uncertainty into structure.

This is not a critique of systems engineering. It is a boundary condition.

Systems engineering asks:

How should the system be structured to meet requirements?

Commitment governance asks:

Have the requirements, site assumptions, resource assumptions, and architectural dependencies earned the right to govern the system?

If the answer is no, engineering excellence may intensify the wrong commitment.

A mission architecture can become coherent around a volatile assumption that remains non-unique. A grid design can become coherent around a point of interconnection whose constraint

INSTRUCTIONAL DRAFT

topology remains unresolved. A repository design can become coherent around assumptions whose long-duration institutional integrity is not yet secure.

The framework therefore sits upstream of architecture hardening. It does not design the system. It asks whether the architecture should be allowed to harden.

4.10 The missing category: threshold governance

The frameworks above are not wrong. They are incomplete at a particular point: the irreversible threshold.

That threshold is where:

- evidence begins to govern action;
- action begins to create exposure;
- exposure begins to create dependency;
- dependency begins to constrain future decisions;
- authority begins to degrade;
- precedent begins to form.

This is the missing category:

Threshold governance.

Threshold governance determines whether commitment may form, and later whether the commitment remains defensible.

It asks:

- Has evidence been over-authorized?
- Has uncertainty lost decision-changing power?
- Is the action reversible in practice, not just in theory?
- Does learning require exposure?
- Does the commitment create dependency?
- Does the institution retain refusal authority?
- Does the commitment create precedent?
- Does the admissible basis survive after entry?

This category is not reducible to risk, readiness, optimization, VOI, resource classification, or assurance.

It is the governance of entry and persistence under irreversible conditions.

INSTRUCTIONAL DRAFT

4.11 Comparative summary

Framework	Primary question	What it handles	What remains unresolved
Risk management	Can risks be mitigated?	Hazards after pathway definition	Whether the pathway should form
Readiness frameworks	Is the capability mature?	Technical and operational maturity	Whether maturity authorizes commitment
Real options	What is the value of waiting or staging?	Flexibility and staged investment	Whether option-preserving actions create exposure
Value of information	Is more information worth obtaining?	Learning value	Whether learning requires irreversible exposure
Adaptive management	Can the system adjust over time?	Monitoring and adjustment	Cases where adjustment becomes non-credible
Resource classification	How confident is the resource evidence?	Resource confidence	Whether resource evidence may govern infrastructure
Planetary protection / assurance	Can mission risk or contamination be controlled?	Safety, reliability, protection	Commitment admissibility and precedent formation
Project finance / diligence	Is the project investable?	Capital, risk, bankability	Whether capital is hardening around unresolved uncertainty
Systems engineering	How should the system be structured?	Requirements, architecture, interfaces	Whether architecture should be allowed to harden

The framework developed here addresses the unresolved category:

whether a commitment may form under current uncertainty, and whether it remains within its admissible basis after formation.

4.12 Section conclusion

Existing frameworks provide essential tools for technical, financial, operational, and institutional decision-making. But disturbance-coupled systems require an additional layer.

That layer is not a better forecast.

It is not an optimization method.

It is not a substitute for engineering, finance, law, regulation, or mission design.

It is a threshold governance framework.

INSTRUCTIONAL DRAFT

Commitment admissibility governs whether entry may be considered.

Commitment integrity governs whether continuation remains defensible.

The next section defines the formal primitives needed to make those conditions precise.

5. Formal Primitives and Structural Logic

The framework requires a set of formal primitives because the central terms are often used loosely in institutional settings. Words such as “commitment,” “reversibility,” “risk,” “evidence,” “authority,” and “uncertainty” can appear clear while hiding the actual decision problem.

This section defines the primitives used in the paper and then formalizes their relationships.

The purpose is not to create mathematical abstraction for its own sake. The purpose is to make threshold reasoning explicit. A commitment cannot be governed if the commitment, threshold, evidence boundary, plausible state space, authority condition, and admissible basis are undefined.

The framework begins from the following claim:

A commitment is not defined by intention alone. It is defined by the exposure created when future refusal becomes harder.

5.1 Commitment

A **commitment** is a decision or sequence of actions that materially reduces the practical ability to refuse, defer, re-site, re-sequence, constrain, or abandon a pathway without disproportionate loss.

Commitment can be created by formal decision, but it can also be created by action.

Formal commitment includes board approval, capital allocation, agency authorization, mission approval, regulatory submission, contract execution, offtake agreement, site acquisition, and launch authorization.

Practical commitment includes repeated access to one site, public designation of a preferred location, infrastructure placement, queue entry, corridor planning, partner reliance, resource-dependent architecture, intrusive verification, disturbance of the target system, and narrative hardening.

A commitment may therefore exist before the institution describes itself as committed.

INSTRUCTIONAL DRAFT

This matters because many exposure-forming actions are framed as preliminary. A preliminary action can still become commitment-bearing if it changes the future decision space.

A useful test is:

After this action, is it still politically, institutionally, financially, operationally, and technically credible to say no?

If the answer becomes materially weaker, commitment is forming.

5.2 Decision Boundary

A **decision boundary** defines what the determination applies to and what it does not apply to.

Without a decision boundary, decisions drift. A determination about bounded exploration can be misused as infrastructure permission. A determination about one site can be transferred to another. A determination about one phase can be applied to capital escalation. A determination about admissibility can be treated as proceed approval.

The decision boundary prevents these errors by specifying the action under review, threshold, evidence boundary, authority boundary, time boundary, domain boundary, non-transferability conditions, and excluded uses.

Decision boundaries are not administrative details. They are part of the governance logic.

Let:

- B_t = decision boundary at threshold time t

A commitment cannot be assessed unless:

$B_t \neq \emptyset$

In prose:

A commitment is not admissible to assess until the boundary of the decision is defined.

5.3 Plausible State Space

The **plausible state space** is the set of materially possible system states consistent with the evidence available at the decision threshold.

INSTRUCTIONAL DRAFT

It is not the set of all imaginable states. It is the set of states that remain credible enough to affect the decision.

A plausible state may describe volatile form, concentration, continuity, accessibility, reservoir structure, containment behavior, mineral continuity, seabed condition, grid constraint topology, orbital traffic behavior, institutional authority condition, public legitimacy, counterparty behavior, or long-duration repository performance.

Let:

- Ω = plausible state space
- $\omega \in \Omega$ = a plausible state consistent with the evidence record
- E_t = evidence available at threshold time t

The plausible state space is defined as:

$\Omega(E_t) = \{\omega : \omega \text{ remains materially consistent with the evidence available at threshold } t\}$

A state is included in Ω only if it remains materially relevant to the commitment decision. The framework does not require consideration of every imaginable state, only those still plausible under the evidence boundary.

The plausible state space matters because admissibility depends not only on the favored interpretation, but on whether other plausible interpretations would change the commitment decision.

A lunar volatile signal may be consistent with several plausible states. If all states support bounded reconnaissance and none supports infrastructure commitment, the decision posture is stable. If one state supports site hardening while another makes site hardening premature, then the plausible state space contains an admissibility flip.

5.4 Exposure

Exposure is the burden created when action makes future refusal, reversal, reconfiguration, re-siting, or abandonment harder.

Exposure can be physical, financial, institutional, political, reputational, regulatory, operational, environmental, architectural, topological, or precedent-forming.

Exposure is not identical to cost. Cost is one form of exposure. Exposure also includes loss of credibility, loss of authority, public expectation, dependency formation, institutional embarrassment, and reduction of future options.

Examples include:

INSTRUCTIONAL DRAFT

Entering an interconnection queue creates exposure to study outcomes, deposits, timing, withdrawal costs, and site-grid coupling.

Drilling creates exposure through capital, permitting, site focus, and physical intervention.

Lunar excavation creates exposure through disturbance, evidence alteration, access repetition, and potential ISRU dependency.

Public announcement creates exposure through expectation and future refusal cost.

Repository siting creates exposure through institutional identity, local reliance, regulatory posture, and long-duration responsibility.

Exposure is the practical material of commitment.

5.5 Reversal Burden and Reversal Capacity

Reversal burden is the cost, difficulty, loss, liability, damage, or structural disruption required to reverse a commitment after it has begun.

Reversal burden may include sunk capital, remediation, re-permitting, engineering redesign, route abandonment, site abandonment, counterparty compensation, loss of queue position, reputation damage, political backlash, regulatory consequences, loss of partner trust, physical irreversibility, precedent reversal, or abandonment of public claims.

Reversal capacity is the institution's practical capacity to absorb the burden of reversal.

It includes financial capacity, technical flexibility, political tolerance, legal flexibility, regulatory room, schedule resilience, authority clarity, partner tolerance, public credibility, operational slack, and alternative pathways.

A commitment can become irreversible not because reversal is impossible in principle, but because the institution lacks capacity to absorb reversal in practice.

Let:

- **C** = proposed commitment
- **D** = Decision Authority
- **R(C,ω,t)** = reversal burden of commitment **C** under plausible state **ω** at time **t**
- **κ(D,t)** = reversal capacity of Decision Authority **D** at time **t**

A commitment crosses an irreversible threshold when:

$$\mathbf{R(C,\omega,t)} > \mathbf{\kappa(D,t)}$$

INSTRUCTIONAL DRAFT

for at least one materially relevant plausible state $\omega \in \Omega(E_t)$.

Equivalently:

$\exists \omega \in \Omega(E_t)$ such that $R(C, \omega, t) > \kappa(D, t)$

If such a state remains plausible, the commitment may be irreversible under unresolved conditions.

The purpose of this notation is not numerical precision. It clarifies the logic. A project may be reversible under favorable states and non-credible to reverse under adverse plausible states. If those adverse states remain unresolved, the threshold may already be governance-relevant.

Irreversible thresholds may occur at site control, queue entry, drilling authorization, injection, excavation, route fixation, fixed infrastructure placement, public primary-site designation, capital escalation, regulatory filing, mission architecture commitment, repository closure, autonomy permissioning, or first-of-kind deployment.

The threshold must be identified before admissibility can be assessed.

5.6 Refusal Credibility

Refusal credibility is the practical ability of the Decision Authority to say no, defer, re-site, re-sequence, constrain, or terminate after an action begins.

Refusal credibility depends on whether refusal remains acceptable under institutional, financial, political, operational, and reputational conditions.

Refusal credibility is preserved when the stop condition is explicit, authority remains clear, alternatives remain viable, capital exposure is bounded, public language remains conditional, counterparties understand the boundary, evidence can still change the pathway, support systems remain non-hardening, and dependency has not exceeded the decision basis.

Refusal credibility degrades when the pathway becomes publicly framed as inevitable, capital is committed in ways that create continuation pressure, counterparties rely on continuation, alternatives become expensive or embarrassing, internal incentives reward progress signaling, regulatory posture becomes difficult to reverse, site, route, or architecture becomes default, or refusal would be interpreted as institutional failure.

Let:

- $F(C, \omega, t)$ = refusal credibility after commitment C under state ω at time t

A simplified condition is:

INSTRUCTIONAL DRAFT

$F(C, \omega, t) = 1$ if $R(C, \omega, t) \leq \kappa(D, t)$

$F(C, \omega, t) = 0$ if $R(C, \omega, t) > \kappa(D, t)$

In institutional application, refusal credibility may be graded rather than binary:

$F \in \{\text{Preserved, Conditional, Fragile, Collapsed}\}$

Refusal credibility is central because a commitment cannot be governed if the ability to refuse has already collapsed.

5.7 Authority Coherence

Authority coherence is the condition in which formal decision rights, practical control, downside responsibility, revocation capacity, and refusal authority remain aligned.

Authority incoherence occurs when the actor with formal authority is no longer the actor controlling practical momentum, bearing downside, or able to stop the pathway.

A board may retain formal authority while capital teams have already structured the transaction. A mission authority may retain formal control while partners have begun relying on the architecture. An agency may retain approval authority while public commitments make reversal non-credible. An investor may bear downside while developers control the sequence. A technical team may create architecture lock-in before the decision authority recognizes the commitment.

Authority coherence is not only legal authority. It is the practical capacity to govern the threshold.

The framework asks:

Can the institution still exercise the authority it claims to possess?

Let:

- $\text{Auth}(D, C, t)$ = authority coherence of Decision Authority **D** over commitment **C** at time **t**

Authority is intact when:

$\text{Auth}(D, C, t) = \text{Coherent}$

If authority has degraded, a commitment may be inadmissible even if evidence appears strong.

INSTRUCTIONAL DRAFT

5.8 Admissibility Flip

An **admissibility flip** occurs when the same proposed commitment produces different admissibility outcomes across plausible states.

The proposed action remains constant. The evidence boundary remains constant. The plausible states differ. The admissibility outcome changes.

Let:

- $A(C, \omega, B_t)$ = admissibility of commitment C under plausible state ω and decision boundary B_t

Where:

$A(C, \omega, B_t) \in \{\text{Admissible}, \text{Inadmissible}\}$

An admissibility flip occurs when:

$\exists \omega_i, \omega_j \in \Omega(E_t)$ such that $A(C, \omega_i, B_t) \neq A(C, \omega_j, B_t)$

This condition means the evidence has not merely left uncertainty about performance. It has left uncertainty about whether the commitment should be allowed to form.

In prose:

A commitment is unstable when one plausible state supports consideration while another plausible state requires deferral or refusal.

An admissibility flip is not merely uncertainty in outcome. It is uncertainty in whether the action should be allowed to form.

5.9 Ignorance Dominance

Ignorance dominance occurs when unresolved plausible states remain capable of changing whether a commitment is admissible.

A concise distinction is:

Uncertainty is decision-relevant when it affects estimates. Ignorance becomes decision-dominant when unresolved plausible states can still change whether a commitment is admissible.

INSTRUCTIONAL DRAFT

Ignorance dominance is active when the plausible state space cannot be collapsed before the threshold, at least one plausible state would materially change the admissibility of commitment, the action required to collapse the state space would itself create the exposure being evaluated, and refusal, re-siting, re-sequencing, or containment may become less credible after action.

Let:

- \mathbf{M} = measurement or verification action intended to reduce uncertainty
- $\mathbf{X}(\mathbf{M})$ = exposure created by measurement action \mathbf{M}
- $\mathbf{T}(\mathbf{C})$ = irreversible threshold associated with commitment \mathbf{C}

Ignorance dominance is active when:

1. The plausible state space remains non-collapsed:

$$|\Omega(\mathbf{E}_t)| > 1$$

2. At least one admissibility flip exists:

$$\exists \omega_i, \omega_j \in \Omega(\mathbf{E}_t): \mathbf{A}(\mathbf{C}, \omega_i, \mathbf{B}_t) \neq \mathbf{A}(\mathbf{C}, \omega_j, \mathbf{B}_t)$$

3. The measurement required to collapse the plausible state space creates commitment-relevant exposure:

$$\mathbf{X}(\mathbf{M}) \geq \mathbf{T}(\mathbf{C})$$

or, more generally:

$$\mathbf{R}(\mathbf{M}, \omega, t) > \kappa(\mathbf{D}, t)$$

for a materially relevant plausible state.

4. Refusal credibility would degrade after measurement or entry:

$$\mathbf{F}(\mathbf{M}, \omega, t+1) < \mathbf{F}(\emptyset, \omega, t)$$

The core condition can be written compactly:

Ignorance Dominance(C) = 1 if unresolved plausible states create an admissibility flip and the action required to resolve them approaches or crosses the same exposure boundary being evaluated.

Ignorance dominance does not mean “we do not know everything.” No frontier system has complete knowledge. It means the remaining ignorance is still structurally capable of changing the decision.

INSTRUCTIONAL DRAFT

5.10 Dependence Topology

Dependence topology is the structure by which early actions create dependencies that later decisions inherit.

It is the map of how one action makes another action easier, more likely, or harder to refuse.

Dependence topology includes site dependencies, route dependencies, corridor dependencies, grid dependencies, resource dependencies, infrastructure dependencies, counterparty dependencies, capital dependencies, regulatory dependencies, communication dependencies, support-system dependencies, authority dependencies, and precedent dependencies.

Commitment often emerges through sequences. An individual step may be bounded, but the sequence may not be.

Example:

Signal → **site preference** → **access repetition** → **support placement** → **disturbance** → **corridor formation** → **resource dependency** → **infrastructure lock-in**

Let:

- $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ represent the commitment-dependence graph
- \mathbf{V} = nodes representing actions, assumptions, assets, institutions, sites, routes, evidence claims, or support systems
- \mathbf{E} = directed edges representing dependence relationships

A directed edge:

$\mathbf{v}_i \rightarrow \mathbf{v}_j$

means that later decision or action \mathbf{v}_j depends on earlier action, assumption, or commitment \mathbf{v}_i .

Let:

- $\mathbf{w}(\mathbf{e})$ = weight of a dependency edge
- $\mathbf{L}(\mathbf{v})$ = lock-in contribution of node \mathbf{v}

A qualitative lock-in index can be represented as:

$$\Lambda(\mathbf{G}, \mathbf{t}) = \Sigma \mathbf{w}(\mathbf{e}) + \Sigma \mathbf{L}(\mathbf{v})$$

The exact values need not be numeric. The structure shows that commitment can emerge through dependency accumulation, not only through one discrete irreversible act.

INSTRUCTIONAL DRAFT

A system becomes topologically commitment-bearing when:

$$\Lambda(G,t+1) > \Lambda(G,t)$$

and the increase reduces refusal credibility or alternative pathway viability.

5.11 Commitment Admissibility

Commitment admissibility is the entry condition under which a proposed irreversible commitment may be considered.

A commitment is admissible only when the evidence, authority, refusal, and dependency conditions are sufficient for the commitment to enter governed consideration.

Admissibility is not approval, execution authorization, optimization, prediction, investment advice, legal advice, engineering validation, regulatory clearance, site selection, or project endorsement.

Admissibility asks a narrower question:

May this commitment be considered under the current evidence and authority boundary?

Formally, strict admissibility requires:

C is admissible under B_t iff $\forall \omega \in \Omega(E_t), A(C,\omega,B_t) = \text{Admissible}$

If there exists a plausible state under which the commitment is inadmissible:

$\exists \omega \in \Omega(E_t)$ such that $A(C,\omega,B_t) = \text{Inadmissible}$

then the commitment fails strict admissibility unless the commitment can be bounded, constrained, or resequenced to remove the inadmissible exposure.

A commitment is admissible when all of the following hold:

1. The decision boundary is defined:

$$B_t \neq \emptyset$$

2. The plausible state space does not contain unresolved admissibility flips, or those flips have been bounded by constraint.
3. Reversal burden remains within reversal capacity across relevant plausible states:

$$\forall \omega \in \Omega(E_t), R(C,\omega,t) \leq \kappa(D,t)$$

INSTRUCTIONAL DRAFT

or, if burden exceeds capacity, the exposure is explicitly accepted, bounded, and authority-coherent.

4. Refusal credibility remains preserved or conditionally preserved:

$F(C, \omega, t+1) \neq \text{Collapsed}$

5. Authority coherence is intact:

$\text{Auth}(D, C, t) = \text{Coherent}$

6. Dependence topology remains within the admissible boundary:

$\Lambda(G, t+1) \leq \Lambda^* _B$

where $\Lambda^* _B$ is the dependence limit allowed by the decision boundary.

7. Precedent transfer is contained:

$P(C, t) \leq P^* _B$

where $P(C, t)$ is precedent exposure and $P^* _B$ is the allowable precedent boundary.

If any condition fails, the commitment is inadmissible or requires deferral, constraint, or refusal depending on where in the sequence the failure occurs.

The output of admissibility is not “proceed.” It is eligibility for further governance.

This is why admissibility must be separated from governance posture.

5.12 Governance Posture

Governance posture is the defensible institutional stance after a commitment has been found admissible to consider.

Admissibility answers whether the commitment may enter governed consideration.

Governance posture answers which stance remains defensible inside that governed space.

The primary posture classifications are:

Proceed-Compatible means proceeding within the defined boundary has not been eliminated by governance constraints.

INSTRUCTIONAL DRAFT

Deferral-Indicated means bounded learning remains the defensible posture.

Refusal-Required means refusal is the only governance-valid posture under the evidence and authority conditions.

This distinction prevents admissibility from being misread as approval.

A commitment may be admissible to consider while deferral remains the only defensible posture.

5.13 Admissible Basis

The **admissible basis** is the set of conditions under which a commitment was allowed to begin or enter governed consideration.

It includes the defined commitment, irreversible threshold, evidence record, plausible state space, dominant uncertainty, active irreversibility regimes, Decision Authority, refusal and revocation conditions, dependency topology, assumptions registry, expiration conditions, reliance perimeter, and precedent containment conditions.

Let:

- $AB(C, t_0)$ = admissible basis of commitment C at entry time t_0

Commitment integrity is evaluated against the admissible basis.

If the commitment continues outside that basis, integrity fails or requires re-evaluation.

5.14 Commitment Integrity

Commitment integrity is the persistence condition under which an existing or emerging commitment remains within its admissible basis as conditions change.

Integrity is maintained when the conditions that made entry defensible remain sufficiently intact.

Integrity degrades when evidence changes, assumptions expire, dependencies accumulate beyond the original boundary, authority degrades, refusal becomes non-credible, support systems harden, precedent expands beyond the defined scope, or the commitment begins to rely on evidence or assumptions not included in the original basis.

Let:

INSTRUCTIONAL DRAFT

- $S(C,t)$ = current state of the commitment at time t
- $I(C,t)$ = commitment integrity state at time t

The commitment retains integrity when:

$$S(C,t) \subseteq AB(C,t_0)$$

That means the current commitment remains inside the evidence, assumption, authority, dependency, irreversibility, and precedent boundary under which it was admitted.

Integrity degrades when:

$$S(C,t) \not\subseteq AB(C,t_0)$$

The integrity output is:

$$I(C,t) \in \{\text{Maintain, Constrain, Re-evaluate, Terminate}\}$$

Where:

Maintain if $S(C,t) \subseteq AB(C,t_0)$

Constrain if $S(C,t)$ is approaching the boundary but can be brought back inside through limits.

Re-evaluate if $AB(C,t_0)$ is no longer reliable because evidence, assumptions, authority, or dependency conditions have materially changed.

Terminate if $S(C,t)$ exceeds the admissible basis and cannot be restored through constraint.

Commitment integrity does not ask whether the original decision was good. It asks whether the commitment remains inside the boundary that made it defensible.

5.15 Precedent Externality

A **precedent externality** occurs when a decision changes the future admissibility environment by becoming something later actors can cite, repeat, normalize, or inherit.

Precedent is not merely memory. It is a form of exposure.

A first action can make later actions appear more acceptable. A bounded pilot can become a general justification. A deferral can be misread as conditional approval. A proceed-compatible posture can be misread as recommendation. A refusal can be misread as a domain-wide rejection.

Let:

INSTRUCTIONAL DRAFT

- $P(C,t)$ = precedent exposure created by commitment C
- $U(C,t)$ = set of future uses or citations of commitment C
- B_t = original decision boundary

Precedent externality occurs when:

$\exists u \in U(C,t)$ such that $u \notin B_t$

In prose:

A decision creates precedent externality when it is cited, repeated, normalized, or transferred beyond the evidence and authority boundary under which it was made.

Precedent containment requires:

$U(C,t) \subseteq B_t$

or else re-screening is required before transfer.

Precedent externality is especially important in frontier systems because early actions may define norms before institutions have mature governance structures.

5.16 Relationship Among Primitives

The primitives form a sequence.

A signal creates interest. Evidence defines a plausible state space. Plausible states may contain ambiguity. Ambiguity becomes decision-dominant when it can cause an admissibility flip. A proposed action creates exposure. Exposure becomes commitment-bearing near an irreversible threshold. The threshold is governed by reversal burden, reversal capacity, and refusal credibility. Authority coherence determines whether refusal can still be exercised. Dependence topology determines whether the action creates inherited constraints. Admissibility determines whether entry may be considered. Governance posture determines which stance remains defensible. Commitment integrity determines whether continuation remains within the admissible basis. Precedent externality determines what the decision may normalize beyond itself.

This sequence can be summarized as:

Evidence → **plausible states** → **ambiguity** → **exposure** → **threshold** → **admissibility** → **posture** → **commitment** → **integrity** → **precedent**

More formally:

INSTRUCTIONAL DRAFT

1. **Evidence defines the plausible state space:**
 $\Omega(E_t) = \{\omega : \omega \text{ remains materially consistent with } E_t\}$
2. **Plausible states reveal admissibility stability or instability:**
 $\exists \omega_i, \omega_j \in \Omega(E_t): A(C, \omega_i, B_t) \neq A(C, \omega_j, B_t)$
3. **Exposure becomes irreversible when reversal burden exceeds capacity:**
 $R(C, \omega, t) > \kappa(D, t)$
4. **Dependence accumulates through topology:**
 $\Lambda(G, t+1) > \Lambda(G, t)$
5. **Admissibility requires bounded exposure, coherent authority, preserved refusal, and contained precedent:**
C admissible only if B_t is defined, no unresolved admissibility flip remains, reversal burden is bounded, authority is coherent, dependencies remain within boundary, and precedent is contained.
6. **Integrity requires the current commitment state to remain inside the admissible basis:**
 $S(C, t) \subseteq AB(C, t_0)$
7. **Precedent control requires future use to remain inside the original boundary:**
 $U(C, t) \subseteq B_t$

5.17 Section Conclusion

The purpose of these primitives is to prevent irreversible decisions from being governed by vague terms.

A commitment is not simply an intention.

Reversibility is not simply theoretical ability to unwind.

Uncertainty is not automatically decision-dominant.

Evidence is not automatically authorization.

Admissibility is not approval.

Proceed-compatible is not a command.

Integrity is not permanence.

Precedent is not neutral memory.

The formal structure does not turn the framework into an optimization model. It clarifies the logical conditions under which a commitment may become inadmissible, posture may become indefensible, integrity may degrade, or precedent may escape its boundary.

INSTRUCTIONAL DRAFT

The next section uses these primitives to develop commitment admissibility as the entry problem.

6. Commitment Admissibility: The Entry Problem

Commitment admissibility is the entry condition for irreversible commitments under persistent uncertainty.

It asks whether a proposed commitment may enter governed consideration before the institution allows evidence, capital, access, infrastructure, public posture, or authority to harden around it.

This is a narrower question than whether the project is desirable, feasible, fundable, technically mature, or strategically important. A project can satisfy those conditions and still be inadmissible if the proposed commitment would harden before uncertainty has lost decision-changing power.

The central question is:

May this commitment be considered under the current evidence and authority boundary?

Commitment admissibility is therefore not approval. It is not authorization. It is not a recommendation to proceed. It is the threshold condition that determines whether a commitment is eligible for governed deliberation at all.

6.1 Why entry requires governance

In reversible systems, the cost of early entry may be limited. An institution can act, learn, reverse, and update without materially damaging future options.

In disturbance-coupled systems, entry changes the decision space.

A site becomes favored.

A corridor becomes default.

Capital begins sequencing.

A queue position creates exposure.

A resource signal becomes an architecture assumption.

A support system begins structuring future access.

An intrusive test changes the evidence baseline.

A public statement creates expectation.

A pilot creates precedent.

The danger is not only that the project might fail. The danger is that the organization may lose the practical ability to refuse before it fully understands what it has committed to.

INSTRUCTIONAL DRAFT

Admissibility governs this entry point.

It asks whether the next action can be allowed to begin without creating exposure that exceeds the evidence, authority, and refusal conditions available at the threshold.

6.2 Admissibility is distinct from feasibility

Feasibility asks whether a project can be made to work.

Admissibility asks whether a commitment should be allowed to harden.

These are different questions.

A lunar ISRU architecture may be technically coherent but inadmissible if the volatile evidence remains unable to distinguish among materially different operating realities.

A geothermal prospect may be technically drillable but inadmissible if the reservoir interpretation remains decision-dominant and drilling would create site, capital, and permitting lock-in.

A grid project may be financially attractive but inadmissible if queue entry would create exposure before the network constraint topology is sufficiently bounded.

A repository may be technically characterized but inadmissible if authority, monitoring, public legitimacy, or long-duration stewardship cannot support the permanence of the commitment.

Feasibility can make action possible. It does not make commitment legitimate.

The separation matters because many institutions are biased toward capability-driven escalation. Once a team can act, the question becomes how to act well. Commitment admissibility restores the prior question: whether the commitment should form.

6.3 Admissibility is distinct from risk tolerance

Risk tolerance asks how much risk an institution is willing to accept.

Commitment admissibility asks whether the decision is structurally eligible for acceptance or refusal.

A high-risk action may be admissible if the risk is understood, bounded, governed, and compatible with refusal authority. A low-probability risk may make a commitment inadmissible if its realization would create non-recoverable exposure and the uncertainty cannot be reduced without crossing the same threshold.

INSTRUCTIONAL DRAFT

The issue is not simply risk magnitude. It is whether unresolved plausible states can still alter the legitimacy of commitment.

This is why the framework does not reduce admissibility to expected value, risk score, or probability threshold.

A small unresolved unknown can dominate admissibility if it controls site selection, corridor geometry, resource dependency, authority retention, or precedent formation.

6.4 Admissibility is distinct from evidence strength

Evidence can be strong and still not commitment-grade.

A signal may be valid for exploration but insufficient for infrastructure.

A model may be useful for ranking but insufficient for site hardening.

A bounded test may support further investigation but not resource dependency.

A regional dataset may support prospectivity but not local construction.

A successful pilot may support learning but not generalization.

The question is not whether evidence exists. The question is what the evidence is allowed to authorize.

This distinction is central in lunar ISRU. Hydrogen indications, thermal context, shadow persistence, radar response, reflectance, and terrain analysis can identify promising regions. They do not by themselves resolve volatile form, concentration, continuity, depth, accessibility, extractability, or construction-scale behavior.

The evidence may be excellent for prospecting. It may be insufficient for site hardening, excavation dependency, or ISRU infrastructure commitment.

Admissibility requires evidence to be matched to the irreversible burden of the proposed action.

6.5 The admissibility set

Admissibility can be understood as a set condition.

Let:

- C be the proposed commitment;
- Ω be the plausible state space;
- ω be a plausible state within Ω ;
- $A(C, \omega)$ be the admissibility condition of commitment C under state ω ;
- B be the defined evidence and authority boundary.

INSTRUCTIONAL DRAFT

A commitment is admissible only if it remains structurally defensible across the relevant plausible states within the boundary B.

This does not mean the commitment must succeed across all states. It means that unresolved states must not still change whether the commitment should be allowed to form.

If the proposed commitment remains defensible under all materially relevant plausible states, then admissibility may be established.

If plausible states imply different admissibility outcomes, then the commitment is unstable.

The framework therefore treats admissibility as a boundary condition, not a score.

6.6 The admissibility flip condition

An admissibility flip occurs when the same proposed commitment produces different admissibility outcomes across plausible states.

This is the core instability condition.

For example, consider a proposed lunar surface infrastructure commitment near a volatile signal.

Under plausible state 1, the signal reflects a concentrated, accessible, construction-relevant volatile deposit. The proposed site commitment may be defensible.

Under plausible state 2, the signal reflects patchy volatile-bearing regolith with poor continuity. The same commitment may require deferral.

Under plausible state 3, the signal reflects diffuse or inaccessible hydrogen that does not justify ISRU dependency. The same commitment may require refusal.

The evidence has not merely left uncertainty about performance. It has left uncertainty about whether commitment should occur.

That is an admissibility flip.

The same structure appears in grid interconnection. Under one plausible network state, queue entry exposes the project to bounded upgrade costs and manageable timing. Under another, queue entry creates cost-allocation exposure, congestion-dominated economics, and withdrawal non-credibility. Under another, the site-grid coupling fails. The same queue-entry decision may be admissible in one state and inadmissible in another.

A commitment should not harden while unresolved plausible states still produce admissibility flips unless the action is bounded in a way that preserves refusal, contains exposure, and prevents over-transfer.

INSTRUCTIONAL DRAFT

6.7 Ignorance dominance

Ignorance dominance occurs when unresolved plausible states still have the power to change whether a commitment is admissible.

It is not the same as uncertainty.

Uncertainty becomes decision-relevant when it affects estimates. Ignorance becomes decision-dominant when unresolved plausible states can still change the legitimacy of commitment.

Ignorance dominance is active when four conditions hold:

1. the plausible state space remains materially non-collapsed;
2. at least one plausible state changes the admissibility of the proposed commitment;
3. the evidence required to collapse the state space cannot be obtained without approaching or crossing the same exposure boundary;
4. refusal, re-siting, re-sequencing, or containment may become less credible after entry.

When ignorance dominance is active, the proposed commitment cannot be treated as ordinary risk acceptance. The institution does not simply face a risky decision. It faces an unstable decision boundary.

The correct posture may be bounded learning, deferral, or refusal, depending on whether uncertainty can be reduced without over-authorizing action.

6.8 Minimum evidence before admissibility

Admissibility requires a minimum evidence condition.

The minimum evidence condition does not demand certainty. It demands that evidence be adequate relative to the irreversible burden of the proposed action.

Before a commitment-bearing action can be considered, the evidence must be sufficient to define:

- the proposed commitment;
- the dominant uncertainty;
- the plausible states still consistent with the evidence;
- the consequences of those states for the commitment;
- the exposure created by action;
- the reversibility of the action in practice;
- the authority needed to refuse, defer, or constrain;
- the conditions under which the decision would be reconsidered.

Evidence is insufficient when it can support interest but not the commitment being proposed.

INSTRUCTIONAL DRAFT

For example, lunar remote sensing may support prospect ranking. It may not support resource-dependent infrastructure. A geophysical anomaly may support further investigation. It may not support drilling commitment. A grid opportunity may support preliminary screening. It may not support queue entry if network constraint ambiguity remains dominant.

The minimum evidence condition is therefore not an evidentiary checklist alone. It is a boundary between what the evidence can support and what the institution is trying to authorize.

6.9 Refusal authority as an admissibility condition

A commitment cannot be admissible if the authority to refuse collapses at entry.

Refusal authority is not an afterthought. It is part of the admissibility condition.

For a commitment to be admissible, the Decision Authority must retain practical capacity to:

- stop;
- defer;
- constrain;
- re-site;
- re-sequence;
- revoke;
- terminate;
- communicate uncertainty;
- absorb reversal burden;
- prevent precedent over-transfer.

If entry itself makes these capacities non-credible, the commitment may be inadmissible even if evidence appears strong.

This is especially important where capital, public posture, counterparty reliance, or institutional alignment creates continuation pressure.

A formal right to refuse is insufficient if practical refusal becomes politically, financially, or institutionally impossible.

6.10 Value of information and the admissibility boundary

Value-of-information reasoning is useful, but it must be bounded by admissibility.

The key question is not only whether additional information is valuable. It is whether the action required to obtain that information is itself commitment-bearing.

If the information can be obtained through bounded, reversible, refusal-preserving action, then deferral or bounded learning may be defensible.

INSTRUCTIONAL DRAFT

If the information can only be obtained by crossing the threshold the information is supposed to justify, then the system may be under ignorance dominance.

In such cases, “more data” is not a neutral recommendation. The data acquisition pathway must itself be governed.

A lunar excavation may provide decisive information about volatile state, but if excavation alters the site, creates corridor dependence, or begins ISRU commitment, the information pathway is commitment-bearing.

A geothermal well may reduce reservoir uncertainty, but if drilling creates capital, permitting, site, and offtake lock-in, it cannot be treated as neutral learning.

An injection test may improve containment understanding while initiating the permanence and liability concerns it seeks to evaluate.

Admissibility therefore asks whether the learning pathway preserves or consumes the conditions required for legitimate commitment.

6.11 Admissibility outcomes

At the formal entry layer, the admissibility condition is binary:

ADMISSIBLE

The proposed commitment may enter governed consideration within the defined evidence and authority boundary.

INADMISSIBLE

The proposed commitment may not enter governed consideration under current evidence, uncertainty, exposure, dependency, or authority conditions.

This binary output is important because admissibility is an eligibility gate.

It does not say the institution should proceed. It says whether the commitment can be considered at all.

If the commitment is inadmissible, the institution should not use governance posture classification to create the appearance of a proceed/defer/refuse choice. The commitment has failed the entry condition.

If the commitment is admissible, the next question becomes which governance posture remains defensible.

That is the role of governance posture classification.

INSTRUCTIONAL DRAFT

6.12 Admissible to consider is not approved to proceed

One of the most important distinctions in the framework is that admissibility does not mean approval.

A commitment can be admissible to consider while deferral remains the defensible posture. It can be admissible to consider while proceeding remains conditional. It can be admissible to consider while refusal remains required under governance posture analysis.

Admissibility only clears the threshold for governed deliberation. It does not choose the posture.

This matters because institutions often convert eligibility into momentum. Once a project is “cleared” for consideration, social and organizational pressure may treat it as implicitly approved.

The framework prevents this by separating:

- **entry eligibility** from **posture defensibility**;
- **admissibility** from **proceeding**;
- **consideration** from **authorization**.

This separation is necessary to preserve decision authority.

6.13 Admissibility under coupled irreversibility

Admissibility becomes more difficult when multiple irreversibility regimes are active at once.

A commitment may involve subsurface disturbance, placement lock-in, topological dependency, authority degradation, and precedent formation simultaneously.

Lunar ISRU is the clearest example. Volatile ambiguity is subsurface. Site preference is placement. Access, power, communications, and logistics are topological. Mission and agency commitment create authority pressure. First resource use creates precedent.

In such a case, admissibility cannot be judged by one regime alone.

A proposed action may appear acceptable as subsurface investigation but inadmissible as placement commitment. It may appear acceptable as site access but inadmissible as precedent. It may appear acceptable as bounded verification but inadmissible if support systems create resource dependency.

The admissibility question must therefore be regime-aware:

Which irreversibility regimes does this action activate, and does any active regime make the proposed commitment inadmissible under the current boundary?

INSTRUCTIONAL DRAFT

If any active regime creates unbounded exposure that cannot be constrained, the commitment may fail admissibility.

6.14 Admissibility and communication

Communication can be commitment-bearing.

Institutions often focus on physical actions, capital actions, or formal approvals. But public language, investor materials, partner statements, agency announcements, and internal narratives can harden commitment before formal execution.

Terms such as “primary site,” “selected corridor,” “go-forward plan,” “strategic resource,” “anchor asset,” or “resource-backed architecture” may reduce future refusal credibility.

Communication can also create precedent. A public claim that a signal is sufficient for infrastructure may normalize future over-authorization of evidence. A public deferral can be misunderstood as delayed approval. A proceed-compatible classification can be misused as endorsement.

Admissibility must therefore consider communication exposure.

A commitment may be inadmissible not because physical action has occurred, but because the proposed communication would degrade refusal authority or create precedent before the evidence has earned that role.

6.15 Boundary conditions and expiration

Admissibility is always boundary-bound.

A commitment may be admissible:

- for one site but not another;
- for bounded verification but not infrastructure placement;
- for internal analysis but not public announcement;
- for one evidence state but not after evidence expires;
- for one authority structure but not after decision rights shift;
- for one phase but not capital escalation.

The admissibility record must therefore define:

- the decision under review;
- the evidence boundary;
- the plausible state space;
- the irreversible threshold;
- the Decision Authority;

INSTRUCTIONAL DRAFT

- the permitted and prohibited interpretations;
- the expiration condition;
- the reconsideration triggers;
- the non-transferability rule.

Without these boundary conditions, admissibility can drift into generalized permission.

A bounded admissibility finding can become a broad approval. A site-specific determination can be transferred to another site. A verification permission can become infrastructure permission. A past decision can become precedent without control.

This is why admissibility requires provenance.

6.16 Section conclusion

Commitment admissibility governs entry into irreversible commitment.

It asks whether the proposed commitment may be considered under the current evidence and authority boundary. It separates feasibility from legitimacy, evidence from authorization, uncertainty from ignorance dominance, and consideration from approval.

A commitment is inadmissible when unresolved plausible states still change the legitimacy of commitment, when evidence is over-authorized, when the learning pathway consumes refusal authority, when reversal burden exceeds reversal capacity, or when active irreversibility regimes create unbounded exposure.

Admissibility is the first major contribution of the framework.

But admissibility alone is not enough.

Once a commitment is admissible to consider, the institution must still classify which posture remains defensible.

That is the subject of the next section.

7. Governance Posture: Between Admissibility and Action

Commitment admissibility determines whether a proposed irreversible commitment may enter governed consideration.

It does not determine that the institution should proceed.

INSTRUCTIONAL DRAFT

This distinction is essential. Without it, an admissibility finding can be misread as approval. A commitment that is admissible to consider may still require deferral. It may remain proceed-compatible only within narrow bounds. It may even require refusal once governance posture is tested against authority, dependency, timing, exposure, and refusal credibility.

This section develops the second decision layer: **governance posture classification**.

The central question is:

If the commitment is admissible to consider, which posture remains governance-defensible?

This is the layer between admissibility and action.

It prevents “may be considered” from becoming “should proceed.”

7.1 Why admissibility needs a posture layer

Admissibility is an entry gate. It evaluates whether the commitment may be brought into governed deliberation under the current evidence and authority boundary.

But entry into deliberation does not resolve the decision.

Once a commitment is admissible to consider, the institution still faces posture selection. It must determine whether proceeding, deferring, or refusing remains defensible under the conditions present at the threshold.

A project may clear admissibility because the commitment is sufficiently defined, the evidence boundary is explicit, refusal authority remains intact, and no active irreversibility regime makes consideration structurally invalid. Yet the correct posture may still be deferral because bounded learning would materially improve the decision without triggering irreversible exposure.

Conversely, a commitment may be admissible to consider because the decision is well-defined and authority remains coherent, but the governance posture may still be refusal because proceeding or deferring would create exposure that cannot be contained.

This distinction matters because institutions often treat procedural eligibility as substantive approval.

Once a project is “allowed into the room,” momentum begins. The fact that the commitment is discussable can become evidence that it is acceptable. The framework prevents this conversion by separating:

- admissibility from posture;

INSTRUCTIONAL DRAFT

- eligibility from action;
- consideration from authorization;
- governance classification from executive decision.

Governance posture exists to preserve this separation.

7.2 Governance posture is not decision authority

Governance posture classification does not replace the traditional decision process.

It does not command the institution to proceed, defer, or refuse. It does not override boards, agencies, mission authorities, investment committees, regulators, technical teams, or legal authorities.

It classifies which posture remains defensible under the evidence, authority, dependency, timing, and irreversibility conditions present at the threshold.

The distinction is precise:

The decision authority decides.

Governance posture classification determines which postures remain defensible for the decision authority to hold.

That distinction protects the framework from becoming an execution authority or decision substitute.

A governance posture classification does not say:

“Proceed.”

It says:

“Proceeding remains compatible with the governance constraints tested.”

It does not say:

“Delay.”

It says:

“Deferral is indicated because bounded learning remains the defensible posture.”

It does not say:

INSTRUCTIONAL DRAFT

“Cancel.”

It says:

“Refusal is required because no posture involving advancement remains governance-valid under the evidence and authority boundary.”

This language matters. “Proceed-compatible,” “deferral-indicated,” and “refusal-required” are posture classifications, not operational commands.

7.3 The three governance postures

The framework recognizes three primary posture classifications.

7.3.1 Proceed-Compatible

Proceed-Compatible means that proceeding within the stated boundary has not been eliminated by the governance constraints tested.

It does not mean the project is approved.

It does not mean execution is authorized.

It does not mean the project is optimal.

It does not mean the outcome is likely to succeed.

It means that, under the defined evidence boundary, authority condition, irreversibility regime, dependency topology, and refusal condition, proceeding remains a defensible posture for the Decision Authority to consider.

Proceed-Compatible requires that:

- the commitment is admissible to consider;
- residual uncertainty no longer controls whether the commitment is legitimate within the defined boundary;
- exposure is bounded or justified by the evidence authority;
- refusal, revocation, or constraint remains possible where required;
- dependencies do not exceed the admissible basis;
- precedent effects are contained;
- the action does not silently expand beyond the boundary reviewed.

Proceed-Compatible is therefore a narrow classification.

INSTRUCTIONAL DRAFT

It is especially important to preserve the phrase “within the stated boundary.” A proceed-compatible posture can be valid for bounded verification but invalid for infrastructure placement. It can be valid for internal site comparison but invalid for public site designation. It can be valid for limited drilling but invalid for capital escalation. It can be valid for one commitment threshold and invalid for the next.

Proceed-Compatible is not general permission.

It is a boundary-specific posture finding.

7.3.2 Deferral-Indicated

Deferral-Indicated means that bounded learning remains the governance-defensible posture, subject to stated controls and prohibited actions.

Deferral is not indecision.

In disturbance-coupled systems, deferral can be the most disciplined posture because it preserves refusal authority while reducing decision-dominant uncertainty.

Deferral-Indicated applies when:

- the commitment is admissible to consider but proceeding would over-authorize the current evidence;
- unresolved uncertainty remains material but can be reduced without crossing the irreversible threshold;
- the action can be resequenced to preserve optionality;
- bounded evidence acquisition can improve the decision without creating unbounded exposure;
- premature advancement would create unnecessary lock-in;
- refusal is not yet required because the pathway may become defensible under changed evidence or structure.

Deferral-Indicated is common in frontier systems because evidence often supports learning before it supports commitment.

A lunar volatile signal may justify bounded reconnaissance while deferring infrastructure placement.

A geothermal prospect may justify additional non-invasive or staged evidence work while deferring drilling commitment.

A grid project may justify further network analysis or alternative POI comparison while deferring queue entry.

INSTRUCTIONAL DRAFT

A critical minerals prospect may justify additional exploration while deferring land, access, offtake, or processing commitment.

The key is that deferral must be bounded. It cannot become a disguised continuation pathway.

A Deferral-Indicated posture should specify:

- what may continue;
- what may not continue;
- what evidence is required;
- what actions would trigger re-review;
- what would convert deferral into refusal;
- what would allow reconsideration for proceed-compatible posture.

Deferral without boundaries can become silent commitment.

7.3.3 Refusal-Required

Refusal-Required means that refusal is the only governance-valid posture under the stated evidence boundary and conditions.

This is not a moral judgment on the project, domain, technology, or institution. It is a structural conclusion about the defined commitment at the defined threshold.

Refusal-Required applies when:

- proceeding would create exposure the evidence cannot justify;
- deferral would not preserve refusal authority;
- bounded learning is unavailable or would itself cross the commitment threshold;
- the commitment depends on unresolved states that remain decision-dominant;
- authority coherence would collapse at entry;
- reversal burden exceeds reversal capacity;
- dependencies would become non-recoverable;
- precedent formation would normalize an inadmissible class of action;
- the proposed action cannot be constrained within an admissible boundary.

Refusal is required when both proceeding and deferral fail as defensible postures.

This is important. Refusal-Required is not simply a stronger form of deferral. It applies when waiting, staging, or further bounded learning cannot repair the structural problem without creating the exposure being evaluated.

INSTRUCTIONAL DRAFT

For example, a lunar architecture that depends on local volatiles before volatile accessibility is bounded may require refusal if the next step would harden infrastructure dependency and no bounded verification path preserves refusal authority.

A grid pathway may require refusal if queue entry, site control, and gen-tie assumptions are already coupled such that entry creates non-recoverable exposure before constraint topology can be bounded.

A repository decision may require refusal if authority continuity, public legitimacy, or long-duration stewardship cannot support the commitment's permanence.

Refusal-Required is one of the most important outputs because it treats non-entry as a valid governance result.

7.4 Posture defensibility conditions

Governance posture is defensible only if the posture remains coherent across several conditions.

These conditions include:

1. evidence authority;
2. uncertainty dominance;
3. exposure containment;
4. refusal credibility;
5. dependency bounds;
6. authority coherence;
7. timing and urgency;
8. communication effects;
9. precedent containment.

Each posture must be tested against these conditions.

7.4.1 Evidence authority

A posture is indefensible if it asks evidence to support more than it can legitimately authorize.

Evidence may justify prospecting but not commitment.

Evidence may justify comparison but not site hardening.

Evidence may justify bounded verification but not infrastructure placement.

Evidence may justify capability confidence but not delegated autonomy near irreversible thresholds.

A Proceed-Compatible posture requires evidence authority to match the proposed commitment.

INSTRUCTIONAL DRAFT

A Deferral-Indicated posture may be appropriate when evidence is meaningful but not commitment-grade.

A Refusal-Required posture may be required when the evidence is being structurally over-authorized and cannot be upgraded without crossing the threshold.

7.4.2 Uncertainty dominance

A posture is indefensible if unresolved uncertainty still controls whether the posture is legitimate.

If plausible states imply materially different posture outcomes, proceeding may be invalid. If uncertainty can be reduced through bounded learning, deferral may be defensible. If uncertainty cannot be reduced without creating the exposure being evaluated, refusal may be required.

The issue is not uncertainty magnitude. It is whether uncertainty remains decision-dominant.

7.4.3 Exposure containment

A posture is defensible only if exposure is contained within the boundary reviewed.

Proceed-Compatible requires that exposure remains bounded, accepted, and within the evidence and authority basis.

Deferral-Indicated requires that the deferral pathway does not itself become commitment-bearing.

Refusal-Required applies when exposure cannot be contained by proceeding or deferring.

Exposure containment must include physical, capital, institutional, topological, authority, and precedent exposure.

7.4.4 Refusal credibility

A posture fails if it destroys the ability to refuse later when refusal remains necessary to governance.

Proceed-Compatible requires that refusal, constraint, or termination remains available where the boundary requires it.

Deferral-Indicated requires that the deferral period preserves real optionality.

Refusal-Required applies when refusal credibility would collapse if the project continued.

Refusal credibility is not symbolic. It must be practical.

INSTRUCTIONAL DRAFT

7.4.5 Dependency bounds

A posture is defensible only if dependencies remain within the admissible basis.

Dependencies can include:

- site dependence;
- corridor dependence;
- grid dependence;
- resource dependence;
- infrastructure dependence;
- counterparty dependence;
- data dependence;
- capital dependence;
- regulatory dependence;
- public narrative dependence.

Proceed-Compatible requires dependency bounds to remain intact.

Deferral-Indicated may require dependency freezes or prohibitions.

Refusal-Required applies when dependencies would exceed the boundary by continuing.

7.4.6 Authority coherence

A posture is invalid if formal authority and practical control diverge so far that the Decision Authority cannot govern the consequence of the posture.

Proceed-Compatible requires coherent authority.

Deferral-Indicated requires the authority to enforce bounded learning and prevent covert escalation.

Refusal-Required applies when only refusal preserves authority coherence.

Authority coherence includes decision rights, veto rights, consequence bearing, revocation capacity, and public communication discipline.

7.4.7 Timing and urgency

Urgency does not automatically justify proceeding.

Timing pressure may be real, but it is a governance variable. It can strengthen or weaken posture defensibility depending on whether delay destroys optionality faster than action creates exposure.

INSTRUCTIONAL DRAFT

A Proceed-Compatible posture may be defensible if delay would create greater irreversible loss and the evidence boundary can support action.

A Deferral-Indicated posture may be defensible if bounded waiting preserves more optionality than it costs.

A Refusal-Required posture may be necessary if urgency is artificial or if urgency is being used to override admissibility discipline.

The framework treats urgency as evidence to be evaluated, not as a command.

7.4.8 Communication effects

A posture can be undermined by how it is communicated.

A proceed-compatible classification can be misread as approval.

A deferral can be misread as pending approval.

A refusal can be misread as rejection of an entire domain.

A bounded action can be misread as strategic commitment.

Communication must therefore preserve the posture boundary.

Proceed-Compatible communication must not imply general authorization.

Deferral-Indicated communication must preserve uncertainty and prevent silent escalation.

Refusal-Required communication must state the defined boundary and prevent over-transfer.

7.4.9 Precedent containment

A posture must consider what it makes easier to cite, repeat, normalize, or inherit.

Proceed-Compatible may create precedent if future actors treat it as broad permission.

Deferral-Indicated may create precedent if future actors treat it as conditional approval.

Refusal-Required may create precedent if future actors treat it as domain rejection rather than threshold-specific refusal.

Posture defensibility therefore requires precedent containment.

A governance posture is not complete until its precedent implications are bounded.

INSTRUCTIONAL DRAFT

7.5 Posture classification across irreversibility regimes

Governance posture should be tested against active irreversibility regimes.

A posture may be defensible in one regime and indefensible in another.

For example, a lunar action may be defensible as bounded subsurface investigation but indefensible as placement commitment if it privileges a site. It may be defensible as scientific reconnaissance but indefensible as precedent if public communication normalizes resource dependency.

The same action must therefore be classified across regimes:

- subsurface;
- placement;
- topology;
- authority;
- precedent;
- coupled irreversibility.

A posture cannot be governance-defensible if it fails in an active regime that cannot be bounded, isolated, or repaired before the threshold is crossed.

This regime-specific posture test prevents superficial classification.

It also explains why some actions are difficult to govern: the problem is not one form of irreversibility but the coupling among several.

7.6 Examples of posture classification

7.6.1 Lunar ISRU

A volatile signal may be strong enough to justify bounded reconnaissance. It may not be strong enough to justify fixed power placement, excavation dependency, or ISRU architecture.

Admissibility may clear for bounded investigation.

The governance posture may be:

Deferral-Indicated for infrastructure commitment.

The rationale:

INSTRUCTIONAL DRAFT

- evidence supports exploration but not resource dependency;
- plausible volatile states still imply different site and excavation decisions;
- disturbance may alter the evidence baseline;
- support systems could create site and corridor lock-in;
- precedent risk is high if a signal becomes an infrastructure anchor.

Proceed-Compatible may apply only to bounded non-hardening investigation.

Refusal-Required may apply to premature ISRU dependency.

7.6.2 Grid interconnection

A project may be admissible to consider entering an interconnection queue if the decision boundary, site, POI, deposit exposure, and withdrawal conditions are defined.

Governance posture may still be:

Deferral-Indicated if constraint-topology ambiguity remains decision-dominant and additional analysis can preserve optionality.

Or:

Refusal-Required if queue entry would create site-grid coupling, capital escalation, and withdrawal non-credibility before upgrade exposure can be bounded.

Proceed-Compatible would require that exposure is bounded, exit remains credible, network ambiguity no longer controls the decision, and site-grid dependencies remain contained.

7.6.3 Critical minerals

A critical minerals prospect may be strategically important and still not justify extraction-pathway commitment.

A Proceed-Compatible posture may apply to bounded exploration.

Deferral-Indicated may apply to land control, processing assumptions, access roads, or offtake alignment until resource continuity, grade, processing, water, power, and closure assumptions become decision-grade.

Refusal-Required may apply if strategic urgency is being used to override evidence adequacy and create irreversible social, environmental, capital, or infrastructure exposure.

7.6.4 Geological disposal

INSTRUCTIONAL DRAFT

A repository commitment may be admissible to consider only under tightly defined evidence, authority, monitoring, and public legitimacy boundaries.

Governance posture may be Proceed-Compatible only if long-duration responsibility, institutional continuity, containment assumptions, and revocation or retrievability logic remain coherent.

Deferral-Indicated may apply if additional characterization can materially improve the basis without creating irreversible public or infrastructural commitment.

Refusal-Required may apply if authority coherence or public legitimacy collapses before entry, even if technical siting evidence remains favorable.

7.7 Posture and institutional responsibility

Governance posture classification preserves institutional responsibility.

It does not remove judgment from the Decision Authority. It makes the boundary of judgment more explicit.

The Decision Authority remains responsible for:

- action;
- inaction;
- reliance;
- communication;
- escalation;
- execution;
- compliance;
- outcomes.

The posture classification states which stance remains defensible under the framework.

An institution may choose to act contrary to a posture classification. But doing so rejects the governance basis of the classification. It cannot later rely on the classification as support for action.

This distinction is essential for institutional defensibility.

7.8 Posture expiration and reclassification

INSTRUCTIONAL DRAFT

Governance posture is not permanent.

A posture classification can expire or require reclassification when:

- new evidence emerges;
- evidence degrades;
- assumptions expire;
- dependencies accumulate;
- authority changes;
- communication alters the public boundary;
- counterparties rely on continuation;
- capital exposure increases;
- physical disturbance occurs;
- precedent risk changes;
- the action changes scope;
- the threshold changes.

A Proceed-Compatible posture may become Deferral-Indicated if new uncertainty becomes decision-dominant.

A Deferral-Indicated posture may become Proceed-Compatible if bounded evidence reduces the relevant ambiguity.

A Deferral-Indicated posture may become Refusal-Required if the deferral pathway itself begins creating exposure.

A Refusal-Required posture may become reconsiderable only if the structural basis for refusal changes, not merely because the institution wants a different answer.

This makes posture classification time-bound, evidence-bound, and authority-bound.

7.9 The posture layer and commitment integrity

Governance posture also establishes the basis for future integrity review.

If a commitment proceeds under a proceed-compatible posture, Instrument C later asks whether the commitment remains within that posture's admissible basis.

If a commitment proceeds despite Deferral-Indicated or Refusal-Required posture, then commitment integrity is compromised at entry.

If a commitment initially defers but continues accumulating exposure through side channels, the deferral posture may fail and require re-evaluation.

INSTRUCTIONAL DRAFT

Thus the posture layer is not only a threshold classification. It becomes part of the record against which future commitment integrity is evaluated.

This is the bridge between entry and persistence.

7.10 Section conclusion

Governance posture classification occupies the space between admissibility and action.

Admissibility asks whether a commitment may be considered.

Governance posture asks which stance remains defensible once consideration is allowed.

The distinction prevents admissibility from becoming approval, preserves the Decision Authority's role, and protects against the institutional tendency to convert eligibility into momentum.

The three posture classifications are:

- Proceed-Compatible;
- Deferral-Indicated;
- Refusal-Required.

Each must be tested against evidence authority, uncertainty dominance, exposure containment, refusal credibility, dependency bounds, authority coherence, timing pressure, communication effects, and precedent containment.

The next section decomposes irreversibility itself into regimes, because posture defensibility depends on how a commitment becomes difficult to reverse.

8. Regimes of Irreversibility

Irreversibility is often treated as a single property: an action is either reversible or it is not. That framing is too coarse for disturbance-coupled exploration systems.

Commitments become difficult to reverse through different structural mechanisms. A project may become irreversible because the subsurface has been disturbed, because a site has become the default location, because a network dependency has formed, because authority to refuse has degraded, or because an early action has created precedent.

These are not the same failure mode.

INSTRUCTIONAL DRAFT

A commitment may be physically reversible but institutionally irreversible. It may be technically reversible but financially non-credible. It may be locally reversible but topologically locked into a larger system. It may be operationally minor but precedent-forming. It may be defensible in one regime of irreversibility and indefensible in another.

This section decomposes irreversibility into five primary regimes:

1. **Subsurface irreversibility**
2. **Placement irreversibility**
3. **Topological irreversibility**
4. **Authority irreversibility**
5. **Precedent irreversibility**

A sixth condition, **coupled irreversibility**, applies when multiple regimes reinforce one another.

The purpose of this decomposition is to make commitment failure diagnosable. It allows an institution to ask not merely whether an action is hard to reverse, but **how** it becomes hard to reverse, through which regime, and at what threshold.

8.1 Why irreversibility requires regimes

In ordinary project language, irreversibility is often associated with physical action: construction, excavation, drilling, injection, deployment, emplacement, or closure. These actions matter, but they are not the only way commitment forms.

A project can become difficult to reverse before physical work begins.

A site can become politically favored.

A corridor can become the assumed route.

A point of interconnection can become the project's organizing geometry.

A resource assumption can become embedded in architecture.

A capital sequence can create continuation pressure.

A public statement can make future refusal costly.

A first demonstration can normalize a future class of activity.

In these cases, irreversibility does not arise from physical permanence alone. It arises from the collapse of practical alternatives.

The framework therefore treats irreversibility as multi-regime. Each regime describes a pathway through which optionality declines and refusal becomes harder.

This has two implications.

INSTRUCTIONAL DRAFT

First, admissibility must be regime-aware. A commitment may be admissible under one regime and inadmissible under another.

Second, commitment integrity must be monitored by regime. A commitment may remain intact physically while failing institutionally, topologically, or precedentially.

8.2 Subsurface irreversibility

Subsurface irreversibility occurs when learning requires disturbance or intervention into an uncertain physical system.

The defining feature is that the act of obtaining knowledge changes the target or creates exposure associated with entering the target.

Examples include:

- lunar trenching or excavation to verify volatile state;
- geothermal drilling to test reservoir productivity;
- critical mineral drilling to confirm continuity and grade;
- carbon storage injection to test containment and pressure behavior;
- repository characterization that creates access pathways and institutional commitment;
- seabed coring or installation preparation that alters baseline conditions.

Subsurface irreversibility is not simply that the ground is disturbed. It is that the disturbance occurs before the system is sufficiently understood, and that the disturbance may change the future decision environment.

In lunar ISRU, excavation or heating may reveal volatile form while altering thermal conditions, mobilizing volatiles, disturbing regolith, or contaminating future measurements. The same action that improves knowledge may degrade the evidentiary baseline.

In geothermal, drilling may reveal reservoir conditions while creating capital exposure, site commitment, permitting posture, and future pressure to continue. The well becomes both an information instrument and a commitment mechanism.

In carbon storage, injection may generate evidence about plume behavior, pressure evolution, and containment while initiating the permanence condition at issue.

Subsurface irreversibility is therefore the regime where the statement “we need more data” becomes governance-sensitive. The question is not whether more data would be useful. It is whether the data can be obtained without creating the exposure the data is meant to justify.

8.2.1 Admissibility implication

INSTRUCTIONAL DRAFT

A subsurface commitment is inadmissible when the dominant uncertainty cannot be reduced without crossing a disturbance threshold that current evidence cannot justify.

This does not prohibit all disturbance. It requires disturbance to be bounded, evidence-generating, and refusal-preserving.

A subsurface action may be admissible if:

- disturbance is limited;
- knowledge gain is material;
- the evidence baseline remains interpretable;
- alternatives remain credible;
- support systems do not harden into dependency;
- termination conditions are explicit;
- the action is not treated as infrastructure permission.

It becomes inadmissible when disturbance converts unresolved ignorance into site, capital, infrastructure, or precedent commitment.

8.3 Placement irreversibility

Placement irreversibility occurs when a location, site, route, corridor, landing zone, access point, or physical geometry becomes difficult to abandon.

The defining feature is spatial fixation.

A project may become committed because one place becomes easier to continue than replace. Placement irreversibility forms when attention, access, capital, support systems, permitting, communications, or public language begin organizing around a location.

Examples include:

- lunar landing site preference;
- primary ISRU site designation;
- mine access road alignment;
- repository siting;
- transmission corridor selection;
- offshore cable landfall;
- gen-tie route fixation;
- seabed installation route;
- drilling pad selection.

INSTRUCTIONAL DRAFT

Placement irreversibility can form before construction. A site may become the project's center of gravity through repeated analysis, mapping, partner conversation, public language, or early access planning.

The danger is that location becomes default before evidence becomes site-grade.

In lunar ISRU, a volatile signal may privilege a site before volatile accessibility, continuity, mechanical behavior, and disturbance response are known. Once access paths, power geometry, communications, and logistics begin to align around that site, re-siting becomes increasingly non-credible.

In grid development, a site may become coupled to a point of interconnection before upgrade exposure and congestion risk are sufficiently bounded.

In geological disposal, a candidate site may become institutionally embedded through characterization, local engagement, regulatory process, and political identity before long-duration integrity is secured.

8.3.1 Admissibility implication

Placement commitment is inadmissible when unresolved uncertainty could still change the site, route, corridor, or geometry that should be selected.

A placement action may be admissible if:

- alternative sites remain viable;
- the site is not publicly or financially hardened;
- evidence supports the spatial scale of commitment;
- access and support systems remain temporary;
- re-siting remains credible;
- the location is not treated as inevitable.

It becomes inadmissible when location begins to govern the project before the evidence has earned authority over location.

8.4 Topological irreversibility

Topological irreversibility occurs when networks, dependencies, queues, corridors, logistics, support systems, or shared infrastructure create structural lock-in.

The defining feature is dependency architecture.

Topological irreversibility is not about one site or action alone. It is about the structure of relationships created among actions, assets, institutions, and future decisions.

INSTRUCTIONAL DRAFT

Examples include:

- grid interconnection queue exposure;
- gen-tie routing;
- shared transmission capacity;
- lunar power and logistics networks;
- communications and navigation anchoring;
- ISRU processing dependency;
- mine-to-processing logistics;
- repository transport and access systems;
- orbital coordination regimes;
- autonomous system permission dependencies.

Topological irreversibility forms when one decision creates a network that later decisions must inherit.

A project may become locked in not because any single node is irreversible, but because the network created among nodes becomes hard to unwind.

In grid interconnection, queue entry can create topological exposure because the project becomes linked to network studies, competing projects, upgrade allocation, congestion, deliverability, gen-tie routing, and withdrawal conditions. The uncertainty is not merely cost. It is constraint topology.

In lunar ISRU, access repetition, power placement, communications, navigation, logistics, and excavation planning create a topology around a resource assumption. The system begins to behave as though the resource interpretation is true before that interpretation is decision-grade.

In critical minerals, a prospect can become topologically locked through roads, water, power, processing, offtake, and tailings assumptions. Even before extraction begins, the infrastructure topology can privilege one resource interpretation.

8.4.1 Admissibility implication

A topological commitment is inadmissible when dependencies begin to form around assumptions that remain decision-dominant.

A topological action may be admissible if:

- dependencies are bounded;
- interfaces remain reversible;
- alternatives remain live;
- network effects are mapped;
- exit remains credible;
- dependency growth is capped;
- no single assumption becomes architecture-governing prematurely.

INSTRUCTIONAL DRAFT

It becomes inadmissible when the network begins to make the assumption true in practice by organizing future decisions around it.

8.5 Authority irreversibility

Authority irreversibility occurs when decision rights, refusal credibility, revocation capacity, or governance coherence degrade after entry.

The defining feature is loss of practical authority.

An institution may retain formal authority while losing the practical ability to use it. This can happen through capital deployment, public commitments, partner reliance, regulatory posture, political expectation, internal incentives, or narrative hardening.

Examples include:

- board alignment that makes reversal embarrassing;
- public announcement of a preferred site;
- partner reliance on continuation;
- capital sequencing that penalizes delay;
- regulatory filing that establishes posture;
- agency commitment to mission architecture;
- offtake agreements;
- strategic-resource framing;
- delegated autonomy near irreversible thresholds.

Authority irreversibility is often invisible because institutions can still point to formal rights. They may say the decision is reversible, optional, non-binding, or preliminary. But if refusal would now be institutionally unacceptable, authority has degraded.

The relevant question is:

Can the Decision Authority still exercise refusal without disproportionate loss?

If not, authority irreversibility has formed.

Authority irreversibility is especially important because it can arise before physical exposure. A public claim, capital process, or partner alignment can make future refusal non-credible even before construction or disturbance.

8.5.1 Admissibility implication

A commitment is inadmissible if entry destroys the authority conditions required to govern it.

INSTRUCTIONAL DRAFT

An authority-sensitive action may be admissible if:

- the Decision Authority is named;
- veto and refusal rights remain practical;
- consequence-bearing is aligned;
- public language remains conditional;
- counterparties understand the boundary;
- capital exposure is capped;
- revocation conditions are explicit.

It becomes inadmissible when the institution cannot credibly say no after entry.

Authority irreversibility also links directly to commitment integrity. A commitment that was admissible at entry can lose integrity if authority later degrades.

8.6 Precedent irreversibility

Precedent irreversibility occurs when an early action changes the future admissibility environment by normalizing a class of commitments.

The defining feature is transfer.

A decision becomes precedent-forming when future actors can cite, repeat, normalize, inherit, or misuse it beyond the original boundary.

Precedent is a form of exposure because it changes what future decisions appear to permit.

Examples include:

- first lunar volatile extraction treated as resource-use permission;
- first lunar site commitment treated as a model for infrastructure sequencing;
- first seabed disturbance treated as acceptable pilot practice;
- first repository authorization treated as proof of institutional feasibility;
- first autonomy delegation near irreversible thresholds treated as operational permission;
- first grid-interconnection posture treated as administrative despite lock-in;
- first critical-minerals project where strategic urgency overrides evidence adequacy.

Precedent irreversibility is particularly powerful in frontier systems because early actions may define norms before mature governance exists.

A bounded action can become generalized.

A local determination can become a domain rule.

A deferral can be misread as future approval.

INSTRUCTIONAL DRAFT

A proceed-compatible posture can be misread as endorsement.
A refusal can be misread as domain rejection.

Precedent therefore requires containment.

The relevant question is:

What could this decision cause future actors to cite, repeat, normalize, or inherit?

This is why the framework treats the “afterlife” of decisions as part of governance.

8.6.1 Admissibility implication

A precedent-forming action is inadmissible when its likely future reuse exceeds the evidence and authority boundary under which it is made.

A precedent-sensitive action may be admissible if:

- the decision boundary is explicit;
- non-transferability is stated;
- public use is controlled;
- evidence authority is not overstated;
- future use requires re-screening;
- precedent implications are recorded;
- the action does not normalize commitment beyond its scope.

It becomes inadmissible when a local action is likely to create a broader rule that the evidence cannot support.

8.7 Coupled irreversibility

Coupled irreversibility occurs when multiple regimes reinforce one another and create system-level lock-in.

This is the most important condition in high-consequence frontier systems.

Many commitments do not fail through a single regime. They fail because regimes interact.

In lunar ISRU:

- volatile uncertainty creates subsurface irreversibility;
- site preference creates placement irreversibility;
- access, power, communications, and logistics create topological irreversibility;
- mission and partner alignment create authority irreversibility;

INSTRUCTIONAL DRAFT

- first resource-dependent infrastructure creates precedent irreversibility.

Together, these regimes can transform a signal into infrastructure lock-in.

The sequence may look like this:

Volatile signal → site preference → access repetition → support placement → disturbance → corridor formation → excavation planning → ISRU dependency → infrastructure lock-in

Each step reinforces the next.

In grid interconnection:

- site preference creates placement irreversibility;
- queue entry and POI selection create topological irreversibility;
- capital and developer expectations create authority irreversibility;
- treatment of queue entry as administrative creates precedent irreversibility.

In geological disposal:

- site characterization creates subsurface and placement commitment;
- repository infrastructure creates topology;
- licensing and public legitimacy create authority irreversibility;
- first authorization creates precedent.

Coupled irreversibility matters because a project may appear manageable when each regime is evaluated separately, while the combined structure is not manageable.

8.7.1 Admissibility implication

A commitment under coupled irreversibility is admissible only if the combined exposure remains bounded.

It is not enough to show that disturbance is bounded if placement is not.

It is not enough to show that placement is bounded if topology is hardening.

It is not enough to show that topology is bounded if refusal authority is collapsing.

It is not enough to show that authority remains formal if precedent is expanding.

Coupled irreversibility requires system-level review.

If any active regime creates unbounded exposure that cannot be isolated, constrained, or reversed, the overall commitment may be inadmissible.

INSTRUCTIONAL DRAFT

8.8 Regime-specific failure patterns

The five regimes produce different failure patterns.

Subsurface failure

The system is disturbed before the dominant subsurface ambiguity is reduced.

Typical failure:

Learning changes the system before learning justifies commitment.

Placement failure

A site, route, or corridor becomes default before the evidence supports spatial fixation.

Typical failure:

Location becomes commitment before location is earned.

Topological failure

A network of dependencies forms around an assumption that remains unresolved.

Typical failure:

Dependency makes the assumption harder to abandon than to continue.

Authority failure

The institution loses practical refusal power before the decision is fully governed.

Typical failure:

The authority to say no collapses while formal authority remains.

Precedent failure

A bounded action becomes generalized beyond its evidence boundary.

Typical failure:

One decision becomes a rule it was never qualified to create.

Coupled failure

INSTRUCTIONAL DRAFT

Multiple regimes reinforce one another until reversal becomes non-credible.

Typical failure:

Small commitments assemble into system-level lock-in.

8.9 Regime-specific governance questions

Each regime has a governing question.

Regime	Governing question
Subsurface	Does learning require changing the system being evaluated?
Placement	Does the action make one site, route, or geometry easier to continue than abandon?
Topology	Does the action create dependencies that later decisions must inherit?
Authority	Does the action weaken the practical ability to refuse, revoke, or re-sequence?
Precedent	Does the action normalize future commitments beyond its boundary?
Coupled	Do multiple regimes reinforce one another faster than evidence or authority can govern them?

These questions should be asked before irreversible thresholds are crossed and after commitments begin to harden.

They provide the diagnostic bridge between admissibility and integrity.

8.10 Relationship to commitment admissibility

Regimes of irreversibility determine what admissibility must test.

A commitment cannot be evaluated only in the abstract. It must be evaluated by the regimes it activates.

If the proposed action activates subsurface irreversibility, admissibility must test whether the evidence supports disturbance.

If it activates placement irreversibility, admissibility must test whether evidence supports spatial fixation.

If it activates topological irreversibility, admissibility must test whether dependencies are bounded.

INSTRUCTIONAL DRAFT

If it activates authority irreversibility, admissibility must test whether refusal authority survives entry.

If it activates precedent irreversibility, admissibility must test whether future transfer can be contained.

If it activates coupled irreversibility, admissibility must test whether the combined system remains governable.

This makes admissibility regime-specific.

8.11 Relationship to governance posture

Governance posture can differ by regime.

A proposed action may be:

- proceed-compatible under subsurface evidence;
- deferral-indicated under placement;
- refusal-required under precedent.

This means posture classification must identify the controlling regime.

For example, bounded lunar trenching may be defensible as subsurface learning but not if it is publicly framed as the first step toward ISRU infrastructure.

A grid queue entry may be manageable financially but indefensible topologically if it couples site control, gen-tie routing, and capital sequencing before constraint ambiguity is bounded.

A critical minerals exploration program may be defensible geologically but indefensible institutionally if strategic-resource language makes future refusal non-credible.

Governance posture is therefore not simply a project-level label. It is a regime-aware classification.

8.12 Relationship to commitment integrity

Commitment integrity can degrade in one regime while remaining intact in another.

A commitment may remain physically bounded but lose authority integrity.

A site may remain technically viable but become topologically over-dependent.

INSTRUCTIONAL DRAFT

A repository may remain geologically stable while institutional authority degrades.
A pilot may remain operationally contained while precedent expands.

Instrumentally, this means that commitment integrity should be evaluated by regime:

- Has subsurface evidence changed?
- Has placement become harder to reverse?
- Has topology accumulated new dependencies?
- Has authority degraded?
- Has precedent expanded?
- Have regimes coupled in ways not present at entry?

A commitment remains within its admissible basis only if the active regimes remain within the bounds under which the commitment was admitted.

8.13 Section conclusion

Irreversibility is not a single condition. It is a set of structural pathways through which refusal, reversal, re-siting, re-sequencing, or abandonment becomes less credible.

The five primary regimes are subsurface, placement, topology, authority, and precedent. Coupled irreversibility occurs when these regimes reinforce one another.

This decomposition matters because commitment can fail in different ways. A project may be technically feasible but placement-inadmissible. It may be evidence-supported but authority-incoherent. It may be locally bounded but precedent-forming. It may be manageable in one regime and indefensible in the coupled system.

Regime analysis therefore provides the diagnostic structure required for both entry and persistence.

The next section develops commitment integrity: the condition under which an existing or emerging commitment remains within its admissible basis as evidence, assumptions, dependencies, authority, and precedent evolve.

9. Commitment Integrity: The Persistence Problem

Commitment admissibility governs entry.

It asks whether a proposed irreversible commitment may be considered under the evidence, uncertainty, authority, dependency, and refusal conditions present at the threshold.

INSTRUCTIONAL DRAFT

But entry is not the end of the governance problem.

A commitment that was admissible at entry may become unstable as conditions change. Evidence may weaken. Assumptions may expire. Dependencies may accumulate. Support systems may harden. Public language may expand the commitment beyond its original boundary. Authority to refuse or revoke may degrade. A bounded commitment may begin functioning as precedent for a broader class of action.

This is the persistence problem.

Commitment integrity is the condition under which an existing or emerging commitment remains within its admissible basis as evidence, assumptions, dependencies, authority, and precedent evolve.

The central question is:

Does the commitment remain defensible within the boundary that made it admissible, or has the structure changed enough that it must be maintained, constrained, re-evaluated, or terminated?

This section develops commitment integrity as the second major contribution of the framework.

9.1 Why admissibility is not enough

Admissibility is time-bound, evidence-bound, authority-bound, and threshold-bound.

A commitment may be admissible at the moment it begins because the evidence supports a bounded action, the uncertainty has been classified, refusal remains credible, dependencies are contained, and the Decision Authority retains coherent control.

Those conditions may not remain true.

A bounded lunar verification action may later begin anchoring site preference, access repetition, support placement, and ISRU dependency.

A geothermal drilling commitment may begin under an explicit reservoir hypothesis, then become capital-locked even after drilling reveals a different system.

A carbon storage project may enter under bounded containment assumptions, then require re-evaluation as plume behavior, pressure response, or monitoring limits diverge from the original basis.

INSTRUCTIONAL DRAFT

A critical minerals project may begin as exploration, then harden into access roads, processing assumptions, water dependency, offtake expectations, and strategic-resource narrative before deposit continuity is adequate.

A repository program may enter under a valid siting and containment basis, then face institutional, monitoring, public legitimacy, or stewardship conditions that undermine continuation.

These are not merely cases where a project becomes riskier. They are cases where the commitment may leave the boundary under which it was admissible.

That is why admissibility requires a persistence counterpart.

9.2 The admissible basis

The **admissible basis** is the set of conditions under which a commitment was allowed to begin or enter governed consideration.

It includes:

- the defined commitment;
- the irreversible threshold;
- the evidence record;
- the evidence freeze or evidence boundary;
- the plausible state space;
- the dominant uncertainty;
- the active irreversibility regimes;
- the Decision Authority;
- the authority boundary;
- the refusal and revocation conditions;
- the dependency topology;
- the assumptions registry;
- the permitted actions;
- the prohibited actions;
- the expiration conditions;
- the reconsideration triggers;
- the reliance perimeter;
- the precedent containment conditions.

Commitment integrity is evaluated against this admissible basis.

The question is not:

Was the original decision correct?

INSTRUCTIONAL DRAFT

The question is:

Does the commitment still remain inside the conditions that made it defensible?

This distinction matters because a commitment can fail integrity without proving that the entry decision was wrong. The world may change. Evidence may improve or degrade. Dependencies may grow. Authority may shift. The original basis may expire.

Integrity review is therefore not a retrospective blame mechanism. It is a persistence test.

9.3 Integrity is distinct from performance

Commitment integrity is not the same as project performance.

A project may perform well while losing integrity. A commitment may produce early success while exceeding its admissible basis.

For example, an early lunar surface operation may succeed technically while normalizing resource dependency beyond the evidence boundary. A critical minerals project may attract capital while embedding weak resource assumptions into infrastructure. A repository may meet early monitoring expectations while institutional continuity or public legitimacy degrades.

Conversely, a project may encounter difficulty while retaining integrity. A bounded verification action may reveal unfavorable evidence, yet remain legitimate because its admissible basis included termination conditions and preserved refusal authority.

Performance asks:

Is the project working?

Integrity asks:

Does the commitment remain defensible within the basis under which it was allowed to form?

That difference is essential in irreversible systems because outcome success can conceal governance failure, and outcome difficulty can coexist with sound threshold governance.

9.4 Integrity degradation pathways

Commitment integrity degrades through several recurring pathways.

INSTRUCTIONAL DRAFT

The five primary pathways are:

1. evidence drift;
2. assumption expiry;
3. dependency accumulation;
4. authority degradation;
5. precedent expansion.

These pathways correspond to the movement of a commitment away from its admissible basis.

9.4.1 Evidence drift

Evidence drift occurs when new observations, measurement degradation, failed assumptions, contamination, or reinterpretation weaken the evidentiary basis that supported entry.

Evidence drift can occur because the system changes, because new evidence emerges, or because the original evidence is found to have been overextended.

In lunar ISRU, evidence drift may occur if intrusive verification reveals that volatile distribution is patchier, deeper, more diffuse, more mechanically difficult, or less accessible than the signal originally suggested. It may also occur if disturbance alters the evidence baseline, making later measurements harder to interpret.

In geothermal systems, drilling may reveal that reservoir permeability, temperature distribution, recharge, stress state, or fault behavior differs from the entry hypothesis.

In carbon storage, plume movement, pressure response, fault behavior, or monitoring performance may diverge from the model used to support injection.

In critical minerals, drilling may reveal discontinuity, grade variability, metallurgical complexity, water constraints, or waste burden inconsistent with the prospectivity narrative.

In geological disposal, new geologic, hydrologic, seismic, geochemical, or institutional evidence may undermine assumptions supporting containment or stewardship.

Evidence drift threatens integrity when the commitment continues to rely on an evidence basis that no longer supports it.

9.4.2 Assumption expiry

INSTRUCTIONAL DRAFT

Assumption expiry occurs when assumptions valid at entry lose validity over time, scale, phase, or context.

Assumptions are often necessary in frontier and infrastructure systems. The problem is not that assumptions exist. The problem is that they can remain embedded after their validity period has expired.

An assumption may expire because:

- the project moves to a new phase;
- the spatial scale changes;
- the operational burden increases;
- the institutional context changes;
- the evidence boundary shifts;
- the assumption was explicitly time-limited;
- the system evolves;
- the decision is transferred to a new site, route, asset, or authority.

For example, a lunar assumption valid for bounded reconnaissance may not remain valid for excavation dependency. A grid assumption valid for screening may not remain valid for queue entry or capital escalation. A mineral prospectivity assumption valid for drilling may not remain valid for processing infrastructure or offtake commitments. A repository assumption valid for siting may not remain valid for closure or intergenerational stewardship.

Assumption expiry threatens integrity when the commitment continues as if the original assumptions still govern a changed decision context.

Integrity requires assumptions to be tracked, bounded, and expired when their authority ends.

9.4.3 Dependency accumulation

Dependency accumulation occurs when support systems, counterparties, capital sequences, infrastructure pathways, logistics, public narratives, or operational routines become harder to unwind than originally bounded.

A commitment may enter with limited dependency. Over time, dependencies can grow.

A lunar verification step may begin with temporary access but later require repeat traverses, power support, communications, logistics, excavation planning, and processing assumptions.

A geothermal project may begin with a drilling commitment but later tie into offtake agreements, road access, power infrastructure, financing, and community expectations.

INSTRUCTIONAL DRAFT

A grid project may begin with queue entry but later accumulate site control, gen-tie routing, study deposits, capital commitments, and partner alignment.

A repository may begin with site characterization but later accumulate access infrastructure, licensing posture, local economic reliance, transport planning, monitoring systems, and public identity.

Dependency accumulation threatens integrity when the actual commitment becomes larger than the commitment that was admitted.

The key question is:

What now depends on this commitment that did not depend on it at entry?

If the answer includes new dependencies that make refusal, re-siting, re-sequencing, or termination non-credible, integrity may be compromised.

9.4.4 Authority degradation

Authority degradation occurs when the institution loses practical ability to refuse, revoke, constrain, re-site, re-sequence, or terminate.

Authority degradation may occur even when formal authority remains.

It may be caused by:

- public commitment;
- capital exposure;
- board alignment;
- partner reliance;
- regulatory posture;
- internal incentives;
- political pressure;
- reputational concern;
- community reliance;
- mission narrative;
- delegated autonomy;
- loss of veto clarity;
- diffusion of decision responsibility.

Authority degradation is especially important because commitment integrity depends on the continued ability to govern the commitment.

A commitment that cannot be stopped cannot be said to remain governed.

INSTRUCTIONAL DRAFT

If the Decision Authority no longer has practical capacity to enforce the boundary, the commitment may exit its admissible basis even if the technical evidence remains favorable.

For example, a lunar program may retain formal authority to revise architecture, but if public and partner commitments have made re-siting non-credible, authority has degraded. A repository authority may retain legal power, but if institutional legitimacy collapses, formal authority may no longer be sufficient. A grid developer may retain a withdrawal option, but if capital and site exposure make withdrawal practically impossible, authority has degraded.

Authority degradation is often the decisive integrity failure.

9.4.5 Precedent expansion

Precedent expansion occurs when an original decision is cited, repeated, normalized, or transferred beyond the evidence and authority boundary under which it was made.

This is one of the most subtle integrity threats.

A bounded action may become a general rule.

A local determination may become domain permission.

A proceed-compatible posture may become public approval.

A deferral may become assumed future approval.

A refusal may become a broader rejection than intended.

A pilot may become precedent for deployment.

Precedent expansion threatens integrity because it changes the meaning and effect of the original commitment. The commitment may remain physically bounded while its institutional consequences expand.

For example, a bounded lunar verification determination may be cited later to justify site hardening, power placement, or ISRU dependency. A seabed pilot may normalize future disturbance. A repository authorization may become a template for future siting even where the evidence and authority conditions differ. A grid queue-entry decision may normalize treatment of queue entry as administrative rather than commitment-bearing.

Integrity requires precedent containment.

A commitment remains within its admissible basis only if its future use remains bounded by its original evidence, authority, and threshold conditions.

9.5 Integrity states

INSTRUCTIONAL DRAFT

Commitment integrity is expressed through four states:

1. **Maintain**
2. **Constrain**
3. **Re-evaluate**
4. **Terminate**

These states classify whether continuation remains defensible.

9.5.1 Maintain

Maintain means the commitment remains within admissible bounds.

The evidence boundary, assumptions, dependency topology, authority structure, refusal credibility, and precedent containment remain consistent with the basis under which the commitment was allowed.

Maintain is appropriate when:

- evidence remains consistent with the original basis;
- assumptions remain valid;
- dependencies remain bounded;
- authority remains coherent;
- refusal or constraint remains credible where required;
- no new trigger requires re-evaluation;
- precedent has not expanded beyond the boundary;
- the commitment remains within the defined scope.

Maintain does not mean the project is guaranteed to succeed. It means the commitment remains governance-defensible within its boundary.

9.5.2 Constrain

Constrain means the commitment may continue only under narrowed scope, added conditions, or explicit guardrails.

Constrain is appropriate when integrity is under pressure but can still be preserved by limiting exposure.

Examples of constraints include:

INSTRUCTIONAL DRAFT

- limiting physical disturbance;
- freezing site expansion;
- prohibiting public commitment language;
- capping capital exposure;
- restricting partner reliance;
- preventing infrastructure hardening;
- limiting the action to bounded verification;
- preserving alternatives;
- requiring additional evidence before escalation;
- maintaining retrievability or revocation capacity;
- preventing precedent transfer.

Constrain is not a soft approval. It is a boundary-preserving action.

It applies when continuation remains possible only if the commitment is prevented from expanding beyond its admissible basis.

9.5.3 Re-evaluate

Re-evaluate means new evidence, changed assumptions, dependency growth, authority degradation, or precedent expansion requires reassessment before further hardening.

Re-evaluate is appropriate when the original basis can no longer be relied upon without a new review, but termination is not yet structurally required.

Re-evaluation may be triggered by:

- new data;
- failed assumption;
- evidence contamination;
- changed site conditions;
- changed authority;
- new capital exposure;
- new counterparty reliance;
- public narrative shift;
- altered regulatory posture;
- unexpected dependency;
- changed plausible state space;
- new precedent use;
- threshold drift.

Re-evaluate is the correct state when the original commitment may still be defensible, but not on the old record.

INSTRUCTIONAL DRAFT

It protects against stale admissibility.

9.5.4 Terminate

Terminate means the commitment no longer satisfies its admissible basis and should not continue within the defined boundary.

Terminate is appropriate when:

- refusal authority has collapsed;
- exit has become non-credible;
- dependencies have exceeded admissible bounds;
- evidence no longer supports the commitment;
- assumptions have failed;
- the commitment now rests on over-authorized evidence;
- continuation creates unbounded exposure;
- precedent expansion cannot be contained;
- the commitment has become structurally indefensible.

Terminate does not necessarily mean the broader domain, project type, or strategic objective is invalid. It means this commitment, under this boundary, no longer remains defensible.

Termination is a valid governance outcome.

9.6 Integrity and the five regimes of irreversibility

Commitment integrity should be evaluated by irreversibility regime.

A commitment may remain intact in one regime and fail in another.

9.6.1 Subsurface integrity

Subsurface integrity asks whether the physical evidence and disturbance boundary remain within the admitted basis.

Questions include:

- Has disturbance altered the system beyond the admitted boundary?
- Has new evidence changed the plausible state space?
- Has measurement contamination degraded future interpretability?
- Has intervention created more exposure than originally bounded?

INSTRUCTIONAL DRAFT

9.6.2 Placement integrity

Placement integrity asks whether site, route, corridor, or geometry remains within the admitted basis.

Questions include:

- Has one location become harder to abandon?
- Have alternatives degraded?
- Has access repetition created spatial lock-in?
- Has public language made the site appear inevitable?

9.6.3 Topological integrity

Topological integrity asks whether dependencies remain bounded.

Questions include:

- Have new support systems formed?
- Has the commitment created network dependencies?
- Has capital, logistics, processing, power, or communications become coupled?
- Has one assumption begun governing multiple future decisions?

9.6.4 Authority integrity

Authority integrity asks whether decision authority remains coherent.

Questions include:

- Can the Decision Authority still refuse?
- Can the commitment still be revoked or constrained?
- Has consequence-bearing diverged from decision control?
- Has public, political, or counterparty pressure made refusal non-credible?

9.6.5 Precedent integrity

Precedent integrity asks whether the decision remains bounded in future use.

Questions include:

- Is the decision being cited beyond its boundary?
- Has a bounded action become a general rule?
- Has public interpretation expanded the commitment?
- Is re-screening required before reuse?

9.6.6 Coupled integrity

INSTRUCTIONAL DRAFT

Coupled integrity asks whether multiple regimes have begun reinforcing one another beyond the admitted basis.

A commitment may require re-evaluation or termination if regimes that were bounded at entry become mutually reinforcing.

For example:

site preference + access repetition + support placement + public language + capital alignment

may create a coupled integrity failure even if no single element appears decisive alone.

9.7 Integrity and time

Commitment integrity is time-sensitive.

Some commitments degrade quickly. Others degrade slowly. Some remain stable for long periods before an assumption expires or an authority condition changes.

The framework therefore treats time as part of the admissible basis.

A commitment may include:

- expiration dates;
- evidence update requirements;
- monitoring intervals;
- assumption review points;
- trigger thresholds;
- authority reconfirmation requirements;
- re-screening conditions.

Time matters differently across domains.

In grid interconnection, integrity may degrade quickly as queue studies, deposits, site control, and capital sequencing evolve.

In lunar infrastructure, integrity may degrade as access repetition, support systems, and architecture assumptions accumulate.

In critical minerals, integrity may degrade as land, offtake, processing, water, and capital commitments accumulate around a prospect.

INSTRUCTIONAL DRAFT

In geological disposal, integrity must be considered across long horizons where institutional memory, monitoring, and public legitimacy may change.

Commitment integrity therefore requires attention not only to state changes, but to time-to-degradation.

9.8 Integrity and evidence freeze

Formal commitment governance often requires an evidence freeze: a defined record of the evidence available at the threshold.

Evidence freeze protects against retrospective justification.

It allows the decision record to be reconstructed and prevents later outcomes from altering the basis on which entry was evaluated.

But evidence freeze does not mean evidence is ignored after entry. It means later evidence must be treated as a change condition.

If new evidence appears, it does not silently modify the original basis. It triggers one of the integrity states:

- maintain if the evidence remains within the basis;
- constrain if the evidence requires narrower action;
- re-evaluate if the basis must be reopened;
- terminate if the commitment no longer satisfies its basis.

Evidence freeze is therefore compatible with learning. It simply prevents learning from being retroactively absorbed into the original justification without governance.

9.9 Integrity and precedent control

Precedent control is part of commitment integrity because decisions have institutional afterlives.

A commitment can leave its admissible basis not only by changing physically or operationally, but by being reused improperly.

A precedent card or equivalent governance record should ask:

- What could this decision make easier to repeat, cite, normalize, or inherit?
- What exactly is the precedent object?

INSTRUCTIONAL DRAFT

- What evidence boundary applies?
- What does not transfer?
- What public communication is permitted?
- What requires re-screening?
- What misuse is most likely?

Precedent control is especially important when the commitment is first-of-kind, public, strategic, or likely to influence future standards.

In lunar ISRU, the first bounded volatile verification could be misused as resource infrastructure permission.

In seabed systems, a pilot disturbance could normalize a broader intervention regime.

In grid interconnection, a queue-entry finding could normalize treatment of queue entry as non-substantive despite lock-in.

In geological disposal, one repository authorization could shape future institutional and regulatory expectations.

Precedent control preserves the boundary of the original decision.

9.10 Integrity and termination

Termination is often treated as failure. In this framework, termination can be a valid integrity-preserving outcome.

A commitment should terminate when continuation would require reliance on a basis that no longer exists.

Termination may preserve value by preventing deeper lock-in. It may protect institutional authority, reduce future liability, prevent precedent expansion, or preserve credibility.

Termination is especially important where sunk cost creates continuation pressure.

A project that cannot terminate when its basis expires is no longer governed. It is merely continuing.

The framework treats termination as a governance function, not an admission of defeat.

9.11 Integrity and institutional memory

INSTRUCTIONAL DRAFT

Commitment integrity requires institutional memory.

A commitment can remain physically stable but lose governance integrity if the institution forgets why it was admitted, what assumptions bounded it, what evidence supported it, what triggers required review, or what uses were prohibited.

This is especially important for long-duration systems.

Geological disposal makes the problem obvious, but the same logic applies in other domains. Large infrastructure, critical minerals, grid assets, seabed systems, orbital regimes, and lunar development can all outlive the teams that initiated them.

Institutional memory requires:

- evidence records;
- assumption registries;
- decision boundaries;
- trigger registers;
- authority records;
- precedent controls;
- change logs;
- review conditions.

Without memory, commitments drift.

Commitment integrity therefore depends on governance records that remain reconstructible.

9.12 Relationship to admissibility and posture

Commitment integrity extends the admissibility and posture record into time.

Admissibility determines whether a commitment may be considered.

Governance posture determines which stance is defensible at the threshold.

Integrity determines whether the commitment remains within that threshold record as conditions change.

If a commitment was inadmissible at entry, integrity is already compromised.

If a commitment proceeded despite Deferral-Indicated or Refusal-Required posture, integrity is compromised at entry.

INSTRUCTIONAL DRAFT

If a commitment was Proceed-Compatible within a defined boundary, integrity asks whether continuation remains inside that boundary.

If a commitment was constrained, integrity asks whether the constraints are being obeyed.

If a commitment was deferred, integrity asks whether deferral is preserving optionality or quietly accumulating exposure.

This relationship is what makes the framework a lifecycle model rather than a one-time screen.

9.13 Integrity failure modes

The major integrity failure modes are:

1. Basis drift

The commitment continues after the evidence or assumptions that supported entry have changed.

2. Scope creep

The commitment expands beyond the defined action or phase.

3. Dependency creep

New dependencies accumulate without re-review.

4. Authority erosion

Decision rights remain formal but refusal becomes non-credible.

5. Communication ratchet

Public or institutional language makes reversal harder.

6. Precedent escape

The decision is cited beyond its boundary.

7. Trigger breach

A defined reconsideration or termination trigger occurs without review.

8. Stale reliance

INSTRUCTIONAL DRAFT

Actors continue relying on a determination after expiration.

9. Coupled escalation

Multiple small changes reinforce one another until the commitment becomes larger than its admitted basis.

These failure modes are diagnosable. That is the point of the integrity concept.

9.14 Section conclusion

Commitment integrity governs persistence.

It recognizes that a commitment that was admissible at entry may become indefensible as evidence changes, assumptions expire, dependencies accumulate, authority degrades, or precedent expands.

Integrity is evaluated against the admissible basis.

The four integrity states are:

- Maintain;
- Constrain;
- Re-evaluate;
- Terminate.

This concept is necessary because frontier and infrastructure commitments do not end at entry. They harden through time.

The next section applies the entry and persistence framework to the forcing case of lunar south polar ISRU, where indirect evidence, persistent ambiguity, disturbance, infrastructure pressure, and precedent formation converge.

10. Forcing Case: Lunar South Polar ISRU

Lunar south polar in-situ resource utilization is the forcing case for commitment admissibility because it concentrates the full structure of disturbance-coupled exploration.

The target condition is inferred indirectly.
Interpretation remains non-unique.
Verification may require disturbance.
Support systems create path dependence.
Infrastructure pressure is high.

INSTRUCTIONAL DRAFT

Precedent implications are significant.

Refusal becomes harder once site, access, power, and resource assumptions begin to align.

The central question is not whether lunar volatiles exist.

The central question is whether volatile evidence has earned the authority to govern irreversible surface commitments.

A volatile signal can justify exploration, ranking, and bounded verification. It does not automatically justify site hardening, repeated access, fixed power placement, excavation dependency, ISRU infrastructure, or precedent-forming resource use.

This section applies the framework through five linked stages:

1. detectability limits;
2. ambiguity persistence;
3. environmental response;
4. architecture under constraint;
5. commitment thresholds.

Together, these stages show how a signal can become a site, how a site can become an architecture, and how an architecture can become a commitment before the underlying uncertainty has ceased to be decision-dominant.

10.1 Why lunar ISRU is the forcing case

The lunar south pole is not merely an example. It is a forcing environment for the theory.

It combines:

- indirect volatile evidence;
- persistent subsurface ambiguity;
- extreme environmental constraint;
- access scarcity;
- illumination and thermal discontinuity;
- operational difficulty;
- high cost of surface intervention;
- limited opportunity for repeated verification;
- infrastructure pressure;
- and high precedent value.

These conditions make ordinary feasibility reasoning insufficient.

INSTRUCTIONAL DRAFT

A site may look promising because multiple signals converge: hydrogen indications, thermal stability, shadow persistence, terrain context, radar or reflectance signatures, and strategic location. But those signals do not automatically determine the physical state, accessibility, concentration, continuity, or construction-scale behavior of volatiles.

The lunar south pole therefore exposes the central admissibility problem:

When does evidence support further exploration, and when does it begin to be over-authorized into infrastructure commitment?

This distinction matters because lunar development will likely require early decisions about landing sites, access paths, power placement, communications, navigation, excavation, processing, and logistics. If those decisions begin hardening around volatile assumptions before the assumptions become decision-grade, infrastructure may form around a state of the subsurface that remains unresolved.

The result is not simply technical risk. It is premature commitment formation.

10.2 Detectability limits

The first problem is detectability.

Available lunar volatile evidence can identify promising regions. It can constrain hypotheses. It can prioritize exploration. It can distinguish stronger candidates from weaker candidates.

But detectability does not equal commitment authority.

Relevant evidence may include:

- hydrogen indications;
- neutron spectroscopy;
- thermal context;
- permanently shadowed region association;
- topography and illumination;
- radar response;
- reflectance data;
- terrain and access constraints;
- prior mission observations;
- analog reasoning.

These data are valuable. They are not neutral. They shape institutional attention and site preference.

But indirect sensing has limits. It may not resolve:

INSTRUCTIONAL DRAFT

- volatile form;
- volatile concentration;
- volatile continuity;
- depth distribution;
- physical accessibility;
- extractability;
- mechanical behavior;
- stratigraphy;
- thermal stability under disturbance;
- construction-scale conditions;
- local heterogeneity;
- operational reliability.

The scale problem is especially important.

A regional signal may support regional interest. It may not support a landing zone, excavation plan, power geometry, traverse corridor, or ISRU dependency. Evidence collected at one spatial or physical scale cannot automatically govern commitments at another scale.

The admissibility implication is direct:

Indirect sensing may justify prospecting and bounded verification. It does not by itself authorize site hardening, excavation dependency, or ISRU infrastructure commitment.

This does not diminish the evidence. It defines its authority.

10.3 Ambiguity persistence

The second problem is ambiguity persistence.

The same volatile signal can remain consistent with multiple operating realities.

A hydrogen indication or volatile-associated signature may correspond to:

1. an architecture-grade deposit;
2. a patchy operational resource;
3. a diffuse or inaccessible signal;
4. mechanically difficult material;
5. a signal that is scientifically meaningful but operationally weak;
6. a local condition that does not generalize to the infrastructure scale.

These possibilities are not merely different scientific interpretations. They imply different commitments.

INSTRUCTIONAL DRAFT

An architecture-grade deposit might support site preference, infrastructure planning, and eventual ISRU dependency if other conditions are satisfied.

A patchy operational resource might support additional reconnaissance but not fixed infrastructure.

A diffuse or inaccessible signal might support scientific interest while failing as an infrastructure anchor.

Mechanically difficult material might require different excavation, power, mobility, and processing assumptions.

A locally meaningful signal might not support broader corridor or architecture decisions.

The ambiguity becomes decision-dominant when these plausible states imply different site, access, excavation, power, logistics, or ISRU decisions.

The central issue is therefore not whether the volatile signal is real. It is whether the signal is still compatible with operating realities that would change the commitment decision.

A project can tolerate ambiguity if all plausible states support the same bounded next step. It cannot treat ambiguity as residual risk if plausible states still imply different commitment thresholds.

In the lunar case, this means:

A signal can support exploration before it supports infrastructure.

That sentence carries much of the framework.

10.4 Environmental response

The third problem is that verification is not neutral.

On the Moon, the actions required to reduce uncertainty may alter the system being evaluated.

Verification may require:

- landing activity;
- repeated traverse;
- drilling;
- coring;
- trenching;
- excavation;

INSTRUCTIONAL DRAFT

- heating;
- regolith handling;
- volatile capture experiments;
- emplacement of instruments;
- power support;
- communications support;
- mobility infrastructure;
- operational staging.

These actions can generate valuable knowledge. But they can also create disturbance.

Disturbance may produce:

- regolith alteration;
- volatile mobilization;
- thermal perturbation;
- measurement contamination;
- compaction or mechanical change;
- plume or exhaust effects;
- local surface modification;
- traffic effects;
- baseline alteration;
- concentration of future activity at the disturbed site.

The governance problem is that the act of learning can become part of the commitment pathway.

If intrusive verification requires support systems, access repetition, or site preparation, then verification may begin structuring future infrastructure. If disturbance changes the evidence baseline, then later measurements may be harder to interpret. If the verification site becomes the focus of future work, then the test has begun to privilege one location.

This creates the core lunar version of disturbance-coupled exploration:

Verification can produce knowledge while altering the physical and institutional conditions under which future knowledge and commitment will be judged.

The admissibility question is therefore not simply whether verification is useful. It is whether the verification pathway preserves refusal authority.

Bounded verification may be admissible when:

- the disturbance is limited;
- the evidence gain is material;
- alternatives remain credible;
- support systems remain temporary;
- site preference does not become infrastructure commitment;

INSTRUCTIONAL DRAFT

- public communication remains conditional;
- termination conditions are explicit;
- the result cannot be over-transferred into ISRU permission.

Verification becomes commitment-bearing when it:

- anchors a site;
- requires repeated access;
- invites support placement;
- alters the evidence baseline;
- creates sunk infrastructure;
- establishes a corridor;
- converts a signal into resource dependency;
- or sets precedent for future extraction.

10.5 Architecture under constraint

The fourth problem is architecture.

Lunar exploration does not occur in isolation. Even limited surface investigation requires architecture: landing, mobility, power, communications, navigation, thermal management, data relay, logistics, and operational sequencing.

These support systems are often treated as enabling infrastructure. But in disturbance-coupled systems, enabling infrastructure can become commitment-forming.

Support systems do not merely allow exploration. They begin to structure future access, dependency, and refusal cost.

Relevant architectural elements include:

- landing site access;
- rover mobility routes;
- communication geometry;
- navigation support;
- power placement;
- thermal management;
- excavation support;
- dust and plume mitigation;
- regolith handling;
- logistics staging;
- repeated traverse corridors;
- equipment pre-positioning;
- data and operations architecture.

INSTRUCTIONAL DRAFT

Each element can privilege one site or operating geometry.

Power placement can favor a specific surface region.
Communication geometry can favor certain access paths.
Repeated rover traverses can become corridors.
Logistics staging can make one location the default.
Excavation support can imply future processing.
Processing assumptions can imply ISRU dependency.

Architecture therefore creates topology.

The issue is not only whether the architecture is coherent. It is whether architecture is hardening faster than understanding.

A lunar support architecture may be technically rational and still premature if it depends on volatile assumptions that remain ambiguous. Once support systems begin forming around one interpretation, the institution may become increasingly committed to that interpretation even if the evidence has not earned that authority.

The governance question is:

Are support systems preserving exploration, or are they beginning to convert exploration into commitment?

This question distinguishes architecture readiness from commitment admissibility.

A system may be architecturally ready to support exploration while still not admissible as infrastructure commitment.

10.6 Surface commitment classes

Lunar commitment can form through multiple surface actions.

Eight surface commitment classes are especially important:

1. **Site hardening**
A candidate site becomes the operational or institutional default.
2. **Landing preparation**
Landing zones, descent/ascent assumptions, or surface approach paths begin shaping future access.
3. **Berm or blast protection emplacement**
Surface modification begins to protect or enable repeated operations.
4. **Excavation or trenching**
Intrusive disturbance changes the regolith and evidence baseline.

INSTRUCTIONAL DRAFT

5. **Corridor establishment**
Repeated traverse or access paths begin structuring mobility and logistics.
6. **Regolith processing dependency**
The architecture begins relying on local material handling or volatile processing.
7. **Fixed infrastructure placement**
Power, communications, navigation, logistics, or processing assets begin anchoring a site.
8. **Subsurface access**
Drilling, coring, trenching, or other intrusive operations create disturbance and site-specific commitment.

These classes are not all equally irreversible. But each can become commitment-bearing depending on scale, timing, evidence basis, support systems, public posture, and dependency formation.

The important point is that lunar commitment may begin before permanent habitation, industrial extraction, or full-scale infrastructure.

It may begin when access repeats.

It may begin when power is placed.

It may begin when the site is publicly named.

It may begin when excavation is planned.

It may begin when ISRU enters the architecture as an assumed dependency.

The framework therefore treats surface actions not by their label, but by their exposure effect.

10.7 Signal-to-lock-in chain

The lunar commitment pathway can be expressed as a signal-to-lock-in chain:

Volatile signal → site preference → access repetition → power/support placement → disturbance → corridor formation → excavation planning → ISRU dependency → infrastructure lock-in

This chain is important because it shows that commitment accumulates.

The volatile signal does not itself create infrastructure. But it creates site preference. Site preference encourages access. Access encourages support. Support enables disturbance. Disturbance concentrates evidence and future work. Repetition creates corridors. Corridors support excavation planning. Excavation planning enables resource dependency. Resource dependency hardens infrastructure.

The chain converts evidence into topology.

INSTRUCTIONAL DRAFT

At each stage, the institution may still describe the work as exploratory, preliminary, or reversible. But the practical ability to refuse may be declining.

The admissibility problem is therefore not located at one obvious moment. It is distributed across the chain.

A governance framework must identify the first point at which the chain becomes commitment-bearing.

In some cases, that point may be primary site designation.

In others, repeated access.

In others, intrusive verification.

In others, fixed power or communications placement.

In others, public use of ISRU-dependent language.

In others, capital or partner alignment around a resource-dependent architecture.

The chain must be governed before it becomes self-justifying.

10.8 Applying the five regimes of irreversibility

Lunar ISRU activates all five regimes of irreversibility.

10.8.1 Subsurface irreversibility

Subsurface irreversibility appears when volatile verification requires disturbance.

Drilling, trenching, excavation, heating, or regolith handling can reduce uncertainty while changing the target condition. The evidence-generating action can also alter the evidence baseline.

The governing question is:

Does learning require changing the volatile-bearing system being evaluated?

10.8.2 Placement irreversibility

Placement irreversibility appears when one site or access geometry becomes the default.

A volatile signal may begin privileging a crater rim, PSR-adjacent zone, access route, or operating region before resource accessibility is decision-grade.

The governing question is:

INSTRUCTIONAL DRAFT

Does this action make one location easier to continue than abandon?

10.8.3 Topological irreversibility

Topological irreversibility appears when access, power, communications, navigation, logistics, excavation, and processing begin reinforcing one another.

The governing question is:

Does the action create dependencies that future lunar decisions must inherit?

10.8.4 Authority irreversibility

Authority irreversibility appears when mission, agency, partner, or public commitments make re-siting, deferral, or refusal less credible.

The governing question is:

Does this action weaken the practical ability to say no later?

10.8.5 Precedent irreversibility

Precedent irreversibility appears because early lunar resource actions may normalize future resource use, infrastructure sequencing, site claims, or operational standards.

The governing question is:

Does this action create a future rule that exceeds the evidence boundary of the current decision?

10.8.6 Coupled lunar irreversibility

The lunar case is most dangerous when these regimes couple.

A volatile signal may begin as evidence. It becomes placement. Placement becomes topology. Topology becomes authority pressure. Authority pressure becomes precedent.

This is why lunar ISRU is the forcing case.

The problem is not one irreversible act. It is the assembly of irreversible commitment from several individually plausible steps.

10.9 Minimum evidence before lunar commitment

INSTRUCTIONAL DRAFT

A lunar commitment is not admissible merely because evidence is strong enough to justify exploration.

Evidence must be adequate relative to the irreversible burden of the proposed action.

Before site hardening, fixed infrastructure, excavation dependency, corridor formation, or ISRU architecture can be considered, evidence must satisfy several conditions.

10.9.1 Multi-modal convergence

Evidence should converge across relevant sensing and observational modes. No single signal should govern infrastructure commitment alone.

10.9.2 Scale match

The evidence must be relevant at the scale of the proposed commitment. Regional evidence cannot automatically support local infrastructure placement.

10.9.3 Ambiguity reduction

Decision-dominant ambiguity must be reduced. Remaining plausible states should not imply materially different site, access, excavation, power, or ISRU decisions.

10.9.4 Disturbance justification

If verification requires disturbance, the expected knowledge gain must justify the disturbance and preserve future interpretability.

10.9.5 Refusal preservation

The action must preserve the ability to stop, re-site, re-sequence, or abandon the pathway.

10.9.6 Dependency containment

Support systems must not create resource dependency before the resource interpretation becomes decision-grade.

10.9.7 Precedent containment

The action must not normalize broader lunar resource commitment beyond the evidence boundary.

These conditions do not demand certainty. They demand that remaining uncertainty no longer govern the commitment.

INSTRUCTIONAL DRAFT

10.10 Applied admissibility logic

The general lunar south polar ISRU posture can be stated as:

Proceed with bounded, non-hardening investigation.

Defer infrastructure commitment.

Refuse premature dependency formation.

This is not an anti-ISRU position.

It is a sequencing claim.

Exploration may proceed without authorizing infrastructure commitment. Bounded verification may be legitimate while site hardening remains premature. Evidence gathering may be valuable while ISRU dependency remains inadmissible.

The admissibility logic is:

- volatile signals can justify prospecting;
- prospecting can justify bounded verification;
- bounded verification can justify further evidence development;
- but none of these automatically authorize infrastructure hardening;
- infrastructure commitment becomes admissible only when remaining uncertainty no longer has the power to change site, access, support, disturbance, resource dependency, or precedent-forming decisions.

This separates learning from commitment.

10.11 Governance posture in the lunar case

After admissibility is evaluated, governance posture must still be classified.

For many lunar ISRU actions, the posture will differ by action class.

Bounded remote analysis

Likely posture:

Proceed-Compatible

Rationale:

INSTRUCTIONAL DRAFT

- no material disturbance;
- alternatives remain open;
- refusal authority preserved;
- evidence development improves decision basis.

Bounded reconnaissance

Likely posture:

Proceed-Compatible or **Deferral-Indicated**, depending on access and support requirements.

Rationale:

- reconnaissance may be justified if it does not create site hardening or dependency;
- posture becomes deferral-indicated if support architecture begins privileging one site.

Intrusive verification

Likely posture:

Deferral-Indicated until disturbance boundaries, evidence gains, and termination conditions are explicit.

Rationale:

- verification can produce knowledge;
- disturbance may alter the system;
- support systems may create dependency;
- action must be bounded.

Fixed power or support placement

Likely posture:

Deferral-Indicated or **Refusal-Required**, depending on evidence and dependency formation.

Rationale:

- support placement can create topological and placement lock-in;
- volatile evidence may not yet justify infrastructure anchoring.

ISRU-dependent architecture

Likely posture:

INSTRUCTIONAL DRAFT

Refusal-Required if volatile accessibility, continuity, concentration, and disturbance response remain decision-dominant.

Rationale:

- resource dependency exceeds the evidence authority;
- infrastructure hardens around unresolved subsurface states;
- precedent risk is high.

This shows why posture must be action-specific. “Lunar ISRU” is not one decision. It is a sequence of commitment thresholds.

10.12 Commitment integrity in the lunar case

Lunar commitment integrity becomes relevant once any bounded action begins.

A lunar action that was initially admissible may lose integrity if it begins exceeding its admissible basis.

Integrity risks include:

Evidence drift

Verification reveals that volatile conditions differ from the original interpretation, but architecture continues as if the original interpretation remains valid.

Assumption expiry

An assumption valid for reconnaissance is carried into excavation planning or infrastructure placement.

Dependency accumulation

Access repetition, power placement, communications, and logistics create site commitment beyond the original scope.

Authority degradation

Mission, partner, or public commitments make re-siting or refusal non-credible.

Precedent expansion

A bounded verification action is cited as justification for broader ISRU infrastructure.

INSTRUCTIONAL DRAFT

The lunar integrity outputs would be:

Maintain if the action remains within its evidence, disturbance, dependency, and authority boundary.

Constrain if the action may continue only under narrower disturbance limits, communication controls, or dependency caps.

Re-evaluate if new volatile evidence, disturbance effects, support hardening, or authority changes alter the admissible basis.

Terminate if the action no longer satisfies its admissible basis or has created unbounded dependency.

This is critical because even a well-governed lunar entry can degrade if exploration architecture begins hardening silently.

10.13 Precedent in lunar ISRU

Lunar ISRU is highly precedent-sensitive.

Early lunar resource actions may shape future norms around:

- volatile evidence;
- site designation;
- resource use;
- infrastructure sequencing;
- surface disturbance;
- excavation;
- ISRU dependency;
- public-private lunar development;
- mission architecture;
- property-like expectations;
- environmental governance;
- autonomy permissions.

A first action can become a reference point.

If indirect volatile evidence is used to justify site hardening, future actors may treat similar evidence as sufficient. If a bounded excavation is publicly framed as ISRU progress, future actors may treat disturbance as normal. If a site becomes publicly known as a resource location, re-siting may become harder. If a proceed-compatible posture is misread as approval, governance boundaries may erode.

INSTRUCTIONAL DRAFT

Precedent control should therefore be part of any lunar determination.

A lunar precedent record should specify:

- what evidence was used;
- what action was allowed;
- what action was not allowed;
- what does not transfer;
- what requires re-screening;
- what public language is prohibited;
- what future use would exceed the decision boundary.

The key rule is:

A bounded lunar evidence action must not become a general lunar infrastructure rule.

10.14 Lunar implications

The framework produces several implications for lunar exploration and development.

First, resource evidence should be classified by authority level. Evidence that supports prospecting should not be used as infrastructure permission.

Second, surface actions should be classified by commitment effect, not by label. A “test,” “demo,” or “reconnaissance” action can be commitment-bearing if it creates site, corridor, support, or precedent exposure.

Third, architecture should not harden faster than understanding. Power, communications, access, and logistics decisions can become commitment mechanisms.

Fourth, disturbance should be treated as both epistemic and structural. It produces knowledge and alters the future decision environment.

Fifth, refusal and deferral should remain legitimate outputs. Not every promising signal should become an infrastructure anchor.

Sixth, lunar governance should be sequence-sensitive. Exploration, verification, site hardening, excavation, resource dependency, and infrastructure commitment are distinct thresholds.

Seventh, precedent should be contained. Early lunar actions may define norms beyond their technical scope.

These implications do not slow lunar development as an end in itself. They protect lunar development from premature commitment formation.

INSTRUCTIONAL DRAFT

10.15 Section conclusion

Lunar south polar ISRU demonstrates the full structure of commitment admissibility.

The evidence is meaningful but indirect.
The ambiguity is persistent and operationally material.
Verification may disturb the target.
Support systems can harden into architecture.
Architecture can produce dependency.
Dependency can create infrastructure lock-in.
Early actions can set precedent.

The governing question is not whether there is ice.

The governing question is whether the evidence has earned the authority to govern irreversible surface commitment.

For current generalized lunar south polar ISRU reasoning, the defensible posture is likely:

bounded exploration may proceed; infrastructure commitment should defer; premature ISRU dependency should be refused.

The next section turns to geological disposal and nuclear waste, where the central problem is not only entry, but the persistence of commitment integrity across deep time.

11. Limiting Case: Geological Disposal and Nuclear Waste

Geological disposal and nuclear waste systems are the limiting case for commitment integrity.

Lunar ISRU concentrates the entry problem: when can exploration evidence justify commitment? Geological disposal concentrates the persistence problem: how does a commitment remain defensible across time horizons longer than ordinary institutional memory?

In repository systems, the commitment is not simply to build, operate, or close a facility. The commitment is to sustain responsibility across deep uncertainty, long-duration containment assumptions, public legitimacy, monitoring limits, institutional continuity, and intergenerational consequence.

The central question is therefore not only:

Is repository entry admissible?

INSTRUCTIONAL DRAFT

It is also:

Can the commitment remain structurally defensible after entry, closure, institutional turnover, and long-duration uncertainty?

This makes geological disposal the canonical integrity case.

11.1 Why geological disposal is the limiting case

Geological disposal compresses several forms of irreversibility into one system.

It involves:

- subsurface uncertainty;
- site and placement commitment;
- long-lived infrastructure;
- transport and access pathways;
- containment assumptions;
- monitoring and retrievability limits;
- public legitimacy;
- regulatory authority;
- institutional stewardship;
- intergenerational responsibility;
- precedent for future disposal regimes.

Unlike many infrastructure commitments, the time horizon of responsibility extends far beyond ordinary project, investment, political, or organizational cycles.

A grid project may face commitment integrity issues over years or decades. A lunar infrastructure program may face them over mission architectures and development cycles. A repository commitment faces them across institutional succession, public trust cycles, knowledge preservation, environmental change, and deep-time containment claims.

This creates a distinct integrity challenge:

The commitment must remain defensible after the original decision-makers, evidence context, institutional memory, and public mandate have changed.

That is why geological disposal is not merely another transfer case. It is the limiting case for persistence.

INSTRUCTIONAL DRAFT

11.2 Entry admissibility in geological disposal

Before integrity can be tested, repository entry must still satisfy admissibility.

The entry problem asks whether the proposed repository commitment may be considered under the evidence and authority conditions present at the threshold.

Relevant entry thresholds include:

- candidate site designation;
- site characterization authorization;
- land or jurisdictional commitment;
- tunnel or access construction;
- licensing submission;
- waste transport pathway selection;
- repository construction;
- emplacement;
- closure;
- post-closure stewardship transition.

Each threshold has a different admissibility burden.

A site may be admissible for further characterization but not for repository construction.

A tunnel may be admissible for investigation but not for emplacement.

A monitoring architecture may be admissible for pre-closure oversight but not sufficient for post-closure confidence.

A regulatory finding may support a defined phase without authorizing later irreversibility.

The repository case therefore reinforces a central framework rule:

Admissibility is phase-specific.

A favorable site indication does not authorize the full disposal pathway. Each irreversible threshold must be governed according to the exposure it creates.

11.3 The admissible basis in disposal systems

For geological disposal, the admissible basis must be especially explicit because later actors may inherit the commitment long after the original basis is forgotten.

INSTRUCTIONAL DRAFT

The admissible basis should include:

- the defined disposal action;
- the repository phase under review;
- the geological evidence boundary;
- containment assumptions;
- hydrologic and geochemical assumptions;
- seismic and structural assumptions;
- engineered barrier assumptions;
- monitoring assumptions;
- retrievability or reversibility assumptions;
- institutional stewardship assumptions;
- public legitimacy conditions;
- regulatory authority conditions;
- transport and access assumptions;
- closure and post-closure transition conditions;
- reconsideration triggers;
- precedent controls.

A repository commitment loses integrity when continuation depends on an assumption that has expired, changed, or exceeded the phase for which it was valid.

For example, an assumption valid during characterization may not be valid after emplacement. An assumption valid before closure may not be valid after closure. A monitoring condition valid during active oversight may not be valid under reduced institutional presence.

This is why repository integrity cannot be treated as a one-time entry decision.

11.4 Subsurface integrity

Geological disposal is fundamentally subsurface, but its subsurface integrity problem differs from ordinary exploration.

In exploration systems, disturbance may reveal a resource. In disposal systems, intervention aims to isolate hazard. The target is not extraction but long-duration containment.

Subsurface integrity asks whether the geological, hydrological, geochemical, thermal, mechanical, and structural assumptions that supported the commitment remain within their admissible bounds.

Key questions include:

- Do the host rock and surrounding formations remain consistent with the containment basis?

INSTRUCTIONAL DRAFT

- Are groundwater pathways, fracture networks, or permeability conditions within the assumed envelope?
- Are geochemical interactions consistent with engineered and natural barrier assumptions?
- Are thermal effects within the modeled range?
- Are seismic, tectonic, or stress conditions within the admissible basis?
- Does characterization remain sufficient for the phase under consideration?
- Does new evidence alter the plausible state space?
- Does uncertainty remain bounded relative to the irreversible threshold being crossed?

Subsurface integrity may be maintained even with residual uncertainty if that uncertainty is bounded within the repository's admissible basis.

It fails or requires re-evaluation when new plausible states emerge that could materially change containment, monitoring, retrievability, closure, or stewardship decisions.

11.5 Placement integrity

Repository siting is a placement commitment of unusually high consequence.

Placement integrity asks whether the selected site remains defensible as the project advances through characterization, construction, emplacement, closure, and post-closure stewardship.

A site may lose placement integrity if:

- new geological evidence undermines the basis for selection;
- transport routes create unexpected exposure;
- local or regional conditions change;
- land-use, political, or community legitimacy conditions degrade;
- access infrastructure creates unanticipated lock-in;
- alternatives were prematurely foreclosed;
- public or institutional narratives make re-siting non-credible before evidence has stabilized.

Placement integrity is not only physical. It includes the social and institutional fact that a repository site becomes difficult to abandon once a host community, region, regulator, and national program begin organizing around it.

A repository site can therefore become placement-irreversible before closure, before emplacement, and sometimes before construction.

That means re-siting credibility must be preserved deliberately during earlier phases.

INSTRUCTIONAL DRAFT

11.6 Topological integrity

A repository is not a point asset. It is a system.

It depends on:

- waste inventory assumptions;
- transport routes;
- access tunnels;
- emplacement layout;
- engineered barriers;
- monitoring systems;
- security systems;
- institutional oversight;
- local infrastructure;
- regulatory processes;
- emergency and retrieval pathways;
- knowledge management systems.

Topological integrity asks whether these dependencies remain bounded and coherent relative to the admissible basis.

A repository may fail topological integrity if the system of dependencies becomes larger or less reversible than originally admitted.

For example:

- transport commitments may harden around a route before public legitimacy or risk governance is stable;
- access infrastructure may make abandonment politically or financially non-credible;
- waste inventory assumptions may change after design commitment;
- monitoring systems may become inadequate for the evolved risk profile;
- retrievability assumptions may be compromised by closure sequencing;
- institutional dependencies may become fragile over time.

The key topological question is:

What begins to depend on the repository commitment, and what does the repository commitment begin to depend on?

If those dependencies expand beyond the original admissible basis, integrity requires constraint or re-evaluation.

INSTRUCTIONAL DRAFT

11.7 Authority integrity

Authority integrity is the central challenge in geological disposal.

The commitment is long-lived, but the institutions that authorize it are not static. Political leadership changes. Agencies reorganize. Public trust shifts. Regulatory standards evolve. Communities change. Technical memory decays. Documents can be lost, reinterpreted, or ignored. Future actors may not share the assumptions of the original decision.

Authority integrity asks whether decision rights, responsibility, refusal capacity, oversight, and institutional memory remain coherent across time.

Key questions include:

- Who has authority to continue, constrain, pause, retrieve, re-evaluate, or terminate?
- Does that authority remain legitimate over time?
- Who bears consequence if assumptions fail?
- Can future institutions reconstruct the original decision basis?
- Are reconsideration triggers visible and enforceable?
- Is public legitimacy part of the admissible basis?
- Does closure remove practical revocation capacity?
- Can authority survive political, financial, or institutional turnover?
- Is responsibility transferred faster than accountability can survive?

Authority integrity can degrade even if the technical containment case remains plausible.

A repository can be geologically strong and institutionally weak.

This is why geological disposal demonstrates the persistence problem so sharply. It requires not only technical containment, but durable governance coherence.

11.8 Precedent integrity

Repository decisions are highly precedent-forming.

A repository authorization can shape:

- future siting processes;
- acceptable evidence standards;
- public consent norms;
- retrievability expectations;
- closure practices;
- monitoring assumptions;

INSTRUCTIONAL DRAFT

- institutional responsibility models;
- intergenerational governance standards;
- treatment of long-duration uncertainty.

Precedent integrity asks whether the repository decision is being used beyond its original evidence and authority boundary.

A specific repository decision should not automatically authorize other sites, other host rocks, other inventories, other communities, other closure conditions, or other regulatory regimes.

The precedent risk is not only that the first decision will be copied. It is that the first decision will normalize the conversion of long-duration uncertainty into institutional confidence too early.

A repository precedent record should therefore specify:

- what decision was made;
- what evidence supported it;
- what assumptions bounded it;
- what does not transfer;
- what future decisions require re-screening;
- what public interpretation is prohibited;
- what would trigger re-evaluation;
- what remains unresolved.

This is the same function as the Precedent Card in the broader framework: it governs the afterlife of a decision.

11.9 Integrity degradation pathways in geological disposal

The general integrity degradation pathways appear strongly in repository systems.

11.9.1 Evidence drift

Evidence drift occurs when new geological, hydrological, geochemical, seismic, thermal, or monitoring evidence changes the plausible state space.

This may not invalidate the repository. But it may require re-evaluation if the new evidence affects containment, retrievability, monitoring, closure, or stewardship.

11.9.2 Assumption expiry

Assumption expiry is pervasive.

INSTRUCTIONAL DRAFT

Assumptions about host rock behavior, engineered barriers, water movement, corrosion, institutional continuity, monitoring duration, public acceptance, and retrievability may be valid only under defined phases.

When a project moves from characterization to construction, or from emplacement to closure, assumptions must not transfer automatically.

11.9.3 Dependency accumulation

Dependency accumulation occurs as repository infrastructure, transport systems, local economic reliance, regulatory posture, waste inventory allocation, and national policy begin depending on continuation.

The more systems depend on the repository, the harder re-evaluation or termination becomes.

11.9.4 Authority degradation

Authority degradation occurs when future institutions lose the practical ability to reconsider, retrieve, constrain, communicate, or accept responsibility.

This may happen through institutional turnover, loss of public legitimacy, fragmented responsibility, or closure decisions that eliminate practical revocation.

11.9.5 Precedent expansion

Precedent expansion occurs when the repository decision becomes generalized into a standard for future disposal programs without preserving its evidence, community, authority, and geological boundary.

This is especially dangerous because repository decisions are rare, high-stakes, and likely to be cited.

11.10 Integrity states in geological disposal

The four integrity states apply as follows.

Maintain

A repository commitment may be maintained when the evidence, assumptions, dependencies, authority, and precedent controls remain within the admissible basis.

Maintain does not mean “safe forever.” It means the commitment remains defensible under its current boundary.

INSTRUCTIONAL DRAFT

Constrain

A repository commitment should be constrained when integrity remains possible only under narrower conditions.

Examples include:

- limiting emplacement scope;
- extending monitoring;
- preserving retrievability;
- restricting closure progression;
- limiting transport expansion;
- requiring additional characterization;
- maintaining public legitimacy conditions;
- preventing precedent transfer;
- capping phase progression.

Constrain preserves integrity by limiting the commitment before it exceeds its basis.

Re-evaluate

A repository commitment requires re-evaluation when new evidence, changed assumptions, authority shifts, dependency growth, or legitimacy changes mean the original basis can no longer be relied upon without reopening the record.

Re-evaluate is especially important at phase transitions:

- characterization to construction;
- construction to emplacement;
- emplacement to closure;
- closure to post-closure stewardship;
- monitoring revision;
- authority transfer.

Terminate

A repository commitment should terminate when it no longer satisfies its admissible basis and cannot be constrained or re-evaluated into defensibility.

Termination may apply to a site, phase, pathway, closure plan, transport route, or authority structure.

Termination does not necessarily reject geological disposal as a category. It rejects the defined commitment under the defined boundary.

INSTRUCTIONAL DRAFT

11.11 Why geological disposal extends the theory

Geological disposal shows that commitment governance cannot stop at entry.

A repository may pass an entry screen and still lose integrity later. It may satisfy technical evidence requirements at one phase and fail authority or precedent requirements at another. It may retain physical containment assumptions while losing public legitimacy or institutional continuity. It may remain locally defensible while becoming dangerous precedent.

This case extends the theory in four ways.

First, it shows that admissibility is not permanent.

Second, it shows that integrity must be evaluated across time and phase.

Third, it shows that authority and institutional memory are not secondary to technical evidence. They are part of the commitment's admissible basis.

Fourth, it shows that precedent must be governed because rare long-duration decisions become reference points for future institutional behavior.

In this sense, geological disposal is not merely an application of the framework. It is the case that forces the framework to include persistence.

11.12 Comparison with lunar ISRU

Lunar ISRU and geological disposal sit at opposite ends of the entry-persistence spectrum.

Lunar ISRU emphasizes entry:

- Is volatile evidence sufficient for site, access, excavation, support, or ISRU commitment?
- Can exploration proceed without becoming infrastructure commitment?
- Can disturbance generate knowledge without creating lock-in?
- Can precedent be contained in a frontier governance environment?

Geological disposal emphasizes persistence:

- Can the commitment remain defensible over time?
- Can assumptions remain bounded across phase transitions?
- Can authority survive institutional change?
- Can public legitimacy and memory persist?
- Can closure occur without eliminating meaningful governance capacity?

INSTRUCTIONAL DRAFT

Both cases share the same structural problem:

Commitment must not exceed the evidence, authority, dependency, and precedent basis that supports it.

The difference is where the pressure is greatest.

Lunar ISRU asks whether commitment should begin.

Geological disposal asks whether commitment can remain valid after it begins.

Together, they establish the two central axes of the theory: admissibility and integrity.

11.13 Section conclusion

Geological disposal and nuclear waste are the limiting case for commitment integrity.

They show that a commitment is not governed simply because entry was admissible. The commitment must remain within its admissible basis as evidence changes, assumptions expire, dependencies accumulate, authority evolves, and precedent expands.

In repository systems, the central governance challenge is long-duration defensibility.

The commitment must survive not only geological uncertainty, but institutional time.

This case demonstrates why the framework requires commitment integrity as a co-equal contribution with commitment admissibility.

The next section turns to critical minerals, where the problem is neither lunar frontier entry nor deep-time persistence alone, but the conversion of prospectivity into extraction, infrastructure, capital, and strategic-resource commitment.

12. Transfer Case: Critical Minerals

Critical minerals demonstrate how prospectivity becomes commitment.

This case differs from lunar ISRU and geological disposal in an important way. Lunar ISRU concentrates the problem of frontier entry. Geological disposal concentrates the problem of long-duration integrity. Critical minerals sit in the middle: evidence begins as geological prospectivity, but quickly becomes entangled with land control, access, permitting, infrastructure, processing, offtake, strategic urgency, capital exposure, and public-interest narratives.

INSTRUCTIONAL DRAFT

The central question is:

Has prospectivity earned the right to govern irreversible extraction pathways?

A mineral indication can be real, strategically important, and worthy of exploration without yet being adequate to anchor extraction commitment.

Critical minerals matter because urgency can overwhelm admissibility discipline. Strategic importance can make projects appear commitment-worthy before the evidence has become decision-grade. Supply-chain anxiety, national security framing, industrial policy, energy transition demand, and capital scarcity can all create pressure to convert a geological signal into a project pathway before the uncertainty has stopped governing the decision.

The governance problem is not whether critical minerals are important. They are.

The governance problem is whether importance is being used to over-authorize evidence.

12.1 Why critical minerals are a transfer case

Critical minerals are a strong transfer case because the same structural sequence appears repeatedly:

prospectivity → land control → access → drilling → resource interpretation → permitting posture → processing assumptions → infrastructure dependency → offtake / capital sequencing → extraction commitment

At the beginning of the sequence, evidence may support exploration. By the middle, the project may begin organizing infrastructure, stakeholders, and capital. By the end, refusal may become difficult even if major uncertainty remains.

The transition from prospect to commitment can occur through many channels:

- lease or land acquisition;
- drilling programs;
- access roads;
- water assumptions;
- power access;
- processing assumptions;
- transport routes;
- waste and tailings planning;
- community engagement;
- permitting pathways;
- offtake discussions;
- strategic-resource designation;

INSTRUCTIONAL DRAFT

- public funding;
- equity or project finance;
- industrial policy narratives.

Each channel can be individually defensible. Together, they can convert uncertainty into lock-in.

This is why critical minerals test the generality of the framework. The system is terrestrial, commercial, strategic, and capital-intensive. It does not have the extreme physical constraints of the Moon or the deep-time burden of nuclear waste. Yet it shares the same commitment structure: evidence is indirect or incomplete, interpretation remains non-unique, verification requires disturbance, and project pathways can harden before uncertainty collapses.

12.2 Prospectivity is not commitment authority

Prospectivity identifies potential. It does not automatically authorize commitment.

A mineral system may show:

- favorable geology;
- geophysical anomaly;
- geochemical signature;
- structural control;
- mineralized intercepts;
- regional trend alignment;
- historical workings;
- analog support;
- strategic commodity relevance;
- infrastructure proximity.

These signals can justify additional exploration. They can justify ranking, staking, sampling, mapping, drilling, or bounded evidence development. They do not by themselves justify irreversible extraction pathways.

The distinction is important because prospectivity is often narratively powerful. A project can become “the lithium asset,” “the rare earths project,” “the domestic supply solution,” or “the strategic minerals corridor” before the deposit has earned the right to govern that identity.

Once that language takes hold, evidence can become subordinate to narrative.

The admissibility question is therefore:

Is the project still an exploration hypothesis, or has the prospect begun functioning as an infrastructure and capital anchor?

INSTRUCTIONAL DRAFT

A prospect becomes commitment-bearing when its assumed future value begins shaping land, access, processing, offtake, power, water, permitting, or capital decisions that would be difficult to reverse if the deposit interpretation changes.

12.3 Ambiguity persistence in critical minerals

Critical mineral systems can exhibit strong ambiguity persistence.

The same exploration evidence may remain consistent with materially different operating realities:

1. **coherent economic deposit;**
2. **geologically interesting but discontinuous mineralization;**
3. **grade-positive but processing-constrained material;**
4. **resource present but infrastructure-burdened project;**
5. **strategically important commodity with weak local project economics;**
6. **environmentally or socially constrained deposit;**
7. **deposit requiring water, power, tailings, or access commitments that exceed its evidence basis.**

These are not merely different technical scenarios. They imply different commitment decisions.

A coherent economic deposit may justify advancing toward feasibility, permitting, infrastructure planning, and capital sequencing.

Discontinuous mineralization may justify more drilling but not project-scale infrastructure.

Processing-constrained material may require metallurgical evidence before offtake or processing assumptions can govern commitment.

A resource with severe water, power, access, or tailings constraints may fail admissibility even if grade is attractive.

A strategically important commodity may still be inadmissible as a project if the deposit cannot support the irreversible pathway being built around it.

Ambiguity becomes decision-dominant when these plausible states imply different land, drilling, access, permitting, processing, water, power, capital, or offtake decisions.

The project may be uncertain in many ways. The key issue is whether remaining ambiguity can still change whether the extraction pathway should be allowed to form.

INSTRUCTIONAL DRAFT

12.4 Disturbance and evidence generation

Critical minerals require disturbance to reduce uncertainty.

This may include:

- trenching;
- drilling;
- bulk sampling;
- access road construction;
- pad preparation;
- geotechnical testing;
- metallurgical sampling;
- water investigation;
- tailings investigation;
- pilot processing;
- site preparation;
- environmental baseline work.

Some of these actions are routine in exploration. But routine does not mean non-commitment-bearing.

A drilling campaign can create physical disturbance, capital exposure, public visibility, community response, permitting posture, and narrative hardening. Access roads can make one site geometry easier to continue than abandon. Metallurgical samples can create processing assumptions. Bulk samples can create downstream partner interest. Environmental baseline work can begin framing a project pathway.

The framework does not argue that these actions should be avoided. It argues that their governance status depends on whether they remain bounded learning or begin functioning as commitment.

A drilling program is more likely admissible when:

- the evidence gain is material;
- disturbance is bounded;
- alternatives remain open;
- public language remains conditional;
- no extraction pathway is implied;
- capital exposure is limited;
- termination conditions are explicit;
- results can still cause refusal or re-sequencing.

A drilling program becomes commitment-bearing when:

- land, access, processing, offtake, and capital begin aligning around an assumed deposit;

INSTRUCTIONAL DRAFT

- public or investor language presents the project as inevitable;
- alternatives become non-credible;
- additional disturbance is justified by prior disturbance rather than new evidence;
- strategic urgency substitutes for evidence adequacy.

The question is not whether exploration should continue. The question is whether exploration is still exploration.

12.5 Land, access, and placement irreversibility

Critical minerals often become committed through land and access before extraction begins.

Placement irreversibility appears when a project's physical geography begins hardening around a deposit interpretation.

Examples include:

- land acquisition;
- lease consolidation;
- access road alignment;
- drilling pad networks;
- processing site assumptions;
- tailings location assumptions;
- water access;
- power connection;
- transport routes;
- port or rail assumptions;
- community or jurisdictional engagement.

A deposit is not only a geological object. It becomes a spatial system.

Once that spatial system begins to form, re-siting or re-sequencing becomes harder. Even if the geological interpretation changes, the project may continue because land, access, permitting, and capital have already concentrated around one pathway.

The admissibility question is:

Has the evidence earned the right to fix this geography?

A project may be admissible for exploration but not for placement commitment. It may be admissible for drilling but not for access-road construction. It may be admissible for land optioning but not land consolidation. It may be admissible for metallurgical testing but not processing-site planning.

INSTRUCTIONAL DRAFT

Placement thresholds must be separated.

12.6 Topological irreversibility in critical minerals

Critical mineral projects are topological systems.

They depend not only on deposit quality, but on relationships among:

- mine site;
- processing route;
- water source;
- power source;
- roads;
- rail;
- ports;
- tailings;
- permitting;
- community acceptance;
- offtake;
- financing;
- downstream supply chain;
- strategic policy support.

A deposit can be geologically promising but topologically weak.

Topological irreversibility appears when these dependencies begin forming before the deposit and project constraints are decision-grade.

Examples:

- offtake expectations before resource confidence;
- processing assumptions before metallurgy is bounded;
- water-dependent plans before water availability is admissible;
- power assumptions before grid or generation dependency is resolved;
- tailings planning before waste behavior is sufficiently understood;
- access roads before project scale is justified;
- public funding before project viability is bounded;
- strategic supply-chain commitments before extraction pathway is defensible.

Topological failure occurs when the project becomes easier to continue than abandon because multiple systems begin depending on the assumption that the deposit will justify them.

This is a key critical minerals pattern:

INSTRUCTIONAL DRAFT

The deposit may still be uncertain while the project system starts behaving as if the deposit is proven.

12.7 Authority irreversibility and strategic urgency

Critical minerals are highly exposed to authority irreversibility because urgency is often politically and economically real.

Governments need supply security.
Manufacturers need feedstock.
Investors want exposure to strategic commodities.
Communities may expect development.
Agencies may want domestic capacity.
Companies may need to signal progress.
Boards may fear missing a strategic window.

These pressures can degrade refusal authority.

A project may begin as an exploration opportunity but become politically framed as necessary. Once that happens, refusal becomes harder. Deferral may be described as weakness. Reconsideration may be interpreted as anti-development. Evidence discipline may be portrayed as obstruction.

This is dangerous because strategic importance can be mistaken for admissibility.

A mineral may be nationally important. The project may still be inadmissible at its current threshold.

Authority irreversibility appears when the institution can no longer distinguish:

We need this commodity

from:

This project has earned irreversible commitment.

Those are different claims.

Commitment governance preserves that distinction.

It allows urgency to be acknowledged without allowing urgency to override evidence.

INSTRUCTIONAL DRAFT

12.8 Precedent irreversibility in critical minerals

Critical minerals can create precedent in several ways.

A project may normalize:

- strategic urgency as sufficient justification for weak evidence;
- accelerated permitting before deposit confidence is adequate;
- infrastructure commitment before processing uncertainty is resolved;
- public funding before refusal conditions are defined;
- community exposure before project viability is bounded;
- environmental disturbance as acceptable because the commodity is important;
- offtake alignment before resource and processing assumptions are decision-grade.

Precedent risk is high because critical minerals sit at the intersection of industrial policy, climate transition, national security, and capital markets.

An early commitment can become a model.

If a government backs one project before evidence is adequate, future projects may cite that action. If investors treat strategic designation as a substitute for admissibility, capital may flow into weakly bounded projects. If permitting accelerates before uncertainty is resolved, institutions may normalize exposure as the cost of urgency.

The precedent question is:

What future commitments does this decision make easier to justify?

A critical minerals determination should therefore specify what does and does not transfer.

A decision to support bounded exploration should not become extraction permission.

A decision to fund characterization should not become infrastructure endorsement.

A decision to recognize strategic importance should not become project approval.

A decision to defer should not become assumed future approval.

A refusal of one project should not become rejection of the commodity class.

Precedent containment is essential.

12.9 Critical minerals and admissibility

The admissibility condition for critical minerals can be stated as:

INSTRUCTIONAL DRAFT

A critical minerals commitment is admissible only when remaining geological, processing, infrastructure, environmental, social, and authority uncertainties no longer have the power to change whether the proposed extraction pathway should form.

This does not mean all uncertainty must be resolved.

It means the uncertainty that remains must not still govern the commitment.

Different thresholds require different evidence.

Exploration-stage admissibility

Prospectivity may be sufficient for mapping, sampling, geophysics, or bounded drilling.

Land-control admissibility

Evidence must justify why the specific land position should become harder to abandon.

Access admissibility

Evidence must justify why roads, pads, or route geometry should form.

Processing admissibility

Metallurgical and mineralogical evidence must justify process assumptions.

Infrastructure admissibility

Resource, access, water, power, tailings, and social-license conditions must be sufficiently bounded.

Capital admissibility

The project must not use strategic urgency to overrun unresolved deposit, processing, environmental, or infrastructure uncertainty.

Extraction-pathway admissibility

The full project pathway must remain defensible across plausible states.

This staged structure prevents a prospect from being treated as a project too early.

12.10 Governance posture in critical minerals

INSTRUCTIONAL DRAFT

Critical minerals often require different posture classifications for different actions.

Bounded exploration

Often **Proceed-Compatible** when disturbance is limited, alternatives remain open, and public language remains conditional.

Expanded drilling

May be **Proceed-Compatible** or **Deferral-Indicated** depending on disturbance, community exposure, capital implications, and whether results can still cause refusal.

Land consolidation

Often **Deferral-Indicated** until deposit continuity, access, processing, and environmental conditions justify placement commitment.

Processing assumptions

Often **Deferral-Indicated** until metallurgical uncertainty is reduced.

Offtake or strategic supply commitments

May be **Deferral-Indicated** or **Refusal-Required** if they create capital and authority lock-in before the resource and processing pathway are decision-grade.

Infrastructure commitment

May be **Refusal-Required** if the project system is hardening around prospectivity rather than evidence sufficient for extraction.

This action-specific posture logic is important because the phrase “critical minerals project” hides multiple commitment thresholds.

12.11 Commitment integrity in critical minerals

Critical minerals also require integrity review after entry.

An admissible exploration or development commitment may lose integrity as the project evolves.

Integrity degradation can occur through:

Evidence drift

INSTRUCTIONAL DRAFT

Drilling reveals weaker continuity, lower grade, more complex mineralogy, or less favorable geotechnical conditions than expected.

Assumption expiry

An assumption valid for exploration is carried into development, processing, or financing.

Dependency accumulation

Roads, water, power, processing, tailings, offtake, and capital commitments begin depending on the project before the evidence basis supports them.

Authority degradation

Strategic narratives, public funding, community expectations, or political pressure make refusal non-credible.

Precedent expansion

A bounded exploration decision is cited as development approval, or strategic designation is treated as permission for accelerated commitment.

Integrity outputs apply:

- **Maintain** if the project remains within its admitted evidence and dependency boundary.
- **Constrain** if continuation requires limits on drilling, capital, land, infrastructure, public language, or offtake.
- **Re-evaluate** if new evidence or changed assumptions undermine the prior basis.
- **Terminate** if the project no longer satisfies the conditions under which it was allowed to continue.

Critical minerals are especially prone to stale reliance, where early prospectivity continues to justify later commitments after the evidence boundary has changed.

12.12 The critical minerals failure mode

The characteristic failure mode is:

Strategic urgency converts prospectivity into infrastructure commitment before the deposit, processing pathway, and dependency structure are decision-grade.

The project may not fail because people ignored uncertainty. It may fail because uncertainty was acknowledged but treated as manageable after commitment.

INSTRUCTIONAL DRAFT

The sequence is often:

1. commodity urgency is real;
2. a prospect appears promising;
3. land and access begin forming;
4. capital and partners align;
5. processing and infrastructure assumptions appear;
6. public and policy narratives reinforce urgency;
7. refusal becomes harder;
8. unresolved deposit or system uncertainty remains decision-dominant.

By the time the uncertainty becomes visible, the project may be too politically, financially, or institutionally committed to stop cleanly.

This is exactly the pattern the framework is designed to identify.

12.13 What the critical minerals case adds to the theory

Critical minerals add four important lessons.

First, they show that **importance does not equal admissibility**. Strategic value can increase pressure to act, but it does not reduce the evidence burden of irreversible commitment.

Second, they show that **prospectivity can become authority**. Geological signals can begin governing land, infrastructure, capital, and policy before becoming commitment-grade.

Third, they show that **urgency can degrade authority**. Institutions may lose refusal credibility because refusal conflicts with strategic narrative.

Fourth, they show that **commitment forms across a project system**, not only at the mine face. Water, power, processing, roads, tailings, offtake, financing, and public policy can all become part of the commitment topology.

This makes critical minerals a strong transfer case from the lunar and repository extremes into terrestrial strategic infrastructure.

12.14 Section conclusion

Critical minerals demonstrate the conversion of prospectivity into commitment.

INSTRUCTIONAL DRAFT

A promising signal can justify exploration without justifying extraction pathways. Strategic importance can justify attention without authorizing infrastructure, capital, processing, offtake, or permitting lock-in.

The admissibility question is whether prospectivity has earned the right to govern irreversible extraction commitments.

The integrity question is whether the project remains within its admissible basis as evidence, dependencies, capital, authority, and precedent evolve.

Critical minerals show why the framework is not limited to lunar or deep-time systems. It applies wherever evidence remains incomplete, verification requires disturbance, and institutional urgency can harden commitments before uncertainty loses decision-changing power.

The next section generalizes this structure across additional domains through a cross-domain transfer matrix.

13. Cross-Domain Transfer Matrix

The framework transfers across domains only where the same decision structure is present.

It is not enough for a system to be uncertain, complex, expensive, or strategically important. The framework applies most strongly where evidence remains incomplete, verification requires action, action creates exposure, and that exposure can harden before uncertainty loses decision-changing power.

The transfer rule is therefore structural:

A domain falls within the framework when five conditions are present:

1. inference is indirect;
2. interpretation remains non-unique;
3. verification requires disturbance, deployment, access, or intervention;
4. intervention creates irreversible exposure;
5. commitment creates path dependence before uncertainty collapses.

This section applies that rule across six additional domains:

- geothermal systems;
- carbon storage;
- seabed and marine infrastructure;
- grid interconnection;
- orbital systems;
- autonomy near irreversible thresholds.

INSTRUCTIONAL DRAFT

These cases differ in technical content. They share a commitment structure.

13.1 Transfer logic

The framework does not transfer by analogy.

It transfers by decision structure.

Lunar volatile exploration, carbon storage, grid interconnection, and autonomous systems are not substantively the same. They do not share the same evidence, physics, institutions, or failure modes. But they can share the same governance problem:

a decision must be made before the system is fully known, and the action required to improve knowledge may itself create hard-to-reverse exposure.

That is the common structure.

The transfer matrix therefore asks the same questions in each domain:

- What is the signal?
- What ambiguity persists?
- What action reveals the uncertainty?
- What exposure does that action create?
- Which irreversibility regimes are activated?
- What is the admissibility question?
- What is the integrity question?

This prevents the framework from becoming rhetorical. It requires every domain to pass through the same structural test.

13.2 Geothermal Systems

13.2.1 Signal and evidence

Geothermal systems are inferred through combinations of geological mapping, geophysics, geochemistry, temperature gradients, structural interpretation, analog fields, heat-flow data, surface manifestations, and exploratory drilling.

INSTRUCTIONAL DRAFT

These data can identify promising systems, but they often do not fully resolve reservoir structure, permeability, recharge, stress regime, sustainable productivity, induced seismicity behavior, or long-term thermal decline before drilling.

The evidence may support prospect ranking. It may not yet support drilling, stimulation, offtake, or infrastructure commitment.

13.2.2 Ambiguity persistence

The same geothermal signal can remain consistent with multiple operating realities:

- productive and sustainable reservoir;
- high-temperature but low-permeability system;
- fault-controlled resource with drilling uncertainty;
- short-lived or compartmentalized reservoir;
- geochemically or mechanically challenging system;
- system requiring stimulation beyond acceptable exposure;
- project whose surface infrastructure would exceed subsurface confidence.

These plausible states imply different commitments.

A productive reservoir may support drilling expansion, offtake, and infrastructure planning. A low-permeability system may require deferral or alternative technology. A fault-controlled system may require additional structural evidence. A mechanically or seismically sensitive system may constrain or refuse stimulation.

Ambiguity becomes decision-dominant when unresolved reservoir states still change whether drilling, stimulation, offtake, or capital commitment is admissible.

13.2.3 Disturbance / exposure action

The key exposure action is drilling.

Drilling generates knowledge. It also creates:

- capital exposure;
- site commitment;
- permitting posture;
- pad and road infrastructure;
- public and stakeholder attention;
- offtake expectations;
- pressure to drill follow-up wells;
- potential physical changes to the system.

INSTRUCTIONAL DRAFT

Stimulation, testing, and production trials can create additional exposure through pressure changes, induced seismicity risk, thermal drawdown, and infrastructure dependency.

13.2.4 Irreversibility regimes

Geothermal systems activate:

- **IR-SUB:** drilling and stimulation disturb the subsurface;
- **IR-PLC:** pads, wells, access roads, and field layout fix location;
- **IR-TOP:** offtake, gathering systems, transmission, and operations create dependency;
- **IR-AUT:** capital, permitting, and stakeholder expectations degrade refusal authority;
- **IR-PRC:** early successful or failed projects can shape future geothermal standards and public tolerance;
- **IR-CPL:** drilling, capital, offtake, and site infrastructure can reinforce one another.

13.2.5 Admissibility question

Has the reservoir interpretation earned the right to govern drilling, stimulation, offtake, or infrastructure commitment?

A geothermal commitment is inadmissible if reservoir ambiguity remains decision-dominant and the action required to resolve it would create exposure that current evidence cannot justify.

13.2.6 Integrity question

Does the geothermal commitment remain within its admissible basis as reservoir behavior resolves?

Integrity may degrade if drilling results diverge from the reservoir model, if stimulation changes exposure, if productivity declines, if induced seismicity assumptions fail, or if offtake and infrastructure dependencies continue despite evidence drift.

Possible integrity outputs:

- **Maintain:** reservoir behavior remains within admitted bounds.
- **Constrain:** limit drilling, stimulation, production, or infrastructure expansion.
- **Re-evaluate:** new reservoir evidence changes the plausible state space.
- **Terminate:** continuation relies on a reservoir basis that no longer holds.

13.3 Carbon Storage

INSTRUCTIONAL DRAFT

13.3.1 Signal and evidence

Carbon storage relies on structural geology, reservoir properties, caprock integrity, fault interpretation, pressure modeling, geochemistry, geomechanics, seismic data, well data, monitoring design, plume modeling, and regulatory assumptions.

These data may support storage potential. They may not fully resolve containment behavior, pressure evolution, fault response, plume migration, induced seismicity, leakage pathways, or monitoring sufficiency before injection.

The system cannot be fully known before the act that initiates the storage condition.

13.3.2 Ambiguity persistence

The same storage evidence can remain consistent with multiple operating realities:

- robust containment and predictable plume behavior;
- pressure-limited storage;
- fault-sensitive storage;
- monitoring-constrained storage;
- containment dependent on assumptions that cannot be directly verified;
- storage requiring infrastructure or liability exposure beyond the evidence basis.

These states imply different injection, monitoring, siting, liability, and permitting decisions.

Ambiguity becomes decision-dominant when unresolved containment or pressure states could change whether injection should proceed, defer, be constrained, or be refused.

13.3.3 Disturbance / exposure action

The key exposure action is injection.

Injection produces knowledge about pressure, plume movement, containment behavior, and monitoring performance. But injection also initiates the very condition being evaluated.

It creates:

- pressure perturbation;
- plume migration;
- monitoring obligation;
- liability exposure;
- regulatory commitment;
- public legitimacy exposure;
- infrastructure dependence;
- permanence claims;

INSTRUCTIONAL DRAFT

- long-duration stewardship.

This makes carbon storage a strong disturbance-coupled system.

13.3.4 Irreversibility regimes

Carbon storage activates:

- **IR-SUB:** injection changes subsurface pressure and fluid state;
- **IR-PLC:** storage site and injection geometry become fixed;
- **IR-TOP:** wells, pipelines, monitoring systems, and emitters create dependency;
- **IR-AUT:** permits, liability, and public commitments shape refusal capacity;
- **IR-PRC:** early storage projects shape standards for future storage regimes;
- **IR-CPL:** injection, monitoring, liability, and infrastructure reinforce one another.

13.3.5 Admissibility question

Is injection admissible before pressure, containment, plume, monitoring, and liability uncertainties are bounded relative to the irreversible burden of storage?

The question is not whether storage is valuable. It is whether this injection commitment is admissible under current evidence and authority conditions.

13.3.6 Integrity question

Does the storage commitment remain within its admissible basis as plume behavior, pressure response, monitoring evidence, liability conditions, and institutional authority evolve?

Integrity may degrade if plume behavior diverges from model expectations, if monitoring becomes inadequate, if pressure response exceeds assumptions, if fault behavior changes the plausible state space, or if liability and authority conditions become incoherent.

Possible integrity outputs:

- **Maintain:** storage behavior remains inside admitted bounds.
- **Constrain:** limit injection rate, volume, pressure, or operational expansion.
- **Re-evaluate:** new monitoring or pressure evidence changes the basis.
- **Terminate:** injection or continuation no longer satisfies the admissible basis.

13.4 Seabed and Marine Infrastructure

INSTRUCTIONAL DRAFT

13.4.1 Signal and evidence

Marine and seabed systems rely on bathymetry, sonar, sub-bottom profiling, sediment sampling, geotechnical surveys, ecological baseline studies, current and wave models, cable route studies, hazard mapping, and permitting evidence.

These data may identify feasible routes or sites. They may not fully resolve seabed response, ecological disturbance, installation behavior, sediment mobility, geohazards, or long-term infrastructure interaction before intervention.

13.4.2 Ambiguity persistence

The same marine evidence can remain consistent with multiple operating realities:

- stable route or installation corridor;
- geotechnically difficult seabed;
- ecologically sensitive area;
- mobile sediment regime;
- installation method that creates greater disturbance than assumed;
- route whose future maintenance burden is underestimated;
- site whose public or regulatory legitimacy is fragile.

These states imply different route, installation, permitting, mitigation, monitoring, and abandonment decisions.

Ambiguity becomes decision-dominant when unresolved seabed or ecological conditions could change whether route fixation, installation, or disturbance is admissible.

13.4.3 Disturbance / exposure action

Exposure actions include:

- seabed coring;
- trenching;
- cable laying;
- anchoring;
- foundation installation;
- dredging;
- route clearance;
- habitat disturbance;
- pilot deployment;
- public permitting posture.

These actions may generate knowledge while changing the physical and ecological baseline.

INSTRUCTIONAL DRAFT

13.4.4 Irreversibility regimes

Marine systems activate:

- **IR-SUB:** seabed disturbance and subsurface sampling;
- **IR-PLC:** route, corridor, landfall, or installation geometry fix placement;
- **IR-TOP:** cables, platforms, ports, grids, and maintenance routes create dependency;
- **IR-AUT:** permitting, public legitimacy, and stakeholder commitments affect refusal;
- **IR-PRC:** first disturbance or pilot projects can normalize future marine intervention;
- **IR-CPL:** route, installation, permitting, and ecological disturbance can reinforce lock-in.

13.4.5 Admissibility question

Should route, landfall, installation, or disturbance commitment harden before seabed, ecological, geotechnical, and permitting uncertainties are bounded?

13.4.6 Integrity question

Does the marine commitment remain within its admissible basis as seabed behavior, environmental response, maintenance needs, public legitimacy, or permitting conditions change?

Integrity may degrade if installation produces unexpected disturbance, if seabed mobility changes route stability, if ecological response exceeds assumptions, or if permitting and public legitimacy shift.

Possible integrity outputs:

- **Maintain:** route and disturbance remain within admitted bounds.
- **Constrain:** limit installation method, footprint, timing, or expansion.
- **Re-evaluate:** new seabed or ecological evidence changes the basis.
- **Terminate:** route or installation no longer remains defensible.

13.5 Grid Interconnection

13.5.1 Signal and evidence

Grid interconnection is not a subsurface disturbance case, but it is a strong topological irreversibility case.

INSTRUCTIONAL DRAFT

Evidence may include site control opportunity, proximity to transmission, point-of-interconnection assumptions, interconnection queue data, congestion patterns, upgrade cost estimates, deliverability analysis, market pricing, utility or ISO/RTO signals, and neighboring project behavior.

These data may suggest an attractive project. They may not fully resolve constraint topology, upgrade exposure, restudy risk, deliverability, curtailment, cost allocation, withdrawal credibility, or neighboring queue dynamics before entry.

13.5.2 Ambiguity persistence

The grid equivalent of subsurface ambiguity is **constraint-topology ambiguity**.

The same site and point-of-interconnection opportunity can remain consistent with multiple network realities:

- bounded upgrade path;
- congestion-dominated pathway;
- restudy-exposed pathway;
- cost-allocation trap;
- deliverability constraint;
- neighboring-project dependency;
- site-grid coupling failure.

These states imply different queue, site, gen-tie, capital, and withdrawal decisions.

Ambiguity becomes decision-dominant when unresolved network states could change whether queue entry, site control, or capital escalation is admissible.

13.5.3 Exposure action

The key exposure actions include:

- queue entry;
- study deposits;
- point-of-interconnection selection;
- site control tied to interconnection assumptions;
- gen-tie routing;
- equipment procurement;
- offtake or financing alignment;
- public project positioning;
- capital escalation.

Queue entry can look administrative. In practice, it can become commitment-bearing.

INSTRUCTIONAL DRAFT

It may create deposits, timelines, study exposure, site-grid coupling, stakeholder alignment, and withdrawal costs. It may also make the project appear more advanced, which can create capital and internal momentum.

13.5.4 Irreversibility regimes

Grid interconnection activates:

- **IR-PLC:** site, POI, and gen-tie geometry fix placement;
- **IR-TOP:** network constraints, queue dynamics, upgrades, and shared capacity create dependency;
- **IR-AUT:** capital, board, developer, and counterparty expectations degrade refusal;
- **IR-PRC:** treating queue entry as administrative can normalize silent commitment;
- **IR-CPL:** site control, queue entry, gen-tie routing, and capital sequencing reinforce one another.

Grid may not activate IR-SUB in the same way as drilling, but it strongly activates topological, placement, authority, and precedent regimes.

13.5.5 Admissibility question

Should queue entry, POI fixation, site control, or capital escalation occur before constraint topology, upgrade exposure, deliverability, and withdrawal credibility are sufficiently bounded?

13.5.6 Integrity question

Does the interconnection commitment remain within its admissible basis as study results, upgrade exposure, congestion, neighboring queue behavior, capital commitments, and withdrawal conditions evolve?

Integrity may degrade if upgrade costs exceed assumptions, restudies change deliverability, gen-tie routing locks in before network reality is clear, or capital continues despite a changed constraint topology.

Possible integrity outputs:

- **Maintain:** interconnection exposure remains bounded.
- **Constrain:** cap deposits, pause capital, preserve alternatives, or limit site-grid coupling.
- **Re-evaluate:** new study results or queue changes alter the basis.
- **Terminate:** queue/site commitment no longer satisfies its admissible basis.

INSTRUCTIONAL DRAFT

13.6 Orbital Systems

13.6.1 Signal and evidence

Orbital systems rely on mission design, orbital mechanics, traffic models, conjunction risk estimates, spectrum coordination, debris models, launch cadence, spacecraft performance, regulatory approvals, and operational assumptions.

These data may support deployment. They may not fully resolve traffic behavior, coordination burden, debris interaction, autonomy requirements, collision avoidance capacity, or reciprocal effects of dense deployment.

13.6.2 Ambiguity persistence

The orbital equivalent is **traffic-regime and dependency-topology ambiguity**.

The same deployment plan may be consistent with:

- manageable coordination burden;
- congestion-sensitive operational regime;
- debris-risk amplification;
- spectrum or coordination constraint;
- autonomy-dependent operations;
- reciprocal vulnerability to other actors;
- precedent that normalizes a crowded orbital regime.

These states imply different deployment, licensing, autonomy, coordination, and scaling decisions.

13.6.3 Exposure action

Exposure actions include:

- launch;
- slot or shell occupation;
- orbital deployment;
- constellation scaling;
- autonomous collision avoidance delegation;
- coordination commitments;
- spectrum reliance;
- public or regulatory commitments.

INSTRUCTIONAL DRAFT

Deployment creates knowledge about operations, but it also enters the shared environment. Once deployed, spacecraft contribute to traffic, collision avoidance burden, debris risk, and coordination expectations.

13.6.4 Irreversibility regimes

Orbital systems activate:

- **IR-PLC:** orbital shell, slot, or trajectory geometry becomes placement;
- **IR-TOP:** traffic, coordination, spectrum, debris, and constellation operations create dependency;
- **IR-AUT:** regulators, operators, and mission authorities may lose practical refusal capacity after launch cadence begins;
- **IR-PRC:** early deployment patterns can normalize orbital congestion and autonomy permissions;
- **IR-CPL:** launch cadence, traffic burden, coordination dependency, and precedent reinforce each other.

13.6.5 Admissibility question

Should deployment, scaling, or autonomy delegation proceed before traffic-regime, coordination, debris, and revocation uncertainties are bounded?

13.6.6 Integrity question

Does the orbital commitment remain within its admissible basis as traffic conditions, debris risk, coordination burden, autonomy behavior, and regulatory context evolve?

Possible integrity outputs:

- **Maintain:** operations remain within admitted traffic and coordination bounds.
- **Constrain:** limit scaling, alter operations, restrict autonomy, or increase oversight.
- **Re-evaluate:** traffic, debris, or coordination evidence changes the basis.
- **Terminate:** continuation or scaling no longer remains defensible.

13.7 Autonomy Near Irreversible Thresholds

13.7.1 Signal and evidence

Autonomous systems rely on simulation, testing, validation, operational history, environmental models, sensor performance, failure-mode analysis, and safety cases.

INSTRUCTIONAL DRAFT

These data may support capability. They may not fully resolve behavior under rare edge cases, degraded conditions, adversarial contexts, communication loss, or morally and physically consequential thresholds.

13.7.2 Ambiguity persistence

The autonomy equivalent is **permission-boundary ambiguity**.

The same performance evidence may be consistent with:

- safe bounded autonomy;
- autonomy requiring human confirmation near thresholds;
- behavior that is reliable in nominal cases but unstable in rare cases;
- autonomy that performs well but lacks revocation clarity;
- autonomy that creates authority delegation beyond what the institution can govern.

These states imply different permission, escalation, oversight, and refusal decisions.

13.7.3 Exposure action

The exposure action is not merely operation. It is permission delegation.

Autonomy may act faster than human oversight can intervene. If the delegated action can disturb, damage, commit, place, extract, inject, collide, or otherwise cross an irreversible threshold, then permission itself becomes commitment-bearing.

13.7.4 Irreversibility regimes

Autonomy near irreversible thresholds activates:

- **IR-SUB:** if autonomous drilling, excavation, injection, or sampling changes a physical system;
- **IR-PLC:** if autonomous placement or routing fixes location;
- **IR-TOP:** if autonomy creates dependency in networks or operations;
- **IR-AUT:** delegation changes decision rights and revocation capacity;
- **IR-PRC:** early autonomous action can normalize future permission regimes;
- **IR-CPL:** autonomous action can couple speed, dependency, authority loss, and precedent.

13.7.5 Admissibility question

Should autonomy be permitted to act near irreversible thresholds before uncertainty, escalation, revocation, and human refusal conditions are bounded?

INSTRUCTIONAL DRAFT

13.7.6 Integrity question

Does the autonomy permission remain within its admissible basis as operational behavior, environmental context, edge cases, authority boundaries, and precedent evolve?

Possible integrity outputs:

- **Maintain:** autonomy remains within bounded permission.
- **Constrain:** narrow action classes, require human confirmation, reduce autonomy envelope.
- **Re-evaluate:** new behavior or context changes the permission basis.
- **Terminate:** autonomy permission no longer remains defensible.

13.8 Cross-domain transfer matrix

Domain	Signal / evidence	Dominant ambiguity	Exposure action	Active regimes	Admissibility question	Integrity question
Lunar ISRU	Hydrogen, thermal, terrain, radar / reflectance	Volatile form, accessibility, continuity	Intrusive verification, support placement, excavation	SUB, PLC, TOP, AUT, PRC, CPL	Has volatile evidence earned infrastructure authority?	Does exploration remain bounded or become site / ISRU lock-in?
Geological disposal	Geological, hydrologic, geochemical, containment evidence	Long-duration containment and institutional stewardship	Siting, access, emplacement, closure	SUB, PLC, TOP, AUT, PRC, CPL	Is each repository phase admissible?	Does the commitment remain defensible across time and authority change?
Critical minerals	Geology, geophysics, geochemistry, drilling	Deposit continuity, processing, infrastructure burden	Drilling, land control, access, offtake, capital	SUB, PLC, TOP, AUT, PRC, CPL	Has prospectivity earned extraction-pathway authority?	Does the project remain within its evidence and dependency basis?

INSTRUCTIONAL DRAFT

Domain	Signal / evidence	Dominant ambiguity	Exposure action	Active regimes	Admissibility question	Integrity question
Geothermal	Geophysics, gradients, geochemistry, structure	Reservoir productivity and behavior	Drilling, stimulation, offtake	SUB, PLC, TOP, AUT, PRC	Has reservoir interpretation earned drilling / infrastructure authority?	Does continuation remain valid as reservoir behavior resolves?
Carbon storage	Storage model, caprock, faults, plume / pressure models	Containment, pressure, plume behavior	Injection	SUB, PLC, TOP, AUT, PRC	Is injection admissible before containment and pressure uncertainty are bounded?	Does storage remain within its basis as plume and pressure evolve?
Seabed / marine	Bathymetry, geotech, ecological baseline	Seabed condition and environmental response	Route fixation, installation, disturbance	SUB, PLC, TOP, AUT, PRC	Should route or installation harden before seabed uncertainty is bounded?	Does the corridor remain defensible as impacts and conditions change?
Grid interconnection	Site, POI, queue, congestion / upgrade signals	Constraint-topology ambiguity	Queue entry, POI fixation, gen-tie, capital	PLC, TOP, AUT, PRC, CPL	Should queue/site commitment form before network constraints are visible?	Does commitment remain defensible as upgrade and congestion resolve?
Orbital systems	Mission design, traffic, debris, coordination evidence	Traffic-regime / dependency-topology ambiguity	Launch, deployment, scaling, autonomy	PLC, TOP, AUT, PRC, CPL	Should deployment proceed before traffic and revocation risks are bounded?	Does deployment remain defensible as traffic and debris conditions evolve?
Autonomy	Testing, simulation, validation,	Permission-boundary ambiguity	Delegated action near irreversible thresholds	SUB, PLC, TOP, AUT,	Should autonomy act before revocation	Does permission remain defensible

INSTRUCTIONAL DRAFT

Domain	Signal / evidence	Dominant ambiguity	Exposure action	Active regimes	Admissibility question	Integrity question
	operational evidence			PRC, CPL	and threshold conditions are bounded?	as behavior and context evolve?

13.9 What the transfer matrix shows

The matrix shows that the framework is not limited to one domain.

It applies wherever three transformations occur.

First, **evidence becomes authority**. A signal, model, test, or estimate begins to justify action beyond its proper evidence level.

Second, **action becomes exposure**. The step taken to learn, position, or advance the project makes future refusal more difficult.

Third, **exposure becomes commitment**. The institution begins inheriting dependencies created by prior action.

The specific technical content differs. Lunar ISRU concerns volatile state. Grid interconnection concerns constraint topology. Geological disposal concerns long-duration containment and stewardship. Autonomy concerns permission boundaries. Carbon storage concerns plume and pressure behavior.

But the governance structure is common:

uncertainty remains, learning requires action, action creates exposure, and exposure hardens into commitment.

13.10 Transfer limits

The framework should not be applied indiscriminately.

It is less useful when:

- evidence is direct and sufficient;
- the action is genuinely reversible;
- learning does not create material exposure;
- alternatives remain intact;

INSTRUCTIONAL DRAFT

- authority remains coherent;
- no dependency forms;
- no precedent is created;
- the decision does not approach an irreversible threshold.

For example, ordinary internal analysis, low-stakes site comparison, reversible desk studies, non-public early modeling, and exploratory conversations without commitment effects may not require commitment governance.

The framework is most relevant when the next step changes the future decision space.

This limit is important because over-application would weaken the theory. Not every decision needs commitment admissibility review. Only decisions approaching irreversible thresholds do.

13.11 Section conclusion

The cross-domain transfer matrix demonstrates that commitment admissibility and commitment integrity are not lunar-only, subsurface-only, or resource-only concepts.

They apply to a class of systems defined by structural coupling:

indirect inference, persistent ambiguity, intervention-dependent learning, exposure creation, and path-dependent commitment.

Lunar ISRU, geological disposal, and critical minerals provide the major case architecture. Geothermal, carbon storage, marine systems, grid interconnection, orbital systems, and autonomy demonstrate transfer.

Across all of them, the core governance questions remain the same:

Should this commitment be allowed to form under current uncertainty?

Does the commitment remain within its admissible basis as conditions change?

The next section turns from transfer to institutional implications: how organizations should preserve evidence boundaries, refusal authority, decision records, and precedent containment.

14. Institutional Implications

Commitment admissibility and commitment integrity are not only analytical concepts. They imply a different institutional posture toward frontier, infrastructure, and exploration decisions.

INSTRUCTIONAL DRAFT

The framework does not require institutions to avoid hard decisions, eliminate uncertainty, or slow all action. It requires institutions to distinguish between evidence that supports learning and evidence that authorizes commitment. It also requires institutions to preserve refusal authority long enough for uncertainty to lose decision-changing power.

The institutional implication is simple:

Organizations must govern the boundary where evidence becomes authorization.

That boundary is often where failure begins.

14.1 Evidence is not authorization

The first institutional implication is that evidence must be classified by what it can legitimately support.

Many organizations treat evidence as cumulative. More evidence increases confidence; increased confidence supports the next step; the next step becomes progress. This logic works when the next step is reversible. It becomes dangerous when the next step is commitment-bearing.

Evidence should be categorized by authority level:

1. **Interest-generating evidence**
Supports attention, inquiry, and early study.
2. **Prospect-ranking evidence**
Supports comparison among candidate sites, routes, resources, or pathways.
3. **Site-comparison evidence**
Supports relative preference, but not necessarily commitment.
4. **Bounded verification evidence**
Supports limited learning actions under constraints.
5. **Design-supporting evidence**
Supports technical assumptions within defined bounds.
6. **Commitment-authorizing evidence**
Supports irreversible or hard-to-reverse commitment within a stated boundary.

The institutional failure occurs when evidence from a lower level is used at a higher level.

A volatile signal becomes site authority.

A geological prospect becomes infrastructure authority.

A queue opportunity becomes capital authority.

A pilot becomes precedent authority.

A model becomes commitment authority.

The framework requires institutions to ask:

INSTRUCTIONAL DRAFT

What is this evidence allowed to authorize?

That question should be asked before the evidence is allowed to shape land, capital, infrastructure, architecture, public posture, or operational dependency.

14.2 Refusal and deferral are valid outcomes

The second implication is that refusal and deferral must be treated as valid governance outcomes.

Many institutions are structurally biased toward continuation. Teams are rewarded for progress. Capital wants deployment. Missions want execution. Agencies want milestones. Developers want permits. Investors want decisions. Public narratives reward momentum.

In that environment, refusal can appear like failure and deferral can appear like indecision.

The framework reverses that assumption.

Refusal can preserve value when commitment would create non-recoverable exposure.
Deferral can preserve authority when uncertainty remains decision-dominant.
Re-evaluation can preserve integrity when the original basis has expired.
Termination can preserve legitimacy when continuation would exceed the admissible boundary.

This is especially important in frontier systems, where the pressure to demonstrate progress can cause evidence to be over-authorized.

A system that cannot refuse is not governing.
A system that cannot defer is not learning.
A system that cannot terminate is not preserving integrity.

Institutions should therefore define refusal conditions before entry, not after failure.

14.3 Decision records must be reconstructible

The third implication is that threshold decisions must be reconstructible from the evidence state at the time of commitment.

This is necessary because outcomes can mislead. A successful project can retroactively make a premature decision appear justified. A failed project can make a defensible decision appear wrong. In irreversible systems, the quality of the threshold decision must be evaluated from the evidence, uncertainty, authority, and reversibility conditions present when the threshold was crossed.

INSTRUCTIONAL DRAFT

A reconstructible decision record should include:

- the decision under review;
- the irreversible threshold;
- the evidence boundary;
- the plausible state space;
- the dominant uncertainty;
- the active irreversibility regimes;
- the Decision Authority;
- the refusal and revocation conditions;
- the dependency topology;
- the permitted and prohibited actions;
- the assumptions registry;
- the reconsideration triggers;
- the expiration conditions;
- the precedent implications.

Without this record, organizations cannot distinguish between a good decision with a bad outcome, a bad decision with a good outcome, and a commitment that lost integrity after entry.

Reconstructibility is therefore not administrative overhead. It is the mechanism by which threshold decision quality remains auditable.

14.4 Evidence freeze and change control

The fourth implication is that formal threshold decisions require evidence discipline.

An evidence freeze is a defined record of the evidence available at the decision threshold. It prevents later information, outcome preference, or institutional pressure from rewriting the basis of the decision.

Evidence freeze does not mean evidence stops mattering. It means later evidence is treated as new evidence, not as retroactive justification.

After evidence freeze:

- new data should not silently modify the original decision basis;
- assumptions should not be changed informally;
- narrative revisions should not alter the threshold record;
- late evidence should trigger reopening, re-evaluation, or a decision change log;
- the original reasoning should remain reconstructible.

This is particularly important when commitments are controversial, expensive, public, or precedent-forming.

INSTRUCTIONAL DRAFT

A project that incorporates new evidence without reopening the decision record may appear adaptive. But it may also be obscuring whether the original commitment remains within its admissible basis.

Change control preserves the distinction between learning and retroactive rationalization.

14.5 Precedent must be governed

The fifth implication is that institutions must govern the afterlife of decisions.

A decision does not end when it is issued. It may be cited, repeated, generalized, misquoted, normalized, or used to justify later commitments.

This is especially true in frontier domains where early actions set norms.

A bounded lunar verification action may later be treated as resource-use permission.

A queue-entry decision may normalize the idea that interconnection is merely administrative.

A repository authorization may shape future siting and closure expectations.

A critical minerals funding decision may normalize strategic urgency as a substitute for evidence adequacy.

An autonomy permission may become a template for future delegated action near irreversible thresholds.

Institutions should therefore record the precedent object.

The relevant questions are:

- What could this decision make easier to repeat, cite, normalize, or inherit?
- What exactly is the precedent?
- What does not transfer?
- What future use requires re-screening?
- What public language is permitted?
- What misuse is most likely?

A Precedent Card or equivalent record prevents a bounded decision from quietly becoming a broader rule.

Precedent control is not bureaucratic caution. It is necessary because first actions can reshape the future admissibility environment.

14.6 Authority must remain aligned with consequence

INSTRUCTIONAL DRAFT

The sixth implication is that institutions must preserve authority coherence.

Formal authority is not enough.

A board may formally approve or reject, while practical commitment has already formed through project teams, counterparties, public language, or capital sequencing. An agency may retain formal discretion while political commitments make refusal impossible. An investor may bear downside while an operator controls escalation. A mission authority may own the decision while partners structure the architecture.

Authority coherence requires alignment among:

- formal decision rights;
- practical control;
- veto power;
- consequence bearing;
- revocation capacity;
- public communication authority;
- responsibility for outcomes.

When these diverge, commitment governance weakens.

Institutions should therefore identify the Decision Authority before commitment hardens. They should also identify who can stop, defer, re-site, re-sequence, constrain, or terminate the commitment.

If no one can credibly stop the pathway, the institution is no longer governing it.

14.7 Decision exposure should be identified early

The seventh implication is that organizations should identify exposure before formal commitment.

Many commitment-bearing actions occur before the institution recognizes them as decisions. Queue entry, site control, public naming, preferred-route language, bounded disturbance, repeated access, access-road planning, partner alignment, or capital pre-positioning may be treated as preparation.

But preparation can create exposure.

Institutions should therefore ask early:

Where does exposure begin?

INSTRUCTIONAL DRAFT

This question should be asked before:

- queue entry;
- site control;
- route fixation;
- drilling;
- excavation;
- injection;
- infrastructure placement;
- support system deployment;
- public designation;
- capital escalation;
- partner reliance;
- resource dependency;
- delegated autonomy.

The purpose is not to halt work. It is to locate the real commitment boundary before it is crossed silently.

14.8 Minimum evidence should be tied to the next threshold

The eighth implication is that evidence requirements should be threshold-specific.

It is not useful to ask whether the evidence is “enough” in general. Evidence is enough only relative to a proposed commitment.

The correct question is:

What must be known before this next irreversible threshold can be considered?

The evidence burden for bounded analysis differs from the evidence burden for site control. The evidence burden for reconnaissance differs from that for excavation. The evidence burden for queue screening differs from that for queue entry. The evidence burden for repository characterization differs from that for closure.

Minimum evidence should therefore be defined for each threshold.

This prevents institutions from using evidence sufficient for one phase to justify a stronger phase.

INSTRUCTIONAL DRAFT

14.9 Governance posture should be separated from executive decision

The ninth implication is that posture classification should not be confused with decision-making authority.

A governance posture classification determines whether proceeding, deferring, or refusing remains defensible under the evidence and authority boundary.

It does not make the decision.

The Decision Authority remains responsible for action, inaction, reliance, communication, execution, and outcomes.

This distinction matters because the framework's role is not to replace boards, agencies, investment committees, mission planners, engineers, regulators, or operators. It is to define whether the available evidence and authority structure can support the posture being considered.

The institutional value lies in preventing structurally indefensible postures from being treated as available options.

14.10 Commitment integrity requires active preservation

The tenth implication is that integrity must be preserved after entry.

A commitment does not remain defensible automatically.

Institutions should track whether:

- evidence has changed;
- assumptions have expired;
- dependencies have accumulated;
- support systems have hardened;
- authority has degraded;
- public language has expanded;
- counterparties have relied on continuation;
- precedent has escaped;
- reconsideration triggers have been crossed.

If any of these occur, the commitment may require constraint, re-evaluation, or termination.

Commitment integrity should be evaluated especially at phase transitions:

INSTRUCTIONAL DRAFT

- from analysis to site control;
- from site control to disturbance;
- from reconnaissance to infrastructure;
- from drilling to development;
- from injection testing to storage commitment;
- from characterization to repository construction;
- from emplacement to closure;
- from pilot to deployment;
- from bounded autonomy to delegated authority.

Integrity failures often occur during transitions because assumptions valid in one phase are carried into another without reauthorization.

14.11 Institutional design implications

The framework suggests several governance design practices.

14.11.1 Decision threshold registers

Institutions should maintain registers of irreversible thresholds. These identify where exposure begins, what evidence is required, who holds authority, and what actions require review.

14.11.2 Assumption registries

Assumptions should be recorded with scope, evidence basis, expiration conditions, and transfer limits.

14.11.3 Commitment trigger registers

Organizations should identify actions that trigger commitment: queue entry, site control, disturbance, capital allocation, public language, infrastructure placement, or partner reliance.

14.11.4 Evidence authority labels

Evidence should be labeled by what it can support: interest, ranking, comparison, verification, design, or commitment.

14.11.5 Refusal conditions

Stop, defer, re-site, re-sequence, constrain, and terminate conditions should be defined before exposure hardens.

14.11.6 Precedent controls

INSTRUCTIONAL DRAFT

Every high-consequence determination should identify what does and does not transfer.

14.11.7 Decision change logs

When evidence, assumptions, authority, or dependency conditions change, the decision record should be updated rather than silently reinterpreted.

These practices convert the theory into institutional memory.

14.12 Implications by institutional type

Different institutions face different versions of the same problem.

14.12.1 Mission agencies

Mission agencies should separate exploration justification from infrastructure authorization. They should define when payloads, traverses, landing sites, power systems, communications, or resource assumptions become commitment-bearing.

14.12.2 Infrastructure developers

Developers should identify when site control, interconnection, route fixation, permitting posture, or capital sequencing begins reducing refusal credibility.

14.12.3 Investors and investment committees

Investors should ask whether capital is being deployed into a pathway whose dominant uncertainty remains unresolved. They should distinguish bankability from admissibility.

14.12.4 Regulators and permitting bodies

Regulators should recognize that early filings, pilot approvals, or conditional permits may create precedent and institutional reliance.

14.12.5 Boards and executive sponsors

Boards should require clarity on decision authority, refusal conditions, and irreversible thresholds before approving capital or public commitments.

14.12.6 Scientific and technical teams

Technical teams should label the authority of evidence and prevent scientific signals from being over-translated into infrastructure decisions.

INSTRUCTIONAL DRAFT

14.12.7 Autonomous systems governance bodies

Autonomy authorities should distinguish capability from permission and define when delegated action may cross irreversible thresholds.

14.13 What the framework does not imply

The framework does not imply that institutions should avoid action until uncertainty disappears.

That would be impossible in frontier systems.

It does not imply that all disturbance is wrong.

It does not imply that all first-of-kind activity should be refused.

It does not imply that strategic urgency is irrelevant.

It does not imply that reversible evidence gathering should stop.

It does not imply that technical expertise is insufficient.

It does not imply that governance should replace engineering, finance, law, regulation, mission design, or operations.

The framework implies something narrower:

When action creates irreversible exposure, the evidence, authority, dependency, and precedent basis of that action must be governed before the action hardens.

This is a discipline of timing and boundary, not paralysis.

14.14 Section conclusion

The institutional implications of commitment admissibility and integrity are substantial.

Organizations must distinguish evidence from authorization. They must preserve refusal and deferral as valid outcomes. They must make threshold decisions reconstructible. They must freeze evidence when formal determinations are made. They must govern precedent. They must align authority with consequence. They must identify exposure early, tie evidence to thresholds, separate posture from decision authority, and preserve integrity after entry.

These practices do not replace existing decision systems. They protect those systems from being asked to optimize commitments that should not yet have hardened.

INSTRUCTIONAL DRAFT

The next section addresses why outcomes cannot validate threshold decisions in irreversible systems, and why decision quality must be reconstructed from the evidence state at the commitment boundary.

15. Why Outcomes Cannot Validate Threshold Decisions

In reversible systems, outcomes can teach directly.

If an action fails, the institution can revise, reverse, adapt, and try again. If an action succeeds, the institution can infer that the decision process was at least practically adequate. Feedback remains available because the system has not permanently foreclosed alternatives.

In irreversible systems, outcomes are less reliable as validators.

A favorable outcome can conceal a defective threshold decision. An unfavorable outcome can obscure a defensible one. Success may reward an inadmissible commitment. Failure may punish a commitment that was defensible under the evidence available at the time.

This creates a central epistemic problem for governance:

The quality of a threshold decision cannot be judged only by what happened afterward. It must be reconstructed from the evidence, authority, uncertainty, and reversibility conditions present when the commitment began.

This section explains why outcome-based validation is insufficient in disturbance-coupled systems and why commitment governance requires threshold provenance.

15.1 The outcome fallacy

The outcome fallacy occurs when a decision is judged primarily by its eventual result rather than by the conditions under which it was made.

In ordinary settings, this may be merely misleading. In irreversible systems, it can be institutionally dangerous.

A successful project can make a premature decision appear justified. The institution may conclude that because the outcome was favorable, the commitment was admissible. But success does not prove that the evidence, authority, and dependency structure at the threshold were sufficient.

INSTRUCTIONAL DRAFT

A failed project can make a defensible decision appear illegitimate. The institution may conclude that because the outcome was unfavorable, the commitment should never have been allowed. But failure does not prove that the threshold decision was inadmissible if the relevant uncertainties were bounded, refusal authority was preserved, and the commitment was defensible across plausible states at the time.

The error is confusing **outcome correctness** with **threshold defensibility**.

The framework rejects that confusion.

A threshold decision must be judged by asking:

- What was known at the time?
- What plausible states remained?
- Which uncertainties were decision-dominant?
- What exposure was created?
- Was reversal still credible?
- Was authority coherent?
- Did the action create dependency or precedent beyond its basis?
- Did the commitment remain within its defined boundary?

The outcome may be relevant evidence for later integrity review. It is not retroactive proof that the original threshold was admissible.

15.2 Success does not prove admissibility

Success can be especially dangerous because it normalizes the pathway that produced it.

A lunar ISRU demonstration may succeed technically. It may extract or detect useful volatile material, validate mobility, or show that a local operation can function. But if the commitment hardened around volatile assumptions that were still decision-dominant at entry, success does not prove that the threshold was admissible. It may simply mean that one favorable plausible state occurred.

A critical minerals project may become profitable. But if strategic urgency caused land, access, processing, water, offtake, and capital commitments to form before the deposit and system constraints were decision-grade, profitability does not prove that the commitment standard was sound. It may instead normalize weak evidence-to-infrastructure translation.

A grid project may survive queue entry and interconnection. But if queue entry was treated as administrative despite unresolved constraint-topology ambiguity, a favorable outcome does not prove the commitment was well-governed. It may encourage future projects to accept similar exposure without the same luck or conditions.

INSTRUCTIONAL DRAFT

A repository may perform well during early monitoring. But early performance does not prove that the long-duration authority, memory, public legitimacy, and stewardship basis remains intact.

Success can therefore create precedent risk.

It can cause institutions to convert a favorable outcome into a general rule:

- “This signal was enough.”
- “This queue entry worked.”
- “This pilot proved the pathway.”
- “This site selection was validated.”
- “This early authorization was justified.”

But the correct governance question is narrower:

Was the decision admissible under the evidence boundary at the time, and does its precedent transfer?

Often, the answer may be that the outcome was favorable but the precedent should not transfer.

15.3 Failure does not prove inadmissibility

Failure can also mislead.

An admissible commitment can fail. Disturbance-coupled systems are uncertain, physically complex, and institutionally difficult. A decision can be made under a defensible evidence boundary and still produce an adverse outcome.

A bounded lunar verification action may reveal that volatiles are inaccessible or mechanically difficult. That does not mean the verification decision was inadmissible if the action was bounded, alternatives remained open, disturbance was controlled, and refusal authority was preserved.

A geothermal well may fail to produce expected flow. That does not automatically mean drilling was inadmissible if the reservoir uncertainty was understood, the evidence justified the test, exposure was bounded, and termination conditions were in place.

A carbon storage project may encounter unexpected pressure response. That does not prove injection was inadmissible if the original commitment included monitoring, constraint, and re-evaluation mechanisms.

INSTRUCTIONAL DRAFT

A critical minerals program may downgrade a resource after drilling. That does not prove exploration was wrong if the drilling was bounded and did not over-authorize infrastructure, offtake, or capital lock-in.

Failure should trigger learning and possibly integrity review. It should not automatically rewrite the entry determination.

The relevant question is:

Did the decision record anticipate that adverse plausible state and preserve the authority to respond?

If yes, failure may be evidence that the governance system worked.

15.4 Outcomes as integrity evidence, not retroactive justification

Outcomes matter, but they matter in the correct place.

They are evidence for commitment integrity, not retroactive validation of entry.

After a commitment begins, new outcomes can reveal whether the commitment remains within its admissible basis. A monitoring result, drilling result, excavation result, queue study, containment observation, or operational event can change the plausible state space.

Those outcomes may support:

- **Maintain**, if the commitment remains within admitted bounds;
- **Constrain**, if continuation remains defensible only under narrower limits;
- **Re-evaluate**, if the original basis can no longer be relied upon without reopening;
- **Terminate**, if the commitment no longer satisfies its admissible basis.

This preserves the role of outcomes without allowing them to distort the original record.

For example, if lunar verification reveals accessible volatiles within the expected range, that may support maintaining or reconsidering a bounded commitment. But it does not automatically validate a prior infrastructure commitment if that commitment exceeded the evidence boundary at entry.

If a storage plume behaves within expected bounds, that may support maintaining the storage commitment. But it does not erase the need to evaluate whether injection was admissible at the time or whether future expansion requires new review.

INSTRUCTIONAL DRAFT

If a grid study produces acceptable upgrade exposure, that may support continued interconnection posture. But it does not prove that queue entry was non-commitment-bearing as a general rule.

Outcomes inform integrity. They do not erase provenance.

15.5 Threshold provenance

Threshold provenance is the reconstructible record of the decision conditions at the moment commitment began.

It is the evidence of whether the commitment was admissible, which posture was defensible, and what boundaries governed the decision.

A threshold provenance record should include:

- the decision under review;
- the commitment threshold;
- the evidence boundary;
- the evidence authority level;
- the plausible state space;
- the dominant uncertainty;
- the active irreversibility regimes;
- the dependence topology;
- the reversal burden;
- the reversal capacity;
- the refusal credibility condition;
- the Decision Authority;
- the authority coherence condition;
- the permitted actions;
- the prohibited actions;
- the assumptions registry;
- the expiration conditions;
- the reconsideration triggers;
- the precedent containment conditions.

Without threshold provenance, institutions cannot distinguish among four cases:

1. a sound decision with a favorable outcome;
2. a sound decision with an unfavorable outcome;
3. an unsound decision with a favorable outcome;
4. an unsound decision with an unfavorable outcome.

INSTRUCTIONAL DRAFT

Only the first and second preserve governance integrity. The third is especially dangerous because it rewards a defective threshold and may become precedent.

15.6 The four outcome-decision combinations

The relationship between threshold defensibility and outcome can be represented as four quadrants.

Threshold decision	Outcome	Interpretation
Defensible	Favorable	Governance and outcome aligned
Defensible	Unfavorable	Sound threshold, adverse state or execution outcome
Defective	Favorable	Lucky success; high precedent risk
Defective	Unfavorable	Governance failure and adverse outcome aligned

The most dangerous quadrant is not always defective decision plus unfavorable outcome. That case is visible.

The more dangerous quadrant may be defective decision plus favorable outcome.

In that quadrant, the institution learns the wrong lesson. It may conclude that the threshold was sound because the outcome was good. It may generalize the decision. It may lower future evidence standards. It may treat luck, favorable state resolution, or external conditions as proof of governance adequacy.

This is how weak precedent forms.

A favorable outcome can turn a bad threshold into a future rule.

That is why outcome success must be separated from threshold admissibility.

15.7 Why irreversible systems amplify outcome bias

Irreversible systems amplify outcome bias because the evidence needed to judge the original decision often becomes obscured after action.

Several mechanisms contribute.

15.7.1 Evidence baseline alteration

INSTRUCTIONAL DRAFT

Disturbance can change the system being evaluated. Later evidence may no longer reflect the pre-commitment state.

Lunar excavation, drilling, injection, seabed disturbance, and repository access can all alter the baseline.

15.7.2 Sunk-cost reinterpretation

Once capital is deployed, institutions reinterpret prior evidence in favor of continuation.

Evidence that was once uncertain becomes described as sufficient because reversing would be costly.

15.7.3 Narrative hardening

Public language, investor materials, agency statements, and project identity can convert uncertain assumptions into institutional facts.

15.7.4 Memory loss

Over time, the original uncertainty, assumptions, and refusal conditions may be forgotten.

This is especially important for long-duration commitments.

15.7.5 Precedent simplification

Future actors may cite only the outcome, not the boundary.

They may say “it worked” without preserving the conditions under which it was allowed.

These mechanisms make provenance necessary.

15.8 Decision quality is evidence-bound and time-bound

A threshold decision must be judged by the record available at the threshold.

It is evidence-bound because only the evidence available at the time could legitimately support the decision.

It is time-bound because the admissibility condition can change as evidence, assumptions, dependencies, and authority conditions evolve.

This does not mean later evidence is ignored. Later evidence belongs to integrity review.

INSTRUCTIONAL DRAFT

The distinction is:

- **Admissibility** is judged at entry under the evidence boundary.
- **Integrity** is judged after entry as the system evolves.
- **Outcome** informs integrity and learning, but does not rewrite entry.

This structure preserves accountability.

Without it, institutions can defend bad threshold decisions using good outcomes or condemn sound threshold decisions using bad outcomes.

Both are errors.

15.9 Implications for scientific and institutional claims

The outcome distinction also affects how scientific and institutional claims should be made.

A successful field action should not automatically be framed as validating a commitment framework. It may validate a measurement, capability, or local condition. It may not validate infrastructure dependency, site hardening, capital sequencing, or future precedent.

Similarly, an unsuccessful field action should not automatically invalidate the underlying exploration rationale. It may reveal that a particular commitment should be constrained, re-evaluated, or terminated.

Claims should therefore be scoped:

- “This evidence supports bounded verification.”
- “This result supports re-evaluation of the site model.”
- “This outcome remains within the admitted basis.”
- “This result does not transfer to infrastructure commitment.”
- “This finding triggers reconsideration before further hardening.”
- “This outcome does not validate the original threshold beyond its boundary.”

Such language preserves scientific integrity and governance discipline.

15.10 Outcomes and the Precedent Card

Because outcomes can be misused, precedent control is essential.

INSTRUCTIONAL DRAFT

A Precedent Card or equivalent record should specify whether an outcome may be cited in future decisions and under what conditions.

It should identify:

- what decision produced the outcome;
- what evidence boundary applied;
- what action was actually authorized;
- what was not authorized;
- what the outcome does and does not prove;
- what future use requires re-screening;
- what public language is allowed;
- what misuse is likely.

For example, a successful lunar bounded verification action may be cited as evidence that the specific verification method produced useful information. It should not automatically be cited as evidence that lunar ISRU infrastructure is admissible.

A successful grid interconnection case may be cited as a project-specific outcome. It should not automatically justify treating queue entry as non-commitment-bearing.

A repository performance milestone may be cited within the monitoring record. It should not automatically validate long-duration integrity beyond the phase reviewed.

Precedent control ensures that outcomes do not escape their decision boundary.

15.11 Section conclusion

Outcomes matter, but they cannot be allowed to govern threshold legitimacy retroactively.

In irreversible systems, success does not prove admissibility. Failure does not prove inadmissibility. Outcomes are evidence for integrity review, not a substitute for threshold provenance.

The quality of a commitment decision must be reconstructed from the evidence, plausible states, irreversibility regimes, refusal credibility, authority coherence, dependency topology, and precedent conditions present when the threshold was crossed.

This is why commitment governance is not prediction.

It is a discipline of evidence-bounded, time-bounded, authority-bounded decision reconstruction.

The next section discusses what the framework adds, where it is limited, and what future work is needed to formalize and apply it across domains.

INSTRUCTIONAL DRAFT

16. Discussion

The framework developed in this paper adds a governance layer for systems in which learning, disturbance, exposure, authority, and commitment formation are structurally coupled.

It does not replace technical, financial, legal, operational, scientific, or regulatory judgment. It addresses a prior and narrower question: whether an irreversible commitment should be allowed to form under the evidence and authority conditions present at the threshold, and whether that commitment remains defensible as conditions change.

The discussion clarifies what the framework contributes, where it is limited, and what further work is needed.

16.1 What the framework adds

The framework adds six distinctions that are often blurred in high-consequence exploration and infrastructure decisions.

16.1.1 Feasibility is not admissibility

A project can be feasible and still inadmissible.

Feasibility asks whether an action can be done. It evaluates capability, engineering, execution, cost, technology, logistics, and operational viability.

Admissibility asks whether the commitment should be allowed to form under the evidence, uncertainty, authority, and reversibility conditions present at the threshold.

This distinction matters because institutions often move from feasibility to momentum. Once a project appears technically possible, the organizational question becomes how to execute it. In disturbance-coupled systems, the prior question remains unresolved: whether the project should be allowed to harden before uncertainty loses decision-changing power.

The framework therefore separates:

Can this be done?

from:

Should this commitment be allowed to form?

That separation is central to governance under irreversibility.

INSTRUCTIONAL DRAFT

16.1.2 Evidence is not authorization

The framework also separates evidence from authorization.

Evidence may be valid, useful, and scientifically meaningful without being commitment-grade.

A signal can support prospecting without supporting infrastructure. A geophysical anomaly can support investigation without supporting drilling commitment. A queue opportunity can support screening without supporting site-grid lock-in. A successful pilot can support learning without authorizing deployment.

The framework asks what the evidence is allowed to support.

This is especially important where evidence is indirect and interpretation remains non-unique. Evidence can become over-authorized when it is used beyond its legitimate decision authority.

The institutional failure is not necessarily misuse of bad evidence. It is often overuse of good evidence.

Commitment governance addresses that failure.

16.1.3 Uncertainty is not ignorance dominance

The framework distinguishes ordinary uncertainty from ignorance dominance.

Uncertainty is expected in frontier systems. No project has complete information before action. The presence of uncertainty alone does not make a commitment inadmissible.

Ignorance becomes decision-dominant when unresolved plausible states can still change whether the commitment is admissible.

This distinction prevents two opposite errors.

The first error is paralysis: treating all uncertainty as disqualifying.

The second error is overconfidence: treating all uncertainty as ordinary risk.

The framework avoids both by asking whether the remaining uncertainty still governs the commitment decision.

If plausible states imply the same bounded action, uncertainty may be acceptable. If plausible states imply different admissibility outcomes, ignorance remains dominant.

INSTRUCTIONAL DRAFT

16.1.4 Exploration is not commitment

The framework separates exploration from commitment.

Exploration may proceed when it preserves refusal authority, limits disturbance, avoids site or architecture hardening, and prevents evidence over-transfer.

Commitment begins when action makes refusal, re-siting, re-sequencing, or abandonment materially harder.

This distinction is important because many systems use exploratory language for commitment-bearing actions.

A “test” may create a corridor.

A “pilot” may create precedent.

A “preliminary route” may harden placement.

A “bounded verification” may anchor future infrastructure.

A “queue entry” may create site-grid commitment.

A “demonstration” may normalize deployment.

The framework therefore evaluates actions by their exposure effect, not their label.

16.1.5 Entry is not persistence

The framework separates commitment admissibility from commitment integrity.

Admissibility governs entry. It asks whether a commitment may be considered under current conditions.

Integrity governs persistence. It asks whether the commitment remains within its admissible basis as evidence changes, assumptions expire, dependencies accumulate, authority degrades, or precedent expands.

This distinction is a major theoretical contribution.

Many frameworks focus on the initial decision or on downstream risk management. This framework treats the commitment as a structure that must remain inside its basis over time.

A commitment can be admissible at entry and lose integrity later.

A commitment can perform well and still exceed its admissible basis.

A commitment can encounter unfavorable evidence and still maintain integrity if it was bounded correctly.

INSTRUCTIONAL DRAFT

This lifecycle view is necessary for long-duration and path-dependent systems.

16.1.6 Outcomes do not validate thresholds

The framework separates outcome from threshold defensibility.

A favorable outcome does not prove that the original commitment was admissible. An unfavorable outcome does not prove that it was inadmissible.

This is important because irreversible systems can reward bad threshold decisions and punish defensible ones. The relevant standard is the evidence, uncertainty, authority, and reversibility condition present at the threshold.

This makes commitment governance a provenance discipline.

It requires decision records that can reconstruct what was known, what remained plausible, what threshold was crossed, what exposure was created, and whether refusal remained credible.

16.2 Relationship to existing domains of practice

The framework is adjacent to several established domains, but it occupies a different layer.

It is adjacent to risk management, but it governs before the pathway becomes defined as the object of risk management.

It is adjacent to systems engineering, but it asks whether requirements, assumptions, and architectures should be allowed to harden.

It is adjacent to resource classification, but it asks whether resource evidence may govern commitment.

It is adjacent to real options, but it asks whether option-preserving actions themselves create exposure.

It is adjacent to value of information, but it asks whether information can be obtained without crossing the threshold it is meant to justify.

It is adjacent to adaptive management, but it asks whether later adjustment remains credible.

It is adjacent to mission assurance, but it asks whether the mission or architecture should be allowed to become commitment-bearing.

INSTRUCTIONAL DRAFT

It is adjacent to project finance, but it asks whether capital is hardening around a pathway before uncertainty has stopped governing the decision.

The framework therefore sits upstream of many established practices.

Its role is not to perform those practices. Its role is to prevent those practices from being asked to optimize commitments that should not yet have formed.

16.3 Generality and limits

The framework is general, but not universal.

It is general because the same commitment structure appears across multiple domains: lunar ISRU, geological disposal, critical minerals, geothermal, carbon storage, marine infrastructure, grid interconnection, orbital systems, and autonomy.

In each case, evidence remains incomplete, interpretation may remain non-unique, learning may require intervention, intervention creates exposure, and exposure can harden into commitment.

But the framework is not universal.

It should not be applied where:

- the action is genuinely reversible;
- learning does not create material exposure;
- evidence is direct and sufficient for the commitment;
- alternatives remain open;
- authority remains coherent;
- dependency does not form;
- precedent does not matter;
- the decision does not approach an irreversible threshold.

Over-application would weaken the framework.

The value of the framework lies in disciplined use at commitment-bearing thresholds.

16.4 The danger of over-formalization

One risk of the framework is that it could become overly formal or procedural.

INSTRUCTIONAL DRAFT

If applied mechanically, it could create the appearance of rigor without improving judgment. It could become another checklist, another approval gate, or another bureaucratic layer. That would defeat the purpose.

The framework is not meant to produce procedural complexity for its own sake.

Its purpose is to preserve judgment before irreversible exposure hardens.

The key is boundary discipline:

- define the decision;
- identify the threshold;
- classify evidence authority;
- test ignorance dominance;
- map dependence topology;
- evaluate refusal credibility;
- define the admissible basis;
- preserve integrity through time;
- contain precedent.

These steps should be proportional to the consequence of the commitment.

A small reversible action does not require the same treatment as a repository, lunar infrastructure commitment, or capital-intensive critical minerals pathway.

The framework must remain threshold-bound.

16.5 The role of judgment

The framework uses formal primitives, but it does not eliminate judgment.

Judgment is required to define the plausible state space, identify the dominant uncertainty, determine whether ambiguity is decision-dominant, assess refusal credibility, classify irreversibility regimes, and determine whether precedent transfer is likely.

The framework does not replace judgment. It structures it.

This is especially important because many of the relevant variables are not purely technical. Authority coherence, public legitimacy, institutional memory, continuation pressure, and precedent expansion cannot be reduced to simple metrics.

They can, however, be made explicit.

The framework's value lies in making hidden judgment visible before the commitment hardens.

INSTRUCTIONAL DRAFT

16.6 Implications for science

For science, the framework clarifies that evidence has authority boundaries.

Scientific evidence can be strong for one purpose and insufficient for another. A measurement may be valid as a detection, weak as a resource classification, and inadmissible as infrastructure authority.

This is not a demotion of science. It protects scientific evidence from being asked to do institutional work it cannot support.

In lunar science, this distinction prevents volatile signals from being prematurely converted into site and ISRU commitments.

In geophysics, it prevents non-unique inversions from being treated as commitment-grade subsurface truths.

In environmental and marine science, it prevents baseline evidence from being overextended into disturbance permission.

In nuclear waste and geological disposal, it prevents technical evidence from being detached from long-duration institutional authority.

The framework therefore supports scientific integrity by preserving the distinction between detection, interpretation, verification, and authorization.

16.7 Implications for infrastructure and capital

For infrastructure and capital, the framework clarifies that capital sequencing is not neutral.

Capital can preserve optionality, but it can also destroy it. It can fund learning, but it can also create continuation pressure. It can discipline projects, but it can also turn unresolved assumptions into institutional commitments.

Investment committees, boards, lenders, insurers, and public funding bodies should therefore ask not only whether a project is attractive or bankable, but whether capital is hardening around a pathway that remains ignorance-dominant.

The framework is particularly relevant where early capital produces:

- site control;

INSTRUCTIONAL DRAFT

- queue entry;
- land acquisition;
- corridor fixation;
- drilling;
- infrastructure pre-positioning;
- partner reliance;
- public expectations;
- or strategic-resource narrative.

In such cases, capital is not merely financing a decision. It is forming part of the decision.

16.8 Implications for frontier governance

For frontier domains, the framework emphasizes precedent.

Early actions in frontier systems often define future norms. A first lunar extraction, first seabed disturbance, first large-scale storage project, first repository authorization, or first autonomy permission can become a reference point.

If these actions are not bounded, they can create governance drift.

Frontier governance therefore requires special attention to precedent containment:

- what was authorized;
- what was not authorized;
- what evidence supported the action;
- what assumptions bounded the decision;
- what does not transfer;
- what future use requires re-screening.

This is especially important because frontier decisions often occur before mature institutional frameworks exist. Early actions can become de facto law, practice, expectation, or standard.

The framework treats precedent as a form of irreversibility.

16.9 Limitations

Several limitations should be acknowledged.

First, the framework is conceptual. It does not yet provide a fully quantified model of admissibility or integrity.

INSTRUCTIONAL DRAFT

Second, the framework depends on high-quality boundary definition. If the decision, threshold, evidence record, or authority condition is poorly defined, the framework cannot function properly.

Third, plausible-state reasoning requires domain expertise. The framework does not generate plausible states independently. It structures how they affect commitment.

Fourth, authority coherence and refusal credibility are partly institutional judgments. They require careful assessment and may be contested.

Fifth, the framework may be difficult to apply in environments where decision authorities are unwilling to accept refusal, deferral, re-evaluation, or termination as legitimate outcomes.

Sixth, the framework is most useful at high-consequence thresholds. It may be excessive for low-consequence or genuinely reversible decisions.

Seventh, commitment integrity requires institutional memory. Without records, assumption tracking, and change control, persistence review becomes difficult.

These limitations define the conditions for responsible use.

16.10 Future work

Future work should develop the framework in several directions.

16.10.1 Formalization of admissibility sets

The admissibility set could be developed more rigorously, including conditions under which plausible-state divergence produces admissibility flips.

16.10.2 Metrics for reversal burden and reversal capacity

Future work could classify reversal burden and reversal capacity across financial, technical, political, regulatory, institutional, and physical dimensions.

16.10.3 Domain-specific threshold taxonomies

Each domain could develop a threshold taxonomy: lunar surface commitments, geothermal drilling thresholds, carbon storage injection thresholds, grid interconnection thresholds, critical minerals project thresholds, and repository phase transitions.

16.10.4 Commitment integrity diagnostics

INSTRUCTIONAL DRAFT

Future work could develop diagnostics for evidence drift, assumption expiry, dependency accumulation, authority degradation, and precedent expansion.

16.10.5 Empirical case reconstruction

Historical cases could be reconstructed using the framework to distinguish outcome success from threshold defensibility.

16.10.6 Governance record design

Research could develop standard structures for evidence freeze, assumption registries, decision change logs, precedent cards, and admissible-basis records.

16.10.7 Relationship to assurance and regulation

The framework could be integrated with mission assurance, infrastructure assurance, environmental review, financial diligence, and regulatory decision processes without replacing them.

16.10.8 Organizational adoption

Future work could study how institutions preserve refusal authority under capital, political, and operational pressure.

16.11 Theoretical contribution

The paper's theoretical contribution can be summarized as follows.

First, it defines a class of systems in which learning and commitment are structurally coupled.

Second, it introduces commitment admissibility as a threshold condition for entry.

Third, it introduces commitment integrity as a persistence condition after entry.

Fourth, it decomposes irreversibility into regimes that can be diagnosed separately and in combination.

Fifth, it identifies ignorance dominance as distinct from uncertainty.

Sixth, it treats precedent as a form of irreversibility.

Seventh, it argues that outcome success cannot retroactively validate threshold decisions.

INSTRUCTIONAL DRAFT

Together, these contributions create a governance framework for frontier systems where the cost of error is not only failure, but entrenchment.

16.12 Practical contribution

The practical contribution is the ability to ask a different class of question before irreversible exposure forms.

Instead of asking only:

- Is the project feasible?
- Is the resource promising?
- Is the risk manageable?
- Is the technology ready?
- Is the project financeable?
- Is the architecture coherent?

the framework asks:

- What commitment is actually forming?
- What evidence is being allowed to authorize it?
- What uncertainty remains decision-dominant?
- What becomes hard to reverse?
- Who can still refuse?
- What dependencies are forming?
- What precedent is being created?
- Does the commitment remain within its basis over time?

These questions do not slow decisions for their own sake. They prevent institutions from losing authority before they realize a decision has been made.

16.13 Section conclusion

The framework contributes a threshold-governance theory for disturbance-coupled exploration and infrastructure systems.

Its central value is not prediction. It is discipline at the boundary where evidence, action, authority, and irreversibility begin to couple.

INSTRUCTIONAL DRAFT

It clarifies that feasibility is not admissibility, evidence is not authorization, uncertainty is not automatically ignorance dominance, exploration is not commitment, entry is not persistence, and outcomes do not validate thresholds.

The final section concludes by returning to the central claim: in frontier systems, the cost of error is not failed execution alone. It is the entrenchment of commitments made before the system was sufficiently understood, and the persistence of commitments after their admissible basis has expired.

next section
Thought for 4s

17. Conclusion: The Cost of Error Is Entrenchment

Exploration and infrastructure systems increasingly operate in conditions where learning, disturbance, exposure, and commitment are structurally coupled.

The target condition is often not directly observable before action. Evidence is indirect. Interpretation remains non-unique. Verification requires intervention. Intervention creates exposure. Exposure creates dependency. Dependency can harden into commitment before uncertainty has lost decision-changing power.

In such systems, the central governance problem is not simply risk.

It is timing.

A project may act too early not because its goal is wrong, not because its technology is impossible, and not because its evidence is meaningless, but because the commitment begins to harden before the evidence, authority, and dependency structure can support it.

That is the failure this paper has sought to name.

17.1 The central claim

The central claim is that some systems require governance before commitment hardens, not merely risk management after action begins.

This paper introduced **commitment admissibility** as the entry condition for irreversible commitments under persistent uncertainty.

Commitment admissibility asks:

May this commitment be considered under the current evidence and authority boundary?

INSTRUCTIONAL DRAFT

It separates feasibility from admissibility, evidence from authorization, exploration from commitment, and uncertainty from ignorance dominance.

A commitment is inadmissible when unresolved plausible states can still change whether the action should proceed, defer, or be refused, and when the action required to resolve that uncertainty would itself create the exposure being evaluated.

This paper also introduced **commitment integrity** as the persistence condition after entry.

Commitment integrity asks:

Does this commitment remain within the admissible basis that made it defensible?

It recognizes that a commitment can become indefensible after it begins. Evidence may drift. Assumptions may expire. Dependencies may accumulate. Authority may degrade. Precedent may expand. The commitment may leave the boundary that once supported it.

Together, admissibility and integrity govern the commitment lifecycle.

Admissibility governs whether entry may be considered.

Integrity governs whether continuation remains defensible.

17.2 What the forcing cases show

Lunar south polar ISRU shows the entry problem in concentrated form.

Volatile evidence can be meaningful without being infrastructure-grade. A hydrogen signal, thermal context, radar response, illumination condition, or terrain advantage can justify prospecting and bounded verification. It does not automatically justify site hardening, access repetition, fixed power placement, excavation dependency, corridor formation, ISRU architecture, or precedent-forming resource use.

The lunar case shows that a signal can become a site, a site can become an architecture, and an architecture can become a commitment before the underlying subsurface state is sufficiently resolved.

The governing question is not whether there is ice.

The governing question is whether the evidence has earned authority over irreversible surface commitment.

Geological disposal and nuclear waste show the persistence problem.

A repository commitment cannot be governed only at entry. It must remain defensible across phase transitions, institutional turnover, public legitimacy, monitoring limits, closure decisions,

INSTRUCTIONAL DRAFT

and intergenerational responsibility. The technical containment case matters, but so do authority continuity, memory, responsibility, and precedent.

The repository case shows that a commitment may lose integrity even if it began with a defensible basis.

Critical minerals show the transfer problem.

A prospect can be real, strategically important, and worthy of exploration while still being premature as the anchor for extraction, land control, processing, water, power, offtake, infrastructure, permitting, and capital commitment. Strategic urgency can convert prospectivity into authority before the evidence has earned that role.

The critical minerals case shows that importance does not equal admissibility.

Together, these cases establish the framework's reach:

- lunar ISRU shows premature entry;
- geological disposal shows long-duration persistence;
- critical minerals show urgency-driven conversion from signal to commitment.

17.3 What the framework changes

The framework changes the institutional question.

Instead of asking only:

Can this project work?

it asks:

What commitment is forming, and has the evidence earned authority over it?

Instead of asking only:

Can the risks be managed?

it asks:

Should this exposure be allowed to form before uncertainty loses decision-changing power?

Instead of asking only:

Is more information valuable?

it asks:

INSTRUCTIONAL DRAFT

Does obtaining that information require crossing the threshold the information is meant to justify?

Instead of asking only:

Did the outcome succeed?

it asks:

Was the threshold defensible under the evidence, authority, and reversibility conditions present at the time?

Instead of asking only:

Was entry justified?

it asks:

Does the commitment still remain within its admissible basis?

These are governance questions, not execution questions.

They sit upstream of engineering, finance, permitting, mission design, regulation, operations, and optimization. They do not replace those systems. They protect them from being asked to optimize a commitment that should not yet have hardened.

17.4 The role of refusal

A central implication of the framework is that refusal is not failure.

Refusal can preserve value.

Deferral can preserve optionality.

Constraint can preserve integrity.

Re-evaluation can prevent stale reliance.

Termination can prevent deeper entrenchment.

Many institutional systems treat these outcomes as negative because they interrupt momentum. But in irreversible systems, momentum can be the failure mode.

A system that cannot refuse is not exercising authority.

A system that cannot defer is not preserving learning.

INSTRUCTIONAL DRAFT

A system that cannot re-evaluate is not preserving integrity.
A system that cannot terminate is no longer governing commitment.

This is especially important in frontier settings, where first actions can become norms, public narratives can harden quickly, and evidence can be over-authorized by strategic urgency.

The framework restores refusal, deferral, constraint, re-evaluation, and termination as legitimate governance outcomes.

17.5 The final contribution

The paper's final contribution is to redefine the cost of error.

In ordinary project terms, the cost of error is often understood as failed execution: cost overrun, technical failure, delay, underperformance, environmental harm, or financial loss.

In disturbance-coupled systems, the deeper cost may be entrenchment.

Entrenchment occurs when a pathway becomes harder to abandon than it was ever qualified to become. It occurs when evidence that supported exploration becomes infrastructure authority. It occurs when a reversible investigation becomes a site anchor. It occurs when a queue entry becomes capital lock-in. It occurs when prospectivity becomes extraction commitment. It occurs when a bounded pilot becomes precedent. It occurs when a commitment persists after its admissible basis has expired.

Entrenchment is dangerous because it can survive success.

A project may succeed locally while lowering the standard for future commitment. A mission may work while normalizing premature infrastructure dependence. A pilot may perform while becoming precedent beyond its boundary. A project may generate value while degrading the discipline required for the next decision.

That is why outcome success is not enough.

The question is whether the commitment was allowed to form under conditions that could defend it, and whether it remains within that boundary as the system evolves.

17.6 Closing statement

Exploration requires action under uncertainty. Infrastructure requires commitment before everything is known. Frontier systems will never wait for complete information.

The framework does not demand certainty.

INSTRUCTIONAL DRAFT

It demands that uncertainty not be converted into irreversible commitment before it loses the power to change the decision.

It demands that evidence not be asked to authorize more than it can support.

It demands that exploration not be allowed to become commitment by accident.

It demands that commitments remain inside the basis that made them defensible.

In frontier systems, the cost of error is not failed execution alone.

It is the entrenchment of commitments made before the system was sufficiently understood, and the persistence of commitments after their admissible basis has expired.