Minimum Evidence Before Lunar Commitment

Governing Surface Hardening Under Persistent Subsurface Ignorance

N. K. Grapsas | Sustainable Exploration.

Purpose

This document supports the poster Irreversible Commitments in Lunar Subsurface Exploration: Ignorance, Disturbance, and Precedent Formation. It defines the decision problem behind early lunar infrastructure commitment and introduces minimum evidence standard for actions that harden the surface, fix operational geometry, or allow volatile assumptions to govern architecture.

1. The Problem

Lunar development is often framed as an architecture problem.

What systems are required? How should mobility, power, communications, logistics, autonomy, resource prospecting, and surface construction fit together? Which capabilities should arrive first? Which subsystems need to be demonstrated before broader operations can begin?

Those questions matter. But they do not exhaust the decision problem.

A lunar architecture can be internally coherent and still commit too early. A surface system can be technically plausible and still depend on subsurface conditions that have not been resolved at the scale of the proposed action. A site can be promising for exploration and still be premature as an infrastructure anchor.

The central issue is the timing of path-dependent commitment under unresolved subsurface uncertainty.

At the lunar south pole, volatile-bearing materials are inferred through indirect evidence. Hydrogen indications, permanently shadowed regions, thermal stability, radar response, reflectance, topographic context, and illumination conditions all help constrain the problem. They support exploration interest. They do not automatically determine where infrastructure should be placed, where power should be fixed, where corridors should form, or where excavation should begin.

The difference matters because lunar development does not become irreversible only at the moment of full construction. Commitment forms earlier. It forms when a site becomes privileged, when repeated access organizes around a location, when a landing zone is prepared, when power is placed, when corridors appear, when excavation begins, when subsurface access alters the environment, and when future missions begin to rely on the same assumptions.

These actions may appear incremental. They are not neutral.

They begin to define the geometry of future development. They influence which measurements become easier or harder to obtain. They shape where future assets can operate. They may disturb or obscure the very evidence needed to decide whether the original commitment was justified.

That is the governing problem.

The question is not whether lunar volatiles exist. The question is whether current evidence is sufficient to permit actions that would allow volatile assumptions to govern infrastructure.

A volatile signal may justify continued investigation. It may justify bounded scouting, additional sensing, or limited intervention designed to improve knowledge. But it does not necessarily justify surface hardening, fixed power placement, corridor formation, excavation geometry, or long-horizon dependency.

The burden of evidence must rise with the burden of the action.

A reconnaissance act can tolerate uncertainty that an infrastructure commitment cannot. A site comparison campaign can proceed under ambiguity that a landing pad cannot. A bounded sampling event may be admissible where a resource-processing dependency is not.

The mistake is to treat all early lunar activity as exploration.

Some exploration remains reversible. Some exploration becomes architecture-forming. Once exploration begins to determine where future operations must occur, it has crossed into commitment.

Minimum evidence before lunar commitment is therefore not a scientific checklist. It is a decision standard.

It asks whether the next action should be allowed to proceed given what remains unresolved and what the action would make difficult to reverse.

2. The Three-Layer Structure

Lunar development should be evaluated across three layers.

3. Architecture Readiness

1. Architecture readiness asks whether the required systems have been identified and whether they fit together in a coherent operational concept.
2. This includes mobility, power, communications, navigation, autonomy, excavation, processing, logistics, thermal support, and human or robotic operations.
3. Architecture readiness is concerned with system design.
4. It asks whether the system could be arranged.

4. Construction Readiness

1. Construction readiness asks whether the system can actually be built, placed, sustained, and operated at the site.
2. This includes terrain, slopes, bearing conditions, regolith behavior, thermal cycles, illumination, access constraints, dust, communication geometry, landing effects, maintenance, excavation behavior, and survivability.
3. Construction readiness is concerned with site and operational feasibility.
4. It asks whether the system could work there.

5. Commitment Admissibility

1. Commitment admissibility asks a different question.
2. Should the next path-dependent action be allowed to proceed under the evidence currently available?
3. This layer is not about whether a concept is attractive. It is not about whether the system can be imagined, engineered, or simulated. It is about whether unresolved uncertainty still has the power to change the legitimacy of the proposed commitment.
4. A system may be architecturally coherent.
5. A site may appear construction-plausible.
6. The next action may still be inadmissible.
7. That gap is the focus of this handout.
8. Architecture readiness is advancing faster than construction readiness. Construction readiness is advancing faster than commitment admissibility. The danger is that an architecture becomes persuasive before the evidence has earned the right to anchor it.

6. Surface Commitment Classes

The following classes identify actions that should be treated as commitment-bearing unless shown otherwise. They are not ranked only by physical scale. Their importance comes from the way they shape future options.

6.1. SC-1 Site Hardening

1. Site hardening occurs when a location begins to receive operational preference through preparation, repeated use, protection, reinforcement, or institutional focus.
2. A site can become hardened before permanent infrastructure is installed. Naming, planning, access rehearsal, repeated mission design, and support

planning can all begin to privilege one location over alternatives.
3. The risk is that a site becomes the center of future development before the evidence has resolved whether it should hold that role.

6.2. SC-2 Landing Pad Preparation

1. Landing pad preparation begins to fix access geometry. It affects plume management, hazard zones, traffic patterns, proximity to assets, and future logistics.
2. A landing zone is not only a place where vehicles arrive. It becomes a spatial anchor. Once prepared, it can pull mobility, power, storage, communications, and operations toward itself.
3. If the landing location is chosen around unresolved resource assumptions, the pad can make those assumptions harder to revisit.

6.3. SC-3 Berm and Blast Protection

1. Berms and blast protection modify the surface in ways that change hazard relationships and operational layout.
2. They can protect assets, but they also define where assets are expected to be, where landing activity is expected to recur, and how future surface traffic may be organized.
3. Because berms alter terrain and future operations, they should not be treated as minor support structures when they reinforce an unverified development geometry.

6.4. SC-4 Excavation and Trenching

1. Excavation and trenching are direct disturbances of the material system being evaluated.
2. They can produce knowledge, but they can also destroy context, mix material, change thermal behavior, alter volatile retention conditions, expose or remove layers, and reshape the site for future use.
3. Excavation is therefore both epistemic and structural. It may help reveal the subsurface, while also changing the conditions under which later measurements must be interpreted.

6.5. SC-5 Corridor Establishment

1. Corridors form when repeated or engineered access begins to organize movement through the landscape.
2. A corridor can begin as a route. It can become a dependency. Once mobility, power, communications, and logistics adapt to it, alternatives may become more costly, less accessible, or institutionally disfavored.
3. Corridor formation is a major source of lock-in because it turns a choice of movement into a pattern of development.

6.6. SC-6 Regolith Processing Dependency

1. Regolith processing dependency arises when architecture begins to assume that local material can support extraction, shielding, construction, oxygen production, feedstock use, or other operational functions.
2. This class is especially important for ISRU.
3. A material may be interesting scientifically and still not be reliable enough for system dependency. Processing assumptions require evidence not only that a material exists, but that it is spatially coherent, mechanically accessible, operationally usable, and available at the scale required.

6.7. SC-7 Fixed Infrastructure Placement

1. Fixed infrastructure includes power, communications, navigation, storage, thermal, processing, excavation, or support assets placed in ways that constrain future action.
2. Fixed assets do not merely support activity. They define where activity is likely to continue.
3. Once fixed assets are placed, future decisions begin to inherit their geometry. This can convert a provisional assumption into an architectural fact.

6.8. SC-8 Subsurface Access

1. Subsurface access includes drilling, sampling, emplacement, trenching, excavation, or any other intervention that enters or disturbs the subsurface.
2. Subsurface access can be necessary for verification. But it must be governed because it can also modify the target environment, alter future measurement conditions, and create pressure for follow-on infrastructure.

3. The question is not whether subsurface access should occur. It is whether the access is bounded verification or the beginning of commitment.

7. Minimum Evidence Before Commitment

Commitment is admissible only when evidence is adequate relative to the commitment burden of the next action.

The evidence burden is not constant. It rises as an action becomes harder to reverse, as it narrows alternatives, as it creates dependencies, and as it affects future measurement.

A lightweight reconnaissance activity may be justified under limited evidence. A fixed power asset, prepared landing zone, corridor, excavation zone, or resource-processing dependency requires stronger justification.

Minimum evidence before commitment should satisfy five conditions.

8. Convergence Across Sensing Modes

Evidence should converge across modes in a way that supports the same operational interpretation.

It is not enough for multiple measurements to suggest a broad volatile possibility. The relevant issue is whether the measurements support the same conclusion at the scale and character of the proposed action.

For example, hydrogen indication, thermal retention context, radar response, and terrain data may all point toward exploration interest. But if they still permit materially different interpretations of distribution, concentration, physical state, or accessibility, they may not justify infrastructure commitment.

9. Scale Match Between Observation and Intervention

The evidence must resolve the scale at which action will occur.

A regional indication cannot automatically justify local infrastructure. A broad signal cannot automatically support a specific excavation plan. A remote sensing pattern cannot automatically determine mechanical accessibility at the scale of wheels, drills, buckets, pads, cables, berms, or support assets.

The evidence must be fit for the commitment being authorized.

10. Reduction of Dominant Subsurface Ambiguity

The remaining ambiguity must no longer be capable of changing the decision.

Uncertainty itself is not disqualifying. Frontier environments always contain uncertainty. The question is whether the uncertainty remains decision dominant.

If multiple plausible subsurface states remain consistent with the evidence, and those states would lead to different infrastructure decisions, then the commitment has not earned release.

11. Knowledge Gain from Disturbance

If disturbance is required, the disturbance must produce decision-relevant information that could not be obtained through less commitment-bearing means.

This prevents an infrastructure act from being mislabeled as investigation.

A bounded verification action should be designed around learning. Its geometry, duration, reversibility, support requirements, and stopping conditions should remain tied to knowledge gain rather than development momentum.

12. Preservation of Future Measurement Conditions

The action must not irreversibly destroy, obscure, or bias the evidence needed for later decisions.

This is critical on the Moon because disturbance can change thermal conditions, expose material, modify surface structure, mix layers, alter access, and shift the meaning of later observations.

A commitment that damages future interpretability may be inadmissible even if it produces some immediate information.

13. Ignorance Dominance Condition

The Ignorance Dominance Condition defines when unresolved uncertainty controls the decision.

A commitment is inadmissible if multiple plausible subsurface states remain consistent with current evidence and produce materially different outcomes.

The threshold is not uncertainty magnitude.

The threshold is outcome sensitivity.

A decision can tolerate large uncertainty if the action remains robust across plausible states. A decision can fail under smaller uncertainty if the unresolved state determines whether the action should occur at all.

In lunar ISRU, this distinction is central.

A volatile indication may be real, but the operational implications may diverge. The same signal may be consistent with a coherent deposit, a patchy resource, a diffuse signal, a mechanically inaccessible concentration, a layer that is unstable under intervention, or a configuration that cannot support the proposed dependency.

Each state implies a different relationship between exploration, infrastructure, and risk.

If one plausible state supports commitment and another invalidates it, authority should not be released merely because the favorable state is attractive.

Material divergence includes loss of resource accessibility, failure of assumed ground behavior, infeasibility of infrastructure placement or operation, stranded assets, irreversible misallocation of power or access systems, and emergence of failure modes that cannot be recovered within mission constraints.

The relevant question is:

Would the decision change if the subsurface reality differed within the range still allowed by current evidence?

If yes, the uncertainty is decision-dominant.

Under decision-dominant uncertainty, commitment must defer to measurement.

14. Architecture Independence Test

The Architecture Independence Test asks whether an action can remain genuinely exploratory, or whether it requires or induces architecture formation.

An action should be treated as commitment-bearing if it requires or induces fixed power emplacement, repeated landing at a fixed site, corridor formation, logistics standardization, communications or navigation anchoring, excavation geometry, surface support dependencies, resource-processing dependency, or institutional precedent for continued use.

The test matters because early exploration can begin to create the architecture it claims only to inform.

A rover traverse may remain exploratory if it preserves site optionality and avoids durable dependency. But if repeated traverses establish preferred routes, support assets, communications relays, or future payload assumptions, the traverse may begin to form a corridor.

A drilling activity may remain bounded verification if it is limited, reversible in consequence, and designed to answer a decision-dominant question. But if the drilling location becomes the assumed center for excavation, power placement, and resource-processing architecture, it has crossed into commitment.

A power asset may be a temporary support element if it remains mobile, bounded, and independent of future site selection. But if it fixes operational geometry, privileges one location, or becomes necessary for subsequent activities, it is architecture-forming.

The Architecture Independence Test prevents premature conversion of investigation into development.

If an action cannot proceed without shaping future architecture, it requires commitment admissibility review.

15. Coupled Thresholds

Lunar development crosses thresholds gradually, but the consequences can accumulate quickly.

The first threshold is reversible investigation.

At this stage, activity is designed to increase knowledge while preserving alternatives. Remote sensing, broad scouting, non-hardening site comparison, and limited observation belong here. The purpose is to improve the evidence state without allowing operational convenience to become commitment.

The second threshold is resource verification.

At this stage, learning begins to require contact, sampling, excavation, drilling, or disturbance. This threshold is unavoidable in many subsurface problems. The Moon cannot be understood fully from orbit. But the need for disturbance does not remove the need for governance.

The third threshold is bounded verification.

Here, intervention is allowed only within limits. The action must be tied to a decision-relevant question. It must have a defined scope. It must avoid permanent support dependency. It must preserve the ability to compare sites. It must avoid becoming the first step of unexamined infrastructure formation.

The fourth threshold is architecture commitment.

At this stage, support systems begin to constrain future pathways. Power, access, mobility, communications, logistics, and surface modification organize around one assumed reality. This is where optionality can begin to collapse.

The fifth threshold is commitment-bearing exploration.

At this stage, site hardening, infrastructure placement, corridor formation, excavation, and precedent formation cause exploration to determine development. The system begins to inherit prior assumptions, not merely test them.

The danger is not that these thresholds exist. Lunar development requires progression.

The danger is crossing them without knowing which threshold has been crossed.

16. Decision Gate

Each proposed lunar surface action should resolve to one of three determinations.

17. Proceed

Proceed means the proposed action is admissible under current evidence.

The action is robust across materially plausible subsurface states. Remaining uncertainty no longer has the power to change whether the action should occur. The action does not create unacceptable lock-in, or the lock-in is justified by evidence adequate to the burden of the commitment.

Proceed does not mean risk is absent.

It means unresolved uncertainty no longer dominates the decision.

18. Defer

Defer means the proposed action is not yet admissible, but the decision may change with additional evidence.

This is the most important determination for lunar development.

Defer is not opposition to ISRU. It is not a rejection of exploration. It is the recognition that current evidence supports further measurement before commitment-bearing action proceeds.

A Defer determination should identify what evidence would matter, what measurement pathways are admissible, and which actions remain reversible.

The purpose of Defer is to preserve optionality while improving the evidence state.

19. Refuse

Refuse means the proposed action is inadmissible.

This determination applies when uncertainty is decision-dominant, outcomes diverge across plausible states, or the action would cross an irreversible threshold without sufficient evidence.

Refuse also applies when the action would destroy critical measurement conditions, create unrecoverable lock-in, or require infrastructure hardening before the relevant state space can be resolved.

Refuse is a valid outcome.

In frontier environments, value is often preserved by commitments that do not proceed.

20. Lunar South Pole Context

The lunar south pole is a concentrated example of the commitment problem.

The region is compelling because several conditions overlap. Permanently shadowed regions may preserve volatiles. Nearby illuminated zones may support power. The terrain creates both opportunity and operational constraint. Resource interest, access difficulty, thermal extremes, communications geometry, and infrastructure need are tightly coupled.

That coupling is exactly why the region requires commitment discipline.

Volatile deposits are inferred through indirect sensing but remain structurally unresolved at the scale of many proposed operations. Hydrogen indications may point toward volatile-bearing material, but the distribution, concentration, physical state, continuity, accessibility, and mechanical context remain uncertain.

These unresolved variables are not secondary details. They can determine whether a site can support excavation, whether power should be placed nearby, whether a corridor should be formed, whether a landing zone should be prepared, whether subsurface access is justified, and whether an ISRU dependency should enter architecture.

The lunar south pole also couples learning and commitment.

Verification may require disturbance. Disturbance may require power, mobility, landing support, communications, and operational staging. Those support systems can begin to harden a site before the underlying resource reality has been resolved.

This creates a structural trap.

To learn enough, a mission may need to act. But the action required for learning may begin to commit the system.

The correct response is not paralysis. It is bounded progression.

Investigation should continue. Measurement should improve. Disturbance may be justified. But the burden of evidence must match the burden of action.

Reversible investigation should be protected.

Bounded verification should be allowed where it reduces decision-dominant uncertainty without creating premature dependency.

Infrastructure commitment should wait until the evidence is sufficient for the decision it would authorize.

21. Worked Decision: Lunar Volatile Signal to Infrastructure Anchor

21.1. Decision Under Consideration

Should a lunar south polar site be allowed to become the primary infrastructure and ISRU architecture anchor based on current volatile evidence?

The relevant question is not whether there is ice.

The relevant question is whether volatile signals should be allowed to govern site selection, power placement, logistics routing, excavation planning, and long-horizon infrastructure dependency.

21.2. Current Signal

Current evidence supports exploration interest.

Hydrogen indications suggest volatile-bearing materials may be present. Permanently shadowed regions provide a plausible preservation environment. Thermal conditions support retention. Radar and reflectance data provide indirect constraints. Terrain and illumination data help define access and operational context.

These signals narrow the problem.

They do not necessarily collapse it.

21.3. Remaining Ambiguity

The same signal can remain consistent with several operational realities.

The site may contain an architecture-grade deposit. It may contain a patchy but usable resource. It may contain diffuse or inaccessible volatile-bearing material. It may contain material that is present but mechanically difficult to extract. It may contain a signal that is locally real but insufficient to support dependency.

Each of these states has different implications for infrastructure.

Some may support commitment. Others would invalidate it.

That divergence is the decision problem.

22. Determination

22.1. Defer.

Not because lunar ISRU is invalid.

Not because volatile exploration should stop.

Defer is warranted because current evidence may not yet justify allowing volatile assumptions to govern irreversible infrastructure commitments.

The admissible path is continued measurement, bounded verification, and site comparison that preserves optionality while reducing decision-dominant uncertainty.

23. What Would Change the Determination?

A Defer determination should not be vague. It should identify the evidence that could change the decision.

A shift toward Proceed would require evidence that reduces the ambiguity that currently controls the commitment.

That evidence would need to demonstrate spatial coherence at the intervention scale, concentration sufficient for the proposed dependency, physical accessibility of the resource, mechanical stability under expected operations, compatibility between resource location and power or access constraints, and limited disturbance effects on future measurement.

It would also need to show that the proposed action does not prematurely harden landing zones, corridors, power placement, or support architecture before the relevant state space has been reduced.

Proceed becomes available only when the proposed action survives the materially plausible subsurface states relevant to the commitment.

Refuse becomes necessary if the action requires infrastructure hardening before the evidence can resolve the state space, or if learning would require crossing a threshold that makes reversal infeasible.

This is the central discipline.

Do not allow the need to learn to become an excuse for unmanaged commitment.

24. Final Statement

The central problem in lunar exploration is not uncertainty.

It is the timing of path-dependent commitment under unresolved uncertainty.

A system is not admissible simply because it can be made to work.

A site is not admissible simply because it is promising.

A volatile signal is not admissible simply because it is real.

In frontier environments, the cost of error is not ordinary iteration. It is entrenchment.

Minimum evidence before lunar commitment means that evidence must be adequate to the burden of the action being authorized.

Reversible investigation should proceed.

Bounded verification should be governed.

Irreversible commitment should wait until the evidence has earned the right to shape architecture.