Compliance · Technical Paper

The Constraints We Put in Orbit, Part III: The Governance Gap Is Filled by Discretion

Part III of III. The governance gap orbit creates is filled by nothing but operator discretion, and the analysts we'd use to check that discretion are paid by the same people whose discretion is in question.

Jennifer McKinney · August 15, 2026 · 10 min read

This is Part III of a three-part series on what orbital AI infrastructure actually escapes, and what it doesn't. Part I examined the heat problem; Part II examined jurisdiction. This part turns to governance and trust.

Movement III: The Governance Gap Is Filled by Discretion, Assessed by the Captured

From "they cannot be reached" to "they cannot be checked"

The first movement located a constraint inside physics. The second located one inside the seams between legal regimes. This one locates the last and worst one inside the institutions we rely on to watch the powerful, and it follows directly from the registry-shopping problem the jurisdiction movement opened.

Once you see that the law governing an orbital asset is the law of the registry state, you can predict the move without speculation, because the maritime world already ran the experiment for a century. Flags of convenience are not hypothetical; they are the dominant structure of global shipping, and they exist precisely because the law that governs the vessel is that of the flag state (UNCTAD Review of Maritime Transport 2024). Space law copied that architecture deliberately (POWER Magazine). So predicting that entities will migrate, newly registered companies and existing ones alike, toward permissive registries and toward states without extradition treaties is not a leap. It is noticing that the same incentive is bolted onto the same legal chassis. The extradition angle sharpens it: the move is not only about which regulator governs the *asset*, it is about whether the humans making the decisions can be *reached*. Those are two different shields, and sophisticated actors stack them. That part of the argument would stand without hedging.

Three mechanisms that look like one

But there is a flattening worth resisting, because it costs analytic precision. The intuition that participation is "gated only by how much you can buy somebody off" bundles three distinct mechanisms that all *look* like buying your way in and are not the same thing. It's easy to go in circles if these get argued as one, so naming them apart is the point.

The first is **capture and corruption** proper, paying or revolving-door-ing your way to favorable treatment. Real, but the least interesting, because it is the most illegal and the most catchable. If this were the only mechanism, the fix would be enforcement.

The second is **legal regime-shopping**, which requires bribing no one. Choosing a permissive registry, structuring through a favorable jurisdiction, exploiting the gap between the data-subject theory and the flag-state theory from the last movement, all of it available to anyone with good lawyers and zero willingness to break a law. This is the more corrosive mechanism precisely because it is clean. You cannot prosecute your way out of it. The arbitrage is the system working as designed, not failing. The fix, if there is one, is harmonization, closing the gaps, not enforcement.

The third is the **capital moat** itself. Even with no corruption and no regime-shopping, the entry ticket to frontier AI plus orbital infrastructure is now measured in the billions, which means the field self-selects down to a handful of actors before ethics ever enters the room. That is not someone being bought off; it is the price of admission excluding everyone but the already-powerful, producing the same outcome, a tiny set of unaccountable players, through an entirely different door. Neither enforcement nor harmonization touches this one; it needs something structural.

The reason the distinction matters: if you misdiagnose which mechanism is operating, you reach for a fix that cannot work. Conflating them is how a serious concern gets dismissed as a grievance.

The trust claim, made structural so it can't be waved away

Now the trust point, engaged rather than dodged, but with a guardrail I want to hold myself to. The empirical observation that the leading actors have repeatedly moved faster than oversight, treating permission as something to seek forgiveness for, is fair and documentable in the aggregate. Where this argument should not go is the slide from "their actions have not earned trust" to "they will never be trusted to do the right thing." The first is an evidence-based read of a pattern. The second is a prediction about future character that, written down, does two things you do not want: it is unfalsifiable, and it converts an analytic claim into a grievance, handing anyone who disagrees an easy reason to discard the rigorous parts.

The stronger move, the one consistent with the whole arc, is not to argue that the actors are bad people. It is to argue that *the structure makes good behavior optional*. We have built a system in which doing the right thing is unenforced and unrewarded, so we are relying entirely on the voluntary virtue of whoever happens to be holding the asset. That claim is far harder to wave away, because it is true regardless of whether any given founder is sincere. Sincerity is not a governance mechanism. A system that requires it is already broken. No one needs to be a villain for the structure to fail; you only need the structure to permit failure and trust that it will not occur.

The watchdog shares the paymaster

And here is the part that makes the trust problem worse than "the operators can't be trusted to self-govern," the part that completes the pattern across all three movements.

The apparatus society normally uses to *independently assess* whether the powerful are behaving, the lawyers, the analysts, the financial press, is, in this domain, financially downstream of the same actors. Recall the digital-flag-state proposal from the last movement: it ran in the financial press and argued for shrinking the applicable-law problem down to a single permissive regime on the explicit ground that compliance friction would cripple the industry (Bloomberg Law). That is the interpretive layer producing the answer operators prefer. And the reason is not individual bad faith; it is that the analytic class generating most of this commentary, law firms with aerospace and finance clients, banks underwriting the launches, the financial outlets whose readership *is* the capital, does not contain a disinterested party. Their objectivity is not compromised by anyone being corrupt. It is compromised by who pays them and what answer preserves the relationship. A reviewer whose firm bills the operator is not going to author the analysis concluding the whole model of the operator is jurisdictionally unworkable, even if that is the true answer, because that analysis has no client.

This argument should stay at the level of structural incentive rather than individual incapacity, because the structural version is both more defensible and more true. There almost certainly exist individuals at those institutions who see the problem clearly and say so privately. The claim that survives scrutiny is not that they are personally incapable of objectivity; it is that the *system filters for the analyses that serve the capital* and suppresses the ones that do not, so what reaches publication is selected regardless of any individual's integrity. The analysis an honest banker might write simply does not get commissioned, or does not get published, or does not get read. Same outcome, sturdier claim, and it cannot be rebutted by pointing to one honest banker. Follow the money is an analytic method, not an accusation.

And here is the test of that claim, the thing that distinguishes a structural diagnosis from a conspiracy theory: the impartial analysis *does* exist, it simply does not sit in the forefront, and the reason it does not is visibility, not merit. the sovereignty-and-equity framing advanced by Arora in Part II is published, peer-respected, and institutionally unconflicted, and yet it surfaces in a nonprofit outlet covering technology outside the West rather than in the venues that shape policy conversation. The independent corroborations of the relocation thesis come from a generalist long-form analyst and an energy-trade space-law commentator, not from the aerospace practice groups whose names carry weight in a regulatory docket. The pattern is not that the unconflicted analysis was refuted; it is that it was never amplified, because no one with a budget had a reason to push it to the front. That is the prediction the structural account makes, and it holds: the captured voices are loud and the unconflicted ones are quiet, and the asymmetry tracks funding rather than rigor. Which is, not incidentally, exactly why doing this analysis without a client to please has a job to do here.

Where the trust lands

The governance gap, taken together, breaks into three interlocking failures: orbit lets the decision-makers relocate beyond easy reach; the gap left behind is filled by nothing but their discretion; and the expertise available to assess that discretion is paid by the same people whose discretion is in question. The watchdog and the watched share a paymaster. That is a more serious indictment than untrustworthy founders, because it means the mechanism society uses to *check* the founders is captured at the source.

This is the third instance of the one pattern. The thermal movement: physics is the constraint the pitch decks wave away, relocated from utility bill to natural law. The jurisdictional movement: sovereignty doesn't subtract in orbit, it stacks, and the seams are where accountability dies, relocated from one regulator to an unenforceable tangle. The institutional movement: the governance gap is filled by discretion, and the interpretive layer is captured at the source, relocated from independent oversight to paid endorsement. Three constraints, one shape. Orbital AI does not escape Earth's constraints, physical, legal, institutional. It relocates them to where they are harder to see and harder to enforce, and calls the relocation progress.

Coda: the honest version of hope

This essay will not end on false comfort, but one calibration is worth offering, because cynicism is as lazy as optimism and this piece should not model either.

The EU AI Act and the recent wave of state-level AI laws are evidence that governance does sometimes catch up (EU AI Act), usually late, usually after harm, but it catches up. So the honest position is neither "the leaders will do the right thing" nor "they never will." It is that the outcome is genuinely unsettled, and that it depends in part on whether people who can see the structure clearly bother to describe it before the window the practitioners estimate, that 18-to-24 months, closes. Which is, not incidentally, the thing this piece is trying to do. A description is not a fix. But you cannot fix what no one has named, and right now the naming is mostly being done by people the operators pay. That a clearer naming might come from somewhere else, a paper written by someone with no client to please, is the modest and defensible version of being pleasantly surprised. The surprise is not that the powerful will suddenly be trustworthy. It is that the structure might yet be described accurately by someone who is not being paid to misdescribe it, and that accurate description is the precondition for everything else.

This concludes the three-part series. Start over at Part I: The Heat Problem, or revisit Part II: Sovereignty Doesn't Subtract in Orbit, It Stacks.

Sources

- European Data Protection Board
- Bloomberg Law
- POWER Magazine
- UNCTAD Review of Maritime Transport 2024
- EU AI Act