

Artemis II, NASA, and the problem with a risk number

David Slater March 24, 2026



Stephen Clark's recent *Ars Technica* piece raises a very interesting issue for anyone who works seriously with risk, assurance, and complex systems.

What matters here is not simply that NASA would not give a number for Artemis II risk, but why they were reluctant to do so.

Artemis II is not a routine transport mission. It is a developmental mission, using a launch system and spacecraft combination with only one prior integrated flight, on a lunar trajectory far beyond the operational envelope of low Earth orbit. NASA describes it as the first crewed flight of SLS and Orion, and an approximately 10-day mission around the Moon and back. That matters, because in a mission like this a single bottom-line probability can easily appear more precise, and more authoritative, than the underlying knowledge really justifies.

There are, broadly, two ways to read NASA's reluctance.

The first is the more charitable, and probably the more correct. The people responsible know perfectly well that probabilistic risk assessment can slip into pseudo-precision when the evidence base is thin and the architecture is still immature. If you have one uncrewed integrated mission, a body of subsystem evidence, many analytical models, and a substantial stack of assumptions, the final percentage is not a discovered truth. It is a constructed estimate. Useful, certainly. But only within limits.

That, I think, was the essence of Lori Glaze's point. Such numbers are often more useful comparatively than absolutely. They may help distinguish between options, or show whether a

design change has improved the position, without honestly supporting a statement such as “the chance of catastrophe is X percent.”

The second reading is organisational and political. NASA knows that once a number is spoken publicly it ceases to be a technical artefact and becomes a headline. It becomes something Congress can weaponize, critics can amplify, managers can be trapped by, and the press can simplify beyond usefulness. A public number also tends to compress the discussion into a single scalar, when the real substance of risk lies elsewhere: in the distribution of uncertainties, the quality of evidence, the strength of barriers, the degree of test coverage, and the residual assumptions still being carried.

So yes, sidestepping the question may look evasive. But it is not irrational.

The recent technical issues make this plain. NASA has had the helium-flow problem to the upper stage, and seal work associated with the core-stage oxygen feed line, alongside continuing attention to Orion’s re-entry and heat-shield questions. None of these are especially meaningful because they nudge some abstract overall number up or down. They matter because they show where the system is still revealing operational sensitivity and where engineering confidence must be earned case by case.

The decision not to run another full fuelling test is a good example. That is not a question that can be settled by waving a risk percentage around. It is an operational judgement under uncertainty. NASA is effectively saying that the successful February wet dress rehearsal, together with the corrective work since then, provides enough confidence that another full test is not worth the delay, wear, schedule disruption, and possible introduction of further complications. That may be right. It may even be the best available judgement. But the important point is that the reasoning matters more than the number.

John Honeycutt’s remarks were, to me, the most revealing. His reference to “failure of imagination” goes to the heart of the matter. In complex, tightly coupled systems, catastrophe often does not arise because somebody forgot to assign a probability. It arises because the system enters a condition, interaction, or pathway that was not properly imagined, tested, or governed. That is a very different problem.

This is why NASA appears caught between two incompatible expectations. On the one hand, the public quite reasonably expects that if astronauts are being sent around the Moon, the agency should be able to say what the risk is. On the other hand, if the mission is genuinely novel, the agency ought also to admit that its numerical confidence has limits. NASA seems to be choosing the second position, but not stating it clearly enough. That is why the answer sounds defensive when the underlying logic is actually quite sound.

A more credible formulation would be something like this: *we do have a probabilistic assessment, but because Artemis II is a developmental mission with limited integrated flight heritage, that figure should not be treated as a precise prediction. What matters more is where the major hazards lie, what evidence supports the mitigations, and which uncertainties remain knowingly accepted.*

That would be a more adult answer.

The wider lesson is not just about NASA. It is about the misuse of quantified risk in first-of-kind, high-consequence ventures. In such settings, risk is not well represented by a single number

pretending to be stable and objective. It is better understood as a structured picture of uncertainties, dependencies, barriers, assumptions, and blind spots.

On that basis, NASA's hesitation makes sense.

But if it wants trust, it cannot merely dodge the question. It has to explain why the number is epistemically fragile, where the principal risk concentrations lie, and what evidence justifies proceeding anyway.

That, to me, is the real story.

Reference

Stephen Clark, "NASA officials sidestepped questions on Artemis II risks—there's a reason why," *Ars Technica*, 14 March 2026.