# When We Mistake Solutions for Objectives

## What Mars Sample Return Reveals About Systems Thinking, Institutional Lock-In, and the Danger of Optimizing the Wrong Question

By Erinn van Wynsberghe
May 29, 2026

NASA's Mars Sample Return mission has long been regarded as one of the most ambitious scientific undertakings of the modern era. The concept was elegant in its simplicity: collect samples from Mars, launch them back into space, rendezvous with another spacecraft, return them to Earth, and allow scientists to study them using the most advanced laboratories ever built.

The mission promised extraordinary scientific value. Yet it also raises an interesting question. What if returning Martian rocks to Earth was never the real objective? At first glance, that may sound absurd. The mission was called ***Mars Sample Return***, after all. Returning samples was the entire point. Or was it? The answer depends on whether we are thinking like project managers or systems engineers.

**The Hidden Error: Baking the Solution Into the Problem**
One of the most common mistakes made by organizations, governments, and institutions is that they inadvertently embed a preferred solution into the problem itself. Once that happens, the solution becomes invisible. Consider two different ways of framing the Mars challenge.

The first framing asks:
> How do we return samples from Mars to Earth?

The second asks:
> How do we maximize scientific knowledge from Martian samples?

The difference may appear subtle, but it fundamentally changes the problem. In the first framing, returning samples becomes the objective. In the second framing, returning samples becomes only one possible solution. The solution space suddenly becomes much larger.

This distinction is not unique to space exploration. It appears repeatedly in public policy, infrastructure planning, defence procurement, climate initiatives, healthcare systems, and organizational strategy. The moment we stop asking what outcome we are trying to achieve and start asking how to implement a chosen approach, we risk optimizing a solution rather than achieving the objective.

**A Fair Counterargument**
To be clear, Mars Sample Return may have served multiple objectives beyond scientific analysis alone. Returning samples could contribute to future exploration architectures, demonstrate technological capabilities, support international partnerships, satisfy planetary protection requirements, or help prepare for eventual human missions to Mars.

The purpose of this article is not to argue that Mars Sample Return was necessarily the wrong choice. The purpose is to illustrate a broader principle: organizations should periodically revisit their underlying

objectives to ensure that a chosen solution remains the best path to achieving them. Whether Mars Sample Return was the correct decision is ultimately less important than the lesson it offers about institutional decision-making.

**Objective Drift**
I call this phenomenon “objective drift”. It occurs when an organization gradually replaces its original goal with the solution that was initially chosen to achieve it.

Over time, discussions shift from:
> What are we trying to accomplish?

to:
> How do we make the current plan succeed?

The distinction seems minor, but it can fundamentally alter decision-making. When objective drift takes hold, alternative approaches become increasingly difficult to see. Entire solution spaces may disappear from consideration, not because they were evaluated and rejected, but because they were never revisited in the first place.

**The Systems Engineering Lens**
Systems engineers are trained to begin with objectives rather than architectures.

Before selecting a solution, they ask a series of questions:
*What are we trying to achieve?*
*Why does that objective matter?*
*What outcomes are we seeking?*
*What constraints exist?*
*What alternative approaches could achieve the same outcome?*

Only after these questions are answered should specific architectures be considered. Applying this thinking to Mars Sample Return produces an interesting result. If the true objective is to generate scientific knowledge from Martian material, then returning samples to Earth is not necessarily the only path. It may not even be the best path. The key point is not whether a different approach is superior. The key point is that alternative approaches should remain visible and available for consideration.

**Alternative: Send the Laboratory, Not the Rocks**
Suppose we revisit the problem from first principles. Rather than asking how to bring samples home, imagine asking how to maximize scientific insight per dollar spent. One possible answer would be to send a sophisticated robotic laboratory to Mars itself. Instead of launching a small cache of rocks back to Earth, a future mission might deliver extensive analytical capabilities directly to the Martian surface. Samples could be collected, processed, examined, and analyzed locally. Only the resulting data would need to be transmitted back to Earth. The mission architecture would be entirely different. The scientific objective, however, would remain the same.

This approach might offer advantages. It could reduce mission complexity. It could eliminate entire stages of the return process. It could allow continuous analysis rather than one-time examination. It could potentially scale over time through additional laboratories.

Perhaps it would prove impractical. Perhaps returning samples would still be the better choice. That is not the point. The point is that the alternative should remain visible until a rigorous comparison is performed. Once a solution becomes embedded in the problem definition, competing ideas often disappear before they are ever evaluated.

**Why Institutions Become Locked In**

If alternative approaches exist, why do organizations continue pursuing legacy architectures? The answer is rarely incompetence. More often, it is the natural consequence of institutional incentives.

Once a large initiative is established, many things begin to form around it: budgets, planning documents, review committees, contractors, career paths, political commitments, public expectations. Over time, changing direction becomes increasingly difficult. Even when technology evolves. Even when costs change. Even when new possibilities emerge. Organizations begin optimizing the chosen architecture rather than periodically reassessing the original objective. This phenomenon appears across virtually every sector. It is one of the most common forms of institutional lock-in.

**Monuments Versus Multipliers**

A related pattern appears in public investment. Institutions often gravitate toward highly visible flagship projects. These projects are easy to communicate. They generate headlines. They create symbolic achievements. They provide clear milestones.

Less visible investments often struggle for attention, even when they generate greater long-term value. A flagship mission can become a monument. An enabling capability can become a multiplier. The distinction matters. Monuments demonstrate achievement. Multipliers increase future capability. Both have value. But institutions frequently underinvest in multipliers because the benefits are less visible and often arrive gradually. This pattern can be observed not only in space exploration, but also in infrastructure, education, healthcare, climate policy, scientific research, and economic development.

**The Broader Policy Lesson**

The most important lesson from Mars Sample Return has little to do with Mars. It concerns how organizations make decisions.

Before committing significant resources to any major initiative, leaders should periodically ask:
*What is the actual objective?*
*What public value are we trying to create?*
*Have we mistaken a solution for the goal?*
*What alternatives have been excluded from consideration?*
*Are we optimizing the mission, or merely optimizing a legacy architecture?*

These questions are uncomfortable. They challenge assumptions. They reopen debates that many stakeholders would prefer to consider settled. Yet they are often the questions that produce the greatest breakthroughs.

**The Wrong Question Problem**

One of the most common institutional failures is not choosing the wrong solution. It is forgetting the original objective. When organizations stop asking “why? and focus exclusively on “how”, they gradually optimize the chosen answer rather than the desired outcome. The result is often an organization that

becomes exceptionally good at executing a plan without ever revisiting whether the plan remains the best path to the objective.

Excellence in execution cannot compensate for ambiguity in purpose. Before asking “How do we solve this problem”, institutions must periodically return to a simpler question “What are we actually trying to achieve?” Whether we are designing climate policy, infrastructure systems, economic incentives, scientific programs, or future missions to Mars, that question may be the most important one we can ask. Only after answering it should we begin debating and designing the solution.