One Objective or None: Why NASA’s Faster, Better, Cheaper Program Was Doomed

In the early 1990s, NASA Administrator Daniel Goldin championed a sweeping reform of how the agency would conduct space exploration, encapsulated in the slogan: *Faster, Better, Cheaper*. Missions would be developed more quickly and on smaller budgets, without sacrificing quality. Between 1992 and 2000, NASA launched dozens of missions under this philosophy, including early successes such as Mars Pathfinder. But the program collapsed amid high-profile failures, most famously the back-to-back losses of the Mars Climate Orbiter and Mars Polar Lander in 1999.

The standard explanations—organizational culture, inadequate testing, political pressure—are all real. But beneath them lies a deeper problem from optimization theory: *a system can be optimized only against a single objective function.* If you care about multiple things, you must fuse them into a single composite goal—something like “successful mission” or “return on investment”—that the organization is actually trying to maximize. FBC did not do this. It named three parameters and elevated all of them to the status of objectives, and that framing was mathematically incoherent from the start.

The Math

An optimization problem needs an objective function: the rule that lets you compare two candidate solutions and decide which is better. That function must be a single scalar quantity (Martins and Ning 2022). Try to optimize against two axes at once—maximize speed *and* minimize cost—and you immediately hit pairs where one option is faster but more expensive and the other slower but cheaper. The framework has nothing to say about which to pick. What you have is not an optimization problem; it is a values problem masquerading as one.

The fix is to combine what you care about into one composite. *Successful science return per dollar*, for example, fuses cost and quality into a single ratio. The trade-offs still exist, but they become *internal* to the objective function—weighted explicitly—rather than left to ad hoc judgment. A team optimizing for that ratio knows immediately that a cheap, fast mission that fails at Mars scores zero. A team told only “faster, better, cheaper” has no such guidance.

Why Single Objective Beats “Pick Two”

A common response to this argument is that the project management triangle already solves it: *fast, good, cheap—pick two.* This is a step in the right direction but does not go far enough. Pick-two tells you which corner of the triangle to face. It does not tell you how far to go along that edge. “Faster and better, accept higher cost” still leaves you choosing among infinitely many points, ranging from a modest speed-and-quality gain at modest extra cost to a dramatic gain at ruinous cost. Pick-two narrows the menu. It does not pick from it.

A single objective function does. Once you commit to maximizing, say, *expected science return per dollar*, every point in the decision space—including those on the Pareto frontier—gets a single score, and the best one wins. The frontier is no longer a menu you must choose from with outside judgment. It becomes the byproduct of an optimization that already has the answer.

This matters in practice in several ways. It makes trade-offs comparable: a project lead facing a real decision—“should I cut two weeks of integration testing to hit the launch window?”—has a clear question with a clear answer rather than a values question to escalate. It exposes hidden disagreements early: when a team says “we want all three” and an executive says “we want all three,” they sound aligned but rarely are. Forcing a conversation about the composite objective surfaces the disagreement while it is still cheap to resolve. It survives delegation: a project of any size involves hundreds of small decisions made by people who will never be in the room with senior leadership, and a single objective gives them a rule they can apply without checking.

Most importantly, it defends against the silent bill-payer problem. Pick-two does not tell you what to do when one of your “picked” parameters starts taking unacceptable damage from pressure on the other. NASA effectively picked faster and cheaper, and quality—nominally not on the list but obviously still required—silently absorbed all the slack until two spacecraft were destroyed. A single composite objective with quality folded into it makes that absorption mathematically visible. You cannot quietly trade away something that is in the numerator of the objective you are trying to maximize.

The honest limitation is that the single-objective approach forces you to do something hard: actually decide what you care about and how to weight it. Pick-two lets you avoid that decision. This is why pick-two is more popular, and why it persists despite being analytically weaker. The trade-offs do not disappear because you decline to define them. They get made anyway, by whoever has the most leverage in the moment, and the cost lands wherever it lands. Pick-two is a description of the problem. A single objective function is an actual solution to it.

What FBC Asked For Instead

FBC named three parameters and demanded improvement across all of them, without ever defining the actual objective—and without even committing to a pick-two. Had the goal been stated as *successful science return per dollar*, “faster” and “cheaper” would have been understood as inputs to be balanced against mission success, not pursued as ends in themselves. A failed mission would have represented infinite cost per unit of science, immediately disqualifying any cost-cutting that materially increased failure risk. The objective function would have done the work of pushing back against reckless trimming. Instead, FBC promised three improvements as if they were independent dials that could all be turned up at once.

How It Failed

The Mars Climate Orbiter is the canonical case. The proximate cause was a unit conversion error—pound-seconds versus newton-seconds—that sent the spacecraft too close to Mars. The deeper question is why a problem this basic was not caught. Post-failure reviews documented that cross-team integration testing and independent review had been pared back as part of cost and schedule reductions. Had the agency been explicitly optimizing for science return per dollar, those reviews would have been recognized as cheap insurance against losing the entire investment. Instead, they looked like budget items to trim, and they were trimmed. Mars Polar Lander, lost the same year, told the same story: another routine failure mode that integration testing was designed to catch, and another set of tests cut without countervailing pressure to preserve them.

What Single-Objective Framing Looks Like

Apollo had a clear composite objective: land a man on the Moon and return him safely before the end of the decade. Schedule and crew safety were built into the success criterion itself—a late landing or a dead astronaut both counted as failure. Cost was an input the program drew on as needed. SpaceX's reusable rocket program has run on a similar principle, with cost-per-kilogram-to-orbit as the de facto objective. Schedule slips and early test failures were tolerated because they served the real goal. In both cases, the organization committed to a single objective it sought to maximize and used that commitment to make consistent trade-offs.

Conclusion

FBC's failure was not that it picked the wrong objective; it was that it never picked one. The slogan flattered the agency and Congress by implying no choice was necessary—that all three parameters could improve indefinitely with sufficient cleverness. The mathematics says otherwise, and Mars proved it right. Any organization that names multiple parameters as goals without consolidating them into a single objective function is setting itself up for the same failure mode: trade-offs made implicitly, by the path of least resistance, with the cost landing on whichever parameter is hardest to measure until it is too late.

Reference

Martins, J., and A. Ning. 2022. *Engineering Design Optimization*. Cambridge University Press. https://mdobook.github.io.