Why Projects Overspend

Idea In Short

Treat every large project as a capital allocation decision before treating it as a delivery challenge. The evidence is clear: cost overruns usually begin long before construction starts or software ships. Leaders should slow the front end, test whether the proposed asset is the right answer to the real business need and force estimates through comparisons with similar projects rather than internal optimism. They should also appoint leaders who have already delivered the same kind of work, break scope into repeatable modules and define clear exit points before sunk costs take over. This article explains why large projects go over budget, where judgment fails and which practical disciplines improve the odds of delivering complex work on time, on budget and with the promised benefits

Large projects rarely go over budget for a single reason. They overspend because flawed assumptions, institutional incentives, technical complexity and delayed course correction reinforce one another over time. What looks like a delivery failure at the end often began as a decision failure at the beginning.

The pattern repeats

Large projects have produced the same warning signs for decades across transport, energy, construction, defense and information technology [I.T.]. A review of megaproject performance discussed by Bent Flyvbjerg and Dan Gardner draws on more than 16,000 projects in 136 countries and shows that 91.5 percent run over budget, behind schedule, or both1. The point is not that every major initiative is doomed. The point is that leaders should treat overrun risk as the default condition unless the project has been designed to resist it.

> There are three kinds of lies: lies, damned lies and cost projections for big projects

That line captures the credibility gap around early estimates. In the source article, the comparison ranges from California's high-speed rail project to nuclear construction and major information technology programs, all cases where the stated budget ceased to describe the actual exposure well before completion. Large initiatives fail in different sectors, but they fail through familiar mechanisms.

The project starts failing before delivery

One of the strongest arguments in the original piece is that most large projects do not become bad during execution. They start bad. Sponsors often define the solution before they have fully clarified the problem, which makes the estimate vulnerable from the first approval memo. If the real objective is to move people or goods between two points, a train may or may not be the best answer. When leaders skip that question, they lock capital into an asset before they have tested whether another design, route, operating model, or policy choice could achieve the same result more cheaply.

The same mistake appears in corporate settings. A company may approve an enterprise platform, a data program, or an operating model redesign because the proposal sounds modern and strategic. Yet the budget may hide unresolved requirements, uncertain interfaces and a weak change plan. Those omissions do not stay small. They resurface later as redesign, vendor disputes, governance churn and requests for additional funding.

Research summarized by the Project Management Institute [P.M.I.] makes the same point in different language:

> misinformation about costs, schedules, benefits and risks undermines project viability, while complexity and unplanned events go unaccounted for in the budget2

That is why the budget overrun often reflects not one shock but an estimate that never represented the full job.

Optimism and incentives distort estimates

Psychology matters, but politics matters just as much. Some teams underestimate cost because they genuinely believe they can outperform precedent. Others do it because realistic numbers would make approval harder. Bent Flyvbjerg has long described this pattern as a mix of optimism bias and strategic misrepresentation, a combination that rewards attractive business cases more than accurate ones3. In practice, the board sees an investable story while the delivery team inherits an unstable baseline.

This dynamic explains why overspending can persist even when smart people manage the work. The estimate may already embed political bargaining, optimistic productivity assumptions, understated contingency and fragile dependencies on permits, suppliers, or stakeholder alignment. Once approval is won, the organization shifts from persuasion to delivery. It then discovers that the number used to authorize the project was never the number required to complete it.

A discussion of sunk costs and tactical underestimation is warranted. These are analogous to making a hole so large that funding must eventually follow. That is not simply colorful rhetoric. It describes a governance failure in which commitment outruns evidence and later choices become more expensive because earlier choices were not challenged hard enough.

Complexity multiplies cost

Large projects are not just scaled-up small projects. They contain more interfaces, more contractual boundaries, more regulatory touchpoints and more ways for one delay to trigger another. The longer the execution period, the more opportunity there is for inflation, scope change, labor shortages, policy shifts, technology obsolescence and leadership turnover to damage the economics. Complexity does not guarantee failure, but it widens the path through which ordinary variance becomes major overrun.

This is especially visible in information technology projects. The source article notes that 18 percent of information technology projects suffer real cost overruns above 50 percent and among those cases the average overrun reaches 447 percent. That distribution matters because it shows a fat-tail risk profile. Many projects may appear manageable for a long time, then a subset deteriorates fast enough to erase the performance of the rest.

A review in The Wall Street Journal also points to inadequate planning and prolonged execution as common drivers of failure, reinforcing the idea that overruns are shaped early

and amplified later. For executives, that means average numbers can be deceptive. What damages the portfolio is not only the median case but the minority of projects that spiral into catastrophic rework, delay and recapitalization.

Why some sectors fail harder

There are three especially troubled categories:

1. high-speed rail
2. nuclear power construction, and
3. information technology

They share a difficult mix of novelty, political visibility, dense stakeholder environments and hard-to-reverse capital commitments. In each case, teams face pressure to commit early while the design, interfaces and external dependencies still carry major uncertainty.

That combination helps explain why some projects attract public confidence at launch and still deteriorate later. High-profile initiatives often begin with symbolic force. They signal ambition, modernity, or national capacity. But symbolism does not reduce engineering difficulty, procurement risk, or integration complexity. In fact, visibility can worsen estimates because sponsors gain more from announcing an attractive number than from presenting a painful one.

A Milken Institute Review discussion of Flyvbjerg and Gardner's work argues that the mindset of political and corporate leaders who launch major undertakings without careful study is often the first suspect in overruns4. That explanation fits both public infrastructure and private transformation programs. The incentive to begin is usually stronger than the incentive to challenge the plan thoroughly.

Planning is a leadership task

The source article's most useful advice is not technical. It is managerial. Leaders need to know where they are going before they decide how to build. That sounds obvious, yet many programs still confuse the asset with the outcome. The train is not the goal. Moving people and freight is the goal. The new software stack is not the goal. Better decisions, lower operating friction, or faster customer response may be the goal.

Once that distinction is clear, planning improves. Teams can compare alternatives, test whether scope supports the stated outcome and cut prestige features that inflate cost without improving the result. This is where reference-class thinking becomes valuable. Instead of asking only what this team hopes to accomplish, leaders ask what comparable projects actually achieved and at what price. Oxford Saïd reporting on Flyvbjerg's work and related reviews repeatedly emphasize comparison with precedent because uniqueness is often overstated even when projects claim to be exceptional.

The discipline here is uncomfortable because it constrains executive storytelling. It replaces aspiration with evidence. Yet without that constraint, budgets become expressions of intent rather than instruments of control.

Think slow, act fast

Flyvbjerg and Gardner's phrase "think slow, act fast" deserves attention because it reverses how troubled projects often behave. Weak projects decide quickly and execute slowly. Strong projects decide carefully and execute with momentum. The difference lies in the quality of definition before major resources are committed.

Thinking slow does not mean indecision. It means pressure-testing the estimate, sequencing permits and procurement realistically, resolving design choices early and validating the operating model that must exist after launch. It also means hiring people who have completed the same kind of work before. The person in charge should have real experience with that exact category of project, not simply seniority or prestige. That point is easy to ignore and expensive to relearn.

When leaders skip this work, the organization pays later through claims, redesign, integration fixes, schedule compression and emergency governance. Those costs often appear operational, but they are the bill for poor strategic preparation.

The Lego principle

Another durable idea from the source text is the "Lego principle":

> break the project into smaller, repeatable components rather than treating it as one giant bespoke undertaking

Modularity improves learning. Each repeated element sharpens forecasting, shortens cycle time and exposes defects before they spread across the whole system.

The contrast between the Empire State Building and the Sydney Opera House serves this point well. The first benefited from disciplined sequencing and repeatable work packages, while the second became a symbol of architectural achievement joined to managerial breakdown. The anecdote matters because most organizations still reward the rhetoric of uniqueness. Yet uniqueness is usually the enemy of reliable cost control.

This principle also travels well into technology and corporate transformation. Modular releases, standard interfaces, repeatable deployment patterns and staged migration reduce the chance that one mistake will contaminate the whole program. Leaders do not remove risk by making work smaller in a cosmetic sense. They remove risk by creating units of work that can be estimated, tested and learned from before full-scale commitment.

Sunk costs trap decision-makers

Few forces are as damaging to project economics as the refusal to revisit a weak commitment. Once a project consumes money, political capital, executive reputation and public attention, cancellation starts to look like failure. That is precisely when rational review becomes harder. Teams keep funding the initiative not because the future case is strong, but because the past cost feels too large to abandon.

This is why the governance model matters as much as the delivery model. Boards and steering committees should define explicit points at which the project can be rescoped, paused, or stopped without stigma. Otherwise every review meeting becomes an exercise in justifying prior decisions. The project then loses the one protection that could still save capital:

the ability to admit that the initial premise was wrong

An Independent Institute review of Flyvbjerg and Gardner's book underscores that strategic misrepresentation and the desire to win approval often shape major projects before execution begins, while poor planning leaves delivery teams to absorb the consequences5. Once sunk costs enter, even obvious warning signs may no longer trigger action.

What executives should do differently

Large projects will always involve uncertainty and no process can remove it. But, leaders can reduce budget failure by changing what gets challenged before approval and what gets monitored after it. There are several disciplines that apply far below the scale of a national megaproject.

1. Define the outcome before locking in the asset or solution
2. Compare estimates against real performance from similar projects rather than internal confidence alone
3. Appoint delivery leaders with direct experience in the same project category
4. Break scope into repeatable modules with clear interfaces
5. Create decision gates where redesign or termination remains a real option
6. Assume that long delivery windows increase exposure and price that risk honestly

These practices do not guarantee success. They do something more useful. They make self-deception harder. In large projects, that is often the difference between a difficult initiative and an uncontrollable one.

- 1Oxford Professor's Latest Book Examines Roots Of Project Failure
- 2What You Should Know About Megaprojects
- 3The Iron Law Of Megaprojects
- 4Making Bigger Better
- 5Book Review: How Big Things Get Done

Summary

Large projects do not usually collapse because teams stop caring about cost. They collapse because executives approve weak assumptions, compress planning and commit to delivery

before the work is truly designed. That pattern shows up across infrastructure, technology, public investment and corporate transformation. The implication is direct. Better outcomes come less from heroic execution than from disciplined choices made before execution begins. Leaders who clarify the real objective, compare against relevant precedents, design modular work and retain the option to stop can reduce overruns without reducing ambition. In major programs, credibility is earned not by promising the impossible but by setting terms that the project can actually meet