"IMPORTANT, TIMELY, INSTRUCTIVE, AND ENTERTAINING."
—Daniel Kahneman, Nobel Prize–winning author of *Thinking, Fast and Slow*

HOW BIG THINGS GET DONE

THE SURPRISING FACTORS THAT DETERMINE THE FATE OF EVERY PROJECT, FROM HOME RENOVATIONS TO SPACE EXPLORATION AND EVERYTHING IN BETWEEN

BENT FLYVBJERG
and DAN GARDNER

HOW BIG THINGS GET DONE

THE SURPRISING FACTORS THAT DETERMINE THE FATE OF EVERY PROJECT, FROM HOME RENOVATIONS TO SPACE EXPLORATION AND EVERYTHING IN BETWEEN

BENT FLYVBJERG
and DAN GARDNER

CURRENCY NEW YORK

Published in the United States by Currency, an imprint of Random House, a division of Penguin Random House LLC, New York.

CURRENCY and its colophon are trademarks of Penguin Random House LLC.

Library of Congress Cataloging-in-Publication Data
Names: Flyvbjerg, Bent, author. | Gardner, Dan, author.
Title: How big things get done / Bent Flyvbjerg and Dan Gardner.
Description: First edition. | New York: Currency, [2023] | Includes bibliographical references and index.
Identifiers: LCCN 2022036817 (print) | LCCN 2022036818 (ebook) | ISBN 9780593239513 (hardcover) | ISBN 9780593239520 (ebook)
Subjects: LCSH: Project management.
Classification: LCC HD69.P75 F58 2023 (print) | LCC HD69.P75 (ebook) | DDC 658.404—dc23/eng/20220830
LC record available at https://lccn.loc.gov/2022036817
LC ebook record available at https://lccn.loc.gov/2022036818

Ebook ISBN 9780593239520

crownpublishing.com

Book design by Fritz Metsch, adapted for ebook

Cover design: Jamie Keenan

ep_prh_6.0_142523652_c0_r1

Contents

INTRODUCTION: CALIFORNIA DREAMIN'

How is a vision turned into a plan that becomes a triumphant new reality?

Let me tell you a story. You may have heard about it, particularly if you live in California. If you do, you're paying for it.

In 2008, Golden State voters were asked to imagine themselves at Union Station in downtown Los Angeles, on board a sleek silver train. Departing the station, the train slips quietly through the urban sprawl and endless traffic jams and accelerates as it enters the open spaces of the Central Valley, until the countryside is racing by in a blur. Breakfast is served. By the time attendants clear coffee cups and plates, the train slows and glides into another station. This is downtown San Francisco. The whole trip took two and a half hours, not much more than the time it would take the average Los Angeleno to drive to the airport, clear security, and get on a plane to queue on the tarmac, waiting for departure. The cost of the train ticket was $86.

The project was called California High-Speed Rail. It would connect two of the world's great cities, along with Silicon Valley, the global capital of high technology. Words such as *visionary* are used too liberally, but this really was visionary. And for a total cost of $33 billion it would be ready to roll by 2020.[1] In a statewide referendum, Californians approved. Work began.

As I write, it is now fourteen years later. Much about the project remains uncertain, but we can be sure that the end result will not be what was promised.

After voters approved the project, construction started at various points along the route, but the project was hit with constant delays. Plans were changed repeatedly. Cost estimates soared, to $43 billion, $68 billion, $77 billion, then almost $83 billion. As I write, the current highest estimate is $100 billion.[2] But the truth is that nobody knows what the full, final cost will be.

In 2019, California's governor announced that the state would complete only part of the route: the 171-mile section between the towns of Merced and Bakersfield, in California's Central Valley, at an estimated cost of $23 billion. But when that inland section is completed, the project will stop. It will be up to some future governor to decide whether to launch the project again and, if so, figure out how to get the roughly $80 billion—or whatever the number will be by then—to extend the tracks and finally connect Los Angeles and San Francisco.[3]

For perspective, consider that the cost of the line between only Merced and Bakersfield is the same as or more than the annual gross domestic product of Honduras, Iceland, and about a hundred other countries. And that money will build the most sophisticated rail line in North America between two towns most people outside California have never heard of. It will be—as critics put it—the "bullet train to nowhere."

How do visions become plans that deliver successful projects? Not like this. An ambitious vision is a wonderful thing. California was bold. It dreamed big. But even with buckets of money, a vision is not enough.

Let me tell you another story. This one is unknown, but I think it gets us closer to the answers we need.

In the early 1990s, Danish officials had an idea. Denmark is a small country with a population less than New York City's, but it is rich and gives a lot of money in foreign aid and wants that money to do good. Few things do more good than education. The Danish officials got together with colleagues from other governments and agreed to fund a school system for the Himalayan nation of Nepal. Twenty thousand schools and classrooms would be built, most of them in the poorest and most remote regions. Work would begin in 1992. It would take twenty years.[4]

The history of foreign aid is littered with boondoggles, and this project could easily have added to the mess. Yet it finished on budget in 2004—eight years ahead of schedule. In the years that followed, educational levels rose across the country, with a long list of positive consequences, particularly a jump in the number of girls in classrooms. The schools even saved lives: When a massive earthquake struck Nepal in 2015, almost nine thousand people died, with many being crushed to death in collapsing buildings. But the schools had been designed to be earthquake proof, as a first. They stood. Today, the Bill & Melinda Gates Foundation uses the project as an exemplar of how to improve health by increasing enrollment in schools, particularly for girls.[5]

I was the planner on that project.[6] At the time, I was pleased with how it turned out, but I didn't think much about it. It was my first big project, and, after all, we only did what we had said we would do: turn a vision into a plan that was delivered as promised.

However, in addition to being a planner, I am an academic, and the more I studied how big projects come together—or fail to—the more I understood that my experience in Nepal was not normal. In fact, it was not remotely normal. As we'll see, the data show that big projects that deliver as promised are rare. Normal looks a lot more like California High-Speed Rail. Average practice is a disaster, best practice an outlier, as I would later point out in my findings about megaproject management.[7]

Why is the track record of big projects so bad? Even more important, what about the rare, tantalizing exceptions? Why do they succeed where so many others fail? Had we just been lucky delivering the schools in Nepal? Or could we do it again? As a professor of planning and management, I've spent many years answering those questions. As a consultant, I've spent many years putting my answers into practice. In this book, I'm putting them into print.

The focus of my work is megaprojects—*very* big projects—and lots of things about that category are special. Navigation of national politics and global bond markets, for example, is not something the average home remodeler has to contend with. But that stuff is for another book. What I'm interested in here are the drivers of project failure and success that are universal. That explains the title. *How Big Things Get Done* is a nod to my expertise in megaprojects, which are big by anyone's standards. But "big" is relative. For average homeowners, a home remodeling can easily be one of the most expensive, complex, challenging projects they ever tackle. Getting it right means as much or more to them as the fate of megaprojects means to corporations and governments. It is absolutely a "big thing."

So what are the universal drivers that make the difference between success and failure?

PSYCHOLOGY AND POWER

One driver is psychology. In any big project—meaning a project that is considered big, complex, ambitious, and risky by those in charge—people think, make judgments, and make decisions. And where there are thinking, judgment, and decisions, psychology is at play; for instance, in the guise of optimism.

Another driver is power. In any big project people and organizations compete for resources and jockey for position. Where there are competition and jockeying, there is power; for instance, that of a CEO or politician pushing through a pet project.

Psychology and power drive projects at all scales, from skyscrapers to kitchen renovations. They are present in projects made of bricks and mortar, bits and bytes, or any other medium. They are found whenever someone is excited by a vision and wants to turn it into a plan and make that plan a reality—whether the vision is to place another jewel in the Manhattan skyline or launch a new business, go to Mars, invent a new product, change an organization, design a program, convene a conference, write a book, host a family wedding, or renovate and transform a home.

With universal drivers at work, we can expect there to be patterns in how projects of all types unfold. And there are. The most common is perfectly illustrated by California's bullet train to nowhere.

The project was approved, and work began in a rush of excitement. But problems soon proliferated. Progress slowed. More problems arose. Things slowed further. The project dragged on and on. I call this pattern "Think fast, act slow," for reasons I'll explain later. It is a hallmark of failed projects.

Successful projects, by contrast, tend to follow the opposite pattern and advance quickly to the finish line. That's how the Nepal schools project unfolded. So did the Hoover Dam, which was completed a little under budget in fewer than five years—two ahead of schedule.[8] Boeing took twenty-eight months to design and build the first of its iconic 747s.[9] Apple hired the first employee to work on what would become the legendary iPod in late January 2001, the project was formally approved in March 2001, and the first iPod was shipped to customers in November 2001.[10] Amazon Prime, the online retailer's enormously successful membership and free shipping program, went from a vague idea to a public announcement between October 2004 and February 2005.[11] The first SMS texting app was developed in just a few weeks.

Then there's the Empire State Building.

A NEW YORK SUCCESS STORY

The vision that became arguably the world's most legendary skyscraper started with a pencil. Who held the pencil depends on which version of the story you trust. In one, it was the architect, William Lamb. In another, it was John J. Raskob, a financial wizard and former General Motors executive. In either case, a pencil was taken from a desk and held vertically, point up. That's what the Empire State Building would be: slim, straight, and stretching higher into the sky than any other building on the planet.[12]

The idea to erect a tower probably came early in 1929 from Al Smith. A lifelong New Yorker and former New York governor, Smith had been the Democratic presidential candidate in the 1928 election. Like most New Yorkers, Smith opposed Prohibition. Most Americans disagreed, and Smith lost to Herbert Hoover. Unemployed, Smith needed a new challenge. He took his idea to Raskob, and they formed Empire State Inc., with Smith acting as the president and face of the corporation and Raskob as its moneyman. They settled on a location—the site of the original Waldorf-Astoria hotel, once the pinnacle of Manhattan luxury—set the parameters of the project, and developed the business plan. They fixed the total budget, including the purchase and demolition of the Waldorf-Astoria, at $50 million ($820 million in 2021 dollars), and scheduled the grand opening for May 1, 1931. They hired Lamb's firm. Someone held up a pencil. At that point, they had eighteen months to go from first sketch to last rivet.

They moved fast because the moment was right. In the late 1920s, New York had overtaken London as the world's most populous metropolis, jazz was hot, stocks were soaring, the economy was booming, and skyscrapers—the thrilling new symbol of prosperous Machine Age America—were leaping up all over Manhattan. Financiers were looking for new projects to back, the more ambitious the better. The Chrysler Building would soon become the tallest of the titans, garnering all the prestige and rental income that went with the title. Raskob, Smith, and Lamb were determined to have their pencil top them all.

In planning the building, Lamb's focus was intensely practical. "The day that [the architect] could sit before his drawing board and make pretty sketches of decidedly uneconomic monuments to himself has gone," he wrote in January 1931. "His scorn of things 'practical' has been replaced by an intense earnestness to make practical necessities the armature upon which he moulds the form of his idea."

Working closely with the project's builders and engineers, Lamb developed designs shaped by the site and the need to stay on budget and schedule. "The adaptation of the design to conditions of use, construction and speed of erection has been kept to the fore throughout the development of the drawings of the Empire State," he wrote. The designs were rigorously tested to ensure that they would work. "Hardly a detail was issued without having been thoroughly analyzed by the builders and their experts and adjusted and changed to meet every foreseen delay."[13]

In a 1931 publication, the corporation boasted that before any work had been done on the construction site "the architects knew exactly how many beams and of what lengths, even how many rivets and bolts would be needed. They knew how many windows Empire State would have, how many blocks of limestone, and of what shapes and sizes, how many tons of aluminum and stainless steel, tons of cement, tons of mortar. Even before it was begun, Empire State was finished entirely—on paper."[14]

The first steam shovel clawed into the Manhattan dirt on March 17, 1930. More than three thousand workers swarmed the site, and construction advanced rapidly, beginning with the steel skeleton thrusting upward, followed by the completed first story. Then the second story. The third. The fourth. Newspapers reported on the skyscraper's rise as if it were a Yankees playoff run.

As workers learned and processes smoothed, progress accelerated. Up went three stories in one week. Four. Four and a half. At the height of construction, the pace hit a story a day.[15] And a little more. "When we were in full swing going up the main tower," Lamb's partner Richmond Shreve recalled, "things clicked with such precision that once we erected fourteen and a half floors in ten working days—steel, concrete, stone and all."[16] That was an era when people marveled at the efficiency of factories churning out cars, and the Empire State designers were inspired to imagine their process as a vertical assembly line—except that "the assembly line did the moving," Shreve explained, while "the finished product stayed in place."[17]

By the time the Empire State Building was officially opened by President Herbert Hoover—exactly as scheduled, on May 1, 1931—it was already a local and national celebrity. Its height was daunting. The efficiency of its construction was legendary. And even though practicality had been at the front of Lamb's mind, the building was unmistakably beautiful. Lamb's drive for efficiency had created a lean, elegant design, and the New York chapter of the American Institute of Architects awarded it the 1931 Medal of Honor.[18] Then, in 1933, King Kong climbed the building on the silver screen while clutching the glamorous Fay Wray, and the Empire State Building became a global star.

The Empire State Building had been estimated to cost $50 million. It acually cost $41 million ($679 million in 2021). That’s 17 percent under budget, or $141 million in 2021 dollars. Construction finished several weeks before the opening ceremony.

I call the pattern followed by the Empire State Building and other successful projects “Think slow, act fast.”

At the start, I asked how a vision is turned into a plan that becomes a triumphant new reality. As we will see, that is the answer: Think slow, act fast.

1

THINK SLOW, ACT FAST

The record of big projects is even worse than it seems. But there is a solution: Speed up by slowing down.

Denmark is a peninsula with islands scattered along its east coast. Danes therefore long ago became experts at operating ferries and building bridges. So it was no surprise, in the late 1980s, when the government announced the Great Belt project. It comprised two bridges, one of which would be the world's longest suspension bridge, to connect two of the bigger islands, including the one with Copenhagen on it. There would also be an underwater tunnel for trains—the second longest in Europe—which would be built by a Danish-led contractor. That was interesting because Danes had little experience boring tunnels. I watched the announcement on the news with my father, who worked in bridge and tunnel construction. "Bad idea," he grumbled. "If I were digging a hole that big, I would hire someone who had done it before."

Things went wrong from the start. First there was a yearlong delay in delivering four giant tunnel-boring machines. Then, as soon as the machines were in the ground, they proved to be flawed and needed redesign, delaying work another five months. Finally, the big machines started slowly chewing their way under the ocean floor.

Up above, the bridge builders brought in a massive oceangoing dredger to prepare their worksite.[1] To do its work, the dredger stabilized itself by lowering giant support legs into the seafloor. When the work was done, the legs were lifted, leaving deep holes. By accident, one of the holes happened to be on the projected path of the tunnel. Neither the bridge builders nor the tunnelers saw the danger.

One day, after a few weeks of boring, one of the four machines was stopped for maintenance. It was about 250 meters (820 feet) out to sea and an assumed 10 meters (33 feet) under the seafloor. Water was seeping into the maintenance area in front of the machine, and a contractor unfamiliar with tunneling hooked up a pump to get the water out. The cables of the pump were trailed through a manhole into the boring machine. Suddenly water started pouring in at a speed indicating a breach of the tunnel. Evacuation was immediate—with no time to remove the pump and cables and close the manhole.

The machine and the whole tunnel flooded. So did a parallel tunnel and the boring machine within it.

Luckily, no one was injured or killed. But the salt water in the tunnel was like acid to its metal and electronics. Engineers on the project told me at the time that it would be cheaper to abandon the tunnel and start again rather than pull out the borers, drain the tunnel, and repair it. But politicians overrode them because an abandoned tunnel would be too embarrassing. Inevitably, the whole project came in very late and way over budget.

This story isn't all that unusual. There are lots more like it in the annals of big projects. But it was that project that nudged me to start a big project of my own—a database of big projects. It continues to grow. In fact, it is now the world's largest of its kind.

And it has a great deal to teach us about what works, what doesn't, and how to do better.

HONEST NUMBERS

After the accident, the recovery, and the eventual completion of the Great Belt bridges and tunnel, everyone agreed that the project had gone badly over budget. But by how much? Management said 29 percent for the whole project. I dug into the data, did my own analysis, and discovered that their number was, shall we say, optimistic. The actual overrun was 55 percent, and 120 percent for the tunnel alone (in real terms, measured from the final investment decision). Still, management kept repeating their number in public, and I kept correcting them, until they did a public opinion poll that showed that the public sided with me. Then they gave up. Later, an official national audit confirmed my numbers, and the case was closed.[2]

That experience taught me that megaproject management may not be a field of what University of Washington public affairs professor Walter Williams called "honest numbers."[3] As simple as it should be in theory to judge projects, in practice it's anything but. In every big project, there are blizzards of numbers generated at different stages by different parties. Finding the right ones—those that are valid and reliable—takes skill and work. Even trained scholars get it wrong.[4] And it doesn't help that big projects involve money, reputations, and politics. Those who have much to lose will spin the numbers, so you cannot trust them. That's not fraud. Or rather, it's not usually fraud; it's human nature. And with so many numbers to choose from, spinning is a lot easier than finding the truth.

This is a serious problem. Projects are promised to be completed by a certain time, at a certain cost, with certain benefits produced as a result—benefits being things such as revenues, savings, passengers moved, or megawatts of electricity generated. So how often do projects deliver as promised? That is the most straightforward question anyone could ask. But when I started to look around in the 1990s, I was stunned to discover that no one could answer it. The data simply hadn't been collected and analyzed. That made no sense when *trillions* of dollars had been spent on the giant projects increasingly being called megaprojects—projects with budgets in excess of $1 billion.

Our database started with transportation projects: the Holland Tunnel in New York; the BART system in San Francisco; the Channel Tunnel in Europe; bridges, tunnels, highways, and railways built throughout the twentieth century. It took five years, but with my team I got 258 projects into the database, making it the biggest of its kind at the time.[5] When we finally began publishing the numbers in 2002, it made waves because nothing like it had been done before.[6] Also, the picture that emerged wasn't pretty.

"Project estimates between 1910 and 1998 were short of the final costs an average of 28 percent," according to *The New York Times,* summarizing our findings. "The biggest errors were in rail projects, which ran, on average, 45 percent over estimated costs [in inflation-adjusted dollars]. Bridges and tunnels were 34 percent over; roads, 20 percent. Nine of 10 estimates were low, the study said."[7] The results for time and benefits were similarly bad.

And these are conservative readings of the data. Measured differently—from an earlier date and including inflation—the numbers are *much* worse.[8]

The global consultancy McKinsey got in touch with me and proposed that we do joint research. Its researchers had started investigating major information technology projects—the biggest of which cost more than $10 billion—and their preliminary numbers were so dismal that they said it would take a big improvement for IT projects to rise to the level of awfulness of transportation projects. I laughed. It seemed impossible that IT could be that bad. But I worked with McKinsey, and indeed we found that IT disasters were even worse than transportation disasters. But otherwise it was a broadly similar story of cost and schedule overruns and benefit shortfalls.[9]

That was startling. Think of a bridge or a tunnel. Now picture the US government's HealthCare.gov website, which was a mess when it first opened as the "Obamacare" enrollment portal. Or imagine the information system used by the National Health Service in the United Kingdom. These IT projects are made of code, not steel and concrete. They would seem to be different from transportation infrastructure in every possible way. So why would their outcomes be statistically so similar, with consistent cost and schedule overruns and benefit shortfalls?

We shifted our research to mega-events such as the Olympic Games and got the same result. Big dams? Same again. Rockets? Defense? Nuclear power? The same. Oil and gas projects? Mining? Same. Even something as common as building museums, concert halls, and skyscrapers fit the pattern. I was astonished.[10]

And the problem wasn't limited to any country or region; we found the same pattern all over the world.[11] The famously efficient Germans have some remarkable examples of bloat and waste, including Berlin's new Brandenburg Airport, which was years delayed and billions of euros over budget, hovering on the verge of bankruptcy only a year after opening in October 2020.[12]

Even Switzerland, the nation of precise clocks and punctual trains, has its share of embarrassing projects; for instance, the Lötschberg Base Tunnel, which was completed late and with a cost overrun of 100 percent.

OVER BUDGET, OVER TIME, OVER AND OVER AGAIN

The pattern was so clear that I started calling it the "Iron Law of Megaprojects": over budget, over time, under benefits, over and over again.[13]

The Iron Law is not a "law" like in Newtonian physics, meaning something that invariably produces the same outcome. I study people. In the social sciences, "laws" are probabilistic (they are in natural science, too, but Isaac Newton didn't pay much attention to that). And the probability that any big project will blow its budget and schedule and deliver disappointing benefits is very high and very reliable.

The database that started with 258 projects now contains more than 16,000 projects from 20-plus different fields in 136 countries on all continents except Antarctica, and it continues to grow. There are some recent and important wrinkles in the numbers, which I'll discuss later, but the general story remains the same: In total, only 8.5 percent of projects hit the mark on both cost and time. And a minuscule 0.5 percent nail cost, time, and benefits. Or to put that another way, 91.5 percent of projects go over budget, over schedule, or both. And 99.5 percent of projects go over budget, over schedule, under benefits, or some combination of these. Doing what you said you would do should be routine, or at least common. But it almost never happens.

Graphically, the Iron Law looks like this:

THE IRON LAW OF PROJECT MANAGEMENT:
"Over Budget, Over Time, Under Benefits, Over and Over Again"

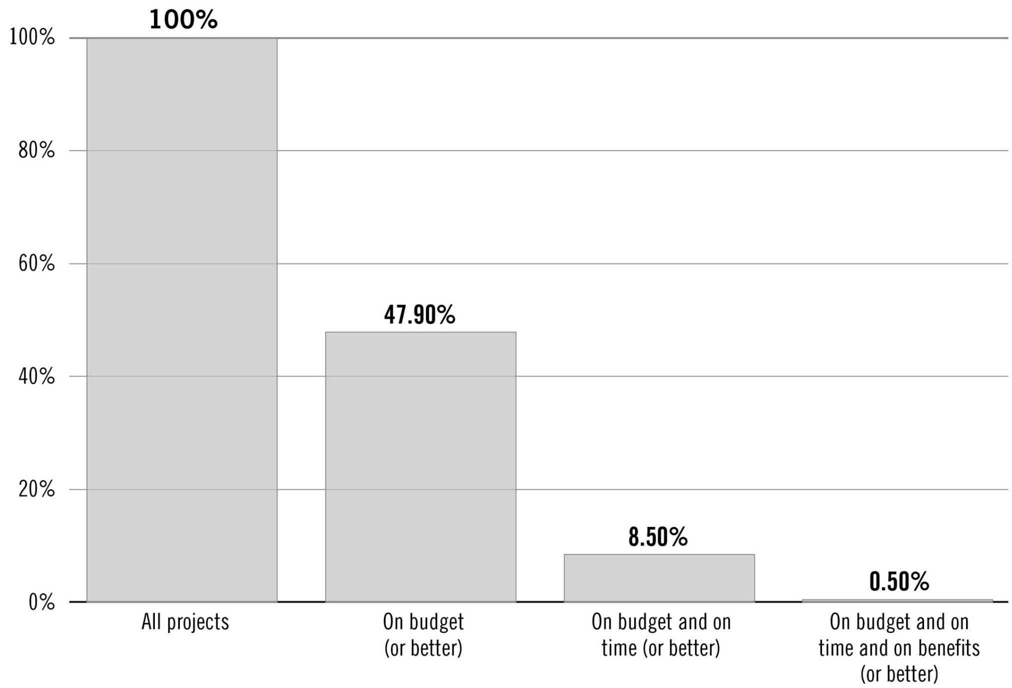

Tellingly, the 0.5 percent of projects that are on budget, time, and benefits are nearly invisible to the naked eye. It's hard to overstate how bad that record is. For anyone contemplating a big project, it is truly depressing. But as grim as those numbers are, they don't tell the full truth—which is *much* worse.

From experience, I know that most people are aware that cost and time overruns are common. They don't know *how* common—they are usually shocked when I show them my numbers—but they definitely know that if they lead a big project, they should consider and protect themselves against overruns, particularly cost overruns. The obvious way to do that is to build a buffer into the budget. You hope it won't be needed, but you'll be covered just in case. How big should that buffer be? Typically, people make it 10 percent or 15 percent.

But let's say you are an unusually cautious person and you are planning the construction of a large building. You put a 20 percent buffer into the budget and think you are now well protected. But then you come across my research and discover that the actual mean cost overrun of a major building project is 62 percent. That is heart-stopping. It may also be project-stopping. But let's say you are the very rare planner who can get your financial backers to cover that risk and still go ahead with the project. You now have an extraordinary 62 percent buffer built into your budget. In the real world, that almost never happens. But you're one of the fortunate few. Are you at last protected? No. In fact, you have *still* drastically underestimated the danger.

That's because you have assumed that if you are hit with a cost overrun, it will be somewhere around the mean—or 62 percent. Why did you assume that? Because it would be true if the cost overruns followed what statisticians call a "normal distribution." That's the famous bell curve, which looks like a bell when graphed. Much of statistics is built upon bell curves—sampling, averages, standard deviations, the law of large numbers, regression to the mean, statistical tests—and it has filtered into the culture and into the popular imagination, where it fits well with how we intuitively grasp risk. In a normal distribution, results are overwhelmingly lumped in the middle and there are very few or no extreme observations at either end—the so-called tails of the distribution. These tails are therefore said to be *thin.*

Height is normally distributed. Depending on where you live, most adult males are around five feet, nine inches (1.75 meters) tall, and the tallest person in the world is only about 1.6 times taller than that.[14]

But the "normal" distribution isn't the only sort of distribution that exists—or even the most common one. So it is not normal in that sense of the word. There are other distributions that are called "fat-tailed" because, compared with normal distributions, they contain far more extreme outcomes in their tails.

Wealth, for example, is fat-tailed. At the time of writing, the wealthiest person in the world is 3,134,707 times wealthier than the average person. If human height followed the same distribution as human wealth, the tallest person in the world would not be 1.6 times taller than the average person; he would be 3,311 miles (5,329 kilometers) tall, meaning that his head would be thirteen times farther into outer space than the International Space Station.[15]

So the critical question is this: Are project outcomes distributed "normally," or do they have fat tails? My database revealed that information technology projects have fat tails. To illustrate, 18 percent of IT projects have cost overruns above 50 percent in real terms. And for those projects the average overrun is 447 percent! That's the *average* in the tail, meaning that many IT projects in the tail have even higher overruns than this. Information technology is *truly* fat-tailed![16] So are nuclear storage projects. And the Olympic Games. And nuclear power plants. And big hydroelectric dams. As are airports, defense projects, big buildings, aerospace projects, tunnels, mining projects, high-speed rail, urban rail, conventional rail, bridges, oil projects, gas projects, and water projects. (See Appendix A.)

In fact, most project types have fat tails. How "fat" their tails are—how many projects fall into the extremes and how extreme those extremes are—does vary. I've cited them in order, from fattest to least fat (but still fat)—or, if you prefer, from most at risk of terrifying overruns to less at risk (but still very much at risk).[17]

There *are* a few project types that do not have fat tails. That's important. I'll explain why, and how we can all make use of this fact, in the last chapter.

But for now the lesson is simple, clear, and scary: Most big projects are not merely at risk of not delivering as promised. Nor are they only at risk of going seriously wrong. They are at risk of going *disastrously* wrong because their risk is fat-tailed. Against that background, it is interesting to note that the project management literature almost completely ignores systematic study of the fat-tailedness of project risk.

What do fat-tailed outcomes look like? Boston's "Big Dig"—replacing an elevated highway with a tunnel, with construction started in 1991—put the city through the wringer for sixteen years and cost more than *triple* what it was supposed to. NASA's James Webb Space Telescope, which is now almost a million miles from Earth, was forecast to take twelve years but required nineteen to complete, while its final cost of $8.8 billion was an astronomical—forgive me—450 percent over budget. Canada's firearms registry, an IT project, went 590 percent over budget. And then there is Scotland's Parliament Building. When it opened in 2004, it was three years late and a bagpipe-exploding 978 percent over budget.

Nassim Nicholas Taleb famously dubbed low-probability, high-consequence events "black swans." Disastrous project outcomes such as these can end careers, sink companies, and inflict a variety of other carnage. They definitely qualify as black swans.

Just look at what a black swan outcome did to Kmart: Responding to competitive pressure from Walmart and Target, it launched two enormous IT projects in 2000. Costs exploded, contributing directly to the company's decision to file for bankruptcy in 2002.[18] Or consider what another IT blowout did to the legendary jeans maker Levi Strauss: Originally forecast to cost $5 million, the project forced the company to take a $200 million loss and show its CIO the door.[19]

There are worse fates for executives. When a troubled nuclear power plant project in South Carolina fell badly behind schedule, the CEO of the company in charge withheld that information from regulators "in an effort to keep the project going," noted a 2021 US Department of Justice press release, which also announced that the executive had been sentenced to two years in federal prison and forced to pay $5.2 million in forfeitures and fines.[20] Black swan outcomes do indeed have consequences for projects and those who lead them.

If you are not a corporate executive or government official, and if the ambitious project you are contemplating is on a much smaller scale than these giants, it may be tempting to think that none of this applies to you. Resist that temptation. My data show that smaller projects are susceptible to fat tails, too. Moreover, fat-tailed distributions, not normal distributions, are typical within complex systems, both natural and human, and we all live and work within increasingly complex systems, which means increasingly interdependent systems. Cities and towns are complex systems. Markets are complex systems. Energy production and distribution are complex systems. Manufacturing and transportation are complex systems. Debt is a complex system. So are viruses. And climate change. And globalization. On and on the list goes. If your project is ambitious and depends on other people and many parts, it is all but certain that your project is embedded in complex systems.

That describes projects of all types and scales, all the way down to home renovations. A few years ago, in a BBC show about renovating historic British properties, one episode featured a London couple who bought a run-down house in the countryside and got a builder to estimate the cost of a complete renovation. He pegged it at $260,000. Eighteen months later, the project was far from done and the couple had already spent $1.3 million.[21] That is the sort of overrun we would expect to find in a fat-tailed distribution. And it is certainly not unique. Later in the book, we'll witness a home renovation in Brooklyn spin wildly out of control and inflict an equally devastating overrun on the unfortunate and unexpecting homeowners.

That London couple were apparently wealthy enough to keep funding the renovation. Similarly, major corporations on the hook for runaway projects may be able to keep things going by borrowing more and more money. Governments can also pile up debt. Or raise taxes. But most ordinary folks and small businesses cannot draw on a big stockpile of wealth, run up debt, or raise taxes. If they start a project that hurtles toward the fat tail of the distribution, they will simply be wiped out, giving them even more reason than a corporate executive or government official to take the danger seriously.

And that starts by understanding what causes project failure.

THE WINDOW OF DOOM

The patterns I mentioned earlier, confirmed by my data, are strong clues: Projects that fail tend to drag on, while those that succeed zip along and finish.

Why is that? Think of the duration of a project as an open window. The longer the duration, the more open the window. The more open the window, the more opportunity for something to crash through and cause trouble, including a big, bad black swan.

What could that black swan be? Almost anything. It could be something dramatic, like an election upset, a stock market collapse, or a pandemic. After Covid-19 emerged in January 2020, projects all over the world—from the 2020 Tokyo Olympics to the release of the James Bond movie *No Time to Die*—were delayed, postponed, or scrapped altogether. Events such as these may be extremely unlikely on any given day, month, or year. But the more time that passes from the decision to do a project to its delivery, the greater their probability.

Notice that these big, dramatic events, which are easily capable of damaging a project so badly that it delivers a black swan outcome, are themselves low probability and high consequence. That is, they are black swans. So a black swan crashing through the window of vulnerability may itself cause a black swan outcome.

But drama isn't necessary for change to batter and bury projects. Even mundane change can do that. Journalists who write biographies of up-and-coming politicians know, for example, that the market for their books depends on the politician continuing to be on the rise when the book is released. Any number of events can change that: a scandal, a lost election; an illness; a death. Even something as simple as the politician getting bored with politics and taking another job would ruin the project. Again, the more time that passes from decision to delivery, the greater the probability of one or more of these events happening. It's even possible that trivial events, in just the wrong circumstances, can have devastating consequences.

It's hard to think of anything more trivial to most people around the world than gusts of wind in the Egyptian desert. Yet on March 23, 2021, it was just such gusts, at just the wrong moment, that pushed the bow of *Ever*

Given, a giant container ship, into a bank of the Suez Canal. The ship got stuck and couldn't be budged for six days, blocking the canal, halting hundreds of ships, freezing an estimated $10 billion in trade each day, and sending shocks rippling through global supply chains.[22] The people and projects who suffered as a result of those supply-chain troubles may never have realized it, but the cause of their trouble was ultimately strong winds in a faraway desert.[23]

A complex systems theorist might describe what happened by saying that the dynamic interdependencies among the parts of the system—the wind, the canal, the ship, and supply chains—created strong nonlinear responses and amplification. In plain English, minor changes combined in a way to produce a disaster. In complex systems, that happens so often that the Yale sociologist Charles Perrow called such events "normal accidents."[24]

Growing complexity and interdependency may make such outcomes more likely in today's world, but they are hardly a new phenomenon. A proverb that originated in the Middle Ages and comes in many forms tells us, "For want of a nail, the shoe was lost. For want of a shoe, the horse was lost. For want of a horse, the rider was lost. For want of a rider, the battle was lost. For want of a battle, the kingdom was lost." This version was published by Benjamin Franklin in 1758, and he introduced it with the warning that "a little neglect may breed great mischief." The key word is *may*. Most nails can be lost without anything bad happening at all. A few such losses will have consequences, but they will be minor, like the loss of one horse or one rider. But sometimes a lost nail may cause something truly terrible.

From the dramatic to the mundane to the trivial, change can rattle or ruin a project—if it occurs during the window of time when the project is ongoing.

Solution? Close the window.

Of course, a project can't be completed instantly, so we can't close the window entirely. But we can make the opening radically smaller by speeding up the project and bringing it to a conclusion faster. That is a main means of reducing risk on any project.

In sum, keep it short!

THE NEED FOR SPEED

How do we get a project done as quickly as possible? The obvious answer—and certainly the most common one—is to set severe timelines, get started right away, and demand that everyone involved work at a furious pace. Drive and ambition are key, goes the conventional wisdom. If experienced observers think a project will take two years, say you will do it in one. Commit to the project, heart and soul, and charge ahead. And in managing others, be fierce. Demand that everything be done yesterday. Like the drummer on a Roman galley preparing to ram a ship, beat the drum at a furious pace.

This thinking is as misguided as it is common. There is a monument to it in Copenhagen.

The Copenhagen Opera House, the home of the Royal Danish Opera, was the vision of Arnold Maersk Mc-Kinney Møller, the CEO and chairman of Maersk, the Danish shipping giant. In the late 1990s, Møller, who was then in his late eighties, decided he wanted a grand building situated prominently at the harborside as his very visible and permanent legacy. And he wanted it designed and built quickly. The queen of Denmark would attend the opening, and Møller had no intention of missing his big night. When Møller asked the architect, Henning Larsen, how long it would take, Larsen said five years. "You'll get four!" Møller curtly responded.[25] With much beating of galley drums, the deadline was met, and Møller and the queen opened the opera house together on January 15, 2005.

But the cost of that haste was terrible, and not only in terms of cost overruns. Larsen was so appalled by the completed building that he wrote a whole book to clear his reputation and explain the confused structure, which he called a "mausoleum."

Haste makes waste.

Even that cost is mild compared to what rushing projects this way can do. In 2021, after an overpass collapsed beneath a metro train in Mexico City, three independent investigations concluded that rushed, shoddy work was to blame. A Norwegian firm hired by the city to conduct an investigation concluded that the tragedy had been caused by “deficiencies in the construction process,” as did a later report released by the attorney general of Mexico City.[26] *The New York Times* did its own investigation and concluded that the city’s insistence that construction be completed before the city’s powerful mayor was due to leave office had been a key contributing cause of the collapse. “The scramble led to a frenzied construction process that began before a master plan had been finalized and produced a metro line with defects from the start,” the *Times* concluded.[27] The collapse of the overpass killed twenty-six people. Haste makes not only waste but tragedy.

MAKE HASTE—SLOWLY

To understand the right way to get a project done quickly, it’s useful to think of a project as being divided into two phases. This is a simplification, but it works: first, planning; second, delivery. The terminology varies by industry—in movies, it’s “development and production”; in architecture, “design and construction”—but the basic idea is the same everywhere: Think first, then do.

A project begins with a vision that is, at best, a vague image of the glorious thing the project will become. Planning is pushing the vision to the point where it is sufficiently researched, analyzed, tested, and detailed that we can be confident we have a reliable road map of the way forward.

Most planning is done with computers, paper, and physical models, meaning that planning is relatively cheap and safe. Barring other time pressures, it’s fine for planning to be slow. Delivery is another matter. Delivery is when serious money is spent and the project becomes vulnerable as a consequence.

Consider a Hollywood director working on a live-action movie project in February 2020. The Covid pandemic is about to arrive. How badly will that hurt the project? The answer depends on what stage the project is in. If the director and her team are writing scripts, drawing storyboards, and scheduling location shoots—if they are planning, in other words—it's a problem but not a disaster. In fact, a lot of the work will probably continue despite the pandemic. But what if, when the pandemic arrives, the director is filming in the streets of New York with a crew of two hundred plus a handful of very expensive movie stars? Or what if the movie is finished but still a month away from its release in theaters that are about to close indefinitely? That's not a problem; it's a disaster.

Planning is a safe harbor. Delivery is venturing across the storm-tossed seas. This is a major reason why, at Pixar—the legendary studio that created *Toy Story, Finding Nemo, The Incredibles, Soul,* and so many other era-defining animated movies—"directors are allowed to spend years in the development phase of a movie," noted Ed Catmull, a co-founder of Pixar. There is a cost associated with exploring ideas, writing scripts, storyboarding images, and doing it all over and over again. But "the costs of iterations are relatively low."[28] And all that good work produces a rich, detailed, tested, and proven plan. When the project moves into the production phase, it will, as a consequence of all that work, be relatively smooth and quick. That's essential, Catmull noted, because production "is where costs explode."

Not only is it safer for planning to be slow, it is *good* for planning to be slow, as the directors at Pixar well know. After all, cultivating ideas and innovations takes time. Spotting the implications of different options and approaches takes more time. Puzzling through complex problems, coming up with solutions, and putting them to the test take still more time. Planning requires thinking—and creative, critical, careful thinking is slow.

Abraham Lincoln is reputed to have said that if he had five minutes to chop down a tree, he'd spend the first three sharpening the ax.[29] That's exactly the right approach for big projects: Put enormous care and effort into planning to ensure that delivery is smooth and swift.

Think slow, act fast: That's the secret of success.

"Think slow, act fast" may not be a new idea. It was on grand display back in 1931, after all, when the Empire State Building raced to the sky. You could even say that the idea goes at least as far back as Rome's first emperor, the mighty Caesar Augustus, whose personal motto was "*Festina lente,*" or "Make haste slowly."

But "Think slow, act fast" is not how big projects are typically done. "Think fast, act slow" is. The track record of big projects unequivocally shows that.

PROJECTS DON'T *GO* WRONG, THEY *START* WRONG

Look at California High-Speed Rail. When it was approved by voters and construction started, there were lots of documents and numbers that may have superficially resembled a plan. But there was no carefully detailed, deeply researched, and thoroughly tested program, which is to say that there was no real plan. Louis Thompson, an expert on transportation projects who chairs the California High-Speed Rail Peer Review Group convened by the California State Legislature, says that what California had in hand when the project got under way could at best be described as a "vision" or an "aspiration."[30] It's no wonder that problems started multiplying and progress slowed to a crawl soon after delivery began.

That is, sadly, typical. On project after project, rushed, superficial planning is followed by a quick start that makes everybody happy because shovels are in the ground. But inevitably, the project crashes into problems that were overlooked or not seriously analyzed and dealt with in planning. People run around trying to fix things. More stuff breaks. There is more running around. I call this the "break-fix cycle." A project that enters it is like a mammoth stuck in a tar pit.

People say that projects "go wrong," which they all too often do. But phrasing it that way is misleading; projects don't *go* wrong so much as they *start* wrong.

This raises an urgent question: If "Think slow, act fast" is the wise approach, why do leaders of big projects so often do the opposite? I'll answer that question in chapter 2.

In chapter 3, I'll look at how to start a project without stumbling into the "Think fast, act slow" tar pit.

People often think that planning is about filling in flowcharts. And too often, it is. But it shouldn't be. In chapter 4, I'll take a close look at what I call "Pixar planning," how the movie studio and others use simulation and iteration to produce a plan that is creative, rigorous, detailed, and reliable—and highly likely to make delivery smooth and swift. I'll use "Pixar planning" as a name and a model for planning, not just at Pixar but for any planning that develops a tested and tried plan; that is, a plan worthy of its name.

In chapter 5, I'll examine the invaluable role of experience in both planning and delivery of big projects—or rather, the invaluable role it could play if it were not so often marginalized, misunderstood, or simply ignored.

In chapter 6, it's on to forecasting. How long will the project take? How much will it cost? Setting the wrong expectations at the start can set you up for failure before you've even started. Fortunately, there is a fix. More fortunately, it's surprisingly easy.

Some people will object to all this emphasis on planning. They believe that big projects, particularly creative projects such as movies, signature architecture, or innovative software, get better results when people take a leap of faith, get started right away, and rely on ingenuity to see them through. In chapter 7, I'll examine this argument in its strongest form—and present the data to prove that it's dead wrong.

But even the best plan won't succeed if it doesn't have a solid team delivering it. So in chapter 8, I'll look at how one giant project successfully drew together thousands of people from hundreds of different organizations with different interests and turned them into a united, determined, effective team that delivered the planned benefits on time and on budget.

In the final chapter, I'll draw on the themes of the previous chapters to explore a concept that brings them all together: modularity. Its potential is huge. Not only can it cut costs, boost quality, and speed things up for a vast array of projects from wedding cakes to subways, it can transform how we build infrastructure—and even help save the world from climate change.

But first we have to answer that question about why projects so often start prematurely. Let me tell you the story of a man in a hurry—and how he almost ruined one of the most beautiful places in the United States.