Per Aspera Industries
Government Use
Portable Document Format
*UNCLASSIFIED

Per Aspera

Per Aspera — 039

A USER'S GUIDE TO "FASTER, BETTER, CHEAPER"

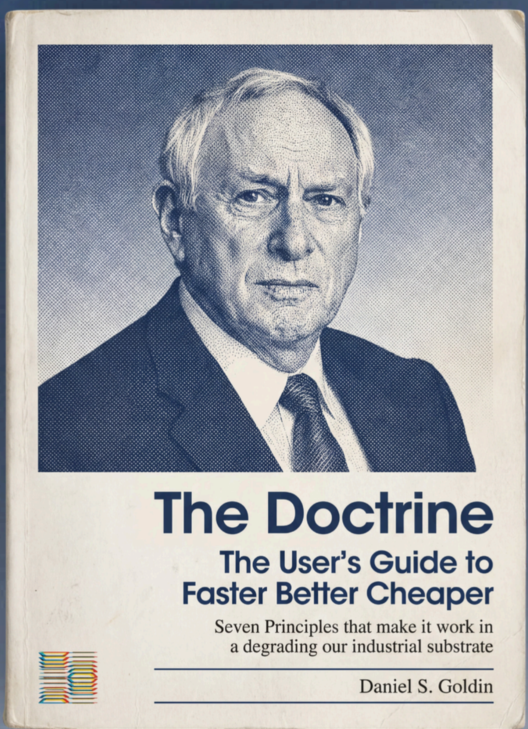

Contributors

Ryan Duffy (Editor in Chief at Per Aspera)
Dan Goldin (Chairman at Per Aspera)
Alan Simon (CEO of North America Closeshoring)

CONTENTS

PROLOGUE: THE FAT STACK OF PAPER

In 1993, Mars Observer was supposed to be the flagship. We'd spent most of a decade building it. The budget had blown past $800M, nearly a 4:1 overrun. And on August 21st, three days before it was supposed to slip into Martian orbit, Observer went silent. We never heard from it again.

After the loss, NASA's new honcho – yours truly – called in the team with a simple request: show me the requirements. They walked in carrying a stack of paper more than 3 feet high. Here was a machine so carefully specified, so relentlessly reviewed, so exhaustively documented that it looked – on paper – like the safest possible bet. It was optimized for compliance, contingency "handling," and passing every gate. Only one problem: somewhere along the way, the mission – arriving and working on Mars – got buried under the requirements.

So I threw the stack out.

Then, I commissioned a replacement with three requirements:

- Land safely on Mars
- Deploy a robot
- Do good science.

Three years later, for a fraction of the cost, Mars Pathfinder bounced onto the surface in airbags, built with the help of a contractor in Delaware better known for spacesuits than interplanetary landings. The robot deployed. **The science was spectacular. Time put the rover on the cover.**


TIME
SPECIAL REPORT
MARS
JULY 4, 1997

This is the crucible from which Faster, Better, Cheaper was born: a way to build things in the real world under real constraints. I ran NASA for nine years as the high priest of this doctrine. I remember FBC as a moment in time when a large institution, briefly, told the truth to itself: the physics is real, the money is finite, and the old way of building hard things will break us.

A lesson learned at high cost. A brief window of clarity. A relapse.

Eventually, the system spit the doctrine out.

Old instincts reasserted themselves after Pathfinder. On the next Mars missions, JPL reintroduced retro rockets and wandered back toward the familiar machinery of precision descent systems and fragile complexity. Both missions failed. The engineers who had bounced a robot across Mars watched their successors return to a method that had already hurt us, and then watched it hurt us again.

If FBC was a virtue, its inverse was a pathology. I've spent the last quarter century watching this pathology spread. Slowly, but surely, it took hold of more institutions until it had become the default operating system for American industry at continental scale.

I'm going to call it SWE: Slow, Worse, Expensive. Symptoms may include different acronyms and different contractors across different agencies and different states; sprawling requirements; timelines that stretch until nobody is accountable, programs that calcify into political entities; and a moral panic every time something fails — even when the failure is exactly what learning looks like in the real world.

We've built a doctrine to outmaneuver this pattern inside one agency.

We're going to need it again, at the scale of a nation.

What follows in this Antimemo is a user's guide to Faster Better Cheaper; seven principles that make it work; signs of SWE degrading our industrial substrate; and a blueprint for rebuilding that substrate in a nation that wants to build again.

SECTION 001

THE DOCTRINE

A lot of historians, armchair pundits, and keyboard warriors have examined Faster, Better, Cheaper from the outside. This is what the playbook actually looks like, as told by the person who created and ran it.

PRINCIPLE 1— START WITH THE PHYSICS

If there's one thing I've learned in my 63 years as an engineer, it is this: you have to understand the physics of the problem before you can simplify it.

You can't compress a schedule for a system you don't understand from first principles. You can't reduce cost on a design whose constraints you haven't internalized. The temptation in any large organization is to "manage" the process via meetings, milestones, made-up paperwork, and contracts.

FBC starts somewhere else. It starts with physics.

PRINCIPLE 2 — GO BIG ONLY WHEN PHYSICS DEMAND IT

Take this as an explicit commandment: "Thou shalt not come forward with a multi-$B program unless the outcome truly demands it."

Some outcomes require scale and exquisite systems. Fine. Pay the price and secure the outcome. But bigness by default creates programs too expensive to fail. They consume all oxygen, crowd out portfolios of other promising solutions, and become politically unsurvivable, which means the institution can't admit problems, because admitting problems would kill the program.

I call these systems Battlestar Galacticas. (You know these when you see 'em. They cost unthinkable sums over their lifetime, blow through tens of billions and well past baseline budgets before fielding anything meaningful, and absorb the political will to try anything else.)

Battlestar Galactica, the capital warship. Specifically, a Colonial Battlestar–class ship.

PRINCIPLE 3 — REQUIRE SIMPLICITY

It wasn't just Mars Observer. I've seen way too many sets of requirements that are printed, stapled, and stacked to the heavens, with every subsystem specified, every interface defined, and every contingency documented. Contrast that with the three requirements for Pathfinder: Land safely, deploy robot, do good science.

"Simplicity of requirements" – or, specifying the mission rather than the method – opens the solution space. Overspecifying costs you money and time. Taken to the extreme, it locks you into decades-long timelines and sole-source dependencies, because you've king-made only one "too-big-to-fail" contractor who can satisfy the full spec.

The tighter the filter, the fewer ideas pass through, and the lower the probability you will land on the best solution.

PRINCIPLE 4 — SIMULATE BEFORE YOU BUILD

Exercise systems digitally before you build them. Test failure modes in simulation where failure costs electrons, not hardware. Run as many iterations as you need before you cut metal. The tools to do this are better now than anything we had at NASA in the '90s or 2000s. High-fidelity simulation tools and physics engines can compress the path from concept to qualification without compromising safety margins, which is the whole trick.

And yet I still walk into factories and see paper, pen, and black plastic notebooks managing operations next to multi-$M machines. That's a technological adoption problem. And given how far our competitors are on world-modeling and sim, it's a national speed problem.

PRINCIPLE 5 — FAIL FAST, FAIL CHEAP, AND LEARN

During my nine years at NASA, we launched 184 robotic spacecraft. Only 16 were designed under FBC principles. Ten succeeded and six failed – a 63% success rate. May be damning if you stop counting there.

But all 16 missions combined cost roughly what the U.S. spent on Cassini alone (~$3.3B in development and primary mission costs) That's ten successful missions for less than one flagship. Success-per-dollar is a more meaningful measurement than success-per-attempt, because attempts are cheap and learning compounds.

We built systems small enough and cheap enough that we could afford to lose some. We tested early, launched early, failed early, and learned what we needed to do before the next mission flew.

You could read breathless stories about those losses, with the media, politicians, and some industry figureheads pounding their chests and crying bloody murder as though someone had died. Nobody had died. Robots had failed. That was the point.

SpaceX is perhaps the best practitioner of this philosophy on the planet, and still to this day has to suffer the consequences and endure the same ritual outrage – from the same types of pundits, politicians, and clueless correspondents – every time a test vehicle breaks apart on camera. The company flew seven Starship integrated flight tests in 21 months: the first exploded, the second broke apart near orbit, and by the fifth, they caught the booster on the launch tower. Each flight iterated on the specific failure mode of the previous one, even if the notion that this is the whole point was lost on the peanut gallery.

Anduril runs a version of the same playbook — self-funded R&D, in-house manufacturing, prototypes fielded on timelines that would still be in the requirements phase at a legacy prime. These companies treat failure as data. The acquisition reflex today is the opposite: delay contact with reality until the system is too big to change and too exquisite to admit error.

PRINCIPLE 6 — RUN A PORTFOLIO, NOT A FLAGSHIP

Would you bet the farm on one exquisite system? No – you'd take many bets, learn quickly, iterate, and scale what works. Here's why:

- If every mission must succeed, then every mission must be expensive.
- If every mission is expensive, you fly fewer.
- If you fly fewer, you learn slower.
- If you learn slower, you get more fragile.
- If you get more fragile, you panic at failure.
- And then you lock yourself into SWE forever.

The way out is to accept that some failures are the price of learning, and to design anti-fragile portfolios so no single loss can kill the program. Each loss generates incremental data that feeds the next mission or a competing approach. Over time, a portfolio will deliver aggregate performance no concentration of a few billion-dollar flagships can match.

Principle 2 still applies: go big when the physics (and the stakes) demand it. And never cut corners on mission-critical systems (such as human spaceflight).

PRINCIPLE 7 — REDUCE PROGRAMS AND PRODUCTS TO SINGLE, SIMPLE METRICS

At NASA, I reduced agency performance down to a simple report card, then compressed it again into two numbers:

- Cycle time (authority-to-proceed to orbit)
- Average spacecraft cost

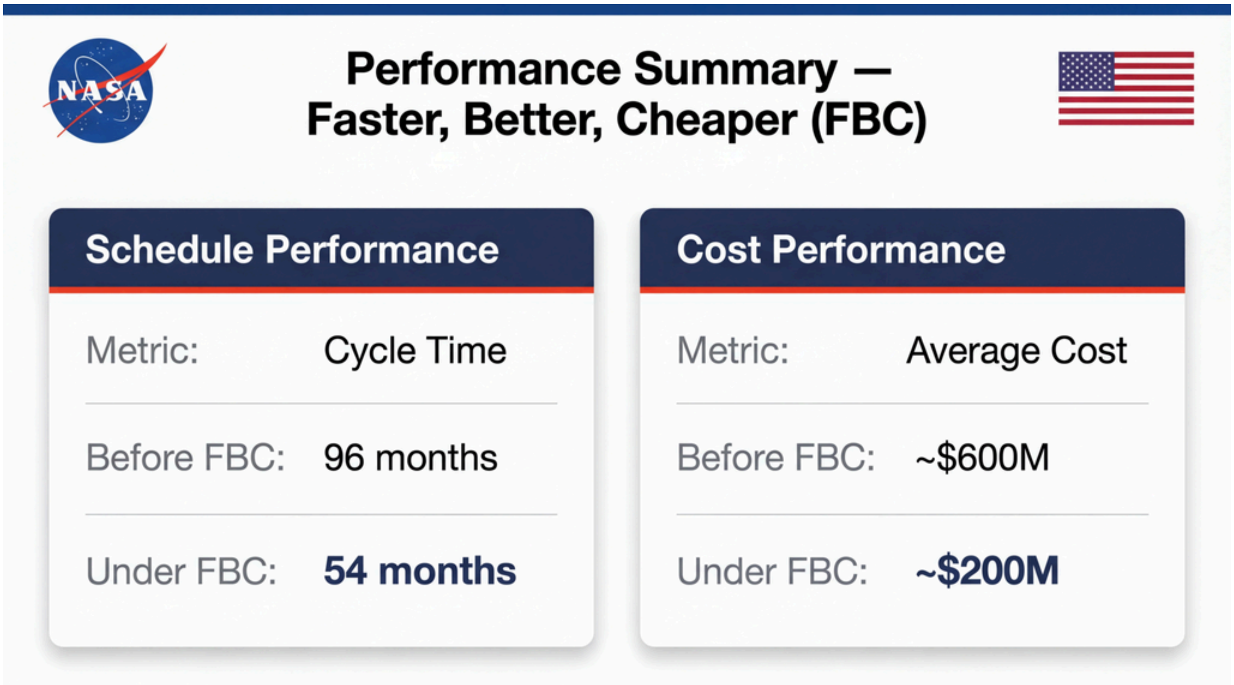

Notice what isn't on that report card: mission success percentage. We lost missions under FBC, and the failures taught us more per dollar than some of the successes.

We did it all on a flat nominal budget to boot. When I arrived, NASA's budget was about $14.3B, and projections had it rising sharply toward $25B within a decade. On my first day I committed to a flat budget for ten years, absorbing inflation without additional funding. Over the decade, this returned ~$40B in "area under the curve" savings to the taxpayer.

To live inside these constraints, you have to change the machine. The civil-service workforce went from about 25,000 to about 17,000 through voluntary buyouts (not forced layoffs). HQ shrunk from 6,000 to ~1,700 people, because more than half of these jobs were a management layer creating more "work," and not new missions.

The same logic of shrinking your mission into a simple, legible set of metrics scales well beyond spacecraft.

I like cycle time for almost any product.

SpaceX, in its first chapter, used the now-famous $/kg (dollars per kilogram) as its north star.

Figure out what the $/kg of your industry is. Pick one or two simple, legible metrics that encode the doctrine you want to run. Then run the organization against that.

SECTION 002

THE SUBSTRATE THAT DISSOLVED

In the last decade of the Cold War, industrial towns like El Segundo ran a complete loop, with proximate, controllable capability from physics to parts. Machine shops doing advanced metallurgy were a short drive away. So were the CEOs of semiconductor fabs you could call directly. TRW, for example, had a local 100,000-sqft. experimental microelectronics facility processing CMOS and gallium arsenide. You could call up Barry Dunbridge, who ran that facility, and say: "Barry, I need 2 dB less noise in my front-end preamp," and he would go off and invent the damn thing that made the system work.

That ecosystem has largely dissolved. Across our industrial clusters, shops shut down, facilities and fabs went dark, and the dense loops dispersed and often disappeared. You cannot run FBC on today's substrate. To get our mojo back, we first have to name the five ways our substrate has degraded – the five channels through which we've installed SWE as the default.

1— SUPPLY CHAINS

First we outsourced steel, industrial chemicals, and advanced materials. Next, it was subassemblies, factory automation, grid-scale batteries, nuclear components, and semiconductors. As we pushed step after step offshore, margins improved on paper, but we lost control over the pace and feasibility of our own work. All of the sudden, the nation's storied manufacturers who remained started to look more like assemblers. FBC requires a portfolio mindset, which requires throughput: parts available, alternatives at the ready, vendors competing, and enough redundancy that a program can swap one path for another without waiting a year. We need to get this throughput back.

2—WORKFORCE

Barry Dunbridge was process control embodied as a human being. FBC relies on a deep bench of that competence: engineers, technicians, machinists, and operators who can move fast without cutting corners on physics or safety. Today, that bench of people who can do the work are aging out faster than they're being replaced.

The U.S. will graduate on the order of 400 petroleum engineers, 450 nuclear engineers, and 300 mining engineers this year, a rounding error compared to China. Credentials ≠ competency, to be sure, but the point stands. How can we seriously tell ourselves we are reshoring energy infrastructure, restarting nuclear reconstruction, and standing up midstream mineral processing with rookie numbers like these?

A thin bench rate-limits you. It forces you to slow down and try to make every shot perfect. You start treating failure as unacceptable. And all of the sudden, you've created the conditions for SWE to take over your organization.

3—ITERATION INFRASTRUCTURE (MACHINE SHOPS, TOOLING EQUIPMENT, ETC.)

In the South Bay era, advanced capability was so clustered that a problem could move from design to fabrication to cut to test, and back, all within a week and one zip code, rarely touching bureaucracy. As that base of shops, tools, and test facilities has depleted and dispersed, "try again" often means sending a file across the ocean and waiting months to get a physical part back. "Try something new" means picking from a TSMC catalog.

This destroys clock time and increases unit cost. And once you start to lose the habit of physical reps, you start demanding certainty upfront: a straight shot to feature-complete. FBC assumes you learn your way to the right answer through many parallel, cheap shots on goal, not with the one mythical "perfect design."

4—COMPUTE ARCHITECTURES

GPUs are a freak of nature, and really, an accident of history. Built for rendering pixels, they turned out to be exceptional matrix multiplication machines, and became the entire substrate of America's AI stack. GPU requirements now drive everything upstream, and are responsible for a non-trivial amount of basis points of GDP growth each quarter. We're planning ~100 GW of new datacenter capacity (e.g., roughly Japan's current electricity use) to feed this single architecture, because we've defined progress as brute-force scaling of this one class of device.

Alternative paradigms could change the physics of the problem: silicon photonic accelerators, analog in-memory compute, optical interconnects, neuromorphic or thermodynamic computing, and so on. It's not that the GPU path is wrong. Transformer inference at scale may genuinely require dense parallel arithmetic. Fine. FBC says pay the price when the physics demand it. But paying the price doesn't mean building big and stopping there. It means scaling what works today while running parallel bets that could break the compute/energy relationship and change unit economics tomorrow. But we're not doing that. Sunk-cost bias and path dependence (which have the signature SWE smell) are crowding out alternative approaches, and relegating them to scientific curiosities.

5—The Design–Production Gap, and Menus Instead of Processes

For a generation, American industry sang itself a comfortable lullaby: design was "strategic" and high-margin, while fabrication was "commodity" and could be done anywhere. By the 2000s, fabless was the orthodoxy. Today, most American engineers go to a catalog. Everyone sends designs to the same foundries, gets the same process design kits, and works inside the same parameters. We traded a culture of invention at the design–production interface, and the seams of process in between, for a culture of menu selection. You can't run FBC off a menu.

SECTION 003

GUESS WHO'S RUNNING FASTER, BETTER, AND CHEAPER?

As America drifts toward Battlestar Galactica mode, China is running something that resembles FBC at national scale. It has a portfolio of many distributed bets, rapid iteration, tolerance for failure, and a "scale the winners" attitude.

On EVs and batteries, Beijing didn't back a single "God car" or "God cell." It let hundreds of subsidized firms fight it out under a protected market umbrella, learning their way to cost and performance through volume, iteration, brutal competition, and eventual consolidation. Plants, pack designs, and integration strategies were tried, copied, discarded, and refined in parallel.

With open-source foundation models, the Chinese technology ecosystem even branded this process: the "Hundred Model War." From 2023 to 2025, hundreds of Chinese foundation models were trained, eventually producing a hundred or so considered substantial enough to matter. The state tolerated obvious duplication and waste in service of building infrastructure, talent, and application diversity. Again: many bets, rapid learning, scale what works, kill what doesn't.

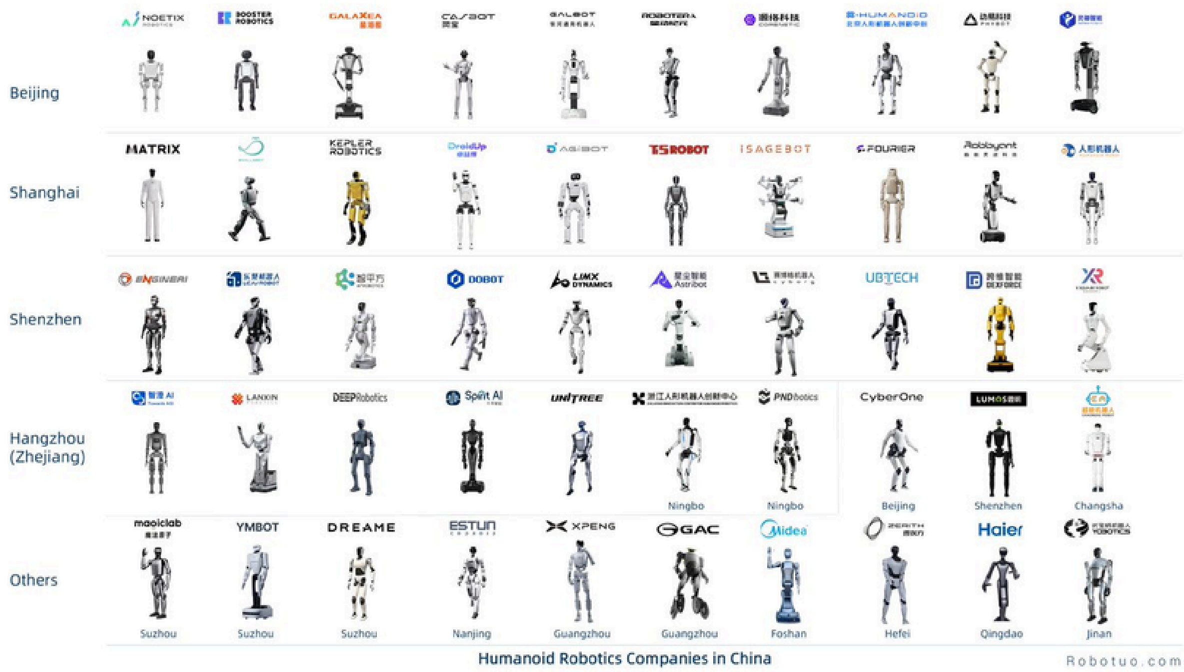

A map of just 50 of the 100+ companies in China developing humanoid robots.

Now China is running the table with this playbook across all of Embodied AI. More than 100 firms across the Mainland are developing humanoid robots, with the leading manufacturers already shipping multiple thousands of units per year. Dozens more offer autonomous forklifts, delivery bots, AGVs, and quadrupeds. In autonomy (specifically L4 passenger vehicles), despite the U.S. holding the architectural crown, China is on pace to finish 2026 with more commercial robotaxis in service than the United States. [See our dispatch from CES for more on this](#).

China has saturated its domestic market with drones, robotaxis, humanoids, and (supposedly) lights-out factories. No one program is singularly brilliant. In fact, on average, they are probably below par. But the sheer volume of attempts generates enough learning to compound. This looks a whole lot like the logic of FBC, except it's being run across entire industries simultaneously. And we're still the ones building the Battlestar Galacticas.

WHAT'S NEXT?

I spent a decade proving this doctrine works inside one agency on a flat budget. The substrate that made it possible — the shops, the workforce, the supply depth, the tight loops from design to test — has since dissolved underneath us. Rebuilding it is a multifaceted problem, but it's solvable, and one that America is uniquely qualified to tackle.

Next week we'll focus on Faster, Better, Cheaper 2.0 – rebuilding the substrate and running the doctrine at the scale of a nation. A sneak peek:

- Five infrastructure layers, each tied to an FBC principle.
- A phased roadmap aligned to the most relevant, executable time unit (American presidential terms).
- A proposal for regional FBC clusters and corridors modeled on what Hsinchu and Shenzhen did in a decade.
- And finally, we'll create a scoreboard that distills the job to be done into a simple set of metrics that we can use to keep ourselves honest as a nation.

So stay tuned for Part II, which will drop right here in one week. And if you're not already subscribed to Per Aspera, [you can do so here](), so you can read it right when it drops.