# From Breakthrough to Breakout

Why the Best Battery Chemistry in the World Won't Matter If You Can't Build the System Around It

**Sinéad O'Sullivan**


## IDEA IN BRIEF

### THE PARADOX

Battery demand is at an all-time high: Grid storage is surging, AI data centres are consuming unprecedented power, and every major economy has declared batteries a strategic priority. Yet the industry is contracting—factory builds are paused, companies are declaring bankruptcies, and capital markets are closing. The problem is that almost nobody can actually make breakthrough batteries reliably, at scale, at consistent quality.

### THE TRAP

The conventional playbook—to prove the chemistry, build a gigafactory, and sell to car companies—will not work. Nine hundred gigawatt-hours of global manufacturing capacity sits unused, while it becomes clear that scaling is not a capital problem, it's a yield problem. And "should we manufacture in China" has become the single highest-stakes question in the industry.

### THE PATH

Get manufacturing yield right before scaling up. Sell to consumer electronics first—the highest-margin, fastest-cycle market. Qualify for US defence contracts in parallel. And if you enter China, build a strict firewall that keeps your best technology at home.


There's a strange thing happening in the battery industry right now, in that demand for batteries is at an all-time high—electricity storage installations surged sixty-six percent in a single year, AI-driven data centres are consuming power at rates that have surprised even bullish forecasters, and announced US manufacturing capacity has doubled since 2022 to over a thousand gigawatt-hours. Every major economy on earth has declared battery manufacturing a strategic priority.

And yet the industry is falling apart. Projects worth billions of dollars have been paused or cancelled amidst the federal EV tax credit expiring in September 2025. Powin, once one of America's top energy storage companies, filed for bankruptcy. Amidst investment drying up. Confidence in this market has, understandably, collapsed.

How can both of these things be true?

The answer is that the battery industry's problem was never demand. It was never even technology—the chemistry coming out of American labs right now is genuinely extraordinary, better than anything China's established manufacturers currently produce. The problem is that almost nobody can actually make these batteries reliably, at scale, at a consistent quality. And that turns out to be a completely different kind of problem than the one most people think they're solving.

For a company with genuinely differentiated chemistry, this paradox is the opportunity.

The weak are exiting the market while the undifferentiated firms are struggling. And the policy environment—a US production tax credit worth thirty-five dollars for every kilowatt-hour produced, import tariffs exceeding eighty-two percent on Chinese batteries, and graphite duties approaching two hundred and twenty percent—has made domestic manufacturing economically viable in ways that were unthinkable three years ago.

But the window to take advantage of this opportunity is narrow, the risks are interconnected, and the conventional playbook—of proving the chemistry, building a gigafactory, and selling to car companies—will not be difficult to enact after 2026.

## THE FACTORY THAT WORKS BEATS THE FACTORY THAT'S BIG

When a battery company announces a new gigafactory, what it is really announcing is a building. What matters is whether the production line inside that building can produce cells that actually work—consistently, predictably, at a defect rate low enough to sell. In manufacturing, this is called yield: the percentage of output that meets specification. And yield is the single thing that separates companies that ship product from companies that burn cash.

The global battery industry currently has about nine hundred gigawatt-hours of announced manufacturing capacity sitting unused. Not because there is no demand, but because the factories cannot reliably produce what they have promised. New chemistries make this harder, not easier—every novel material and cell design requires new production techniques, new quality testing, new ways of catching failures before they reach customers. None of this can be bought off the shelf. Rather, it must be learned, production run by production run.

The counterintuitive implication is the following: that the smartest thing a battery company can do right now is not build a gigafactory. It is actually to build a smaller production line, get the yield right, and prove that it can manufacture consistently before scaling up. The US production tax credit is paid on output, not ambition. Therefore, it rewards batteries that roll off a functioning production line—not press releases about future gigafactories. And that incentive is politically contingent; it exists within a shifting legislative landscape that cannot be assumed to last indefinitely. Every month spent pursuing scale before achieving manufacturing consistency is a month in which potential credit revenue goes unclaimed, inside a policy window that may close without warning.

Prioritising yield before capacity is therefore not a sign of caution or conservatism. It is a strategic sequencing choice that separates companies that generate revenue from those that generate announcements.

## SELL TO THE RIGHT CUSTOMER FIRST

Most battery startups want to sell to car companies. It is the obvious move—electric vehicles are the biggest battery market in the world. It is also a trap.

However, automotive customers take three to five years to qualify a new battery supplier; they are extremely price-sensitive, especially since the US EV tax credit expired; and they already have relationships with massive Chinese manufacturers who can undercut almost anyone on cost. A startup entering that market is bringing a knife to a price war.

Consumer electronics—smartphones, wearables, power tools, drones, portable medical devices—is a far better starting point. These customers pay two to three times more per unit of battery capacity than car companies do, and they make purchasing decisions in twelve to eighteen months, not five years. Further, they care about performance in ways their end users can actually feel: a phone that charges in half the time, a medical device that lasts twice as long. That is the kind of advantage a breakthrough chemistry can sell.

## MARKET ENTRY SEQUENCING

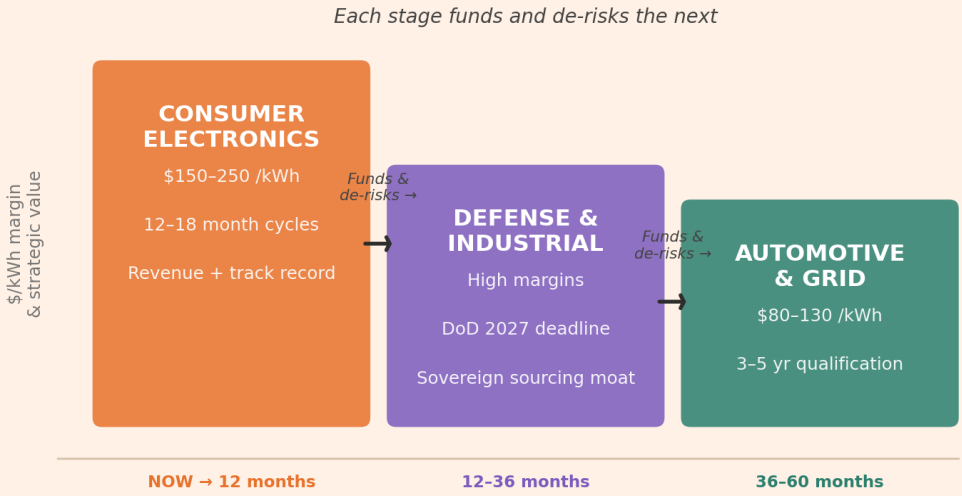


*The optimal sequence: each stage funds and de-risks the next. Start where margins are highest and cycles are fastest, then work toward volume.*

The logic is sequential: Start with consumer electronics to build revenue, manufacturing credibility, and a track record of actually delivering. Then move into defence and industrial applications—the US military faces a 2027 deadline to eliminate Chinese-sourced materials from its batteries, and a domestically manufactured cell with proven reliability is exactly what procurement officers are looking for. Defence revenue is high-margin, long-duration, and provides a strategic moat that no Chinese competitor can replicate in the US market.

Then, and only then, approach automotive and grid storage—from a position of established manufacturing track record rather than unproven promise. Each stage funds and de-risks the next.

*Nine hundred gigawatt-hours of capacity exists on paper. What doesn't exist on paper is a manufacturing process that works.*

## THE CHINA QUESTION

This is where it gets genuinely difficult.

China is the world's largest battery market. The biggest smartphone companies design and assemble there. The deepest pool of battery engineering talent works there and the most mature manufacturing supply chains run through there. For a company selling into consumer electronics, it is very hard to avoid China entirely.

It is also, right now, the most dangerous place in the world to put proprietary technology.

**The case for building there is real.** China's consumer electronics supply chain is unmatched. The major smartphone manufacturers—including those selling into global markets—design and assemble in China. A battery company targeting this market cannot credibly serve its full customer base without manufacturing nearby. Local production also avoids the tariff exposure that makes exporting finished cells to Chinese customers prohibitively

expensive. And access to China's battery engineering talent and established component suppliers can accelerate learning in ways that simply are not replicable in the US today.

**The case against China is existential.** In late 2025, Beijing classified lithium battery technologies as dual-use items—meaning the government can decide what knowledge, equipment, and people are allowed to leave the country. Separately, China now requires government approval before any company can transfer key battery manufacturing techniques abroad. For a US company operating in China, this means any production process developed there could be subject to forced sharing, mandatory licensing, or outright seizure under national security rules.

And it cuts both ways. The US Department of Defense has blacklisted CATL, China's largest battery manufacturer, and over a hundred other Chinese companies. An American startup with significant Chinese operations risks being locked out of US defence contracts, government-backed automotive supply chains, and politically sensitive infrastructure projects—which happen to be some of the most valuable and defensible markets available.

**The answer is structured separation.** Set up a Chinese entity that is legally, operationally, and technologically distinct from the US parent. Different people, different intellectual property, different manufacturing processes. The Chinese operation uses an older, less advanced version of the company's technology—good enough to compete in China's consumer electronics market, but not the cutting-edge chemistry that makes the company valuable. The crown jewels stay in the US.

### THE STRUCTURED SEPARATION MODEL

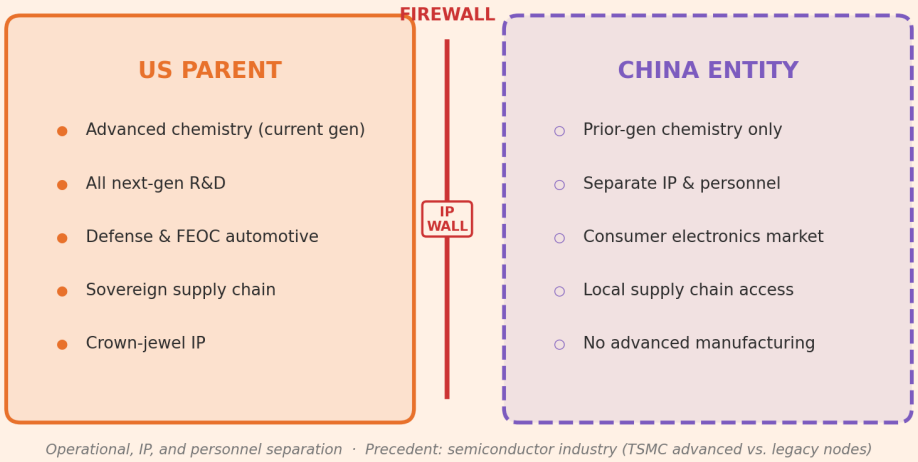


*The structured separation model: completely distinct entities on each side of the firewall. This approach has precedents in the semiconductor industry, where companies maintain strict separation between advanced and legacy production across jurisdictions.*

This model is not risk-free—nothing involving China is. But it is the only approach that does not force a binary choice between the world's largest market and the intellectual property that makes the company worth building.

## THESE RISKS DON'T COME ONE AT A TIME

The risks facing a scaling battery company are not independent. They interact and compound. A deterioration in US–China relations simultaneously makes it more dangerous to operate in China, disrupts the supply of raw materials your factory needs, and accelerates the political pressure to manufacture domestically. A company that treats these as separate items on a risk checklist will be perpetually surprised. One that understands how they connect can plan ahead.

The single most important risk is the most boring-sounding one: *yield failure*. If you cannot make batteries consistently on a production line,

everything else falls apart. You cannot earn the production tax credit. You cannot meet customer delivery timelines. You cannot build the track record that defence and automotive buyers need before they will commit. Yield failure is upstream of everything.

The second-order risk is supply chain concentration. China controls roughly eighty percent of the world's cathode material processing, over ninety percent of anode materials, and ninety-four percent of lithium iron phosphate battery production. A company that builds supply chains for critical materials outside of China—particularly graphite, where Chinese control exceeds ninety percent and export restrictions are already in effect—holds an advantage that goes far beyond unit economics. It becomes qualification-proof: the only supplier that defence and regulated-industry buyers can actually use.

## THE RISK MAP

| RISK | WHAT IT MEANS | SEVERITY | WHEN |
|---|---|---|---|
| **Can't make it at scale** | Lab results don't translate to factory-floor consistency | Critical | 6–18 months |
| **Policy reversal** | The US production tax credit gets cut, restructured, or expires | High | 12–24 months |
| **China takes your IP** | Technology developed in China gets forcibly shared or seized under national security rules | High | Ongoing |
| **Materials cut off** | China restricts exports of the raw materials your factory needs | High | Already happening |
| **Tariff escalation** | New tariffs on components make your costs unpredictable | Medium–High | 2026–2027 |
| **Chinese price war** | Chinese incumbents slash prices to undercut your entry markets | Medium | 2–3 years |
| **Slow customer sign-off** | Customers take longer to qualify your battery than planned | Medium | 18–36 months |
| **Funding dries up** | Investors retreat from hardware startups | Medium | Ongoing |

*The critical insight: many of these risks move together. When US–China relations deteriorate, IP risk, supply chain disruption, and policy acceleration all intensify at the same time.*

## THE WINDOW

The hype era of the battery industry is over. The period of announcement-driven euphoria, easy capital, and gigafactory press releases has given way to something harder and more interesting: an industrial era that rewards execution over ambition, working factories over announced ones, and strategic clarity over opportunistic expansion.

For a company with breakthrough chemistry, this is not a hostile environment. In fact, it is precisely the environment in which genuine technical differentiation becomes decisive. But a better battery does not, by itself, build a business. It has to be paired with manufacturing discipline, the right market in the right order, supply chains that do not run through your competitor, and a clear-eyed answer to the China question.

The window is open. The policy tailwind is real. The demand is structural. But the clock is political, and no one controls it. The companies that act with precision in the next eighteen months will define the industry for the next decade. The ones that mistake announcements for execution will join the nine hundred gigawatt-hours of capacity that exists on paper and nowhere else.

*The frameworks in this article are developed formally—with dynamic models, quantified architecture gaps, and cross-domain evidence—in two companion papers by Sinéad O'Sullivan. Copies available on request: s@sinead.co*