Why Geoeconomic Strategy Keeps Failing

The Architecture Problem Companies Miss

Sinéad O'Sullivan

IDEA IN BRIEF

THE PARADOX

Companies have never been better equipped to understand geopolitics. Consider that Chief Geopolitical Officers, political intelligence units, and scenario planning functions are now standard at major multinationals; yet their strategic bets keep failing anyway: friend-shoring stalls, supply chain diversification from China does not materialize at scale, new market entries underperform. The standard post-mortem blames bad forecasting or slow execution, but that diagnosis is wrong.

THE ARGUMENT

The problem is trying to execute strategies into markets whose supporting architectures do not yet exist. Technical readiness—the capability to produce the good or service—has outrun both market readiness and institutional readiness, hence companies are making sound strategic bets on structurally incomplete markets. The gap is not political; it is architectural.

THE IMPLICATION

There are only three rational responses to an architecture lag problem: build the missing architecture, wait for it to mature, or identify the actors who are actually building it and partner there. Most corporate geoeconomic strategy deploys none of them.

Companies have invested heavily in understanding geopolitics. The past decade saw the creation of an entire industry: chief geopolitical officers, political risk subscriptions, scenario planning workshops, and dedicated intelligence units modeled on government practice. McKinsey, BlackRock, and Goldman Sachs now publish geopolitical outlooks. The World Economic Forum has a geopolitical risk index. By any measure, corporate awareness of the political environment has never been higher.

And yet the strategies that awareness is supposed to inform keep failing. Friend-shoring—the reorientation of supply chains toward geopolitically aligned partners—has not materialized at the scale governments and executives announced. CHIPS Act subsidies have attracted fab announcements but not the ecosystem density that makes semiconductor manufacturing competitive. European hydrogen strategies have produced roadmaps but not markets. These examples point to the fact that the gap between strategic intention and operational reality is not closing.

The conventional diagnosis is familiar: companies underestimated political complexity, moved too slowly, or failed to align incentives across partners. The remedies proposed are familiar too: better political intelligence, faster decision-making, more coordinated government policy. These diagnoses are not wrong exactly, but they are incomplete in a way that is important to outcomes: they treat a *structural* problem as an *execution* problem.

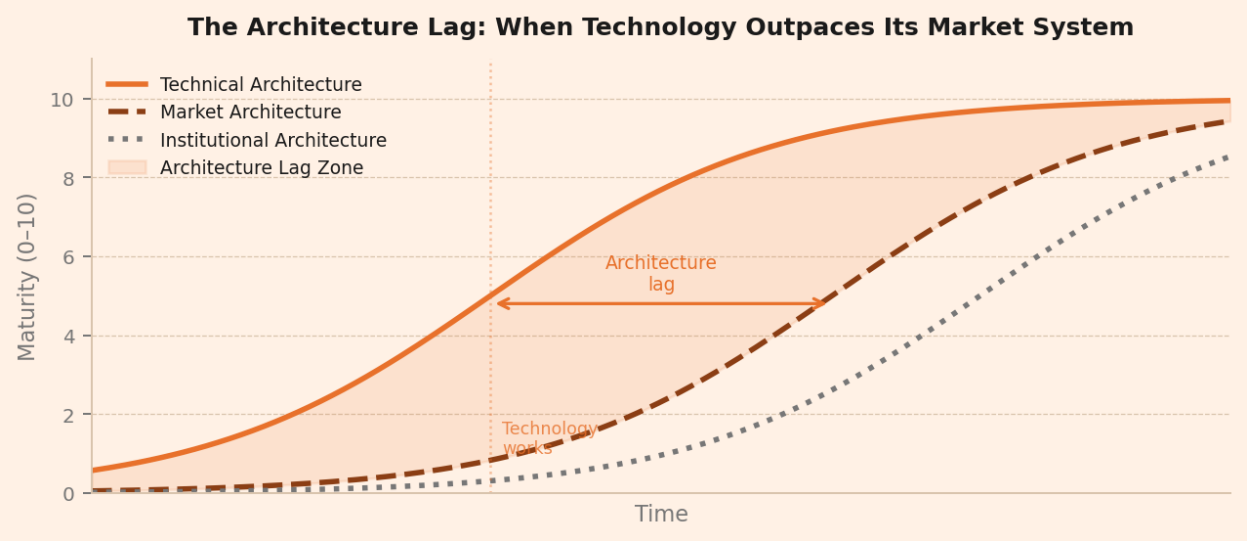

Figure 1. When technical maturity outpaces market and institutional readiness, the result is an architecture lag—a period in which the technology works but the economic system needed to support it does not yet exist.

The Actual Mechanism

Market formation is not a single event. It requires the co-evolution of three distinct architectures: technical, market, and institutional. *Technical architecture* encompasses the engineering systems, production processes, and supply chains that make a product physically possible. *Market architecture* encompasses pricing mechanisms, demand signals, financing instruments, and competitive structures that make it economically viable. *Institutional architecture* encompasses the regulatory frameworks, standards bodies, liability rules, and public agencies that make it legally operable and socially legitimate.

These three architectures do not mature at the same rate: technical architecture often moves faster—it is the product of focused R&D investment and engineering ingenuity. Conversely, market and institutional architectures are slower because they require the coordination of many actors with different interests: regulators, financiers, insurers, customers, and competing producers. So, when technical maturity significantly outpaces the other two, the result is an *architecture lag*.

Architecture lag is not the same as early-stage technology, which exists within a functional market system, even if it is small or nascent. A technology experiencing architecture lag, on the other hand, exists within a market system that is structurally *incomplete*—the supporting mechanisms are absent or immature, not just young. The distinction matters because the remedies are different. Early technologies need time and capital, whereas technologies in architecture lag need deliberate institutional construction.

"The problem is not misreading the geopolitical environment. It is trying to execute sound strategies into markets whose supporting architectures do not yet exist."

Three Cases

Friend-shoring is the clearest current example. The political logic is sound: reducing dependence on Chinese manufacturing by redirecting supply chains to Vietnam, India, Mexico, and other aligned partners; and the technical capability exists on both sides in that Vietnamese and Indian manufacturers can both produce a wide range of goods. The question though is whether the market and institutional architectures needed to support reorientation at scale are present.

Indeed, they are not. The supplier ecosystems that make Chinese manufacturing competitive—dense

networks of component producers, specialized logistics, trained workforces, and embedded quality systems—took decades to construct and reflect sustained institutional investment. Replicating them requires not just capital but the coordinated buildup of standards bodies, technical training institutions, financing instruments for smaller suppliers, and cross-border regulatory recognition. These are not problems that procurement decisions solve.

The CHIPS Act presents the same structure at the level of national industrial policy. The United States can subsidize the construction of advanced semiconductor fabs, and it has: consider that TSMC, Samsung, and Intel have all announced major US investments. But fab construction is only one layer of the technical architecture. An advanced semiconductor ecosystem requires proximity to materials suppliers, specialized equipment manufacturers, process engineers, and university research programs calibrated to industry need. The institutional architecture—the workforce pipeline, the research linkages, the supplier development programs—lags years behind the concrete.

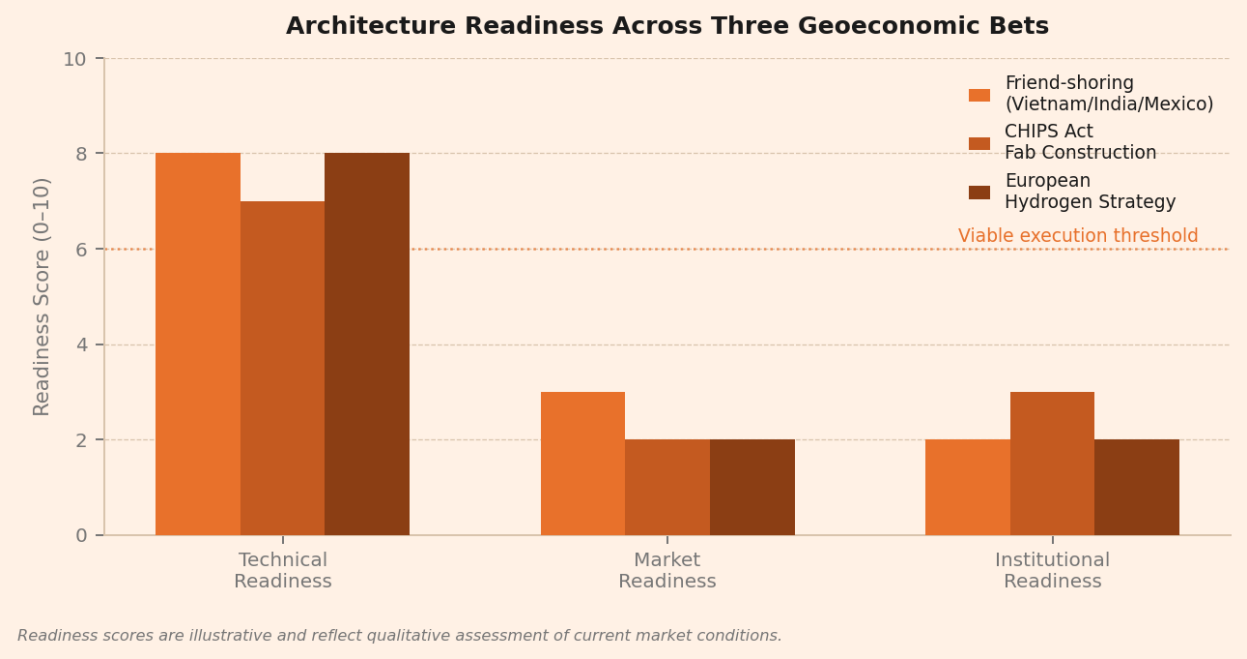

Figure 2. Architecture readiness across three major geoeconomic bets. In each case, technical readiness substantially exceeds market and institutional readiness, producing conditions in which announced strategies cannot achieve their stated objectives within the timeframes governments and companies have set.

European hydrogen strategy illustrates the same dynamic at the level of energy industrial policy. The technical case for hydrogen as an industrial decarbonization pathway is robust: green hydrogen produced from renewable electricity can substitute for fossil fuels in steel, cement, chemicals, and heavy transport. The technology works, however the problem is that the market architecture—stable long-term offtake contracts, transparent pricing benchmarks, cross-border transport infrastructure, and standardized safety and purity specifications—does not yet exist at the scale needed for industrial buyers to commit capital. Without those signals, the investment in production capacity cannot be justified, which means the cost reductions that would make hydrogen competitive cannot materialize.

Three Rational Responses

If the problem is architectural rather than informational, the remedies must be architectural, too. There are three responses that are actually commensurate with the problem, and most corporate geoeconomic strategy deploys none of them.

The first is to build the missing architecture. This means treating market and institutional infrastructure as assets to be constructed, not conditions to be

waited for. A company that needs a regional hydrogen market can participate in the consortia that establish pricing mechanisms, push regulators to accelerate standards development, and structure long-term offtake agreements that give other producers the demand signal they need to invest. This is expensive and slow, and it typically requires collective action across competitors. It is also the only strategy that actually solves the problem.

The second is to wait for the architecture to mature. This is strategically passive but intellectually honest. A company that recognizes the architecture lag in, say, US semiconductor manufacturing can avoid over-investing in a supply chain reconfiguration that will not be viable for a decade, and instead position itself to move quickly when the ecosystem density required for competitiveness is actually present. The mistake, however, is to treat strategic patience as failure.

The third is to identify the actors who are building the missing architecture and partner with them directly. In every architecture lag there are institutions—usually combinations of government agencies, development finance institutions, and first-mover firms—that are doing the slow, unglamorous work of constructing the market and institutional frameworks. Figure 3 maps this for the CHIPS Act: of the fifteen components across three architecture layers, only four are operational. The rest are either under construction or not yet begun. These actors are undervalued by capital markets precisely because their contribution is infrastructural rather than product-based.

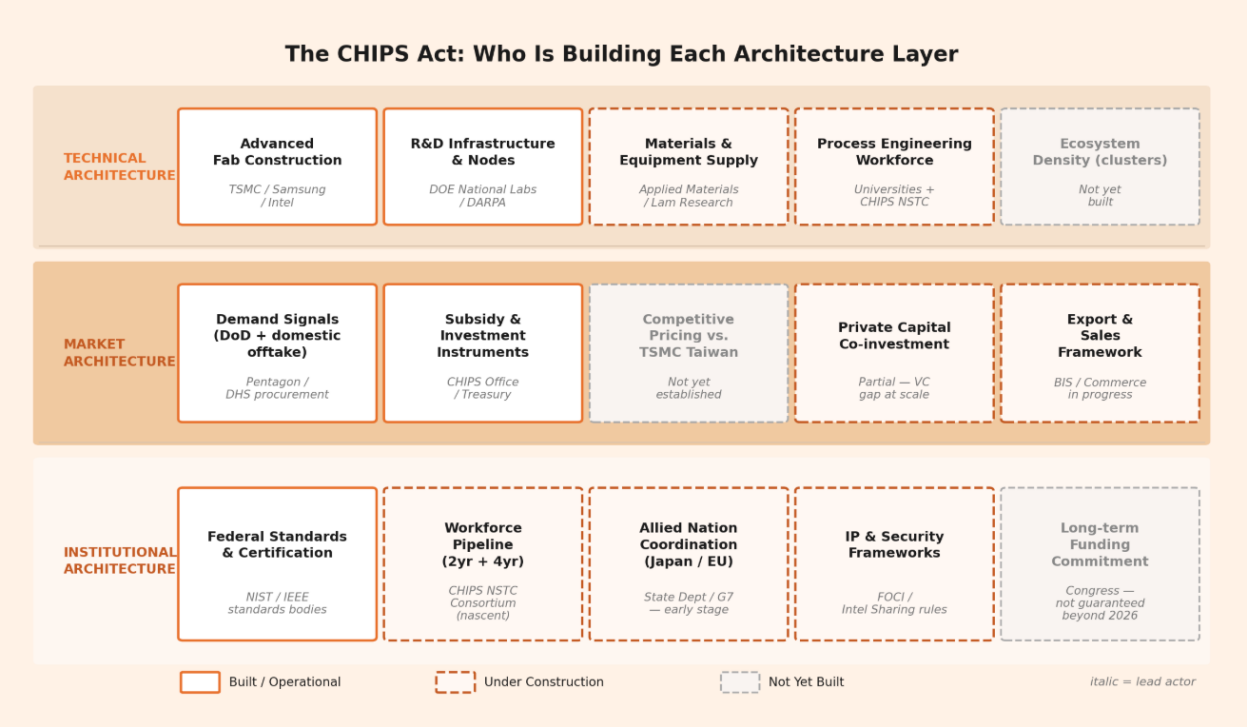

Figure 3. The CHIPS Act: who is building each component of the three architecture layers. Solid borders = built or operational; dashed orange = under active construction; dashed gray = not yet built.

What This Means for Governments

The dominant policy response to architecture lag is to fund the technology and wait for markets to follow. This is not wrong, but it is incomplete. Technology subsidies without parallel investment in market and institutional architecture produce fabs without ecosystems, hydrogen plants without offtake, and supply chain announcements without supplier networks. The public investment lands, but the strategic objective does not materialize.

The more effective intervention is to treat market and institutional construction as explicit policy outputs, not assumed byproducts of technology funding. That means investing in standards bodies before they are needed, designing financing instruments that solve the coordination problems private capital cannot, and building regulatory frameworks on a timeline that precedes commercial deployment rather than following it. It also means being honest about timelines: architecture construction takes longer than election cycles, and strategies that depend on it should be designed accordingly.

The CHIPS Act is instructive here. Of the fifteen architecture components mapped in Figure 3, the two funded most directly—fab construction and R&D infrastructure—are operational. The ones most dependent on deliberate institutional construction—workforce pipelines, allied coordination, long-term funding commitment—remain partial or absent. The next phase of industrial policy requires a shift in emphasis from building things to building the systems that make things viable.

What This Means for Boards

The right question for a board evaluating geoeconomic strategy is not "what is our political risk exposure?" That question produces a risk register, not a strategy. The right question is: which of our strategic bets are contingent on market architectures that do not yet exist, and what is our theory about when and how those architectures will be constructed?

This reframing has practical implications:

It changes the investment thesis for market entry: the relevant uncertainty is not political stability but architectural completeness.

It changes the partner selection logic: the most valuable partners are not those with the best political access but those with the deepest role in building the market infrastructure.

It changes the timeline: strategies contingent on immature architectures need longer horizons and more explicit staging than conventional strategic planning allows.

It also changes what political intelligence is actually for. Better forecasting of government policy is useful, but it is secondary to understanding the institutional construction process—who is building the standards, who is designing the financing instruments, who is negotiating the regulatory frameworks. Those are the actors whose decisions will determine whether a given architecture matures on a useful timeline. They are almost never the ones whose statements appear in geopolitical risk reports.

"The most valuable partners are not those with the best political access. They are those with the deepest role in building the market infrastructure that does not yet exist."

What This Means for Investors

Architecture lag creates a systematic mispricing of risk. Capital markets discount the actors doing institutional construction because their contribution is infrastructural and its value is hard to measure on a product revenue basis. A standards body, a workforce consortium, a government-backed offtake facility—none of these look like compelling investments in isolation. But they are the *rate-limiting inputs for entire industries*. The investors who recognize this earliest hold the most strategically valuable positions when architectures mature.

The practical implication is that investment theses built on architecture readiness rather than technology readiness will systematically outperform in sectors with large architecture lags. The question is not whether the technology works, because it usually does by the time capital is being deployed at scale. The question is whether the market and institutional architecture is far enough along to support commercial returns on a realistic timeline. Answering that question requires a different kind of due diligence: less focused on engineering specs and patent portfolios, more focused on who is building the pricing mechanisms, the offtake structures, and the regulatory frameworks.

The investment of the past decade in geopolitical risk capacity was not wasted, and understanding the political environment is necessary. But it is not sufficient, and the gap between the two has become visible in the consistent underperformance of strategies that were politically well-informed but architecturally naive. Closing that gap requires a different kind of analysis: less focused on political forecasting, more focused on the structural conditions that determine whether a given market can actually form.

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