Freedom to Operate

A Board Framework for Governing Strategic Freedom of Action

By Reinout Schotman
June 2026

Executive Summary

In September 2025 a number of companies in the Gulf discovered that they could no longer use one of the world's leading artificial intelligence models. No government had sanctioned them. No contract had been breached. Nothing about their own conduct had changed. Anthropic, a private company in San Francisco had revised its terms of service, and any firm more than half-owned by investors from China, Russia, Iran or North Korea was now outside the line, wherever in the world it operated and whichever cloud it ran on. The asset these firms wanted to buy was for sale. The permission to buy it had been withdrawn.

The rules had moved. The firms had not.

This is the shift that this framework exists to address. For most of the past forty years boards have worked on the assumption that interdependence breeds restraint: if everyone depends on everyone else, no one can afford to weaponise the relationship. The assumption has quietly inverted. Dependence has itself become an instrument of power. Global supply chains, payment systems, cloud platforms and AI models did not flatten the world so much as wire it, and wiring produces hubs. Whoever controls a hub controls more than their own conduct. They control the choices available to everyone who must pass through it.[1]

The governing question of this framework is concrete: under what conditions does an organisation remain free to pursue its strategy? That is not a question about ownership. A firm can own its strategy, its products, its customers and the infrastructure they run on and still find that the freedom to act on that ownership rests with another party. Ownership determines what a firm possesses. Permission determines what it may do with it. A board that grasps this distinction has stopped measuring its security by the size of its balance sheet and started measuring it by the number of parties who can say no.

Freedom to Operate is not sovereignty and not independence. It does not ask a firm to eliminate its dependencies, which is neither possible nor wise. It asks something more demanding: that a board knows which of its freedoms its strategy cannot survive without, who is in a position to constrain them, what would follow if they did, and what the board intends to do about it before the question is forced.

A framework earns a place at the board table only if it answers three questions in order. The first is *so what*: why the ground has moved, and why ownership no longer settles the matter. The second is *what now*: how to weigh each dependency, choose a stance

[1] Henry Farrell and Abraham L. Newman, 'Weaponized Interdependence: How Global Economic Networks Shape State Coercion', *International Security*, 44(1), 2019; and *Underground Empire: How America Weaponized the World Economy* (Henry Holt, 2023). The broader intellectual lineage runs from Robert O. Keohane and Joseph S. Nye, *Power and Interdependence* (Little, Brown, 1977), who distinguished sensitivity interdependence from vulnerability interdependence, a distinction directly relevant to the irreversibility analysis in Part III.

towards it, and take the decisions that make the stance real. The third is the one most boards never reach in the open, *are we comfortable with this*: which dependencies the firm chooses to live with, deliberately, and whether it can say why. Answering all three is a discipline. Answering only the first is an education. This paper is about the difference, and the difference is a board responsibility, not a risk exercise.

Part I. Strategic Intent: Which Freedoms Matter?

Consider a single dependency, a reliance on one cloud region, and three firms that share it. For a real-time payments platform the dependency is existential, because an outage stops the business within minutes. For a manufacturer running overnight batch jobs it is a manageable inconvenience. For a firm that has not yet moved off its own servers it is irrelevant. The dependency is identical in all three cases. Its strategic meaning is not.

This is why a framework built on dependency cannot begin with a list of dependencies. It has to begin with strategy, because strategy is what gives a dependency its weight. The first question is not what we rely on but which freedoms our strategy cannot do without. These are concrete and worth naming plainly:

- the freedom to serve the markets we have chosen
- the freedom to ship the products we make
- the freedom to obtain the inputs those products require
- the freedom to fund operations and growth
- the freedom to attract and keep the people who do the work
- the freedom to use data
- the freedom to transact

Two of these deserve a further word, because they are less self-evident than the rest. The freedom to use data is now a distinct strategic freedom rather than a subset of obtaining inputs: cross-border data flow is increasingly gated by jurisdiction-specific rules that have nothing to do with the data's commercial value. The freedom to transact is the least self-evidently strategic of the seven, and for that reason the most exposed. Payment rails, correspondent banking and clearing systems are treated as background infrastructure until access to them is withdrawn, at which point a firm discovers that its ability to be paid was never fully its own.

A strategy is, in part, a claim about which of these freedoms the firm intends to exercise and how far. Until that claim is explicit, no assessment of dependency can separate the threats that would end the strategy from the irritations that would merely test it.[2]

The term itself has a useful lineage. Freedom to operate began in intellectual property law, describing whether a product could be sold without infringing another party's rights. It has since entered the language of states, in debates over technological sovereignty and strategic autonomy, and the language of trading blocs. The European

[2] On strategy as the prior act of deciding what business the firm is in, see Peter F. Drucker, *Management: Tasks, Responsibilities, Practices* (Harper and Row, 1973).

Union built its single market on four founding freedoms, the movement of goods, services, capital and persons, and in 2018 formally added data as a fifth, extending the same logic to a resource the original four had no need to name.[3] We bring the term to the level of the firm, because the logic that now governs nations and trading blocs governs companies in miniature. As all three, firms, states and unions, come to depend on technologies, ecosystems and jurisdictions they do not control, all three face the same question about where their freedom of action actually resides.

For a board, the implication is procedural and strict. Before the firm maps a single dependency, it should agree which freedoms are existential to the strategy it has chosen. That agreement becomes the filter for everything that follows. A dependency that touches none of the named freedoms does not deserve board attention, however large it looms on an operational risk register. A dependency that touches one of them deserves the full analysis set out below. Strategy decides what counts.

Part II. Strategic Enablers: What Makes Those Freedoms Possible?

A firm that has built its business on customer relationships is often surprised to discover how little of that relationship it actually controls. The search that brings a customer to its door belongs to Google. The device the customer uses belongs to Apple. The payment that completes the transaction belongs to Visa. The infrastructure that processes it belongs to AWS. The firm owns the business. Others own the pathways through which the business operates. Ownership and control have separated, and the separation is structural rather than accidental.

Behind every strategic freedom sits a set of enablers, and the question that matters is not how many there are but under whose control they fall. Some enablers a firm owns. Some it rents, from a cloud provider, a landlord, a lender. Some it neither owns nor rents but simply relies on, because they belong to an ecosystem or an institution that was there before the firm arrived. The categories are familiar; what distinguishes this analysis is the second axis laid across them, the axis of control:

- **Technology**: owned, licensed, or accessed through a provider whose terms govern what you may build
- **Capital**: held, or supplied by lenders and investors under their conditions and jurisdiction
- **Infrastructure**: operated, leased, or depended upon as a shared utility whose governance sits elsewhere
- **Market access**: direct, or mediated by platforms and distributors whose algorithms decide your reach

[3] Regulation (EU) 2018/1807 on a framework for the free flow of non-personal data, adopted by the European Parliament in 2018 and described on adoption as establishing data as the fifth freedom of the single market, alongside the free movement of goods, services, capital and persons. See also Enrico Letta, *Much More Than a Market* (report to the European Council, April 2024), proposing a further extension of the same concept to research, innovation and knowledge.

- **Ecosystems**: standards, suppliers and complementors a firm joins but does not govern
- **Talent**: present, mobile, and never owned at all

The boundary of the firm has always been drawn by the choice between making and buying.[4] What has changed is that the choice is no longer settled by cost alone. It is increasingly settled by who controls access to what is bought, and under whose jurisdiction that control sits.

Talent deserves separate treatment, because it behaves unlike every other enabler. Capital can be raised in advance and held. Inventory can be stockpiled. Software can be placed in escrow against a supplier's collapse. Talent can do none of these. It cannot be accumulated against a shortage, it cannot be duplicated, and it cannot be locked away for safekeeping. It moves towards opportunity rather than towards ownership. A firm whose strategy depends on access to scarce expertise has placed a bet on the continued openness of the borders, national and organisational, across which that expertise travels. Narrow those borders and the strategy is tested at its foundation, with no amount of capital able to substitute for the people who no longer arrive.[5]

Part III. Strategic Dependencies and Control Points: Who Holds Power Over the Enablers?

A firm that buys a critical component from a single supplier tends to file that supplier under procurement. The filing misses what the relationship has become. The supplier who owns the only source of a critical input, or the platform that sits between a company and its customers, is not merely a counterparty. It holds a measure of control over the dependent firm's freedom of action, and the size of that control is a question worth answering deliberately rather than discovering under pressure.[6]

For each enabler that supports a named freedom, four probes establish where power sits and how durable it is:

- **Ownership**: who actually owns the thing the firm relies on?

---

[4] Ronald H. Coase, 'The Nature of the Firm', *Economica*, 4(16), 1937; and on asset specificity and the cost of losing access to specialised inputs, Oliver E. Williamson, *The Economic Institutions of Capitalism* (Free Press, 1985).

[5] This treatment of capital as an enabler others may govern is consistent with our argument elsewhere that capital is rarely the binding constraint on ambition in the Gulf context. The two claims describe different things: a firm may hold ample capital and still find access to it shaped by who will lend, in which currency, and under whose regulatory reach.

[6] The view of dependency as a power relationship draws on resource dependence theory: Jeffrey Pfeffer and Gerald R. Salancik, *The External Control of Organizations* (Harper and Row, 1978).

- **Jurisdiction**: who writes the rules under which it operates? This is often a different party from the owner, and often the more consequential of the two.
- **Replaceability**: could the dependency be substituted immediately, within months, within years, or not at all?
- **Irreversibility**: has the dependency hardened into a structure? Early success draws adoption, adoption sets a standard, the standard raises the cost of switching, and the rising cost of switching deepens the dependency.[7]

These four probes are, in substance, a map of the firm's control points seen from the position of whoever holds them. Our own work on strategic control points examines the constructive side of this logic, where the prize is to become the node an ecosystem depends on. Here we simply turn the lens around. The control point that is an asset when the firm holds it is a liability when someone else does, and the analysis is identical in either direction. The deeper principle is that power often lies less in the ability to act than in the ability to constrain another's choices.[8] The party who can deny access need do nothing at all. The credible possibility that they might is already a form of control.

Mapping dependencies this way does not yet tell a board what to do. It tells the board where it is exposed and how deeply, which is the groundwork for everything that follows. It also produces the bridge to the next part of the argument. A dependency, once mapped, creates exposure. That exposure turns strategic at the precise moment another actor can deny access. That is the moment a dependency stops being a commercial fact and becomes a question of permission.

Part IV. Strategic Permission: How Networks Become Instruments of Power

Return to September 2025. The company that closed the gate was Anthropic, the developer of the Claude models, and it did so by redrawing the boundary of who may use its technology around ownership and jurisdiction rather than around conduct or price.[9] None of the firms affected had violated their agreements. None had changed their behaviour. The rule had moved, and they had not. A Gulf company with significant Chinese ownership could find a frontier capability withdrawn not by a regulator but by

[7] On increasing returns and lock-in, W. Brian Arthur, *Increasing Returns and Path Dependence in the Economy* (University of Michigan Press, 1994).

[8] The formulation of power as the capacity to constrain rather than to act is a paraphrase of the argument in Thomas C. Schelling, *The Strategy of Conflict* (Harvard University Press, 1960), not a direct quotation.

[9] Anthropic, 'Updating restrictions of sales to unsupported regions', 5 September 2025, reported by Bloomberg and the Financial Times on the same date. The revised terms prohibit access by organisations more than 50 per cent owned, directly or indirectly, by entities headquartered in restricted jurisdictions including China, Russia, Iran and North Korea, regardless of where the subsidiary operates and including access through third-party cloud providers such as AWS.

a paragraph in a private firm's terms of service, and could find it withdrawn even when running on a third-party cloud. The decision reached across borders that no single government's writ would have reached.

This is the clearest current example of what we call involuntary coercion infrastructure: organisations that become instruments of state power, sometimes under compulsion and sometimes by conviction, because a government finds it easier to control a provider than to control the provider's users. The label needs one honest qualification. Anthropic was not coerced into its decision; it has argued publicly for stronger controls and acted ahead of any legal requirement to do so. From the position of the firm that loses access, the provider's motive is irrelevant. Whether the gate was closed under orders or out of conviction, the dependent firm experiences the same thing: a freedom it assumed it possessed turns out to have been a permission someone else could revoke.

The structural version of the same dynamic sits one level up. In January 2025 the United States extended its export controls beyond AI chips to the AI models themselves, and placed most of the world, including the Gulf, in a tier whose access depended on conditions rather than on open trade.[10] The specific rule was later rescinded and replaced, and access for the region was widened rather than narrowed. Yet the structure survived the policy. Later in the year the United States approved advanced chip exports to the region's national AI champions on the condition that they accept security and reporting requirements overseen by a US agency.[11] Whether the gate opens wider or narrower in any given year, the gate exists.

This points to a broader condition that deserves its own name. Much of what the Gulf, and indeed many economies, has built in recent years rests on capabilities owned and governed elsewhere. The AI models, the cloud infrastructure, the financial networks and the semiconductor supply chains that underpin the region's ambitions were designed, built and remain controlled by parties operating under a different jurisdiction. We call this borrowed freedom: the condition of an organisation, or a state, that exercises strategic capability it does not itself control and cannot itself guarantee. The United Arab Emirates has pursued its AI ambition with real seriousness, accepted demanding conditions and navigated this landscape with notable skill. The point is precisely that even a capable and well-governed state finds its freedom of action running through permissions it does not hold. We offer this as a structural observation. For a board, the

[10] United States Bureau of Industry and Security, *Framework for Artificial Intelligence Diffusion*, January 2025, extending export control to advanced AI model weights and establishing tiered country access. The rule was rescinded and replaced under the succeeding administration in May 2025, with access for the Gulf widened rather than narrowed.

[11] United States Commerce Department approval of advanced chip exports to the region's national AI champions, including the UAE's G42, conditioned on security and reporting requirements overseen by the Bureau of Industry and Security, reported November 2025. G42 had earlier realigned its vendor and ownership relationships away from Chinese technology, a step widely read as a precondition of continued access to US capability.

question is whether it knows which of its own freedoms are borrowed, and what it would do if the lender revised the terms.

Two asymmetries make this more dangerous than ordinary commercial dependence. The first is the difference between import control and export control. A restriction on what a firm may buy creates friction; it raises costs and slows the firm down, but the firm survives it. A restriction on whether the firm may continue to sell, ship or operate is a different category of threat entirely, because it can remove the foundation of the strategy rather than merely tax it. The second asymmetry is the one boards trained in competitive analysis are least prepared to recognise. A firm is an economic actor and behaves as one: it seeks profit and can be bargained with on those terms. A state is a political actor, and it may accept a real economic loss to secure a political end. A company can negotiate with a supplier, because both sides want the deal to happen. It cannot negotiate with a national interest, because the other side may treat the firm's loss as the objective rather than the price.

This is where the most influential framework in strategy reaches its limit. Michael Porter's analysis of competitive forces, including the bargaining power of suppliers and buyers, assumes that the actors involved are economically rational and reachable through commercial incentives.[12] The assumption holds for competitors and counterparties. It breaks the moment a political actor enters, because that actor is not playing the commercial game and cannot be priced within it. Freedom to Operate begins where Porter's forces end. Where Porter explains the contest among parties who all want to make money, this framework addresses the party who may be willing to lose money to make a point.

The claim that permission now matters more than ownership is a strong one, and honesty requires naming the conditions under which it would be false. It is weakest where a firm operates entirely within a single stable jurisdiction that also controls the enablers it depends on, because there ownership and permission coincide. It is weaker for enablers that are genuinely fungible, where a denied source can be replaced before the denial does lasting harm. And it is weaker in industries where no node has yet concentrated enough to act as a choke point. A board should test the claim against its own circumstances rather than accept it wholesale. The argument is not that ownership has ceased to matter, but that for a growing class of strategically important enablers it has stopped being sufficient. Ownership determines what a firm possesses. Permission determines what it may do with it.

Part V. Strategic Exposure: How Dangerous Is This Dependency?

A board that has worked this far holds a great deal of diagnosis and no way to act on it. It knows which freedoms matter, what enables them, who controls those enablers and how permission is exercised. What it lacks is a means of telling a dangerous dependency from a tolerable one, so that scarce attention and capital flow to the exposures that could end the strategy rather than to whichever risk happens to be loudest in the room.

[12] Michael E. Porter, *Competitive Strategy* (Free Press, 1980).

Strategic exposure is that means. It adds no new information. It combines what the earlier parts established into one assessment a board can use.

Three dimensions do the work, and the discipline lies in resisting a fourth for the sake of apparent rigour. The first is **criticality**: how essential is the freedom this dependency supports. Criticality functions less as a dimension to be scored than as a filter to be passed. A dependency that supports no freedom the board has named existential does not enter the assessment at all, which keeps the exercise focused and closes the loop back to where the framework began. The second is **replaceability over time**, which folds in what an earlier view treated as a separate axis. Irreversibility is not distinct from replaceability; it is replaceability seen over a longer horizon, the recognition that a dependency easy to replace today may be impossible to replace once lock-in has compounded. The third, and the dominant one, is **permission concentration**: how many parties could deny access, and whether that power sits with one actor or is spread across many. Where permission is concentrated, it overrides whatever comfort the other two dimensions might otherwise offer.

We hold this instrument at two resolutions, because a foundational analysis and a boardroom conversation need different things from it. At full resolution, each qualifying dependency is scored on the three dimensions to produce a profile, and active governance is triggered wherever criticality is high, replaceability over time is low, and permission is concentrated. This scored profile belongs in the firm's own analytical work. For the boardroom, the same logic reduces to a simpler mapping between exposure profile and posture, set out in the Appendix. The argument travels without the matrix. The matrix is a tool, and it belongs in the toolkit rather than the essay.

Part VI. Strategic Posture: How Do We Respond?

A posture is the firm's chosen stance towards a dependency, and the value of the exposure assessment is that the stance should follow from the profile rather than from habit. Five postures cover the field, each belonging to a recognisable region of the exposure map:

- **Accept.** Where a dependency is replaceable and the power to deny it is widely distributed, the right move is to accept it consciously and on the record, rather than spend scarce attention reducing an exposure that can do little harm.
- **Diversify.** Where a dependency is replaceable but permission is concentrated in one party, widen the field of suppliers or routes until no single party holds the gate.
- **Insulate.** Where a dependency is hard to replace but permission is widely distributed, the threat is disruption rather than denial. Build the buffers, contractual protections and contingency that let the firm absorb a shock without losing the freedom behind it.
- **Influence.** Where exposure is high and the firm has standing to shape the party that holds the gate, invest in the relationships, regulatory engagement and reciprocal dependence that make denial less likely.

- **Control.** Where exposure is highest, integrate the dependency, pre-commit resources against its loss, or move to occupy the control point oneself. This is the posture that connects most directly to our work on strategic control points, where the strongest position is to become the node others depend on rather than to remain dependent on someone else's.

The objective in choosing among these is not self-sufficiency, which is rarely achievable and seldom worth its cost, but optionality, the preservation of room to act when circumstances change. These five postures are not a menu from which a board picks the one it prefers. They are the coherent responses to particular profiles, and the test of a board's seriousness is whether it can explain why a given exposure warrants the posture assigned to it, and at what cost. Every posture spends something. Diversification sacrifices the economies of concentration. Insulation ties up capital in buffers that earn nothing until they are needed. Influence and control are expensive and slow. A posture chosen without reckoning its cost is an aspiration, not a decision.

Part VII. Board Actions: What Must the Board Itself Own?

A common failure looks like diligence. A board identifies a serious dependency, asks management to monitor it and prepare a contingency plan, notes the matter in the minutes, and moves on. It has done something, and it has not done the thing that was its own to do. Monitoring, planning and migration are management's work. The board's work is narrower and cannot be delegated, because it is the part of the decision that carries the accountability.

A single test sorts what belongs to the board from what does not: can the board alone authorise or own it? Setting the firm's appetite for concentration risk passes, because risk appetite is a board prerogative. Recording an accepted dependency in the firm's governance passes. Placing the authority to escalate when permission is threatened, before it is threatened, passes. Fixing the cadence at which exposures are re-examined, and treating that cadence as a board duty rather than a management courtesy, passes. Monitoring the news does not. This is, incidentally, where Porter's framework ends and this one begins. Porter names the forces a firm faces but says nothing about the governance that must answer to them.

The board's actions follow the exposure profiles, which keeps them concrete. For a dependency that is critical and held under concentrated permission, the board sets a stated limit on acceptable concentration, mandates a diversification effort without designing it, and names who holds escalation authority should access be withdrawn. For a dependency that has become irreversible, the board records the lock-in as an explicit, dated acceptance rather than an unexamined default, sets the burden of proof required before any further irreversible commitment, and reserves the strategic capital an eventual exit would demand. For a dependency that runs to a political actor, the board separates political risk from operational risk in its own oversight structure, defines the escalation protocol and assigns its ownership at board level, and names in advance the jurisdictions on which the firm will not allow itself to become dependent, whatever the commercial attraction.

In each case the board is not doing management's job. It is doing the part that cannot be delegated, because accountability cannot be separated from the decision that creates it.

Part VIII. The Conscious Acceptance Decision: Are We Comfortable With This?

Boards rarely regret the dependencies they consciously accepted. They regret the ones they inherited without noticing.

The distinction matters because Freedom to Operate is not the elimination of dependency. It is the conscious acceptance of dependency, and that acceptance belongs to the board rather than to the risk committee, the technology function or the procurement team, because the accountability is strategic and sits with the people who answer for the strategy as a whole.

A board may decide, entirely legitimately, that the firm will depend on a single cloud provider for the great majority of its operations. The decision is not wrong because the share is large. It is wrong only if the board cannot say why. An acceptance that cannot be defended in hindsight was never consciously made. It was a position the firm drifted into and later discovered it held. The difference between the two is whether the reasoning was explicit at the time and would survive being read back a year later. A dependency inherited without noticing is not a strategic choice. It is a vulnerability waiting for the moment when the rules move.

Part IX. Fundamental Principles

Stripped to its load-bearing claims, the framework rests on a small number of principles, each intended to stand on its own.

1. Freedom is strategy-dependent; the same dependency means different things to different strategies.
2. Dependencies are unavoidable. Unmanaged dependencies are a choice.
3. Ownership determines what you possess. Permission determines what you may do with it.
4. Import control creates friction. Export control creates existential risk.
5. Not all dependencies are replaceable, and not all replaceability lasts.
6. Irreversibility is usually invisible until it is too late to reverse.
7. Strategic relevance attracts political attention.
8. Resilience is not independence. It is optionality.
9. Power is asymmetric, because economic actors seek profit while political actors may accept loss.
10. Freedom to Operate is not a risk exercise. It is a board responsibility.

The principle that anchors the rest concerns structural power: those who control the frameworks within which others operate do not merely compete with them, they shape the choices available to everyone else.[13] A board that understands its Freedom to Operate is, in the end, a board that has decided not to let those choices be made for it.

Board Questions

The framework becomes a working instrument when a board can take itself through it. The following questions follow the same path the argument took. They are not a checklist to complete but a way of finding where drift may already have replaced decision.

- Which freedoms are existential to our strategy?
- What do those freedoms rest on, and under whose control?
- Who can deny us permission, whether a government, a technology vendor, a financial institution or a platform owner?
- Which of our dependencies have already become irreversible?
- Are we dealing with economic actors or political actors, and do we understand the difference?
- What happens if permission is withdrawn tomorrow: can we still serve customers, ship products, reach our data, pay suppliers and keep our people?
- Which posture have we deliberately chosen for each material exposure?
- Which dependencies are we consciously willing to accept, with reasons we would be content to read back a year from now?

And this brings us to visualizing the full Freedom to Operate framework.

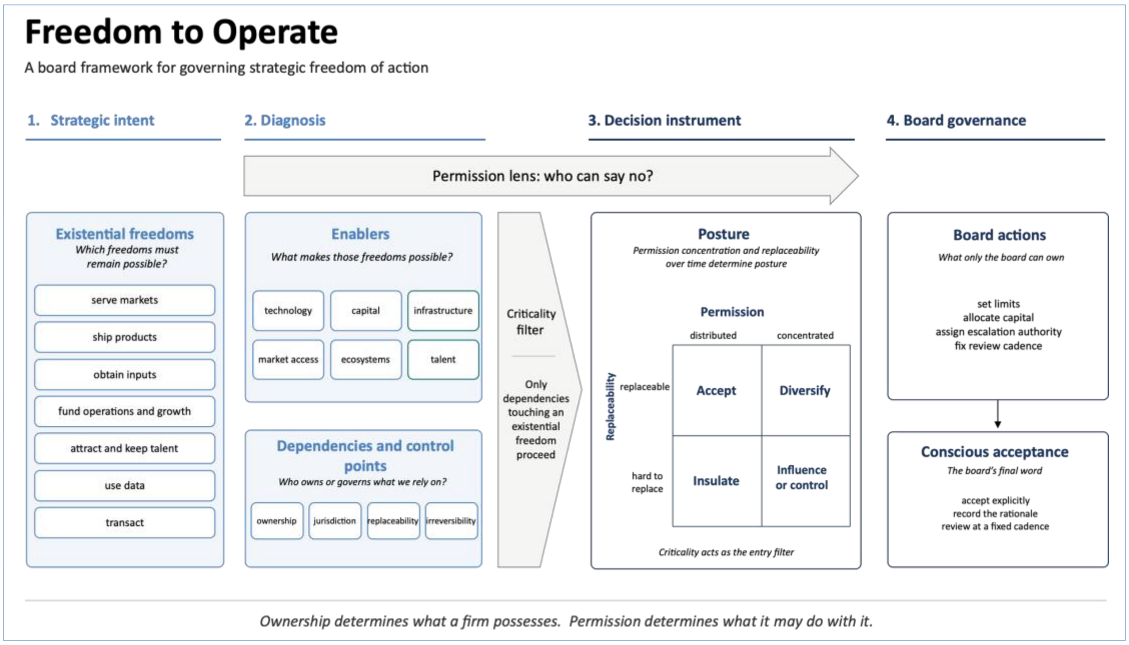

Figure 1: The full Freedom to Operate framework

[13] Susan Strange, *States and Markets* (Pinter, 1988), on structural power exercised through control over security, production, finance and knowledge.

Closing Reflection

Strategy defines where a firm wants to go. Freedom to Operate determines whether it is allowed to get there. A board cannot eliminate uncertainty, cannot rid itself of dependency, and cannot forecast geopolitics. It can do something more useful. It can decide which of its freedoms matter, understand who is in a position to constrain them, and prepare while preparation is still cheap, rather than discovering the constraint at the moment it is imposed.

Read against the three questions, the argument resolves into a single line of thought. The world now turns on permission rather than ownership. The board's answer is a deliberately chosen posture, made real by the acts only the board can perform. And the final test is whether the board can say, consciously and aloud, which dependencies it accepts and why. That last answer is not a risk decision, an operational decision or a procurement decision. It is a board responsibility, and it is the difference between a firm that has secured its freedom of action and one that has merely assumed it.

Appendix. The Exposure Map

The two-by-two set out below reduces the scored-profile analysis to its two most decisive questions: how concentrated is the permission, and how replaceable is the dependency. It is a communication tool and a prompt for discussion, not a substitute for the full assessment. Criticality has already done its work as the filter that determines which dependencies enter the analysis at all. Irreversibility lives inside the time horizon of the replaceability axis.

	Permission concentrated	*Permission distributed*
Replaceable	Diversify	Accept
Hard to replace	Influence or Control	Insulate

A dependency in the bottom-left quadrant, hard to replace and held under concentrated permission, is the existential case. It requires the highest-priority posture and the most demanding board actions. A dependency in the top-right quadrant, replaceable and under distributed permission, can be accepted without further governance burden. The other two quadrants call for Insulate and Diversify respectively, each with its own cost and its own timeline. The posture follows from the profile. The matrix makes that visible.