

Policy Brief No. 233 – March 2026

Realist Innovation Diplomacy: Managing Knowledge Flows Under Rivalry

Mark Daley

Key Points

- Openness in science and technology is a policy variable, not an axiom. In an era of systemic rivalry, states must choose where to collaborate and where to constrain, and the choice should follow from explicit reasoning about what is at stake.
- A two-dimensional map organizes instrument choice: domain salience (Δ) indexes how strongly advances affect relative power; chokepoint intensity (\times) captures how concentrated control is along supply, compute or standards chains.
- Verification capacity is the hinge variable. As the ability to detect leakage improves, trusted coalitions can expand without proportionate risk, making investment in secure enclaves, audit protocols and credible sanctions a strategic priority.
- Proportionality requires diagnostics. Controls should specify, at enactment, the observable indicators that would trigger review, and policy makers should defend the open commons in domains where closure yields no security dividend.

Introduction

States are building walls and bridges simultaneously. In the same policy cycle, governments impose export controls on advanced semiconductors, screen foreign investments in artificial intelligence (AI) firms and restrict researcher access to sensitive facilities, while also fast-tracking talent mobility within alliances, standing up joint research and development (R&D) accelerators and funding trusted data enclaves for coalition partners. The apparent contradiction dissolves once one sees the pattern: this is not incoherence but strategy.

The familiar triadic slogans of science diplomacy — “science in diplomacy,” “science for diplomacy” and “diplomacy for science” — offer little purchase here (The Royal Society 2010). They describe roles and aspirations, not a logic of instrument choice. They presume openness as an unalloyed good and treat closure as deviation. Yet in an era of systemic rivalry, openness is not a default; it is one policy variable among several, whose value turns on what is at stake and who controls the bottlenecks.

About the Author

Mark Daley is the chief artificial intelligence (AI) officer at Western University in London, Ontario, Canada, and a professor in the department of computer science with cross-appointments in five other departments as well as with the Rotman Institute of Philosophy and the Western Institute for Neuroscience. He is also a faculty affiliate of the Vector Institute in Toronto.

Mark's current work focuses on AI and science and innovation diplomacy. In 2025, he completed diplomatic training at France's Institut national du service public and the Académie diplomatique et consulaire and helped lead the Research 7+ summit of Group of Seven research leaders. His policy research on realist innovation diplomacy and AI diplomacy has been cited in outlets including *Politico* and *The Globe and Mail*.

Mark was named in *Maclean's* Power List 2024 of the top 100 Canadians shaping the country and in Constellation Research's AI150, a list of the 150 top global executives leading AI transformation efforts. In October 2024, he was appointed the Natural Sciences and Engineering Research Council of Canada Scholar in Residence in AI.

He has previously served as vice-president (research) at the Canadian Institute for Advanced Research, and as chief digital information officer, special advisor to the president and associate vice-president (research) at Western.

This policy brief advances a realist account of what might be called “knowledge statecraft”: the purposeful design of knowledge flows (for example, research collaboration, data and computation access, talent mobility, standards participation) to shape the distribution of capabilities and bargaining leverage. The core claim is simple: instrument choice in innovation diplomacy follows two complementary logics: *aggregation* (clubs, alliance filtering, preferential openness) and *constraint* (export controls, screening, standards

tactics). Their relative appeal varies systematically with two parameters: how much a domain matters for relative power, and how concentrated the chokepoints are that govern access. Here a “chokepoint” is defined as the ability to control, or at least mediate, access to something of value.

The pay-off is practical. Rather than ad hoc responses to each new dual-use controversy, decision makers can locate domains on a two-dimensional map and select instrument bundles accordingly. The map does not resolve every hard case, but it clarifies *why* cases are hard and *what* should move if a chosen bundle is working. In competitive orders, that discipline — proportionate, revisable and measurable — is how one preserves the epistemic commons that science requires while navigating the security imperatives that rivalry imposes.

The Domain Salience–Chokepoint Intensity Framework

Two parameters structure the strategic environment for any scientific or technological domain: domain salience and chokepoint intensity. Together they define a policy map (see Figure 1) that links domain characteristics to instrument choice.

Domain Salience

Domain salience indexes how strongly marginal advances in a field translate into relative capabilities and bargaining power. High domain salience domains are those where leadership confers meaningful military, economic or coercive advantage and where leakage to rivals carries correspondingly high costs. Leading-edge compute and semiconductor tooling sit squarely in this category: the capability overhang from frontier AI systems, the concentration of advanced chip fabrication and the entanglement of these technologies with defence applications all elevate domain salience. Undersea systems and quantum sensing exhibit similar profiles.

Low domain salience domains, by contrast, are those where advances diffuse broadly, confer few positional advantages or generate

externalities (for example, climate science, pandemic preparedness) that make cooperation the dominant strategy regardless of rivalry. Much of basic astrophysics and Earth observation falls here. The key intuition is that domain salience is not a fixed property of a field; it varies with the state of technology, the proximity of military applications and the structure of competition. A domain can traverse the map as context changes.

Chokepoint Intensity

Chokepoint intensity captures how concentrated control is along supply, compute, data or standards chains (among others). High chokepoint intensity domains feature narrow vendor bases, rare tools, critical intellectual property clusters or committee gatekeeping that creates leverage over access. Semiconductor manufacturing equipment, where a handful of firms supply irreplaceable lithography systems, exemplifies extreme chokepoint intensity (Farrell and Newman 2019; Miller 2022). Standards bodies for wireless communications (viz., third-generation partnership project [3GPP]) and AI management systems (International Organization for Standardization/International Electrotechnical Commission [ISO/IEC 2023] Joint Technical Committee 1/Subcommittee 42) represent a different species of chokepoint: here, leverage flows through editorial control, conformance testing and the ability to set defaults that others are incentivized to follow.

Low chokepoint intensity domains are those where supply is diffuse, alternatives are plentiful and no single actor can credibly deny access. Open-source software ecosystems, broadly published scientific literature and commodity cloud compute (modulo concentration concerns) approximate this condition. The practical significance of chokepoint intensity is that it determines *which* instruments can bite. Denial is effective only where real chokepoints exist; standards tactics pay off only where committees can truly gatekeep.

Two Logics: Aggregation and Constraint

The salience-chokepoint map organizes two complementary logics of knowledge statecraft.

Aggregation collects and aligns. It operates through exclusive clubs, alliance-filtered collaboration, joint R&D infrastructure, trusted data enclaves

and preferential mobility channels. The goal is to raise within-bloc capability faster than rivals can match; pooling talent, harmonizing tooling, sharing evaluation infrastructure and reducing the alignment costs that are otherwise a source of friction for coalition science. Aggregation becomes more attractive as domain salience rises (there is something meaningful to protect) and as verification capacity improves (clubs can confidently expand without risking leakage).

Constraint narrows and denies. It operates through export controls, investment screening, participation limits and standards tactics that fix rules advantaging the coalition or degrading rival access. Constraint becomes necessary where chokepoints allow denial to bind (chokepoint intensity is sharp) and where domain salience is high enough that the security benefits outweigh the innovation drag of closure. The two logics are complements, not substitutes: effective strategy typically combines aggregation within trusted coalitions with constraint at the coalition boundary.

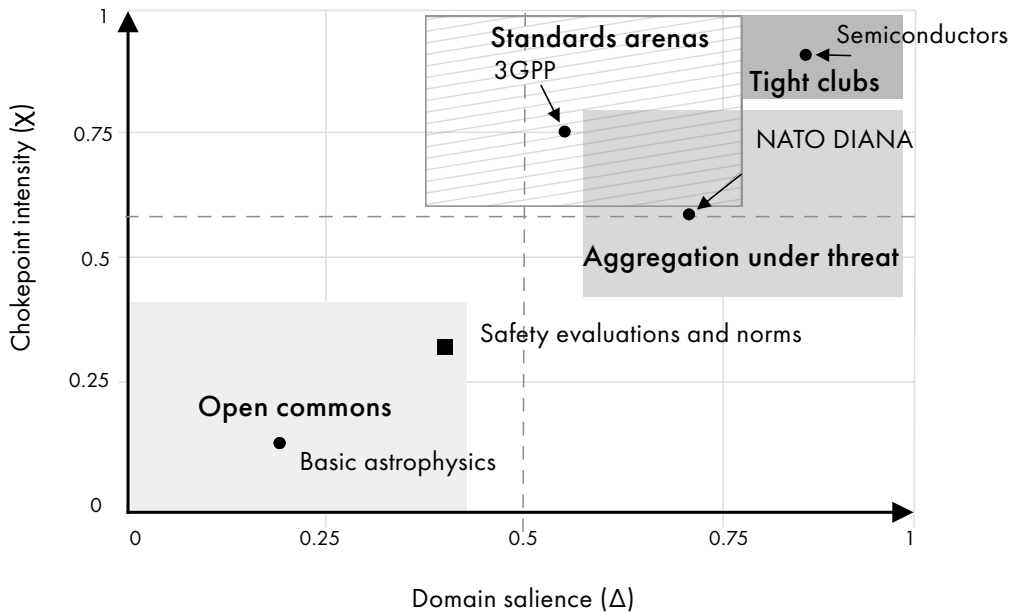
The Quadrant Map

Figure 1 locates familiar policy arenas on the domain salience-chokepoint intensity plane and suggests the instrument tilt appropriate to each region.

High Domain Salience, High Chokepoint Intensity (Upper-Right Quadrant)

Targeted denial and tight clubs dominate. Semiconductor tooling and advanced compute controls exemplify this logic: cross-bloc flows are curtailed while within-alliance collaboration intensifies. Verification capacity (audits, secure enclaves, credible sanctions) is the hinge variable that determines how large clubs can grow. Standards tactics at interfaces (testing protocols, compliance regimes) complement outright denial.

Figure 1: Policy Instrument Regions on the Domain Salience–Chokepoint Intensity Map



Source: Author.

Note: Shading indicates typical tilts: open commons at low salience and weak chokepoints; aggregation under threat as salience rises; standards arenas where chokepoints create rule leverage; and constraint with tight clubs when both parameters are high.

High Domain Salience, Moderate Chokepoint Intensity (Upper-Edge Region)

Aggregation under threat takes precedence. AUKUS¹ Pillar II and the North Atlantic Treaty Organization’s (NATO’s) Defence Innovation Accelerator for the North Atlantic (DIANA) instantiate this pattern: coalition R&D, shared test ranges, trusted computation and mobility pipelines, with denial reserved for acute leakage risks. Investment in reducing alignment costs (for example, harmonized tooling, joint training, interoperable evaluation) expands the feasible scope of collaboration.

Moderate-to-High Chokepoint Intensity, Varying Domain Salience (Right Edge)

Standards arenas function as power platforms. Where market or protocol concentration creates gatekeeping, rule setting can project influence even when denial leaks. Editorial capacity, reference artifacts (test suites, conformance labs, benchmark

data sets) and procedural presence matter more than blunt exclusion. AI governance standards and telecommunications specifications illustrate the stakes (ISO/IEC 2023).

Low Domain Salience, Low Chokepoint Intensity (Lower-Left Quadrant)

The open commons dominates. Denial is unnecessary and counterproductive; clubs are superfluous. Light governance (reproducibility norms, curation, safety baselines, shared infrastructure) addresses coordination failures without restricting flows. Basic astrophysics and much of climate science occupy this space. Crucially, the open commons is not a residual category but a positive choice, appropriate where rivalry does not demand closure.

The map is not static. Domains drift as technology matures, suppliers enter or exit and security perceptions shift. Strategy must therefore be dynamic: placements should be stated with explicit proxies, reviewed periodically and revised when the underlying parameters move.

¹ AUKUS is a trilateral security partnership between Australia, the United Kingdom and the United States.

Policy Levers by Quadrant

The domain salience-chokepoint intensity map is not merely diagnostic; it organizes a menu of instruments. What follows surveys the lever set appropriate to each region, with attention to the conditions under which each instrument is likely to succeed or backfire.

Where Denial Bites: High Domain Salience, High Chokepoint Intensity

When domain salience is acute and chokepoints are sharp, constraint instruments come to the fore. Export controls, deemed-export rules, investment screening and end-use restrictions can materially slow rival capability acquisition, provided they are designed with discipline. The pathologies of denial are well known: over-broad scope chills domestic innovation, drives substitution and alienates allies whose firms bear asymmetric costs. Effective denial therefore requires narrow, published criteria that make the security rationale legible; time-limited licences with explicit review triggers; and investment in verification capacity — secure enclaves, transaction logging and credible sanctions for evasion — so that the boundary of the club can expand as trust accumulates.

The October 2022 US semiconductor controls illustrate both the logic and the risk (Bureau of Industry and Security 2022; Allen 2022). By targeting advanced lithography, high-bandwidth memory and frontier compute, the rules exploited genuine chokepoints. Yet the unilateral character of the initial package imposed coordination costs on allied suppliers and invited acceleration of indigenous substitution efforts. Subsequent rounds of allied alignment — the Dutch and Japanese restrictions of 2023 — partially corrected the imbalance (Second Chamber of the States General 2023). This is a reminder that denial without coalition discipline can be a wasting asset.

Where Aggregation Pays: High Domain Salience, Moderate Chokepoint Intensity

When salience is high but chokepoints are diffuse, the strategic premium shifts to aggregation: pooling talent, harmonizing tooling, sharing evaluation infrastructure and reducing the alignment costs

that otherwise cause friction in coalition science. AUKUS Pillar II, as well as the NATO Science and Technology Organization and DIANA, instantiate this logic, standing up joint accelerators, trusted test ranges and mobility pipelines that raise within-bloc capability faster than rivals can match.

Two variables govern how far aggregation can extend. Verification capacity determines how confident partners can be that sensitive knowledge will not leak; as verification capacity rises, through technical controls, audit regimes and credible enforcement, clubs can admit new members without proportionate risk. Alignment cost captures the friction of joint work: incompatible security classifications, divergent ethics review, mismatched compute stacks, visa delays, and procurement or acquisition inefficiencies. Investment in reducing alignment cost (for example, mutual recognition agreements, common evaluation tooling, harmonized training curricula) expands the feasible scope of collaboration. The policy implication is that aggregation is not a static capacity but a design problem: build the verification and reduce the friction, and the coalition's knowledge base compounds.

Where Standards Are the Battleground

Along the upper edge of the map, where chokepoint intensity is elevated but domain salience varies, standards arenas function as power platforms. Control flows less through denial than through rule setting: who drafts the text, which test suites become normative and whose conformance labs certify compliance and interoperability. AI management-system standards (ISO/IEC 42001), telecommunications specifications (3GPP) and emerging compute-governance frameworks all exhibit this dynamic.

Effective standards statecraft requires sustained presence: delegations with technical depth, editorial stamina and procedural fluency. It requires investment in reference artifacts (for example, benchmark data sets, evaluation protocols, interoperability test beds) that instantiate preferred norms before committees convene. And it requires coalition coordination to prevent fragmentation into rival blocs, each with incompatible conformance regimes. The goal is not exclusion but default setting: shaping the rules so that others find compliance and participation easier than defection.

Where the Commons Should Prevail

In the lower-left quadrant (low domain salience, low chokepoint intensity), an open scientific commons remains the dominant strategy. Denial is unnecessary (no security dividend justifies the cost) and clubs are superfluous (diffuse supply means there is little to gatekeep). Basic astrophysics, much of climate science and large swathes of biomedical research occupy this space.

The policy imperative here is protection, not restriction. Reproducibility norms, open-access mandates, curation infrastructure and mobility guarantees preserve the public-good character of science. The chief risk is security inflation: the temptation to extend controls into low domain salience domains on speculative threat scenarios, thereby chilling collaboration without meaningful security gain. Policy makers should resist domain creep, insist on explicit domain salience assessments before expanding restriction and treat the commons as an asset to be defended rather than a residual to be tolerated.

Recommendations

Five priorities follow from the framework advanced here.

- **Map domains explicitly:** Governments should maintain a living inventory of scientific and technological domains, scored against stated proxies for domain salience and chokepoint intensity. The proxies need not be precise and reasonable people will disagree on weights, but they must be transparent and revisable. An explicit map disciplines instrument choice, exposes inconsistencies (why is domain A subject to controls while domain B, with similar domain salience, is not?) and creates a shared vocabulary for interagency and allied coordination. The alternative is ad hoc response to each new dual-use controversy, with predictable incoherence.
- **Invest in verification capacity:** Verification is the hinge variable that determines how large trusted coalitions can grow. Secure research enclaves, tamper-evident logging, audit protocols and credible sanctions for evasion all raise verification capacity and

thereby expand the set of partners that can be admitted to high domain salience collaboration without proportionate leakage risk. This is infrastructure, not rhetoric; it requires sustained funding, technical standards and institutional homes. The pay-off is a larger, more capable coalition knowledge base.

- **Reduce alignment costs:** Even willing allies face friction — incompatible security classifications, divergent ethics review, mismatched compute stacks and visa delays. Each friction point raises alignment costs and shrinks the feasible scope of joint work. Mutual recognition agreements, harmonized training curricula, common evaluation tooling and streamlined mobility channels are the practical agenda. The goal is to make coalition science easier to do than to avoid.
- **Tie controls to diagnostics:** Export controls and screening regimes should specify, at enactment, the observable indicators that would trigger review — evidence of successful substitution, measurable harm to allied innovation and shifts in the underlying domain salience or chokepoint intensity. Pre-registering these diagnostics commits policy makers to revisit restrictions rather than allowing them to ossify. Proportionality is not a one-time judgment but a continuing obligation.
- **Protect the commons:** In low domain salience, low chokepoint intensity domains, the open scientific commons is not a concession to idealism but the strategically correct posture. Policy makers should resist security inflation; require explicit salience assessments before extending restrictions; and defend reproducibility norms, open-access mandates and researcher mobility as public goods. The commons is an asset; it should be governed, not abandoned.

Conclusion

The argument of this brief is easily stated: in an era of systemic rivalry, openness is a policy variable, not an axiom. Instrument choice in innovation diplomacy follows two complementary logics — aggregation and constraint — whose relative appeal varies with domain salience and chokepoint intensity. A two-dimensional map linking these parameters to lever bundles can discipline what would otherwise be ad hoc, reactive and inconsistent policy.

Nothing in this framework licenses indiscriminate closure. Realism about knowledge politics is not a warrant for security inflation, for treating every scientific exchange as a leakage vector, or for extending denial into domains where the security dividend is speculative and the innovation cost is real. The framework's value lies precisely in its capacity to distinguish: to say here denial is warranted and there the commons should prevail, with explicit reasoning that can be challenged and revised.

Knowledge statecraft will not recede as a policy domain. The technologies that most exercise strategists (for example, AI, advanced compute, autonomous systems, bioengineering) are artifacts of organized inquiry, and the networks that produce them are simultaneously assets to be cultivated and vulnerabilities to be managed. How states navigate this tension will shape the character of international order for decades. The choice is not between openness and closure but between disciplined strategy and drift.

Works Cited

- Allen, Gregory C. 2022. "Choking Off China's Access to the Future of AI." Center for Strategic & International Studies. October. https://csis-website-prod.s3.amazonaws.com/s3fs-public/2023-04/221011_Allen_China_AccessToAI.pdf.
- Bureau of Industry and Security. 2022. "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification." *Federal Register* 87 (197). www.govinfo.gov/app/details/FR-2022-10-13/2022-21658.
- Farrell, Henry and Abraham L. Newman. 2019. "Weaponized Interdependence: How Global Economic Networks Shape State Coercion." *International Security* 44 (1): 42–79. https://doi.org/10.1162/isec_a_00351.
- ISO/IEC. 2023. *ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system*. Geneva, Switzerland: International Organization for Standardization. www.iso.org/standard/42001.
- Miller, Chris. 2022. *Chip War: The Fight for the World's Most Critical Technology*. New York, NY: Scribner.
- Standing Committee on Foreign Trade Development Cooperation. 2023. "Announcement of upcoming export control measures for advanced semiconductor production equipment." Government letter, Government of the Netherlands, March 8. www.tweedekamer.nl/kamerstukken/brieven_regering/detail?id=2023Z04037&did=2023D09406.
- The Royal Society. 2010. *New frontiers in science diplomacy: Navigating the changing balance of power*. London, UK: The Royal Society. www.aas.org/sites/default/files/New_Frontiers.pdf.

About CIGI

The Centre for International Governance Innovation (CIGI) is an independent, non-partisan think tank whose peer-reviewed research and trusted analysis influence policy makers to innovate. Our global network of multidisciplinary researchers and strategic partnerships provide policy solutions for the digital era with one goal: to improve people's lives everywhere. Headquartered in Waterloo, Canada, CIGI has received support from the Government of Canada, the Government of Ontario and founder Jim Balsillie.

À propos du CIGI

Le Centre pour l'innovation dans la gouvernance internationale (CIGI) est un groupe de réflexion indépendant et non partisan dont les recherches évaluées par des pairs et les analyses fiables incitent les décideurs à innover. Grâce à son réseau mondial de chercheurs pluridisciplinaires et de partenariats stratégiques, le CIGI offre des solutions politiques adaptées à l'ère numérique dans le seul but d'améliorer la vie des gens du monde entier. Le CIGI, dont le siège se trouve à Waterloo, au Canada, bénéficie du soutien du gouvernement du Canada, du gouvernement de l'Ontario et de son fondateur, Jim Balsillie.

Credits

Research Director, **Transformative Technologies** Tracey Forrest

Program Manager Grace Wright

Publications Editor Christine Robertson

Graphic Designer Sepideh Shomali

Copyright © 2026 by the Centre for International Governance Innovation

The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the Centre for International Governance Innovation or its Board of Directors.

For publications enquiries, please contact publications@cigionline.org.



The text of this work is licensed under CC BY 4.0. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

For reuse or distribution, please include this copyright notice. This work may contain content (including but not limited to graphics, charts and photographs) used or reproduced under licence or with permission from third parties.

Permission to reproduce this content must be obtained from third parties directly.

Centre for International Governance Innovation and CIGI are registered trademarks.

67 Erb Street West
Waterloo, ON, Canada N2L 6C2
cigionline.org

