

Coercion Without an Author

*How Chokepoint Control Lets a State Compel Without Ordering Anything, and Why
Attribution Rules Cannot See It*

Christopher I. V. Farmer

Independent researcher, Geneva, Switzerland

6 August 2026

Suggested citation. Christopher I. V. Farmer, 'Coercion Without an Author' (2026) preprint.

Status. Preprint. Submitted for publication; not yet peer reviewed.

Keywords. economic coercion; attribution; state responsibility; countermeasures; retorsion; chokepoints; weaponised interdependence; evidential burden

Licence. Creative Commons Attribution 4.0 International (CC BY 4.0).

Contents

Coercion Without an Author	2
How Chokepoint Control Lets a State Compel Without Ordering Anything, and Why Attribution Rules Cannot See It	2
Abstract	2
Note on sources and method	3
Table of Cases	4
United Kingdom	4
International courts and tribunals	4
World Trade Organization	4
Investment tribunals	4
Table of Legislation and Instruments	4
Treaties and international instruments	4
European Union	5
Part I — Introduction	5
1. The problem stated	5
2. Onto what foundation	5
3. What this paper does	5
4. What is new here, and what is not	6
Part II — Coercion Without an Author	8
5. Weaponised interdependence and the premise the remedial literature does not state	8
6. What has actually changed	9
6A. The contrast case: what attributable coercion costs	10
6B. A worked case: Lithuania	10
7. What the attribution literature already knows	13
8. Why the law cannot reach it	14
8A. Deniability through corporate structure	19
9. The pricing consequence	20
Part III — Measuring What Cannot Be Insured	22
10. Why measurement is the only remaining response	22
11. Data and method	22
12. Results	23
13. What this measurement cannot see	26
13A. Why a tariff-line measure understates network exposure	26
14. What the measurement is for	28
Part IV — The Verification–Attribution Gap, Formally	28
15. Why an argument in this form is needed	28
16. The setting	28
17. Two channels of response, and only one of them is constrained	29
18. What the comparative statics separate	31
19. Non-attributive counter-leverage, correctly stated	31
19B. Where this sits in the signalling literature	34
20. Objections	35
Part V — What Follows	36
21. The design problem	36
22. Proposal two: an effects-based evidential mechanism for economic coercion	36
23. The English courts have already been here	42

24. Enforcement, and why a rule nobody enforces is worse than none	43
Part VI — Conclusion	44
25. Reflective analysis	44
26. What would change the analysis	45
27. Final remarks	45
Bibliography	46
Books	46
Contributions to edited collections	46
Journal articles	46
Reports, working papers and official documents	49
Appendix — Verification apparatus	50
What the apparatus cannot do, established by six classes of failure	50

Coercion Without an Author

How Chokepoint Control Lets a State Compel Without Ordering Anything, and Why Attribution Rules Cannot See It

Christopher I. V. Farmer

Abstract

Economic coercion is priced, deterred and remedied only when it can be attributed. Every instrument the international legal order possesses for responding to it — the prohibition on intervention, the secondary rules of state responsibility, countermeasures, the security exceptions of the trading system, and the domestic sanctions architecture built on top of them — presupposes a provable author. This paper argues that the presupposition no longer holds, measures what remains when it fails, and finds that the last consolation the literature offers is not supported by the data.

The doctrinal argument is that six routes exist and five fail at the same point. The prohibition on intervention fails because purpose cannot be shown; state responsibility fails because control cannot be shown to the standard *Bosnian Genocide* preserved; the security exception fails because motive cannot be tested, and its post-2019 jurisprudence consists of one adopted panel report and four appealed into a vacancy; countermeasures fail because the predicate cannot be established to a standard that makes the response safe; and non-violation nullification, which requires no wrongfulness at all, still requires the measure to be shown to have been applied by the respondent, which is the half of the problem that matters. Only creeping expropriation reasons from effect to authorship, and it fails on standing and forum rather than on proof.

The paper's worked case is a pair. In *China — Measures Concerning Trade in Goods and Services* the European Union brought a complaint over the treatment of Lithuanian trade, suspended the panel twice within days of the lapse deadline, and withdrew it on 1 December 2025 without a report or a finding. In *China — Enforcement of Intellectual Property Rights* the same complainant established an **unwritten** Chinese policy from a pattern of individual acts and obtained a binding award. The difference is not written against unwritten. It is whether the conduct leaves a documentary trace in the hands of someone other than the acting state. Judicial orders do; customs discretion does not. That criterion is stated ex ante, it is falsifiable, and it locates the boundary of the paper's own claim.

The formal argument yields a result that is narrower and more uncomfortable than deterrence failure. Where the attainable probability of attribution falls below the point at which a proportionate countermeasure can deter, the countermeasure channel closes — it needs a provable wrongful act and is capped by proportionality — but untargeted retorsion does not, because lawful unfriendly conduct requires no attribution and no

predicate. **Retaliation does not stop; it migrates to instruments that cannot distinguish the responsible from the rest.** The point is not that the cap is low but that it is *legal*: a chosen rule with a lawful substitute permanently beside it. The European Union’s anti-coercion regulation is that structure in legislative form — a determination of coercion as the trigger, non-performance of international obligations permitted only “insofar as” wrongfulness obtains, and every listed measure drafted so that it need not be treaty-inconsistent.

The empirical work is original and its headline finding cuts against the paper’s own framing. Bilateral trade in 2024 at six-digit tariff-line resolution — 11,250,411 flows covering USD 22.94 trillion — is reduced to 323,929 importer-product cells and scored for concentration, world supply grip and the evidenced availability of substitutes. Severe chokepoints in production inputs are **rare**: 1,533 cells of 233,666, carrying 0.77 per cent of input trade by value, an order of magnitude consistent with the European Commission’s narrower 2021 screen and here extended to the whole of world trade. Where they exist, one state is the leading supplier in 73.3 per cent of them.

The second empirical result is negative and it removes the standard consolation. If deterrence survives below the threshold at all, it survives through mutual exposure — a target holding counter-leverage over the state it depends on can respond without alleging anything. That reciprocity is not there. The obvious test has no power: 21 of the 23 states China depends on are dependent on China in return, but 142 of the 152 states holding any severe dependency are, so coincidence predicts 21.5. Excluding China and testing the remainder against a **degree-preserving** null gives 38 reciprocated dependencies against 37.5 expected. Magnitudes point the same way: across the eighteen mutual pairs that admit of a ratio the median imbalance is fifteen to one, and two are within a factor of two. Reciprocal severe dependency occurs at exactly the rate the degree distribution predicts. The mutual-hostage reading of interdependence is not supported at tariff-line resolution.

The paper closes with an evidential mechanism: where a measurable interruption coincides with a demonstrable grip on world supply and with the published absence of a verifiable ordinary explanation, a tribunal may draw an inference of fact in the manner *Corfu Channel* permits where the respondent holds the material. It shifts an evidential burden, not the legal burden and not the standard of proof, and it is drafted against the passage in *Corfu Channel* which says that exclusive control alone does neither.

A companion paper applies the same cost structure to markets rather than to law. The two share a structure and not a cause, and neither depends on the other.

Keywords: economic coercion; attribution; state responsibility; countermeasures; retorsion; chokepoints; weaponised interdependence; evidential burden

Note on sources and method

Every authority in this paper has been checked against a primary or authoritative record before citation. Journal literature was resolved against the Crossref registry and accepted only where title, first author and year agreed; English authority was checked against a citation index of 75,837 judgments harvested from Find Case Law under a computational-analysis permission granted by The National Archives; the judgments of the International Court of Justice, the reports and awards of WTO panels and arbitrators, the awards of investment tribunals and the instruments of the European Union were checked individually against the official text. Anything that could not be verified was removed rather than approximated.

This is stated not as a courtesy but because the paper’s subject makes it material. A paper arguing that the cost of proof is the binding constraint would be self-refuting if it declined to pay that cost. The apparatus is set out in the Appendix, together with the classes of failure that survived it — including two successive misreadings, in opposite directions, of the single instrument closest to the paper’s subject. They are reported because a paper making this argument should be judged partly on what its own verification missed.

The derivation code for the empirical work in Part II, the registers, and the verification tooling are published in full and run offline. Every figure reported is reproducible from the stated sources by the stated method.

Table of Cases

United Kingdom

Bank Mellat v HM Treasury (No 2) [2013] UKSC 39, [2014] AC 700
Barton v Morris [2023] UKSC 3, [2023] AC 684
Belhaj v Straw [2017] UKSC 3, [2017] AC 964
Shagang Shipping Co Ltd v HNA Group Co Ltd [2020] UKSC 34, [2020] 1 WLR 3549
URS Corporation Ltd v BDW Trading Ltd [2025] UKSC 21
Youssef v Secretary of State for Foreign and Commonwealth Affairs [2016] UKSC 3, [2016] AC 1457

International courts and tribunals

Air Service Agreement of 27 March 1946 between the United States of America and France (1978) 18 RIAA 417
Application of the Convention on the Prevention and Punishment of the Crime of Genocide (Bosnia and Herzegovina v Serbia and Montenegro) [2007] ICJ Rep 43
Corfu Channel (United Kingdom v Albania) (Merits) [1949] ICJ Rep 4
Gabčíkovo-Nagymaros Project (Hungary/Slovakia) [1997] ICJ Rep 7
Military and Paramilitary Activities in and against Nicaragua (Nicaragua v United States of America) (Merits) [1986] ICJ Rep 14
Prosecutor v Tadić (Judgment) IT-94-1-A (15 July 1999)

World Trade Organization

China — Enforcement of Intellectual Property Rights WT/DS611/ARB25
China — Measures Concerning Trade in Goods and Services WT/DS610
China — Measures Related to the Exportation of Rare Earths, Tungsten and Molybdenum WT/DS431/AB/R
China — Measures Related to the Exportation of Various Raw Materials WT/DS394/AB/R
European Communities — Measures Affecting Asbestos and Asbestos-Containing Products WT/DS135/AB/R
Japan — Measures Affecting Consumer Photographic Film and Paper WT/DS44/R
Russia — Measures Concerning Traffic in Transit WT/DS512/R
Saudi Arabia — Measures Concerning the Protection of Intellectual Property Rights WT/DS567/R
United States — Certain Measures on Steel and Aluminium Products WT/DS544/R, WT/DS552/R, WT/DS556/R, WT/DS564/R
United States — Origin Marking Requirement WT/DS597/R

Investment tribunals

Generation Ukraine Inc v Ukraine ICSID Case No ARB/00/9, Award (16 September 2003)
Metalclad Corporation v United Mexican States ICSID Case No ARB(AF)/97/1, Award (30 August 2000)
Siemens AG v Argentine Republic ICSID Case No ARB/02/8, Award (17 January 2007)
Técnicas Medioambientales Tecmed SA v United Mexican States ICSID Case No ARB(AF)/00/2, Award (29 May 2003)

Table of Legislation and Instruments

Treaties and international instruments

Agreement on the Application of Sanitary and Phytosanitary Measures (WTO) 1994, arts 2.2, 5.1, 5.2
General Agreement on Tariffs and Trade 1994, art XXI
General Agreement on Tariffs and Trade 1994, art XXIII:1(b)
International Law Commission, Draft Articles on Responsibility of States for Internationally Wrongful Acts, in Report of the International Law Commission on the Work of its Fifty-Third Session (2001) UN Doc

European Union

Regulation (EU) 2023/2675 of the European Parliament and of the Council of 22 November 2023 on the protection of the Union and its Member States from economic coercion by third countries [2023] OJ L2023/2675

Each instrument listed is cited in a footnote for the provision identified there, and no provision is relied on for a proposition it does not state.

Part I — Introduction

1. The problem stated

The cost of applying economic pressure has fallen and the cost of proving who applied it has not.

Identifying which twelve tariff lines, delayed simultaneously, would halt an adversary’s automotive sector within a quarter was until recently the work of a specialist team over months. It is now the work of an afternoon, on public data, and it leaves no institutional trace. Meanwhile the instruments of pressure have become deniable: a customs clearance is slow, a sanitary finding is raised, a licence is not refused but not granted, guidance is given to domestic firms and never written down. Each act is individually explicable and individually contestable. Collectively they are an instrument of policy, and proving that is a different and much harder thing than observing it.

This paper is about what a body of law does when that gap opens. Its answer is not that the law fails, which would be too simple, and not that it fails to deter, which would be too strong. It is that **the response migrates**: from instruments that require an identified author to instruments that do not, and therefore from instruments that discriminate between the responsible and everyone else to instruments that cannot.

2. Onto what foundation

The prohibition on intervention requires coercion bearing on matters a state is permitted to decide freely, and *Nicaragua* held that the United States embargo, whatever else it was, did not breach it.¹ Onto that foundation the last two decades have built an elaborate architecture of unilateral measures — designations, export controls, secondary sanctions, investment screening — every element of which is triggered by an identified actor doing an identified thing. The architecture works, when it works, because attribution is possible: a designated person is named, a controlled item is listed, a transaction is traced. Where authorship cannot be established, the designation architecture does not degrade gracefully. It does not engage at all.

What engages instead is the subject of Part IV.

3. What this paper does

Part II does three things. It sets out the premise the remedial literature does not state, positioning the argument against the work on informal sanctions which has already stated the descriptive half of it; it offers the paper’s worked case, the paired disputes over Lithuanian trade and over Chinese anti-suit injunctions, which between them locate the boundary of the claim more precisely than either does alone; and it surveys six doctrinal routes and shows where each fails.

Part III is the empirical contribution: an exhaustive census of supply chokepoints across the whole of world trade at tariff-line resolution, and a test of whether the mutual exposure on which the standard consolation depends actually exists.

¹ *Military and Paramilitary Activities in and against Nicaragua (Nicaragua v United States of America) (Merits)* [1986] ICJ Rep 14 [244]–[245] (non-intervention); [275]–[276] and [292(11)] (Treaty of Friendship, Commerce and Navigation, art XIX).

Part IV states the argument formally: two constraints, their intersection, and the conditions under which it is empty.

Part V proposes an evidential mechanism and states, at length, the objection that comes closest to defeating it.

Part VI concludes, and names the three developments that would require the analysis to be revised.

4. What is new here, and what is not

4.1 What is not new

That attribution is hard, probabilistic and politically produced, and that deterrence suffers accordingly, is the settled conclusion of two decades of work on cyber conflict.² **That a low detection probability combined with a capped penalty can make deterrence infeasible** is the judgment-proof problem.³ **That where the predicate for a countermeasure cannot be shown, retorsion becomes the primary available response, and that it carries no proportionality limit, no notification requirement and no predicate,** is Kosseff's, and Ruys records that states deliberately leave the characterisation of their measures open in order to avoid the countermeasures conditions.⁴ **That punishing the innocent is the price of deterrence under imperfect monitoring** is Green and Porter, and Abreu, Pearce and Stacchetti.⁵ **That economic coercion is now often informal and deniable, and that informality is what lets the coercing state avoid legal challenge,** is Lim and Ferguson's, on the Korean and Australian episodes.⁶ **That trade dependence is asymmetric rather than mutual, and that the asymmetry is where coercive leverage lives,** is Hirschman's, and Keohane and Nye's after him.⁷ **That bilateral dependency can be mapped at six-digit resolution across the whole of trade** has been done by the European Commission, by the CEPPI, by Baldwin, Freeman and Theodorakopoulos, and by the McKinsey Global Institute.⁸ **That reciprocity in a network must be measured against a null model** is Garlaschelli and Loffredo's, and the degree-preserving family of nulls for world trade specifically is Squartini, Fagiolo and Garlaschelli's; Lee and colleagues have already reported that the dependency subnetwork of world trade is less mutual than its null.⁹

Priority is claimed for none of these.

4.2 What is offered as new

First, the two-channel substitution result (Part IV §17.2). The descriptive core is Kosseff's and is conceded. What is offered is analytical: that the constraint on the first channel is a *chosen legal rule* rather

²See Part III §12.

³Steven Shavell, 'The judgment proof problem' (1986) 6 *International Review of Law and Economics* 45.

⁴Jeff Kosseff, 'Retorsion as a Response to Ongoing Malign Cyber Operations' in T Jančárková and others (eds), *12th International Conference on Cyber Conflict: 20/20 Vision — The Next Decade* (NATO CCDCOE Publications 2020) 9; Tom Ruys, 'Sanctions, Retorsions and Countermeasures: Concepts and International Legal Framework' in Larissa van den Herik (ed), *Research Handbook on UN Sanctions and International Law* (Edward Elgar 2017) 19.

⁵Edward J Green and Robert H Porter, 'Noncooperative Collusion under Imperfect Price Information' (1984) 52 *Econometrica* 87; Dilip Abreu, David Pearce and Ennio Stacchetti, 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring' (1990) 58 *Econometrica* 1041.

⁶Darren J Lim and Victor Ferguson, 'Informal economic sanctions: the political economy of Chinese coercion during the THAAD dispute' (2022) 29 *Review of International Political Economy* 1525.

⁷Albert O Hirschman, *National Power and the Structure of Foreign Trade* (University of California Press 1945); Robert O Keohane and Joseph S Nye, *Power and Interdependence: World Politics in Transition* (Little, Brown 1977) ch 1.

⁸European Commission, 'Strategic dependencies and capacities' (Commission Staff Working Document) SWD(2021) 352 final; Vincent Vicard and Pauline Wibaux, 'EU Strategic Dependencies: A Long View' (CEPII Policy Brief No 2023-41, June 2023); Richard Baldwin, Rebecca Freeman and Angelos Theodorakopoulos, 'Hidden Exposure: Measuring US Supply Chain Reliance' (2023) 2023(2) *Brookings Papers on Economic Activity* 79; McKinsey Global Institute, *The complication of concentration in global trade* (January 2023).

⁹Diego Garlaschelli and Maria I Loffredo, 'Patterns of Link Reciprocity in Directed Networks' (2004) 93 *Physical Review Letters* 268701; Tiziano Squartini, Giorgio Fagiolo and Diego Garlaschelli, 'Randomizing world trade. I. A binary network analysis' (2011) 84 *Physical Review E* 046117, and 'Randomizing world trade. II. A weighted network analysis' (2011) 84 *Physical Review E* 046118; Mi Jin Lee and others, 'Uncovering hidden dependency in weighted networks via information entropy' (2021) 3 *Physical Review Research* 043136.

than a wealth or technology constraint, that a lawful and untargeted substitute therefore sits permanently alongside it, and that the consequence is that falling attribution quality does not stop retaliation but strips it of legal form. Nor is the redirection without domestic analogue, as an earlier draft asserted: Shavell’s comparison of liability with regulation is an account of enforcement migrating when liability fails against a judgment-proof party.¹⁰ What differs is the identity of the substitute — a public regulator domestically, the injured party itself internationally — and that difference is the contribution.

Second, the negative reciprocity result (Part IV §19.1). The mutual-dependence reading underwrites a live argument in the security literature: Cha’s case for collective resilience computes the counter-dependence of the states China has coerced and builds a deterrence proposal on it.¹¹ The result here does not refute that proposal — Cha argues for a *coalition* pooling counter-leverage precisely because no single state holds enough, and a negative dyadic result argues for pooling. What it refutes is the weaker and much more widely held claim that mutual exposure is a general property of interdependence on which an individual state may rely. Reciprocal severe dependency occurs at exactly the rate a degree-preserving null predicts, and the magnitudes are asymmetric by a median factor of fifteen. The direction confirms Hirschman; what is new is that the proposition has been measured across the whole of world trade rather than asserted, and that the measurement is reproducible.

The route to it is also offered, because it is instructive. Two positive versions were computed and discarded — one defeated by a base rate, one by a degree distribution — and each was found only because the previous one had been. Negative results of this kind are not usually reported, because the positive version is available to anyone who stops testing when the number looks right.

Third, the census specification (Part III). The negative finding — that severe chokepoints are rare — is not a discovery and agrees with the Commission’s own screen. What is offered is the specification: the restriction to production inputs by end-use classification, which removes the pistachios-and-wristwatches problem that dominates concentration rankings; an alternative-supply term which is a residual supply index transposed from market-power screening;¹² the decomposition at §9.1 separating cells that are severe because substitution is physically unavailable from those that are severe because the leader dominates world supply, together with the demonstration that one limb of the composite cannot bind at all; and an instrument that runs offline and reproduces every figure.

Fourth, the documentary-trace criterion (Part II §6B.1). The distinction usually drawn is between formal instruments and administrative discretion. The paired disputes show that to be too coarse: an *unwritten* policy was successfully established and an adjudicated finding obtained, while a written-and-unwritten mixture was abandoned without one. The operative variable is whether the conduct leaves a documentary trace in the hands of a party other than the acting state. That criterion is stated *ex ante*, it explains both outcomes, and it identifies what the mechanism in Part V would have to be built on.

4.3 What would falsify each claim

The **two-channel result** fails if states in fact confine themselves to countermeasures when facing deniable pressure — that is, if the observed response to unattributable coercion is legal claim rather than reciprocal unfriendliness.

The **reciprocity result**, being negative, fails in the opposite direction from the others: it fails if recomputation across several reference years, or at a finer nomenclature, or on a trade-weighted rather than count-based measure, shows reciprocal severe dependency exceeding its degree-preserving expectation. That is a real possibility and the instrument is published so that it can be tested.

The **census finding** fails if a finer nomenclature reveals that sub-heading dependencies are dense enough to reverse the incidence result, which Part III §10 concedes is possible and would be a valuable correction.

¹⁰Steven Shavell, ‘Liability for Harm versus Regulation of Safety’ (1984) 13 *The Journal of Legal Studies* 357.

¹¹Victor D Cha, ‘Collective Resilience: Deterring China’s Weaponization of Economic Interdependence’ (2023) 48 *International Security* 91.

¹²Anjali Sheffrin, ‘Predicting Market Power Using the Residual Supply Index’ (FERC Market Monitoring Workshop, California Independent System Operator, 3–4 December 2002).

The **documentary-trace criterion** fails on an instance of coercion that left no external record and was nonetheless established to an adjudicated finding, or on an instance that left an ample external record and could not be.

Part II — Coercion Without an Author

5. Weaponised interdependence and the premise the remedial literature does not state

Farrell and Newman’s account of weaponised interdependence is the standard treatment and the right place to begin.¹³ Their argument is that the global economy is not a flat network of exchange but a set of hub-and-spoke structures — in payments, in information, in finance — whose asymmetries allow the state with jurisdiction over the hub to exercise two capabilities: a *panopticon* effect, seeing flows others cannot, and a *chokepoint* effect, denying access to flows others depend on.

The analysis is correct and has been productive. But it contains a premise that is rarely stated in that literature because it has rarely needed stating: that the exercise of the chokepoint is **observable and attributable**. The entire apparatus of response assumes it. A state that has been denied access to a payments network knows it has been denied, knows by whom, and can respond — by retaliation, by litigation, by building an alternative. Mulder’s history of the interwar economic weapon rests on the same assumption: the blockade was an announced instrument of policy, and its announcement was part of its coercive force.¹⁴

The sanctions literature inherits the premise wholesale. Hufbauer and his collaborators, Pape, Baldwin, Drezner, Morgan, Bapat and Kirshner disagree about whether sanctions work, about what “work” means, and about how selection bias distorts the observed success rate.¹⁵ They agree, without discussion, that a sanction is a thing which has been imposed by an identified state on an identified target for a stated reason. Drezner’s “hidden hand” is hidden in the sense that the *threat* often achieves the result without the imposition; it is not hidden in the sense that the coercing state is unknown.¹⁶

The efficacy debate itself illustrates how much rests on the premise. Wagner’s analysis of interdependence and bargaining power establishes that economic leverage translates into political influence only under specifiable conditions, and Bapat and Morgan’s reconsideration of multilateral against unilateral measures shows that the coalition structure changes the outcome.¹⁷ Norrlof’s power analysis of dollar hegemony explains why one particular hub confers leverage the others do not.¹⁸ The contested claims about the effect of corporate withdrawal from Russia after 2022 are a live instance: the disagreement is about magnitude and measurement, and it is *possible* only because the withdrawals were announced.¹⁹

One literature has seen the premise and rejected it, and this paper’s claim must be positioned against it rather than announced over it. The work on *informal* economic sanctions — Lim and Ferguson on Chinese coercion of Korea during the THAAD dispute, Ferguson, Waldron and Lim on the Australian measures — takes non-announcement and deniability as the *defining* features of the conduct it studies, and observes that

¹³Henry Farrell and Abraham L Newman, ‘Weaponized Interdependence: How Global Economic Networks Shape State Coercion’ (2019) 44 *International Security* 42.

¹⁴Nicholas Mulder, *The Economic Weapon: The Rise of Sanctions as a Tool of Modern War* (Yale University Press 2022).

¹⁵Gary Clyde Hufbauer, Jeffrey J Schott, Kimberly Ann Elliott and Barbara Oegg, *Economic Sanctions Reconsidered* (3rd edn, Peterson Institute for International Economics 2007); Robert A Pape, ‘Why Economic Sanctions Do Not Work’ (1997) 22 *International Security* 90; David A Baldwin, ‘The Sanctions Debate and the Logic of Choice’ (2000) 24 *International Security* 80; T Clifton Morgan, Navin Bapat and Valentin Krustev, ‘The Threat and Imposition of Economic Sanctions, 1971–2000’ (2009) 26 *Conflict Management and Peace Science* 92; Navin A Bapat and T Clifton Morgan, ‘Multilateral Versus Unilateral Sanctions Reconsidered’ (2009) 53 *International Studies Quarterly* 1075; Jonathan Kirshner, ‘The microfoundations of economic sanctions’ (1997) 6 *Security Studies* 32.

¹⁶Daniel W Drezner, ‘The Hidden Hand of Economic Coercion’ (2003) 57 *International Organization* 643.

¹⁷R Harrison Wagner, ‘Economic interdependence, bargaining power, and political influence’ (1988) 42 *International Organization* 461; Bapat and Morgan (n 15).

¹⁸Carla Norrlof, ‘Dollar hegemony: A power analysis’ (2014) 21 *Review of International Political Economy* 1042.

¹⁹Jeffrey Sonnenfeld and others, ‘Business Retreats and Sanctions Are Crippling the Russian Economy’ (2022) working paper. The findings are contested and are cited here for the existence of the measurement rather than for its conclusions.

informality is what allows the coercing state to avoid both WTO challenge and a formal legal response.²⁰ That literature has therefore already made the descriptive observation, and any claim that the premise is unstated is wrong.

What that literature asks is whether deniable coercion *works* — whether it achieves concessions, and at what cost to the coercer. What it does not ask is what deniability does to the *remedial architecture*: whether the institutions that allocate consequence continue to function, what happens to deterrence when the predicate for every formal response is unavailable, and where the response goes instead. That is the question of Movement II, and it is a different question from the efficacy question. The contribution here is not the observation that coercion has become deniable. It is the demonstration that the remedy set has a threshold, that the threshold is reached, and that what lies beyond it is not silence but a second channel with none of the first channel's disciplines.

This paper's claim is therefore that the premise is now doing far more work than the *remedial* literature can bear, and that the reason is a change in the technology of both coercion and concealment.

6. What has actually changed

Three things, none of which individually is remarkable, and which together are.

First, chokepoint identification has been democratised. The empirical exercise in Part III of this paper — mapping every importer-product pair in world trade, scoring supply concentration and world supply grip, and identifying where substitution is obstructed — was performed on a personal computer from a licence-free dataset in a working day. Twenty years ago that capability sat inside a small number of national statistical and intelligence agencies. It now sits with anyone who wants it, including the coercing party, the coerced party, and every commercial actor between them. The strategic significance is not that chokepoints exist; it is that their location is now common knowledge, which changes the game from one of private information to one of coordination.

Second, the instruments of pressure have become deniable. Coercion no longer requires a public instrument. It can be effected through customs delay, sanitary and phytosanitary findings, “voluntary” guidance to domestic firms, licensing slow-walking, informal consumer boycott, and the withdrawal of an export licence for reasons stated as environmental or administrative. Each of these is individually explicable, individually contestable, and collectively devastating. None of them announces itself. Cormac and Aldrich's analysis of implausible deniability in covert action identifies the same structure in the security domain: the point of modern deniability is not that the sponsor is unknown, but that it cannot be *proved*, which is a different and more useful property.²¹ Carson's work on covert intervention shows why states prefer this even when everyone suspects: unprovable action permits both sides to avoid the escalation that public attribution would compel.²²

Third, and most consequentially, the analytical labour of designing such a campaign has collapsed in cost. Identifying which twelve tariff lines, delayed simultaneously, would halt an adversary's automotive sector within a quarter; determining which of a target's suppliers are themselves dependent on inputs the coercing state controls; modelling the second-order effects through a production network — these were, until very recently, the work of a specialist team over months. They are now the work of an afternoon, and they leave no institutional trace.

The asymmetry that results is exactly the one established in Part II. The cost of *conducting* economic coercion has fallen sharply. The cost of *proving* that economic coercion has been conducted has not fallen

²⁰Darren J Lim and Victor Ferguson, ‘Informal economic sanctions: the political economy of Chinese coercion during the THAAD dispute’ (2022) 29 *Review of International Political Economy* 1525; Victor Ferguson, Scott Waldron and Darren J Lim, ‘Market adjustments to import sanctions: lessons from Chinese restrictions on Australian trade, 2020–21’ (2023) 30 *Review of International Political Economy* 1255.

²¹Rory Cormac and Richard J Aldrich, ‘Grey is the new black: covert action and implausible deniability’ (2018) 94 *International Affairs* 477.

²²Austin Carson, ‘Facing Off and Saving Face: Covert Intervention and Escalation Management in the Korean War’ (2015) 70 *International Organization* 103; Michael Poznansky, ‘Feigning Compliance: Covert Action and International Law’ (2019) 63 *International Studies Quarterly* 72.

and has arguably risen, because plausible alternative explanations for any given measure are now cheaper to generate too.

6A. The contrast case: what attributable coercion costs

Before accepting that unattributable coercion is analytically distinct, it is worth establishing that the attributable kind is well understood. It is, and unusually precisely, because the 2018–2019 tariff episode gave economists an identified shock with a known author.

The findings are consistent across independent teams. Amiti, Redding and Weinstein find near-complete pass-through of the tariffs into domestic prices, with the incidence falling on importers and consumers rather than on foreign exporters.²³ Fajgelbaum and colleagues reach the same conclusion by a different route and quantify the aggregate real-income loss.²⁴ Bown’s account places the episode in the forty-year history of special protection and documents its scale.²⁵ Caldara and colleagues isolate the effect of the *uncertainty* the episode generated, distinct from the tariffs themselves, and find it measurably depressed investment.²⁶

Two things about this literature matter here.

First, the precision is entirely a function of attribution. Every one of these studies identifies its effect by knowing exactly which measures were imposed, by whom, on what date, against which tariff lines. The econometrics is only as good as that identification, and the identification was free: the measures were announced, gazetted and litigated. Remove the announcement and none of these papers can be written.

Second, the sanctions literature has reached the same standard for the same reason. The Global Sanctions Data Base codes thousands of episodes precisely because sanctions are declared instruments.²⁷ Working from it and comparable sources, researchers have measured effects on growth, on life expectancy, and on the specific economic interests of targeted elites.²⁸ Whatever one concludes about efficacy, the *measurement* is possible because the author is known.

The contrast is the argument. Where coercion is attributable, we can price it, litigate it, insure against it and study it; where it is not, we can do none of those things, and the difference is not one of degree. This is why the shift toward deniable instruments described at §6 is not a marginal change in tactics. It removes the identification on which every one of these capabilities rests.

6B. A worked case: Lithuania

The mechanism described at §6 is not hypothetical, and a paper that asserted it without an instance would deserve the objection. The clearest documented episode is the treatment of Lithuanian trade with China from late 2021.

Following Lithuania’s decision to permit the opening of a Taiwanese representative office under that name, importers of goods originating in Lithuania, transiting Lithuanian ports, or otherwise connected to Lithuania began to encounter difficulties obtaining customs clearance in China. The measures reported included the rejection of Lithuanian consignments by customs authorities, restrictions affecting multinational companies incorporating Lithuanian inputs into products manufactured elsewhere, and import bans on alcohol, beef, dairy, logs and peat, several of them expressed in sanitary and phytosanitary terms.

²³Mary Amiti, Stephen J Redding and David E Weinstein, ‘The Impact of the 2018 Tariffs on Prices and Welfare’ (2019) 33 *Journal of Economic Perspectives* 187.

²⁴Pablo D Fajgelbaum and others, ‘The Return to Protectionism’ (2020) 135 *Quarterly Journal of Economics* 1.

²⁵Chad P Bown, ‘The 2018 US-China trade conflict after forty years of special protection’ (2019) 12 *China Economic Journal* 109.

²⁶Dario Caldara and others, ‘The economic effects of trade policy uncertainty’ (2020) 109 *Journal of Monetary Economics* 38.

²⁷Gabriel Felbermayr and others, ‘The global sanctions data base’ (2020) 129 *European Economic Review* 103561.

²⁸Mirko Draca and others, ‘On Target? Sanctions and the Economic Interests of Elite Policymakers in Iran’ (2022) 133 *The Economic Journal* 159; Jerg Gutmann, Matthias Neuenkirch and Florian Neumeier, ‘Sanctioned to Death? The Impact of Economic Sanctions on Life Expectancy and its Gender Gap’ (2021) 57 *The Journal of Development Studies* 139; Matthias Neuenkirch and Florian Neumeier, ‘The impact of UN and US economic sanctions on GDP growth’ (2015) 40 *European Journal of Political Economy* 110.

Every feature this paper has attributed to deniable coercion is present. **No instrument was published.** The measures were administrative, individually explicable, and individually contestable: a customs clearance is delayed, a phytosanitary concern is raised, a shipment is found non-compliant. **The pressure was extended to third parties** — companies with no connection to the dispute discovered that using a Lithuanian component impeded their own access, which converts the target from one state into every firm with a supply relationship to it. And **the stated reasons were not the operative ones**, but proving that requires access to the deliberations of a foreign administration.

The European Union requested consultations on 27 January 2022, requested a panel on 7 December 2022, and the Dispute Settlement Body established one on 27 January 2023.²⁹ The full procedural history is now available and it is more instructive than the outcome an earlier draft of this paper anticipated:

Date	Step
27 January 2022	Consultations requested
7 December 2022	Panel requested
27 January 2023	Panel established
18 April 2023	Panel composed
25 January 2024	Panel suspends work at the European Union's request under DSU art 12.12, for an indeterminate period
24 January 2025	Work resumes, one day before the twelve-month lapse of the panel's authority
27 January 2025	Suspended again at the European Union's request
1 December 2025	The European Union notifies the Dispute Settlement Body that it no longer considers it necessary to pursue the complaint, "considering the key objectives behind this dispute have been met and relevant trade has resumed"

No panel report was ever issued. No finding was ever made. The dispute was suspended twice, on each occasion within days of the deadline at which the authority for the panel's establishment would have lapsed under DSU art 12.12, and then withdrawn.³⁰

An earlier draft of this paper treated the bringing of the dispute as refuting the strongest form of its claim. On the complete history the concession was too generous, and the case now cuts three ways.

It confirms the mechanism. The measures took the form the analysis predicts — administrative, deniable, and extended to third parties precisely because third-party pressure is harder to characterise as a measure against the target at all.

It is the clearest available instance of the timing result. The measures took effect in late 2021. Four years later the proceeding was abandoned on the express ground that the trade had resumed. That is not a vindication of the system; it is the system being overtaken by the facts it was invoked to correct. The remedy did not arrive slowly. It did not arrive, and by the time it could have arrived nobody needed it. Whether the resumption of trade was caused by the pendency of the dispute is precisely what cannot be established, which is itself the paper's point in miniature: the counterfactual is unobservable, so the instrument's contribution to the outcome can be neither proved nor priced.

²⁹ *China — Measures Concerning Trade in Goods and Services*, WT/DS610. Consultations requested 27 January 2022; panel requested 7 December 2022; panel established by the Dispute Settlement Body 27 January 2023; panel composed 18 April 2023.

³⁰ *China — Measures Concerning Trade in Goods and Services*, WT/DS610: Request for Consultations by the European Union (27 January 2022) WT/DS610/1; Request for the Establishment of a Panel (7 December 2022) WT/DS610/5; Communication from the Panel (25 January 2024) WT/DS610/12 (suspension under DSU art 12.12); Communication from the Panel (29 January 2025) WT/DS610/13 (resumption 24 January 2025 and further suspension effective 27 January 2025); Communication from the European Union (1 December 2025) WT/DS610/14 (complaint not pursued). Understanding on Rules and Procedures Governing the Settlement of Disputes, art 12.12 (authority for establishment of the panel lapses where work is suspended for more than twelve months).

It nonetheless preserves the narrower claim rather than the stronger one. A dispute *was* brought, and the pendency of a WTO complaint is not nothing — it is a cost imposed on the respondent and a signal to third states, and this paper should not pretend otherwise. What the history establishes is not that deniable coercion is beyond legal reach, but that the reach is contingent on the complainant’s continued willingness to prosecute a claim whose evidentiary demands are severe and whose political cost rises as the underlying harm recedes. The European Union’s stated reason for withdrawal cuts against the paper on one point and for it on another: it suggests the pressure abated, and it demonstrates that when it abated the legal channel was abandoned without a ruling that might have disciplined the next episode. No precedent was created. The next state to do this faces exactly the law that existed in 2021.

A second and older instance is worth naming because it cuts the other way. Formal Chinese export restrictions on rare earths, tungsten and molybdenum *were* successfully challenged, precisely because they took the form of published quotas and duties rather than administrative discretion.³¹ That contrast is the whole argument of §6A in miniature: the same underlying objective is reachable when pursued through an instrument and substantially unreachable when pursued through a discretion.

6B.1 The counter-case, and what it actually shows

There is a third instance which is the strongest objection this paper faces, and it involves the same complainant and the same respondent. In *China — Enforcement of Intellectual Property Rights* the European Union challenged an **unwritten** Chinese policy under which courts issued worldwide anti-suit injunctions against holders of standard-essential patents, restraining them on pain of heavy fines from litigating non-Chinese patents anywhere but China. On 21 July 2025 an appeal arbitrator constituted under DSU art 25 found the unwritten policy inconsistent with the agreement on trade-related intellectual property rights. China stated at the Dispute Settlement Body that it had withdrawn the policy, referring to a notice of its Supreme People’s Court of 23 September 2025. The Commission recorded that statement on 1 April 2026 and expressly reserved its position, noting that *because* the policy is unwritten it would continue to monitor the situation to ensure full compliance and take further action if necessary.³² Compliance is therefore asserted by the respondent and unverified by the complainant, for exactly the reason this paper gives: there is no instrument to inspect, so withdrawal cannot be confirmed by reading anything. The thesis operates on the counter-case at the point of remedy even where it fails at the point of liability.

That is an unwritten, discretionary, deniable measure, established as a policy and adjudicated to a binding result, with an assertion of compliance the complainant has not been able to verify — and it happened while the Appellate Body was inoperative. It appears to refute both the doctrinal claim of this Part and the institutional claim at §8.3, and the appearance must be taken seriously rather than distinguished away.

It does not refute them, and the reason identifies the operative variable more precisely than the Lithuania case does. The Chinese anti-suit injunction policy was unwritten. Its *instances* were not. Each injunction was a judicial order, issued by an identified court, addressed to identified parties and supported by reasons, and each was therefore documented in the hands of the patent holder against whom it was made. The complainant established the existence of an unwritten policy by assembling a pattern from those instances — which is exactly the “series of facts linked together and leading logically to a single conclusion” that *Corfu Channel* contemplates and that §22 proposes to systematise.

It should be stressed that the instances were *not* on any public register, and the arbitrator so found: China was held to have acted inconsistently with the transparency obligations of the intellectual property agreement by failing to publish one of the decisions and by declining to supply information about them when asked.³³

³¹ *China — Measures Related to the Exportation of Rare Earths, Tungsten and Molybdenum* — Reports of the Appellate Body (7 August 2014) WT/DS431/AB/R, WT/DS432/AB/R, WT/DS433/AB/R, adopted 29 August 2014; *China — Measures Related to the Exportation of Various Raw Materials* — Reports of the Appellate Body (30 January 2012) WT/DS394/AB/R, WT/DS395/AB/R, WT/DS398/AB/R, adopted 22 February 2012.

³² *China — Enforcement of Intellectual Property Rights* — Award of the Arbitrator under DSU art 25 (21 July 2025) WT/DS611/ARB25 and WT/DS611/ARB25/Add.1; and see the statement of China recorded in the minutes of the Dispute Settlement Body, WT/DSB/M505, referring to the notice of the Supreme People’s Court of 23 September 2025.

³³ DS611 (n 32): findings of inconsistency with the Agreement on Trade-Related Aspects of Intellectual Property Rights, art 63.1 (failure to publish the anti-suit injunction decision in *Xiaomi v InterDigital* together with the reconsideration decision) and art 63.3, first sentence (failure to be prepared to supply information on the decisions on request).

That finding is not an aside. It confirms that what made the policy reachable was not publication but the existence of a documentary trace **in the hands of a party other than the acting state** — here, the served addressees, who could produce what they had been served with. The measure was deniable in form and fully documented in operation, and the documentation was in private hands throughout.

The Lithuanian measures were the opposite. A consignment held at a border generates no reasoned order, no addressee of record, and no published instance. There is nothing to assemble. The distinction that matters is therefore not written against unwritten, and §6A's framing of it as instrument against discretion is too coarse. **It is whether the conduct leaves a documentary trace in the hands of anyone other than the acting state.** Judicial acts do. Customs discretion does not.

Three consequences follow, and the paper is better for them. The claim of this Part must be narrowed accordingly: unwritten measures are not beyond legal reach, and measures leaving no external record substantially are. §8.3's account of the appellate vacuum must be qualified: the interim arbitration route under DSU art 25 produced a binding award against a major economy and an assertion of compliance in response, which is a considerably better record than the unadopted panel reports suggest even though the assertion remains unverified. And **DS611 is a worked precedent for the technique proposed at §22** — inference of a state policy from a documented pattern of individually defensible acts — which places that proposal within existing practice rather than beyond it.

What DS610 and DS611 establish together is therefore not that the system fails, but where its boundary runs. The same complainant, against the same respondent, in overlapping years, obtained a binding finding against conduct that left a documentary trace in private hands, and abandoned without a finding a complaint against conduct that left none.

7. What the attribution literature already knows

The cyber domain has had two decades of experience with exactly this problem and its conclusions transfer.

Rid and Buchanan's analysis established that attribution is not a technical fact discovered by forensics but a process of judgement under uncertainty, operating at several levels — tactical, operational, strategic — and producing conclusions that are always probabilistic.³⁴ Lindsay's assessment of deterrence against cyberattack reaches the harder conclusion: the attribution problem is not merely a delay in the deterrent response but a structural impediment to it, because a deterrent threat conditioned on an unreliable identification is not credible.³⁵ Gartzke's argument that cyberwar is a myth rests on the same foundation from the other direction: capabilities that cannot be publicly and credibly demonstrated cannot perform the signalling function that coercion requires.³⁶

Nye's treatment of deterrence and dissuasion in cyberspace is the most constructive and, for that reason, the most instructive.³⁷ He sets out four mechanisms — the threat of punishment, denial by defence, entanglement, and normative taboo — and does not abandon punishment. What he does observe is that three of the four are not strongly hindered by the attribution problem. The reading offered here goes one step further than Nye does: if the three attribution-independent mechanisms are the ones that remain serviceable, then a portfolio weighted towards them is a retreat from punishment-based deterrence whether or not it is presented as one. That characterisation is this paper's and not Nye's, and it is contestable — a portfolio may be a portfolio rather than a retreat. But the direction of the drift is not in dispute, and it is the same direction Part IV derives formally. Borghard and Lonergan reach a parallel conclusion about coercion in cyberspace: the domain is poorly suited to coercive bargaining precisely because the signalling requirements cannot be met.³⁸

³⁴Thomas Rid and Ben Buchanan, 'Attributing Cyber Attacks' (2015) 38 *Journal of Strategic Studies* 4.

³⁵Jon R Lindsay, 'Tipping the scales: the attribution problem and the feasibility of deterrence against cyberattack' (2015) 1 *Journal of Cybersecurity* tyv003.

³⁶Erik Gartzke, 'The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth' (2013) 38 *International Security* 41.

³⁷Joseph S Nye Jr, 'Deterrence and Dissuasion in Cyberspace' (2017) 41 *International Security* 44.

³⁸Erica D Borghard and Shawn W Lonergan, 'The Logic of Coercion in Cyberspace' (2017) 26 *Security Studies* 452.

Egloff's work adds the political economy. Attribution is not merely difficult; it is *produced*, by particular actors, with particular incentives, and its public assertion is a strategic act rather than a report of a finding.³⁹ Where the attributing party is also the retaliating party, the epistemic and the political are not separable, and every target of an attribution claim knows it.

The doctrinal literature has arrived independently at the same impasse. Roscini's treatment of evidentiary issues in cyber-related state responsibility disputes documents the absence of any settled standard of proof, the difficulty of discharging whatever standard is applied, and the resulting incentive for states to avoid formal dispute settlement altogether.⁴⁰ Tsagourias identifies the same gap in the self-defence context.⁴¹ Efrony and Shany's survey of state practice after *Tallinn Manual 2.0* finds that states have largely declined to invoke the framework in specific cases, which is a revealed preference about its usability rather than about its content.⁴²

The conclusion the cyber literature has reached, and which the economic-coercion literature has not yet reached, is this: **where attribution cost is high, the operative constraint on an aggressor is not the adjudicable law. It is the target's own exposure.** Part IV states that precisely, and states it more narrowly than the cyber literature does: what closes is the countermeasure channel, not retaliation as such, and what replaces deterrence where exposure is reduced is *prevention*, which §25 distinguishes.

8. Why the law cannot reach it

Six doctrinal routes exist. Five fail for a reason connected to proof rather than to substance; the sixth fails on standing and forum, and is set out with the others because its technique is the one §22 borrows. An earlier draft of this Part described four and called the list exhaustive; §8.3B sets out the two it omitted, both of which are better suited to the problem than any of the original four.

8.1 Non-intervention

The prohibition on intervention requires coercion bearing on matters in which a state is permitted to decide freely. *Nicaragua* is the governing authority and it is not helpful to the victim of economic pressure.⁴³ The Court held the United States embargo to be a breach of the bilateral treaty of friendship, commerce and navigation, but declined to find a breach of the customary principle of non-intervention. Jamnejad and Wood's analysis of the principle concludes that its content in the economic sphere remains substantially undetermined, and that state practice does not support a general prohibition on economic pressure.⁴⁴

So the starting position, before any question of proof arises, is that most economic coercion is not unlawful. That alone is significant. But note what would be required even to run the argument in a case where the coercion crossed whatever threshold might exist: the applicant would have to establish the measure, its coercive purpose, and its authorship. The first is observable. The second and third are exactly what a deniable campaign is designed to obscure.

8.2 State responsibility and attribution

Where the conduct is that of a non-state actor — a state-owned enterprise acting on informal guidance, a customs authority applying discretion, domestic firms responding to a signal — attribution must be

³⁹Florian J Egloff, 'Public attribution of cyber intrusions' (2020) 6 *Journal of Cybersecurity* tyaa012; Florian J Egloff and Myriam Dunn Cavelty, 'Attribution and Knowledge Creation Assemblages in Cybersecurity Politics' (2021) 7 *Journal of Cybersecurity* tyab002.

⁴⁰Marco Roscini, 'Evidentiary Issues in International Disputes Related to State Responsibility for Cyber Operations' in *Cyber War: Law and Ethics for Virtual Conflicts* (Oxford University Press 2015) 215.

⁴¹Nicholas Tsagourias, 'Cyber attacks, self-defence and the problem of attribution' (2012) 17 *Journal of Conflict and Security Law* 229.

⁴²Dan Efrony and Yuval Shany, 'A Rule Book on the Shelf? Tallinn Manual 2.0 on Cyberoperations and Subsequent State Practice' (2018) 112 *American Journal of International Law* 583.

⁴³*Nicaragua* (n 6) paras 245–249.

⁴⁴Maziar Jamnejad and Michael Wood, 'The Principle of Non-intervention' (2009) 22 *Leiden Journal of International Law* 345.

established under the secondary rules of state responsibility, codified by the International Law Commission.⁴⁵ Four of those rules could in principle reach the conduct described at §6, and the difficulty is different in each.

Article 4 attributes the conduct of organs of the state. A customs administration is such an organ, and where a coercive campaign is executed through customs discretion this is the article that bites. The obstacle is not attribution but characterisation: an individual clearance decision is attributable and lawful, and what is unlawful — if anything is — is the pattern, which no single attributable act evidences.

Article 5 attributes the conduct of entities empowered to exercise elements of governmental authority. It reaches a state-owned enterprise only in respect of the governmental functions it is empowered to exercise, not its commercial decisions. A state-owned trading company declining to buy is, on its face, exercising no governmental authority at all.

Article 8 attributes conduct carried out on the instructions of, or under the direction or control of, the state. This is the article that would have to do the work, and the threshold set for it is the one the paper's argument turns on. In *Nicaragua* the International Court of Justice required "effective control of the military or paramilitary operations in the course of which the alleged violations were committed", and held that general financing, organisation, training and equipping did not suffice.⁴⁶ The Appeals Chamber in *Tadić* proposed a lower "overall control" test for organised armed groups.⁴⁷ In *Bosnian Genocide* the Court declined to adopt it for purposes of state responsibility, holding that the overall-control test "stretches too far, almost to breaking point, the connection which must exist between the conduct of a State's organs and its international responsibility".⁴⁸ The threshold is therefore high, it has been tested, and it has been deliberately kept where it is.

Article 11 attributes conduct that a state acknowledges and adopts as its own. It is the one route that does not require proof of the relationship at the time of the act — but it requires the state to say so, and no state pursuing deniable coercion will.

Talmon's analysis of the responsibility of outside powers for the acts of entities they support demonstrates how demanding effective control is in application.⁴⁹ Milanović anticipated the problem the Court would face and Cassese's comparison of the *Nicaragua* and *Tadić* tests documents a body of law that has repeatedly declined to lower the threshold.⁵⁰

That is a defensible position when the alternative is holding states liable for conduct they did not direct; the Court's reasoning in *Bosnian Genocide* is not timid but principled. It is nonetheless a catastrophic position when the conduct in question is designed from the outset to sit below the threshold — and designing it that way is now cheap. The rules of attribution were built to prevent states being held responsible for what they did not do. They are being used to ensure states are not held responsible for what they did.

8.3 The security exception

A state facing coercion dressed as regulation might challenge the measure in the trading system. The respondent will invoke the security exception, which permits a Member to take action it considers necessary for the protection of its essential security interests, taken in time of war or other emergency in international relations.⁵¹ The panel in *Russia — Traffic in Transit* held the exception justiciable, which was reported as

⁴⁵International Law Commission, 'Draft Articles on Responsibility of States for Internationally Wrongful Acts with Commentaries' in *Report of the International Law Commission on the Work of its Fifty-Third Session* (23 April–1 June and 2 July–10 August 2001) UN Doc A/56/10, arts 4, 5, 8 and 11 and commentaries; annexed to UNGA Res 56/83 (12 December 2001) UN Doc A/RES/56/83.

⁴⁶*Military and Paramilitary Activities in and against Nicaragua (Nicaragua v United States of America) (Merits)* [1986] ICJ Rep 14 [115].

⁴⁷*Prosecutor v Tadić* (Judgment) IT-94-1-A (15 July 1999) [116]–[145].

⁴⁸*Application of the Convention on the Prevention and Punishment of the Crime of Genocide (Bosnia and Herzegovina v Serbia and Montenegro)* [2007] ICJ Rep 43 [406]; and see [400]–[407].

⁴⁹Stefan Talmon, 'The Responsibility of Outside Powers for Acts of Secessionist Entities' (2009) 58 *International and Comparative Law Quarterly* 493.

⁵⁰Marko Milanović, 'State Responsibility for Genocide' (2006) 17 *European Journal of International Law* 553; Antonio Cassese, 'The Nicaragua and Tadić Tests Revisited in Light of the ICJ Judgment on Genocide in Bosnia' (2007) 18 *European Journal of International Law* 649.

⁵¹General Agreement on Tariffs and Trade 1994, art XXI(b)(iii).

a significant assertion of jurisdiction.⁵² Read carefully, the standard adopted concedes what matters: the panel will not review the state’s designation of its essential security interests beyond good faith, and will ask only whether the measures are “not implausible” as measures to protect them.

This is not doctrinal deference. It is evidential surrender, and it is rational: a panel has no means of establishing a state’s actual motive, and a standard requiring it to do so would be unworkable. Voon’s assessment identifies the same tension.⁵³ Pinchis-Paulsen’s history of the negotiation of the exception shows that the drafters understood the problem and chose ambiguity deliberately.⁵⁴ The result is that a measure with a stated regulatory purpose and a real coercive purpose is, for practical purposes, unreviewable.

That said, the position did not stop in 2019, and an account treating a single panel report as settled doctrine would be incomplete. The panel in *Traffic in Transit* did not concede everything: it held that the existence of an “emergency in international relations” is an objective requirement which the panel determines for itself, reserving to the invoking state only the characterisation of its essential security interests.⁵⁵ Subsequent panels applied that objective element and did not invariably accept the invocation — in *Saudi Arabia — Intellectual Property Rights*, in the four *United States — Steel and Aluminium* proceedings, and in *United States — Origin Marking (Hong Kong, China)*.⁵⁶ The correct characterisation is therefore that the *objective* limb is reviewable and the *subjective* limb is not.

But that qualification must itself be qualified, and the qualification is fatal to it. *Traffic in Transit* was not appealed and was adopted. Every subsequent security-exception panel report was appealed into a non-functioning Appellate Body and none was ever adopted. The Appellate Body has been unable to hear appeals since December 2019, when the number of serving members fell below the quorum of three; the United States appealed the *Steel and Aluminium* and *Origin Marking* reports on 26 January 2023 into that vacancy, where they remain; and the *Saudi Arabia* appeal was suspended and the complaint ultimately not pursued. The Multi-Party Interim Appeal Arbitration Arrangement under DSU art 25 supplies an interim appellate function among its participants, and no outcome restoring the Appellate Body has followed.⁵⁷ That mechanism should not be written off: §6B.1 records that an arbitration under art 25 produced a binding award against China in July 2025, in response to which China asserted withdrawal of the policy. What it cannot do is bind a Member which has not joined it — and the United States has not — which is why the security-exception reports remain where they are.

The doctrinal position on the most important exception in the trading system therefore rests on one adopted panel report and four of no legal status whatever. That is not a footnote to the argument; it is the argument in a second register. A body of law developed by reports that were appealed into a void is a body of law that exists in the commentary and not in the system, and a respondent invoking art XXI today knows that the objective limb the panels articulated has never survived appellate scrutiny and cannot be made to. The evidentiary constraint identified above is compounded by an institutional one: even where a panel *can* establish the objective element, the finding can be neutralised by a notice of appeal costing nothing.

8.3A The standard of proof, where one is articulated at all

A further difficulty sits underneath every route that turns on proof. An earlier draft of this section asserted that international adjudication has no settled standard of proof. That is wrong, and the error mattered, because it allowed the argument to avoid the authority that most directly obstructs it.

⁵² *Russia — Measures Concerning Traffic in Transit* (n 7).

⁵³ Tania Voon, ‘The Security Exception In WTO Law: Entering a New Era’ (2019) 113 *AJIL Unbound* 45.

⁵⁴ Mona Pinchis-Paulsen, ‘Trade Multilateralism and U.S. National Security: The Making of the GATT Security Exceptions’ (2020) 41 *Michigan Journal of International Law* 109.

⁵⁵ *Russia — Traffic in Transit* (n 7) [7.71]–[7.77], [7.101]–[7.102], [7.138].

⁵⁶ *Saudi Arabia — Measures Concerning the Protection of Intellectual Property Rights* — Report of the Panel (16 June 2020) WT/DS567/R; *United States — Certain Measures on Steel and Aluminium Products* — Reports of the Panel (9 December 2022) WT/DS544/R, WT/DS552/R, WT/DS556/R, WT/DS564/R; *United States — Origin Marking Requirement* — Report of the Panel (21 December 2022) WT/DS597/R.

⁵⁷ Notices of appeal were filed by the United States on 26 January 2023 in DS544, DS552, DS556, DS564 and DS597; the appeal in DS567 was suspended and the complaint was not subsequently pursued. On the interim appellate mechanism see Understanding on Rules and Procedures Governing the Settlement of Disputes, art 25, and the Multi-Party Interim Appeal Arbitration Arrangement, JOB/DSB/1/Add.12 (30 April 2020).

For the precise problem described here — proof of state responsibility from circumstantial evidence where the respondent controls the material — the International Court of Justice has spoken twice, and the standard it set is high *for the category of case before it*. That qualification matters and an earlier draft omitted it. In *Corfu Channel* the Court admitted circumstantial evidence and held that a victim state denied access to the territory where the facts occurred “should be allowed a more liberal recourse to inferences of fact and circumstantial evidence”; such inference is “of special weight when it is based on a series of facts linked together and leading logically to a single conclusion”, and is admissible “provided that they leave no room for reasonable doubt”.⁵⁸ Two distinct standards must be kept apart here, and an earlier draft of this paper ran them together. The requirement of “a degree of certainty” attaches to the separate allegation that Albania had **connived** in the mine-laying — a charge the Court described as of exceptional gravity and rejected. The “no room for reasonable doubt” formulation is the standard the Court applied to the ordinary circumstantial inference of Albania’s **knowledge**, and on that standard it found *for* the applicant. In *Bosnian Genocide* the Court restated the point: claims against a state involving charges of exceptional gravity “must be proved by evidence that is fully conclusive”.⁵⁹

Two things follow, and the second is more favourable to this paper than the first suggests.

The *first* half of *Corfu Channel* is the best authority in existence for the mechanism proposed at §22, and it is judicial rather than legislative: the Court itself recognises that exclusive control by the respondent justifies a relaxation of what the applicant must show. The *second* half is the constraint that defeats that mechanism if it is framed as a presumption rather than as an inference, and §22.2 is now framed accordingly.

The “**fully conclusive**” standard in *Bosnian Genocide* is expressly attached to charges of exceptional gravity, and there is no authority applying it to a campaign of customs delay and sanitary findings. A state pleading economic coercion can therefore argue that it is not engaged. What that state cannot escape is the “**no room for reasonable doubt**” requirement governing circumstantial inference generally, which *Corfu Channel* applied to an ordinary question of knowledge. An earlier draft of this paper treated both formulations as confined to exceptional gravity and drew comfort from it. That was wrong, and the comfort was misplaced.

The position that survives is narrower and still useful. A claimant is not required to meet a standard reserved for allegations of connivance in mass casualty. It *is* required to assemble facts which, taken together, admit of no reasonable alternative explanation — which is precisely why limbs (a) to (c) of the mechanism at §22.2 are drafted as a conjunction including the *absence of a published ordinary explanation*, rather than as a presumption arising from market position. The doctrinal opening is not that the standard is low. It is that the Court has said inference can satisfy it, and has done so on facts thinner than those the instrument in Part III can supply.

This paper does not assume the point in its own favour. The analysis in Part IV holds on either standard, because $\bar{}$ is bounded below one by the evidence available rather than by the standard applied to it.

What Green’s analysis of the Court’s self-defence practice establishes is not the absence of a standard but its *inconsistent application* between cases without acknowledgment, so that a party cannot reliably predict what it must establish.⁶⁰ That is a different and narrower complaint, and it is the one the evidence supports. Bartels’s analysis of the consequences of withdrawing from provisional application illustrates the same phenomenon in a distinct setting: procedural posture frequently determines substantive outcome in ways the substantive rules do not disclose.⁶¹

The operative position is therefore worse for the claimant than an absence of standards would be, though for a narrower reason than an earlier draft gave. An unstated standard can be argued for. A stated requirement that circumstantial inference admit of no reasonable alternative, applied to conduct engineered to supply

⁵⁸ *Corfu Channel (United Kingdom v Albania) (Merits)* [1949] ICJ Rep 4, 17–18.

⁵⁹ *Bosnian Genocide* (n 48) [209].

⁶⁰ James A Green, ‘Fluctuating Evidentiary Standards for Self-Defence in the International Court of Justice’ (2009) 58 *International and Comparative Law Quarterly* 163.

⁶¹ Lorand Bartels, ‘Withdrawing Provisional Application of Treaties: Has the EU Made a Mistake?’ (2012) 1 *Cambridge Journal of International and Comparative Law* 112.

alternatives cheaply, cannot — because the conduct is designed to manufacture precisely the reasonable alternative that defeats the inference.

8.3B The two routes an earlier draft omitted

The two omissions were material because both are, on their face, better suited to the problem than any of the original four.

Non-violation nullification or impairment. GATT art XXIII:1(b) permits a complaint where a benefit accruing under the agreement is nullified or impaired by “the application by another contracting party of any measure, whether or not it conflicts with the provisions of this Agreement”.⁶² The complainant must show a measure, a benefit, and nullification or impairment of that benefit through frustration of a legitimate expectation reasonably held at the time of the concession.⁶³ The remedy is limited: under DSU art 26.1(b) there is no obligation to withdraw the measure, only to make a mutually satisfactory adjustment.⁶⁴ The cause of action has been read narrowly and used sparingly, and panels have described it as an exceptional remedy to be approached with caution.⁶⁵

The relationship to this paper’s proposal must be stated precisely, because an earlier draft over-conceded it. Art XXIII:1(b) is effects-based as to **wrongfulness**: it dispenses with the need to show a breach. It is not effects-based as to **authorship**. The complainant must still establish “the application by another contracting party” of a measure — which is the attribution requirement, unmodified, and it is the half of the problem this paper is about. The existing effects-based cause of action therefore solves the wrong half. That is a sharper point than the concession it replaces, and it is the answer to the objection that §22 proposes something the system already has.

Indirect and creeping expropriation. Investment treaty practice contains the most developed body of law anywhere on inferring interference from *effects*. The doctrine of creeping expropriation exists precisely because a state can destroy the value of an investment through a sequence of individually defensible administrative acts, and tribunals have developed tests looking to cumulative effect rather than to any single instrument or declared purpose.⁶⁶ *Metalclad* treated the effect on the investment as decisive irrespective of declared purpose; *Tecmed* set out the cumulative-effect analysis; *Generation Ukraine* insisted on a sustained pattern rather than isolated acts; *Siemens* confirmed that a creeping expropriation is composed of acts which individually would not qualify.⁶⁷ That is the same evidential problem this paper describes, addressed by a related technique, in a forum where it is routine.

Three limits keep this from disposing of the question. Investment arbitration protects an investor with a qualifying investment under a treaty, not a state whose economy has been squeezed, so the protected interest is narrower than the harm considered here. The conduct threshold is applied to the respondent’s *own* regulatory acts within its own territory, not to conduct it may have induced in third parties or in its own firms acting commercially. And the doctrine operates where the host state’s authorship of the measures is not in dispute — the argument is about their characterisation, not about who took them.

Both routes nonetheless establish something the paper must concede: **effects-based reasoning is not novel in international economic law**. What is not available anywhere in these bodies of law is effects-based reasoning about *authorship*, which is what §22 addresses and what its claim to novelty is now confined to.

⁶²General Agreement on Tariffs and Trade 1994, art XXIII:1(b).

⁶³*Japan — Measures Affecting Consumer Photographic Film and Paper* — Report of the Panel (31 March 1998) WT/DS44/R, [10.35]–[10.82], adopted 22 April 1998.

⁶⁴Understanding on Rules and Procedures Governing the Settlement of Disputes, art 26.1(b).

⁶⁵*European Communities — Measures Affecting Asbestos and Asbestos-Containing Products* — Report of the Appellate Body (12 March 2001) WT/DS135/AB/R, [186], adopted 5 April 2001.

⁶⁶Rudolf Dolzer and Christoph Schreuer, *Principles of International Investment Law* (2nd edn, Oxford University Press 2012) ch 6.

⁶⁷*Metalclad Corporation v United Mexican States*, ICSID Case No ARB(AF)/97/1, Award (30 August 2000) [103]; *Técnicas Medioambientales Tecmed SA v United Mexican States*, ICSID Case No ARB(AF)/00/2, Award (29 May 2003) [114]–[116]; *Generation Ukraine Inc v Ukraine*, ICSID Case No ARB/00/9, Award (16 September 2003) [20.22]; *Siemens AG v Argentine Republic*, ICSID Case No ARB/02/8, Award (17 January 2007) [263].

8.4 Countermeasures

A countermeasure is conduct that would otherwise be unlawful, its wrongfulness precluded because it is taken in response to an internationally wrongful act.⁶⁸ It may be taken only against the state responsible for that act, only to induce compliance, and only while the wrongful act continues; it must be proportionate, and it must be terminated when the obligation is complied with.⁶⁹ The International Court of Justice applied those conditions in *Gabčíkovo-Nagymaros*, holding Variant C unlawful for want of proportionality notwithstanding the prior breach; the *Air Services* tribunal had earlier accepted approximate rather than arithmetic commensurability, and both are the origin of the “of order unity” bound used at Part IV §17.1.⁷⁰

The predicate is the difficulty. Countermeasures require a prior internationally wrongful act **attributable** to the target state, and a state that takes them on an attribution that cannot be sustained has itself committed a wrong and forfeited the preclusion. A state responding to unattributable coercion with countermeasures is therefore betting its own legal position on an attribution it cannot prove. Rational states do not take that bet. What they do instead is the subject of Part IV §17.2, and it is not that they do nothing.

8.5 The common structure

Five of the six routes fail at the same point. Non-intervention fails because purpose cannot be shown. State responsibility fails because control cannot be shown. The security exception fails because motive cannot be tested. Countermeasures fail because the predicate cannot be established to a standard that makes the response safe. Non-violation nullification fails because, although it needs no wrongfulness, it still needs the measure to be shown to have been applied by the respondent. Creeping expropriation fails not on proof but on standing and forum: it reasons from effect, but only for an investor with a qualifying investment, and only about acts whose authorship is not in dispute.

Five of these are not separate doctrinal problems. They are one evidentiary problem wearing five costumes. The sixth, creeping expropriation, is different in kind: it fails not because the applicant cannot prove what happened but because the wrong applicant is in the wrong forum, which is why §22 borrows its technique rather than its cause of action. The doctrines are, in each case, substantively reasonable. What has changed is not the law’s content but the cost of satisfying it, and a rule whose satisfaction has become prohibitively expensive is not a rule that has been weakened. It is a rule that has been repealed without anyone voting on it.

8A. Deniability through corporate structure

A separate literature, rarely joined to this one, establishes that the corporate form itself supplies deniability at low cost.

Zucman’s reconciliation of international investment positions demonstrated that a substantial stock of wealth is unaccounted for in the aggregate statistics, held through structures that the reporting system does not see.⁷¹ Damgaard, Elkjaer and Johannesen show that a large fraction of recorded foreign direct investment is phantom — passing through entities with no substantive operations — so that the ownership recorded in the data is not the ownership that exists.⁷² Alstadsæter, Johannesen and Zucman quantify the distributional consequence.⁷³ Johannesen and Zucman’s assessment of the G20 transparency programme found that it relocated activity rather than eliminating it.⁷⁴

⁶⁸ILC Articles on State Responsibility (n 45) art 22 and commentary.

⁶⁹ibid arts 49, 51, 52 and 53 and commentaries.

⁷⁰*Gabčíkovo-Nagymaros Project (Hungary/Slovakia)* [1997] ICJ Rep 7 [83]–[87]; *Air Service Agreement of 27 March 1946 between the United States of America and France* (1978) 18 RIAA 417 [83], [90]–[96].

⁷¹Gabriel Zucman, ‘The Missing Wealth of Nations: Are Europe and the U.S. net Debtors or net Creditors?’ (2013) 128 *Quarterly Journal of Economics* 1321.

⁷²Jannick Damgaard, Thomas Elkjaer and Niels Johannesen, ‘What is real and what is not in the global FDI network?’ (2024) 140 *Journal of International Money and Finance* 102971.

⁷³Annette Alstadsæter, Niels Johannesen and Gabriel Zucman, ‘Tax Evasion and Inequality’ (2019) 109 *American Economic Review* 2073.

⁷⁴Niels Johannesen and Gabriel Zucman, ‘The End of Bank Secrecy? An Evaluation of the G20 Tax Haven Crackdown’ (2014) 6 *American Economic Journal: Economic Policy* 65.

Findley, Nielson and Sharman’s field experiment is the most direct evidence and the most damning. Soliciting the incorporation of anonymous shell companies from providers across many jurisdictions, they found widespread willingness to supply them in breach of the applicable standards, and — counter-intuitively — that compliance was no better in the jurisdictions with the strictest formal rules.⁷⁵ The rules were not the constraint; the probability of the breach being established was.

The bearing on this paper is direct. Coercion exercised through a commercial actor requires, for attribution, that the actor be linked to the state directing it. That linkage runs through corporate ownership, and the evidence above establishes that corporate ownership is systematically obscurable at a cost far below the cost of penetrating the obscurity. The state responsibility threshold discussed at §8.2 is demanding in law; in practice the demand is discharged against a record that has been constructed not to answer it.

This is also why the instrument published alongside this paper resolves control chains from the accounting-consolidation register rather than from any beneficial-ownership claim. The register records what entities were obliged to file, which is a weaker fact than ownership and a verifiable one — and the difference between those two is the whole subject of this paper.

9. The pricing consequence

The economic corollary is direct. Its first step is not new — Biener, Eling and Wirfs apply verifiability as an insurability criterion and find cyber risk failing it, and the redrafting of war and state-backed-attack exclusions across the cyber market after the *NotPetya* coverage litigation is that finding operating commercially, and the redrafted wordings turn expressly on attribution. What is offered here is the chain that follows from it, and specifically the last two steps, which convert a familiar underwriting difficulty into a claim about prices.

A risk is insurable, in the standard formulation, when the loss is definable, the event is fortuitous, the probability is estimable, and — critically — the occurrence is *verifiable*, so that the insurer can distinguish a covered loss from an uncovered one and from a fabricated one. Biener, Eling and Wirfs’s analysis of cyber risk against these criteria found insurability problems across several dimensions, with information asymmetry and the difficulty of loss quantification prominent among them.⁷⁶ Eling and Wirfs’s later measurement work confirms the difficulty of even establishing what such events cost.⁷⁷

Unattributable economic coercion fails the verifiability criterion completely. Consider a manufacturer whose sole viable supplier of a critical input is located in a state with which its own government is in dispute, and whose supply is interrupted. Was that coercion, a commercial decision, a regulatory action, a logistics failure, or a genuine shortage? The answer determines whether a political risk policy responds. The answer is unknowable at reasonable cost. An insurer that writes the risk therefore writes an unbounded and unmeasurable exposure; an insurer that excludes it leaves the risk with the manufacturer; and in practice policies contain exclusions broad enough to make the second outcome the normal one.

The result is not that coercion risk is mispriced in the ordinary sense of being wrongly estimated. It is that **there is no price**, because there is no transferable contract. The risk sits, unhedged and largely unmeasured, on the balance sheets of firms that in many cases have not identified that they hold it.

9.1 Why the market cannot price it either

The insurance argument has a market counterpart, and it is stronger than the insurance one because it does not depend on any contractual technicality.

Prices aggregate information through trade, and the microstructure literature is precise about the conditions. Glosten and Milgrom show that the spread a market maker sets reflects the probability of facing a better-

⁷⁵Michael G Findley, Daniel L Nielson and Jason C Sharman, ‘Causes of Noncompliance with International Law: A Field Experiment on Anonymous Incorporation’ (2015) 59 *American Journal of Political Science* 146.

⁷⁶Christian Biener, Martin Eling and Jan Hendrik Wirfs, ‘Insurability of Cyber Risk: An Empirical Analysis’ (2015) 40 *The Geneva Papers on Risk and Insurance — Issues and Practice* 131.

⁷⁷Martin Eling and Jan Wirfs, ‘What are the actual costs of cyber risk events?’ (2019) 272 *European Journal of Operational Research* 1109.

informed counterparty; Kyle models the rate at which private information is impounded into price through the order flow.⁷⁸ Easley, Kiefer and O'Hara demonstrate the consequence for thinly traded assets: where informed trade is rare relative to noise, the price adjusts slowly and incompletely.⁷⁹ The common requirement is that *somebody* holds information which trading can reveal.

Unattributable coercion fails that requirement at the root. There is no party who knows that a given disruption was coercive, because the fact is not merely private — it is unestablished. A market cannot aggregate information that does not exist in anyone's hands, and no volume of trading produces it.

What markets do instead is price the *observable proxy*, and there is now good evidence that they do this well. Baker, Bloom and Davis's economic policy uncertainty index, Caldara and Iacoviello's geopolitical risk index, and Hassan and colleagues' firm-level political risk measures all extract usable signals from text, and all are shown to predict investment and volatility.⁸⁰ Bloom's work on uncertainty shocks and Pástor and Veronesi's model of policy uncertainty establish the transmission.⁸¹ Bianchi, Ludvigson and Ma document the macroeconomic consequences of belief distortions.⁸²

But notice what these measures track: *announced* policy, *reported* geopolitical events, *disclosed* firm-level discussion. They are measures of the attributable. A campaign conducted through customs discretion and licensing delay generates no announcement, appears in no event count, and is discussed in no earnings call as coercion because the firm experiencing it does not know that is what it is. The best available instruments for pricing political risk are, by construction, blind to the category this paper is concerned with.

The behavioural literature adds a caution against assuming the shortfall is corrected elsewhere. Barberis's survey of psychology-based asset pricing establishes that investors under-react to information that is diffuse and hard to interpret, which is a fair description of a supply interruption of unknown cause.⁸³ Under-reaction and ambiguity aversion pull in opposite directions, and which dominates is an empirical question this paper cannot resolve; what can be said is that neither produces a price that reflects the underlying probability, because the underlying probability is not known to anyone.

The formal position is worse still. Where the probability distribution over causes is itself unknown — not merely unknown in realisation but unspecified — the decision problem is one of ambiguity rather than risk, and the standard apparatus does not apply. Gilboa and Schmeidler's maxmin representation and Hansen and Sargent's robust-control approach both describe agents who cannot form a unique prior and who therefore optimise against the worst case.⁸⁴ That is a reasonable description of a manufacturer facing a supply interruption of unknown origin, and it has a specific consequence: an ambiguity-averse agent overweights the adverse interpretation, which means the observable behaviour of firms facing this class of risk will be excessive precaution rather than mispricing. The cost does not vanish. It appears as over-hedging, redundant inventory and foregone specialisation, none of which is recorded anywhere as a cost of coercion.

This is where the empirical work becomes necessary. If the risk cannot be transferred, it must be measured and managed structurally, and measurement requires knowing where the exposures actually are. Part III supplies that measurement.

⁷⁸Lawrence R Glosten and Paul R Milgrom, 'Bid, ask and transaction prices in a specialist market with heterogeneously informed traders' (1985) 14 *Journal of Financial Economics* 71; Albert S Kyle, 'Continuous Auctions and Insider Trading' (1985) 53 *Econometrica* 1315.

⁷⁹David Easley and others, 'Liquidity, Information, and Infrequently Traded Stocks' (1996) 51 *The Journal of Finance* 1405.

⁸⁰Scott R Baker, Nicholas Bloom and Steven J Davis, 'Measuring Economic Policy Uncertainty' (2016) 131 *Quarterly Journal of Economics* 1593; Dario Caldara and Matteo Iacoviello, 'Measuring Geopolitical Risk' (2022) 112 *American Economic Review* 1194; Tarek A Hassan and others, 'Firm-Level Political Risk: Measurement and Effects' (2019) 134 *Quarterly Journal of Economics* 2135; Tarek A Hassan and others, 'Sources and Transmission of Country Risk' (2024) 91 *Review of Economic Studies* 2307.

⁸¹Nicholas Bloom, 'The Impact of Uncertainty Shocks' (2009) 77 *Econometrica* 623; Luboš Pástor and Pietro Veronesi, 'Uncertainty about Government Policy and Stock Prices' (2012) 67 *The Journal of Finance* 1219.

⁸²Francesco Bianchi, Sydney C Ludvigson and Sirui Ma, 'Belief Distortions and Macroeconomic Fluctuations' (2022) 112 *American Economic Review* 2269.

⁸³Nicholas Barberis, 'Psychology-Based Models of Asset Prices and Trading Volume' (2018) *Handbook of Behavioral Economics* 79.

⁸⁴Itzhak Gilboa and David Schmeidler, 'Maxmin expected utility with non-unique prior' (1989) 18 *Journal of Mathematical Economics* 141; Lars Peter Hansen and Thomas J Sargent, 'Robust Control and Model Uncertainty' (2001) 91 *American Economic Review* 60.

Part III — Measuring What Cannot Be Insured

10. Why measurement is the only remaining response

Part II concluded that unattributable economic coercion cannot be reliably litigated or insured, and Part IV adds that the countermeasure channel through which it might be deterred is closed. That conclusion sounds like counsel of despair and is not. It is a statement about which instrument remains, and the instrument that remains is structural: an exposure that has been identified can be reduced, and an exposure that has been reduced cannot be exploited whatever the state of the law.

That reframing has a cost. It requires knowing where the exposures actually are, at a resolution fine enough to act on and across the whole of trade rather than the handful of commodities that attract commentary. Most published work on supply concentration is either case-study — rare earths, semiconductors, pharmaceuticals — or aggregate, reporting national import concentration in a way that averages away the very cells that matter. Neither is usable by a firm or a ministry deciding what to do next.

This Part sets out an exhaustive measurement, its method, its limits, and its results. The instrument that produced it is published in full, runs offline, and reports the same numbers to anyone who runs it.

11. Data and method

The source is the BACI database maintained by the CEPII, version 202601, in the 2022 Harmonized System nomenclature, reference year 2024.⁸⁵ BACI reconciles the mirror flows reported by importers and exporters into a single harmonised bilateral series at six-digit tariff-line resolution; it is distributed under the Etalab Open Licence 2.0. The reference year contains **11,250,411 bilateral flows** totalling **USD 22.94 trillion**.

Flows are aggregated into (importer, product) cells. A cell is retained where the importer purchased at least USD 1 million of the product in the year: below that threshold supplier shares are dominated by rounding and single consignments, and a reported hundred-per-cent dependency is noise presented as a finding. **323,929 cells** survive.

Four measures are computed for each cell.

CR1 is the share held by the largest single supplier. **HHI** is the Herfindahl–Hirschman index of supplier shares. These describe the buyer’s own purchasing and are the measures on which most published concentration work stops.

GRIP is the leading supplier’s share of *world* exports of the product. This is the measure that matters, and its omission is the principal defect in concentration-based accounts of supply risk. A buyer taking one hundred per cent of a good from a state holding four per cent of world supply can re-source within a quarter. A buyer taking sixty per cent from a state holding eighty-five per cent of world supply cannot re-source at any price. The buyer’s own purchasing pattern is uninformative about which situation obtains.

ASR, the alternative supply ratio, is the exportable capacity of the rest of the world excluding the leader, expressed as a multiple of what this importer buys, counting only capacity actually evidenced by exports in the reference year. This is a transposition rather than an invention, and the lineage should be named: it is the *residual supply index* used as the standard pivotal-supplier screen in electricity market monitoring since 2002 — the capacity of all suppliers but the largest, as a fraction of demand, with the largest supplier pivotal where the ratio falls below one.⁸⁶ Two things change in the transposition. The denominator is this importer’s purchases rather than market demand, because the question is whether *this* buyer can be re-supplied rather than whether the market can. And the threshold is set at three rather than near unity as in electricity, where a residual supply ratio a little above one is treated as sufficient, because re-sourcing across borders requires displacing other buyers rather than dispatching idle plant. Idle capacity, mothballed

⁸⁵Guillaume Gaulier and Soledad Zignago, ‘BACI: International Trade Database at the Product-Level’ (2010) CEPII Working Paper 2010-23. Version 202601 released 22 January 2026, distributed under the Etalab Open Licence 2.0.

⁸⁶Anjali Sheffrin, ‘Predicting Market Power Using the Residual Supply Index’ (FERC Market Monitoring Workshop, California Independent System Operator, 3–4 December 2002).

capacity and capacity that could be built are excluded deliberately: coercion operates on the timescale of the first and not the second.

Exposure is then the product of a concentration term and a difficulty-of-substitution term:

$$E = \text{CR1} \times \max(\text{GRIP}, 1 - \min(\text{ASR}/3, 1), 0.6 \times \text{RCR})$$

where RCR is the concentration of the alternative supplier set. The difficulty term takes the *worst* of three ways substitution can fail — global dominance by the leader, insufficient residual capacity anywhere, and alternatives that are themselves concentrated — because these are not compensating conditions. A buyer facing any one of them is stuck, and a weighted average would allow a comfortable reading on one to purchase a dangerous reading on another. That property is what makes most composite risk indices unusable for decision, and avoiding it is worth the loss of elegance.

An earlier specification wrote the first term as $\max(\text{CR1}, \text{HHI})$, on the reasoning that either measure of concentration might bind. It cannot. On shares summing to one, $\sum s^2 - (\max s)(\sum s) = \text{CR1}$, so the Herfindahl index never exceeds the leading share and the maximum is always CR1; the two coincide only where a single supplier holds the whole market. This was confirmed across all 323,929 cells, in none of which HHI exceeded CR1. The redundancy was harmless to every result but is recorded because a composite whose first term silently reduces to one of its arguments should not be presented as a choice between two. HHI is retained in the published cell data as a descriptive field and does no work in the score.

End use is classified against the United Nations Broad Economic Categories, revision 5, whose third digit distinguishes intermediate consumption, gross fixed capital formation, and final consumption. The distinction is not cosmetic. Ranked on concentration alone, Germany’s most “exposed” imports in the reference year are pistachios, cut flowers and Swiss wristwatches: all genuinely single-sourced, none of them coercion risk. What makes a dependency coercible is that it enters a production function. **233,666 cells — 72.1 per cent of the total — are production inputs**, and the results below are restricted to them.

12. Results

12.1 Coercive leverage is rare

The exposure distribution across production-input cells is severely right-skewed.

Exposure band	Cells	Share
0.0–0.1	142,511	61.0%
0.1–0.2	50,246	21.5%
0.2–0.3	19,491	8.3%
0.3–0.4	9,529	4.1%
0.4–0.5	5,392	2.3%
0.5–0.6	3,261	1.4%
0.6–0.7	1,703	0.7%
0.7–0.8	943	0.4%
0.8–0.9	501	0.2%
0.9–1.0	89	0.04%

More than four-fifths of all production-input dependencies score below 0.2. **1,533 cells — 0.66 per cent — score at or above 0.70**, and a further 4,964 fall between 0.50 and 0.70.

The 0.70 band must not be over-read, and an earlier draft of this paper over-read it. Because the difficulty term is a maximum over three distinct failure modes, the composite does not say *which* mode drives any given cell. Decomposing the 1,533:

Binding failure mode	Severe cells
Leader’s share of world exports (GRIP)	1,367
Residual world capacity insufficient (ASR)	166
Alternative suppliers themselves concentrated (RCR)	0

Where two limbs are exactly equal at the published precision the cell is assigned to the capacity limb, which affects one cell of the 1,533; assigning ties the other way gives 1,368 and 165.

The strong reading — that no combination of alternative suppliers could physically cover the loss within the reference year — is therefore true of **166 cells, not 1,533**. In the other 1,367 the world outside the leader could in principle supply this buyer’s volume; what makes the cell severe is that the leader holds most of world exports, so re-sourcing means bidding for a thin residual against every other buyer displaced at the same time. That is a real and different vulnerability — it is the one that binds in a general denial rather than a bilateral one — but it is not physical impossibility, and the two should not be reported under one sentence. The third limb records zero, and that is a property of the specification rather than a finding. Since CR1 = 1, a cell scores 0.70 or above only if the difficulty term does, and the third limb can reach 0.70 only if RCR = 1.167, which is impossible for a concentration index bounded above by one. **The alternative-supplier concentration limb cannot bind in any severe cell by construction.** It is retained because it discriminates in the middle of the distribution, where it does bind, but it contributes nothing to the severe set and an earlier draft wrongly read its zero as evidence about the world. If that limb is intended to be capable of driving a cell into the severe band, its coefficient would have to exceed 0.70; it is left at 0.6 deliberately, because an alternative supplier set that is concentrated is a weaker form of exposure than a leader who dominates world supply, and should not on its own be called severe.

The aggregate value at stake is correspondingly small. Production-input cells carry USD 16.29 trillion of trade; the severe cells carry USD 126 billion, or **0.77 per cent**. Measured by value, economic coercion through supply denial has a strikingly narrow footprint.

That finding cuts against a substantial body of political commentary and should be stated plainly: **the world economy is not, in general, a lattice of chokepoints**. The overwhelming majority of traded inputs have deep and diversified supply. Accounts of interdependence as pervasive vulnerability are not supported by the data at tariff-line resolution.

12.2 Where it does exist, it is held by one state

The rarity of leverage makes its distribution more important, not less. Across the 1,533 severe production-input cells worldwide, the leading supplier is:

Leading supplier	Severe cells	Share
China	1,123	73.3%
United States	69	4.5%
India	38	2.5%
Brazil	35	2.3%
Netherlands	34	2.2%
Canada	28	1.8%
Bangladesh	28	1.8%
South Africa	19	1.2%
Türkiye	16	1.0%
Indonesia	14	0.9%

Forty-eight states appear as the leading supplier in at least one severe cell. One of them accounts for nearly three-quarters of them.

This is the paper’s central empirical finding and it requires care in interpretation. It is **not** a finding that China exercises coercion, and nothing in trade data could establish that. It is a finding about *structural position*: in 73.3 per cent of the world’s severe production-input chokepoints, the state that would have to act for the chokepoint to close is the same one. The measure is descriptive of capability, not of intent, and Part IV argues that in a world where intent cannot be proved, capability is the only variable that can be governed.

The broader distribution across grip is consistent. The median leading-supplier share of world exports across production-input cells is 0.121; the ninetieth percentile is 0.414; the ninety-ninth is 0.730. In **14,970 cells** — **6.4 per cent** — the leading supplier holds at least half of world exports.

12.3 The largest holder of leverage is itself substantially exposed

The realist observation, and the one most often missed in policy discussion, is that the structure is not one-sided.

Ranked by the count of severe production-input dependencies, the most exposed economies are India (42), the United States (42), Vietnam (40), **China (35)**, Korea (33), Brazil (31), Japan (30), Russia (29), Indonesia (28) and Egypt (27).

China holds three-quarters of the world’s severe chokepoints and is simultaneously the fourth most exposed economy by the same measure. Measured by value at severe risk within its own production inputs above USD 100 million, China is the *most* exposed of the major economies:

Economy	Input lines >\$100m	Severe	Value at severe risk	Share of filtered imports
China	952	15	USD 34.2bn	1.99%
United States	1,532	22	USD 18.4bn	0.88%
India	624	7	USD 3.5bn	0.66%
Canada	645	6	USD 2.3bn	0.73%
Korea	525	2	USD 2.0bn	0.48%
United Kingdom	661	2	USD 1.1bn	0.27%
Japan	552	5	USD 1.0bn	0.23%
Italy	686	2	USD 0.5bn	0.15%
Germany	1,148	3	USD 0.4bn	0.05%
France	792	1	USD 0.1bn	0.03%

The largest single severe dependencies in world trade illustrate the point:

Importer	Product	Value	Exposure	Leading supplier
China	Ferro-nickel	USD 14.2bn	0.87	Indonesia, 92%
Vietnam	Flat panel display modules	USD 10.5bn	0.90	Korea, 97%
China	Chromium ores and concentrates	USD 5.5bn	0.82	South Africa, 87%
United States	Coniferous S-P-F sawn wood	USD 4.2bn	0.94	Canada, 98%
China	Cobalt mattes and intermediates	USD 3.2bn	0.82	DR Congo, 93%
Belgium	Medicaments containing hormones	USD 2.8bn	0.90	Ireland, 100%
China	Flat panel display modules	USD 2.3bn	0.76	Korea, 90%
United States	Oriented strand board	USD 2.2bn	0.71	Canada, 93%
United States	Railway goods vans and wagons	USD 2.2bn	0.90	Mexico, 95%

Four of the nine largest severe dependencies in world trade are Chinese. Three are metallurgical inputs on which the industrial base depends; the fourth is flat panel display modules from Korea, which is the same

dependency Vietnam holds at the head of the list and is not a raw material at all. That is a materially different picture from the one in which China holds leverage and others hold vulnerability. It is tempting to read it as a mutual-hostage structure supporting a stability result, and Part IV §19.1 tests that reading directly. It does not survive: the reciprocity is no greater than the degree distribution predicts, and the magnitudes are overwhelmingly asymmetric. What §12.3 establishes is that the largest holder of leverage is also substantially exposed. What it does not establish, and what §19.1 shows the data will not support, is that it is exposed *to the states it holds leverage over* in amounts that would restrain it.

Germany's position deserves a note of its own. On this measure Germany is the *least* exposed major economy — three severe lines out of 1,148, 0.05 per cent of filtered import value — which is at odds with a substantial commentary describing German industry as acutely dependent. The discrepancy is explained by the end-use filter and by resolution: Germany's widely-discussed dependencies are concentrated in final goods and in sub-heading specificities that HS six-digit cannot see. That is a limit of the instrument and is stated as such below, not argued away.

13. What this measurement cannot see

Four limits. Two understate exposure; the third cuts both ways and is flagged as such; the fourth overstates it.

The nomenclature is the ceiling. HS six-digit is the finest resolution in the source. Where the decisive dependency sits inside a heading — extreme-ultraviolet lithography inside the semiconductor-machinery line, a single qualified grade inside a chemical line, one approved supplier inside an aerospace fastener category — the instrument cannot see it and will under-read. It does not guess, and the reader should treat every result as a lower bound.

Origin is customs origin. Material refined in one state and shipped from another appears as coming from the second. Where refining is more concentrated than shipping — the normal case in minerals — measured concentration understates real concentration. The rare-earth line is the clearest example: measured grip on the compounds heading is far lower than the concentration of the refining capacity that produces them.

Capacity is evidenced, not theoretical. Alternative supply counts only what a state actually exported. This is deliberate but it cuts both ways: it excludes capacity that could be brought online in a crisis, and so overstates difficulty in the medium term while correctly capturing it in the short.

Everything contractual and everything held is invisible. Strategic reserves, long-term offtake agreements, qualified second sources held on retainer, and engineering changes that design the input out will all change the real answer, and none appears in trade data.

A fifth limit is analytical rather than empirical, and is the most important. **Concentration is not consequence.** The end-use filter distinguishes inputs from consumption, which is coarse. It does not weight by criticality, because trade data contains no information about what breaks when a line stops. A USD 5.5 billion chromium dependency and a USD 5.5 billion dependency in a substitutable industrial chemical score alike and are not alike. Deciding which matters is the analyst's job and this instrument does not pretend to do it.

13A. Why a tariff-line measure understates network exposure

The measurement in this Part is bilateral and product-level. Real economies are networks, and a substantial literature establishes that shocks in networks do not stay where they land.

Acemoglu and colleagues demonstrated that idiosyncratic shocks to individual sectors do not wash out in the aggregate where the input-output structure is sufficiently asymmetric — a small number of sectors are disproportionately connected, and disturbances to them propagate.⁸⁷ Barrot and Sauvagnat established the microeconomic counterpart using natural disasters as an instrument: firms suffer measurable losses when a supplier of a *specific* input is disrupted, and the effect is concentrated precisely where the input cannot

⁸⁷Daron Acemoglu and others, 'The Network Origins of Aggregate Fluctuations' (2012) 80 *Econometrica* 1977.

readily be substituted.⁸⁸ Carvalho and colleagues, exploiting the Great East Japan Earthquake, traced the propagation through both upstream and downstream links and found substantial aggregate consequences.⁸⁹ Bonadio and colleagues quantified the transmission through global value chains during the pandemic.⁹⁰ Elliott, Golub and Leduc model the endogenous formation of supply networks and show that firms' privately optimal sourcing decisions produce a network that is more fragile than the socially efficient one — fragility is an equilibrium outcome, not an oversight.⁹¹

The financial-network literature supplies the formal machinery for the same phenomenon and reached it earlier. Elliott, Golub and Jackson, Acemoglu, Ozdaglar and Tahbaz-Salehi, Gai and Kapadia, and Glasserman and Young all establish that the relationship between connectivity and systemic risk is non-monotonic: dense connection dissipates small shocks and amplifies large ones.⁹² There is no reason to expect production networks to behave differently, and the empirical work above suggests they do not.

Two implications follow for the figures reported at §12, and both cut in the same direction.

The measure captures first-order exposure only. A country whose direct imports of a good are diversified may nonetheless depend on it entirely, because its diversified suppliers all depend on the same upstream source. This is the transshipment problem of §13 stated in network terms, and it means the true exposure is bounded below by the measured exposure and can substantially exceed it.

The value figures understate consequence by more than they understate incidence. Severe cells carry 0.77 per cent of production-input trade by value. The output at risk is not that value; it is the output that stops when the input stops, which the propagation literature suggests is larger by an order of magnitude or more in the concentrated cases. A USD 5.5 billion chromium dependency is not a USD 5.5 billion exposure; it is an exposure to whatever fraction of stainless steel production depends on it.

The critical-minerals literature illustrates the point with the material this instrument is worst at seeing. Gulley, Nassar and Xun map the overlap between Chinese and United States dependence on the same resources and find substantial common exposure.⁹³ Mancheri and colleagues trace the effects of Chinese policy on rare-earth supply-chain resilience and document adjustment that occurred over years rather than quarters.⁹⁴ Riddle and colleagues model the market dynamics under agent-based assumptions, and Gaustad and colleagues assess the extent to which circular-economy strategies can substitute for primary supply.⁹⁵ The common finding is that adjustment is possible and slow, which is precisely the parameter the ASR measure treats as binary.

None of this is a reason to prefer a network measure. Input-output tables at the resolution required do not exist for most economies, are years out of date where they do, and cannot be reconciled across borders at tariff-line granularity. The bilateral measure is what can actually be computed for every state and every product, which is why it is the one published here. But it should be read as a floor, and §12's finding that severe chokepoints are rare should be read as a statement about *direct* dependency and not about the total

⁸⁸Jean-Noël Barrot and Julien Sauvagnat, 'Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks' (2016) 131 *Quarterly Journal of Economics* 1543.

⁸⁹Vasco M Carvalho and others, 'Supply Chain Disruptions: Evidence from the Great East Japan Earthquake' (2021) 136 *Quarterly Journal of Economics* 1255.

⁹⁰Barthélémy Bonadio and others, 'Global supply chains in the pandemic' (2021) 133 *Journal of International Economics* 103534.

⁹¹Matthew Elliott, Benjamin Golub and Matthew V Leduc, 'Supply Network Formation and Fragility' (2022) 112 *American Economic Review* 2701.

⁹²Matthew Elliott, Benjamin Golub and Matthew O Jackson, 'Financial Networks and Contagion' (2014) 104 *American Economic Review* 3115; Daron Acemoglu, Asuman Ozdaglar and Alireza Tahbaz-Salehi, 'Systemic Risk and Stability in Financial Networks' (2015) 105 *American Economic Review* 564; Prasanna Gai and Sujit Kapadia, 'Contagion in financial networks' (2010) 466 *Proceedings of the Royal Society A* 2401; Paul Glasserman and H Peyton Young, 'How likely is contagion in financial networks?' (2015) 50 *Journal of Banking & Finance* 383.

⁹³Andrew L Gulley, Nedal T Nassar and Sean Xun, 'China, the United States, and competition for resources that enable emerging technologies' (2018) 115 *Proceedings of the National Academy of Sciences* 4111.

⁹⁴Nabeel A Mancheri and others, 'Effect of Chinese policies on rare earth supply chain resilience' (2019) 142 *Resources, Conservation and Recycling* 101.

⁹⁵Matthew Riddle and others, 'Global critical materials markets: An agent-based modeling approach' (2015) 45 *Resources Policy* 307; Gabrielle Gaustad and others, 'Circular economy strategies for mitigating critical material supply issues' (2018) 135 *Resources, Conservation and Recycling* 24.

exposure of an economy to disruption in another.

14. What the measurement is for

The purpose is not to rank states or to supply a grievance. It is to make a specific argument available: that where a legal remedy is unavailable because authorship cannot be proved, the residual policy instrument is the structure of dependency itself, and that structure is now measurable exhaustively, cheaply, and by anyone.

That has an uncomfortable implication which Part IV develops formally. If exposure can be measured by the potential target, it can be measured equally well by the potential coercer — the data is the same and the licence is free. The measurement does not create the vulnerability; it removes the informational asymmetry that previously protected some parties by accident. Whether that is stabilising or destabilising is not obvious, and is precisely the question a signalling model is suited to answer.

Part IV — The Verification–Attribution Gap, Formally

15. Why an argument in this form is needed

Parts II to V establish a pattern: in credence markets and in economic coercion alike, the cost of doing something has collapsed while the cost of proving who did it has not, and the institutions that allocate consequence depend on the second cost being low. That is an observation, not a result. It does not tell us whether the failure is gradual — so that better forensics, more disclosure and firmer doctrine would restore the position incrementally — or structural, so that they would not.

The distinction determines the policy. If deterrence degrades smoothly as attribution becomes harder, the right response is to invest in attribution. If it fails at a threshold, that investment is wasted beyond a point and the resources belong elsewhere.

What follows is a feasibility argument, not a game. There is no signal here in the technical sense — the coercing state's action is unobserved and is not chosen to convey anything — and no equilibrium is solved for. Two constraints are written down, their intersection is examined, and the conditions under which it is empty are identified. That is a modest apparatus and it is the honest one; presenting it as a solved signalling model would claim more than is delivered.

16. The setting

C is a state holding a supply position in a traded input. **T** is a state dependent on that input, with exposure e as measured in Part III over a dependent line of value V .

Supply disruptions occur for reasons unconnected to policy — commercial decisions, logistics failures, regulatory action taken for its stated reason, shortages — with probability q . Call these *innocent* disruptions. They are common, and q is not small. **C** may also deny supply, obtaining a benefit B , being the value of whatever concession the pressure is intended to extract **net of C's own cost of denying** — the lost export revenue, the reputational cost with third parties, and the cost of any substitution its own producers must make. Defining B net matters and an earlier draft did not. Gross of that cost the thresholds below are harder to satisfy than they should be, and the treatment would be asymmetric, because §19 imposes a cost condition on **T's** counter-threat and none on **C's** original act. Netting it makes conditions (1) and (4) easier to satisfy, which runs against this paper's negative conclusions and is the reason for stating it.

T observes a disruption but not its cause. This is the entire problem: the observable is identical whether nature or **C** produced it, and Part II showed that making it identical is now cheap.

T may invest in attribution and thereby identify the cause correctly with probability at most $\bar{\cdot} < 1$. The bound is not a modelling convenience but the settled finding of the attribution literature: attribution is

a probabilistic judgement produced under uncertainty, never a fact recovered by forensics, and it does not converge on certainty however much is spent.⁹⁶

17. Two channels of response, and only one of them is constrained

The literature on deterrence under imperfect detection begins with Becker: what governs behaviour is the expected sanction, so a low probability of detection can be offset by a high penalty.⁹⁷ Deterrence requires

$$\bar{c} \cdot P \geq B \quad (1)$$

where P is the punishment imposed. Everything turns on what bounds P , and here the analysis must be more careful than a domestic-law analogy allows, because international law offers **two** response channels with entirely different constraints.

17.1 Countermeasures, which are capped

A countermeasure is a measure that would otherwise be unlawful, rendered lawful by a prior internationally wrongful act. It must be proportionate — commensurate with the injury suffered, taking into account the gravity of the wrongful act and the rights in question — and a disproportionate response is itself wrongful.⁹⁸ So on this channel

$$P \leq L, \text{ where } L \text{ is the injury and } \bar{c} \text{ is of order unity} \quad (2)$$

Combining (1) and (2), deterrence through countermeasures is feasible only where

$$\bar{c} \geq B / (\bar{c} \cdot L) \quad (3)$$

For $\bar{c}$ below that value the two constraints are jointly unsatisfiable: **no proportionate countermeasure deters**. And countermeasures carry a second and independent difficulty. They require a *prior wrongful act attributable to T's counterparty* as their legal predicate, so a state that resorts to them on an attribution it cannot prove has committed a wrong itself. On this channel attribution enters twice — once through $\bar{c}$ in (1), and once as a condition of lawfulness — which is why it is the channel that fails first.

Two qualifications keep (3) honest. The injury for proportionality purposes is not simply the disrupted trade value: it includes the non-material injury to the right to decide freely, and the practice accepts approximate rather than arithmetic commensurability. And where the response is collective, P scales with the number of responding states. Both loosen (2) without removing it.

17.2 Retorsion, which is not

Retorsion is the channel the analysis cannot ignore, and its omission would invalidate the result. Retorsion is unfriendly conduct that is nonetheless lawful: raising a tariff within bound rates, closing public procurement to suppliers of a given nationality, tightening export-licence policy, withdrawing trade preferences, expelling officials. Some instruments of this kind are targeted at an identified person and some are not, and §17.2 turns on the difference; the examples here are the untargeted ones. It requires **no prior wrongful act, no attribution, and no proportionality** in the countermeasures sense. It is not constrained by (2) at all.

That matters enormously, and it means condition (3) does not describe the whole position. A distinction must be drawn here that an earlier draft elided, because as drafted it contradicted Part I §3. Measures of the kind described there — designations, export controls, investment screening — are retorsion as to their *legal predicate*: they need no prior wrongful act and no attribution in the ILC sense. They are not free of attribution as to their *operational trigger*: a designation must name someone, a control must list an item, and both require the administering state to identify a person or a consignment. So §3's statement stands — where authorship cannot be established the designation architecture does not engage — and what §17.2

⁹⁶Rid and Buchanan (n 34); Lindsay (n 35); Egloff (n 39).

⁹⁷Gary S Becker, 'Crime and Punishment: An Economic Approach' (1968) 76 *Journal of Political Economy* 169.

⁹⁸International Law Commission, Articles on Responsibility of States for Internationally Wrongful Acts (2001) UN Doc A/56/10, art 51.

adds is that the *untargeted* forms of retorsion do not share that limitation. A general tariff, a procurement exclusion by nationality, a withdrawal of preferences and a licensing slowdown require no identification of anyone. They are the measures that remain available when the targeted ones cannot fire, and they are correspondingly blunt: they fall on the whole of a state's trade rather than on the actor responsible for anything. That is the migration, stated precisely: not from law to lawlessness, but from instruments that discriminate between actors to instruments that cannot.

So the correct statement is not that deterrence fails. It is narrower and, on reflection, more interesting:

Where τ is low, a deterrent can be delivered only by retorsion, because the countermeasure channel is closed by (3). Retaliation therefore survives in a form that is unadjudicated and reciprocally available. Whether *deterrence* survives with it is not established here: that depends on the retorsion channel's own feasibility condition, which is set out below and is not derived.

Three consequences follow, and none is comfortable.

Retorsion has no proportionality discipline, so it has no natural stopping point. The constraint that (2) imposes is not merely a limit on the responder; it is the mechanism by which a response is recognised as a response rather than as a fresh injury. Remove it and each party's retaliation is, to the other, an unprovoked act calling for its own answer.

Retorsion carries no legitimacy premium. A countermeasure asserts a legal claim and invites the other party to answer it; that is its function. Retorsion asserts nothing and so generates no record, no forum, and no possibility of the dispute being resolved rather than continued. It converts a legal disagreement into an open-ended exchange.

Retorsion is available to both sides on identical terms. A target that responds to deniable coercion with lawful unfriendly measures has not deterred; it has opened a second front on which the coercing state may act with exactly the freedom it enjoyed on the first.

A fourth point limits the first three, and it is required by the evidence. Retorsion is not bounded by (2), but it is not bounded by nothing. It is bounded by the retaliator's own cost — a tariff is paid by the state that imposes it — and by escalation risk, neither of which this argument models. Introduce either and the retorsion channel acquires its own feasibility condition, which may also fail. What can be said without a model is that its bound is a matter of self-interest rather than of law, and self-interested bounds are not adjudicable, not stated in advance, and not the same for a large state as for a small one.

The conclusion is therefore not that law fails to deter, and it is narrower than an earlier draft claimed. It is that **as attribution cost rises, retaliation migrates from instruments that require an identified author to instruments that do not** — from a regime of claims, predicates and proportionality to one of reciprocal unfriendliness with no terminating rule and no discrimination between the responsible and the rest.

One qualification keeps that honest. Migration is a claim about the *legal form* of retaliation and not about its volume: whether the quantity of deterred coercion rises or falls depends on the retorsion channel's own feasibility condition, which is not derived here. What the argument establishes is that falling attribution quality does not stop retaliation, it strips it of legal form.

The structure has since been enacted, and the way it was enacted qualifies the first two consequences above rather than confirming them. The European Union's anti-coercion regulation, examined at §22.0, conditions its response measures on a determination of coercion rather than on a wrongful act, permits escalation to the non-performance of international obligations only "insofar as" wrongfulness obtains, and drafts every measure in its annex so that it need not be treaty-inconsistent. That is the two-channel structure of this section in legislative form: a default channel requiring no predicate, and an upper channel gated by exactly the proof that Part II shows is ordinarily unavailable. A legislature drafting against this problem, with every incentive to solve it, produced an instrument that concedes it.

But it also disciplined the lower channel, and that must be reported because it cuts against the argument.

Article 11(1) requires Union response measures to be proportionate and not to exceed the level of injury, taking into account the gravity of the coercion — which is the countermeasures proportionality formula, imported by choice onto a channel that international law does not require to be capped at all. Article 12 supplies review, amendment, suspension and termination; the determination is a published implementing act; and an offer to negotiate is required before the measures take effect, engagement having already been required before adoption.⁹⁹ So the claim that the surviving channel has no proportionality discipline and no stopping rule is a claim about **general international law**, not about every instrument built on it. Where a state or a bloc chooses to impose those disciplines on itself, they exist.

That is the more interesting finding. A legislature under no obligation to cap its lower channel capped it anyway, supplied it with a terminating rule, and attached it to a published determination. The unbounded character of retorsion is evidently recognised as a defect by the actors best placed to exploit it — which is evidence that the defect is real, and evidence that it is remediable by design rather than by doctrine. It is also the strongest available argument for the proposal at §22: if the discipline is going to be supplied unilaterally, it is better supplied by a public and computable trigger than by a political determination.

Whether the migrated state of affairs is *worse* than deterrence failing outright is a question this paper cannot answer in the welfare terms it has set itself, and an earlier draft asserted it without argument. Unbounded reciprocal retorsion may well deter more episodes than a closed countermeasure channel, while imposing deadweight loss on both parties and on third states; which effect dominates is not obvious and is not derived here. What can be said, and is a claim about legality rather than about welfare, is that the migrated regime retains the appearance of a functioning response while removing every feature that made the response a legal one — the predicate, the forum, the proportionality limit and the stopping rule.

17.3 The discontinuity, and what it is not

Condition (3) is a threshold in $\bar{c}$ for a single dyad and a single episode: below the critical value the feasible set is empty, above it it is not. It should not be read as implying that the aggregate quantity of deterred coercion falls discontinuously. With heterogeneity in B across episodes and states — which certainly exists — the *measure* of episodes satisfying (3) falls smoothly as $\bar{c}$ declines, because different episodes cross their own thresholds at different points. The discontinuity is a property of the individual case; the aggregate is a distribution over cases. An earlier formulation of this Part elided the two, and the correction weakens the rhetorical claim without affecting the analytical one.

18. What the comparative statics separate

Instruments that raise expected punishment through the countermeasure channel fail below the threshold. Better forensics raise $\bar{c}$ toward its bound but cannot pass it. Firmer doctrine raises the *legal* availability of a response without raising the *evidential* ability to establish its predicate.

Instruments that operate through retorsion remain *available*, at the cost identified at §17.2. Whether they deter is a separate question this argument does not answer. They are available, they are lawful, and they will be used. What they will not do is produce a stable equilibrium or a resolved dispute.

Instruments that reduce the aggressor’s benefit work without either difficulty. Diversification, substitution, stockholding and qualification of second sources reduce B directly and depend on proving nothing. Note that this is *prevention* rather than deterrence — it removes the motive rather than raising the expected cost — and the paper has elsewhere conflated the two. The distinction matters because prevention is unilaterally available and deterrence is not.

19. Non-attributive counter-leverage, correctly stated

There is a further channel, and Part III’s evidence locates it.

⁹⁹Regulation (EU) 2023/2675, arts 11(1), 12 and 8(7)(b).

Suppose T also holds a supply position on which C depends. T can then respond to a disruption not by punishing C — which requires attribution — but by reducing its own supply to C. This requires no allegation and no legal predicate; it is the exercise of a commercial freedom.

An earlier version of this Part stated the resulting condition as $B - B$, where B is the loss C suffers. **That was wrong**, and the error is instructive because it is the same error the paper diagnoses elsewhere: it made attribution disappear by arithmetic rather than by argument.

If T withholds supply on *any* disruption, C's expected payoff from denial is $B - B$ and from restraint is $-q \cdot B$, because nature disrupts anyway and T cannot tell the difference. Deterrence therefore requires

$$B \cdot (1 - q) \geq B \quad (4)$$

The q term is not a detail. It says that unconditional retaliation has deterrent value only to the extent that C's action *changes* the probability of suffering it. Where T cannot distinguish denial from innocent disruption, the deterrent is discounted by exactly the frequency of innocent disruption — and q was stipulated at §16 to be substantial.

Two further conditions must hold and were previously unstated. T's own cost of withholding supply must be low enough that the threat is credible, since T is also injured by the interruption it creates. And T must be able to *commit*: a threat to retaliate on every disruption is, at the moment of a disruption T believes innocent, one T would prefer not to execute.

That commitment requirement is the analytically important point, and it generalises — though less strongly than an earlier draft asserted. **A party that can commit to punishing without proof does not defeat the attribution problem; it converts it into a base-rate problem, and recovers a deterrent scaled by the informativeness of the observable — at the price of punishing the innocent.** Since §16 stipulates that q is not small, the recovered deterrent is substantially degraded even under perfect commitment, and it vanishes entirely as q approaches one. Condition (4) is one instance of that principle rather than a special channel.

The principle is not new and the paper should not present it as such. It is the central result of the literature on repeated games under imperfect public monitoring: Green and Porter show that sustaining cooperation when the monitoring signal is noisy requires punishment phases triggered by observations that may be entirely innocent, and Abreu, Pearce and Stacchetti generalise the construction.¹⁰⁰ Punishing the innocent as the price of deterrence without attribution was established there forty years ago. What is different in the present setting, and what Part IV adds, is that the constraint on the punishment is a *legal* cap rather than a wealth or technology constraint, and that a lawful uncapped substitute exists alongside it — which is the two-channel structure at §17 and has no analogue in the industrial-organisation setting.

Condition (4) should therefore **not** be identified with Proposal Two. An earlier draft said it was “precisely what Proposal Two implements”, and that is wrong twice over. Proposal Two as reframed at §22.2 is an inference of fact shifting an evidential burden, not a commitment to respond without proof, and §22.2 makes rebuttal by a bona fide state cheap — which is the opposite of a commitment device. Condition (4) describes what a *state* might do unilaterally. Proposal Two describes what a *tribunal* might infer. The two are separate mechanisms operating on separate actors, and conflating them concealed the fact that the paper's model prices a cost its proposal does not incur.

19.1 Is the mutuality dyadic? A test

Condition (4) requires that T holds leverage over **C specifically**. An aggregate observation that C is exposed to *somebody* establishes nothing. Part III §12.3 reported that the state leading 73.3 per cent of severe production-input chokepoints is itself heavily exposed, but its severe exposures are in ferro-nickel, chromium and cobalt intermediates — sourced from Indonesia, South Africa and the Democratic Republic

¹⁰⁰Edward J Green and Robert H Porter, ‘Noncooperative Collusion under Imperfect Price Information’ (1984) 52 *Econometrica* 87; Dilip Abreu, David Pearce and Ennio Stacchetti, ‘Toward a Theory of Discounted Repeated Games with Imperfect Monitoring’ (1990) 58 *Econometrica* 1041.

of the Congo, which are not obviously the states over which it holds leverage. On its face that is an aggregate coincidence and not a mutual hostage.

The question is empirical and the instrument answers it. Taking every severe production-input cell, and asking for each state exposed to a given leader whether that state also leads a severe cell in which the leader is the dependent importer, the result is as follows.

States with at least one severe production-input dependency	152
States severely exposed to China	142
States China is severely exposed to	23
Of those 23, also severely exposed to China	21
Distinct mutual pairs worldwide	40

The China result is worthless and must be discarded. Twenty-one of twenty-three is 91.3 per cent, which looks decisive until it is set against the base rate. Of the 152 states holding any severe production-input dependency at all, 142 — 93.4 per cent — are severely exposed to China. A randomly chosen set of 23 such states would be expected to contain 21.5 that are exposed to China. The observed value is 21. On a hypergeometric test the probability of a result at least this reciprocal under the null of no dyadic structure is 0.82. The test does not merely fail to reach significance; it returns very slightly *less* mutuality than coincidence predicts. Because China leads 73.3 per cent of the world’s severe chokepoints, virtually any state with an exposure of any kind is exposed to China, and no statement about China’s dyads can discriminate between mutual hostage and arithmetic. An earlier draft reported the 21 of 23 as confirmation of condition (4). It confirms nothing.

Removing China does not rescue it, and the way it fails is instructive. The natural repair is to delete the dominant node and test whether reciprocity in the remainder exceeds chance. Excluding China as both importer and supplier leaves 108 states holding 315 directed severe dependencies, of which 38 are reciprocated, forming 19 mutual pairs. Against a uniform random graph of the same density — 315 edges among 108×107 ordered pairs, a density of 0.027 — the expected number of reciprocated edges is 8.6, and 38 observed is **4.4 times** the chance rate. A draft of this section reported that figure as a real finding.

It is not, and the reason is the same defect as before at smaller scale. A uniform-density null assumes every state is equally likely to be an endpoint. This network is nothing like that: in the eighteen interpretable pairs the United States appears in seven, India in six, the Netherlands in four and South Africa in three. Where degrees are that unequal, reciprocated edges arise from the degree sequence alone, because a state that is an endpoint of many edges is correspondingly likely to be an endpoint of the reverse of any given one. The correct null holds the degree sequence fixed. Rewiring the observed network 200 times by degree-preserving double-edge swaps gives a mean of **37.5** reciprocated edges against **38** observed — a ratio of 1.01, with the observed value at the 50th percentile of the null distribution.

There is no reciprocity above chance in the severe-dependency network. Mutual severe dependency occurs at exactly the rate the degree sequence predicts, and nothing dyadic is happening.

That is a negative result and it is the honest one. It survives two attempts to obtain the positive version, and both attempts failed for the same reason — a null model that did not represent the structure already visible in the data. The first failure was a base rate, the second a degree distribution. Reporting only the second, or only the first, would have concealed how easily this measurement produces a confirmation of whatever it is pointed at.

Magnitude confirms the negative reading rather than softening it. Condition (4) is a magnitude condition, $B(1 - q) > B$, which requires the counter-leverage to exceed the leverage by a margin that widens as q rises. One substitution has to be declared before the test is read. B is the concession C extracts and B is the loss C suffers from T ’s counter-denial; the data measures neither, and what is used below as a proxy for both is the *value at severe risk* in each direction. That is injury, not benefit, and the two are the same object only if each party’s gain from coercing is proportional to the other’s exposure. The proxy is

used because nothing better exists at this resolution, and it is tolerable here for two reasons: the primary result is the count-based reciprocity test, which does not use B or B at all, and the magnitude asymmetries reported below are of an order — a median of fifteen to one — far outside any plausible proportionality error. Measuring value at severe risk in both directions across the eighteen mutual pairs that admit of a ratio — the nineteenth involves an aggregated reporting territory and is excluded — the median asymmetry is **15 to 1** and only **two of the eighteen** are within a factor of two. The largest are Canada–United States (USD 2.0bn against USD 7.5bn) and Mexico–United States (USD 1.0bn against USD 4.2bn); at the other extreme Brazil–India stands at 395 to 1. So of the pairs where reciprocity exists at all, almost none is reciprocal in the sense condition (4) needs, and the two that come closest — Malaysia–India at 1.2 to 1 and Japan–India at 1.5 to 1 — are the pairs where a symmetric standoff is *least* consequential.

The same is true of China’s own dyads once they are examined for size rather than existence. China holds USD 34.8 billion of production inputs at severe risk from the twenty-three states it depends on — the figure at Part III §12.3 is USD 34.2 billion because that table is restricted to import lines above USD 100 million, whereas this one counts every severe cell — against USD 18.4 billion held at risk by those states from China; but in thirteen of the twenty-three the counterpart is the more exposed, and in the other ten China is more exposed, in four cases by more than an order of magnitude — Indonesia 18 to 1, New Zealand 17 to 1, South Africa 12 to 1 and the Democratic Republic of the Congo 92 to 1 — with Namibia and Finland holding no reciprocal exposure at all. There is no mutual-hostage equilibrium here. There is a set of highly asymmetric bilateral relationships in which, in a handful of cases, the structurally weaker party holds the larger claim.

What this does to condition (4). It does not falsify it; a sufficient condition that is rarely satisfied is not a false condition. What it establishes is that **the condition is rarely satisfied**, which matters more for policy than for theory. Deterrence through mutual exposure is not a general feature of the trading system that a state can rely on. It is available in a small number of asymmetric relationships and, on the paper’s own algebra, it is available precisely to the *more* exposed party rather than to the pair — because (4) requires $B > B/(1 - q)$, so it is the state with the larger counter-leverage that can make the threat, not two states with comparable positions. An earlier draft had this backwards, describing the near-symmetric dyads as those where the condition “holds most strongly”. They are the dyads where it fails.

Three limits must be stated. The 142 and the 23 count different populations — cells in which a state leads against cells in which it depends — and are not comparable as measures of “exposure”; they are reported together only to establish the base rate that defeats the China test. The customs-origin qualification at Part III §13 bites hardest here and applies to more than one case: Indonesian ferro-nickel substantially reflects Chinese-owned smelting capacity located abroad, and Congolese cobalt intermediates substantially reflect Chinese-owned upstream assets. Applied consistently that removes two of the three largest Chinese exposures, leaving South African chromium as the one major dependency that is unambiguously on another state’s decisions — which weakens the mutual-hostage reading of China’s position considerably and is stated here rather than confined to a footnote. And the whole exercise rests on a single reference year, in a measure known to churn: applying the Commission’s bottom-up screen to two decades of trade data, Vicard and Wibaux find that one dependent product in five identified in 2018 was no longer so identified in 2019, and that close to half of those identified in 2014 had dropped out five years later.¹⁰¹ The mutual pairs reported above are a snapshot and should be treated as one until the calculation has been run across several years, which the published instrument makes possible and which §4.3 names as the falsification test. That the exposure structure moves this fast is not only a caveat on the measurement; it is a substantive result for §26, because an equilibrium resting on mutual exposure inherits the volatility of the thing it rests on.

19B. Where this sits in the signalling literature

The argument above is not a signalling model, but the informational problem it describes is the one that literature has spent forty years on, and the relationship is worth stating precisely — including where it does not hold.

Jervis’s security dilemma is the foundational case: where the same action is consistent with defensive and

¹⁰¹Vincent Vicard and Pauline Wibaux, ‘EU Strategic Dependencies: A Long View’ (CEPII Policy Brief No 2023-41, June 2023) https://cepii.fr/PDF_PUB/pb/2023/pb2023-41.pdf accessed 4 August 2026.

offensive intent, the observer cannot infer intent from the action, and the resulting uncertainty is itself a source of conflict.¹⁰² Fearon's rationalist explanation of war supplies the bargaining structure — conflict occurs because private information about resolve cannot be credibly communicated and parties have incentives to misrepresent it.¹⁰³ The situation here is that structure with the informational problem moved: the uncertain party is not unsure how far the other will go, but whether the other did anything at all.

That move matters, and it is where the analogy breaks. In a signalling model the informed party *wants*, at least sometimes, to be believed, and the literature is about how it makes itself credible — Fearon's tying hands and sinking costs, Sartori's reputational account, Weeks's audience costs.¹⁰⁴ Every one of those mechanisms operates on a sender who has chosen to communicate. Deniable coercion is the opposite case: the acting party's objective is that no inference be drawn at all. Slantchev's analysis of feigned weakness is the closest antecedent, showing that states manipulate what the other side can infer whenever doing so is cheap.¹⁰⁵

The reputational equilibria of Kreps and Wilson and of Milgrom and Roberts are the most instructive failure.¹⁰⁶ Both show that a player facing a sequence of opponents will invest in a reputation for toughness where the investment is observed and remembered. Reputation requires conduct to be linked to a name. Where an action cannot be attributed, no reputation forms and repeated play does not discipline — which is the conclusion the companion paper reaches about credence markets by an entirely different route. The convergence is not coincidental: both are consequences of a missing link between an act and its author, which is the general structure this paper is describing and which Townsend formalised for a principal who must pay a fixed cost to learn the truth of a report.¹⁰⁷

Buchanan's cybersecurity dilemma completes the picture and sharpens the discomfort.¹⁰⁸ Measures taken for defence are indistinguishable from preparations for attack, so defensive investment itself generates insecurity. Applied here: a state that maps its dependencies is doing exactly what a state planning coercion would do, with the same data. The instrument in Part III serves a ministry assessing its own vulnerability and one designing a campaign against another's equally well. That is an uncomfortable property of the work and it is stated rather than avoided, because the alternative — restricting the analysis — is not available. The data is public, the method is elementary, and the party most likely to have run it already is the one with the most resources.

20. Objections

Bounded attribution is doing all the work. It is, and it is the assumption best supported by evidence. If τ were near unity neither the model nor the problem would be interesting.¹⁰⁹

This is Becker with extra steps. The reply must be more careful than the objection deserves, because bounded sanctions are not new. Shavell's judgment-proof problem is the standard case of a wrongdoer whose assets cap the recoverable penalty, and it produces exactly the structure derived at (3): where p is low and f is capped, the product may be insufficient at any feasible enforcement level.¹¹⁰ The contribution here is not the structure but two things about its application. First, the cap in the international setting arises from a *legal* rule rather than from the wrongdoer's wealth, which means it is chosen and could be changed. Second — and this is what §17.2 adds — the presence of a parallel channel that does not require an identified author means the binding cap does not suppress retaliation but *redirects* it. That redirection does have a domestic

¹⁰²Robert Jervis, 'Cooperation under the Security Dilemma' (1978) 30 *World Politics* 167.

¹⁰³James D Fearon, 'Rationalist explanations for war' (1995) 49 *International Organization* 379.

¹⁰⁴Anne E Sartori, 'The Might of the Pen: A Reputational Theory of Communication in International Disputes' (2002) 56 *International Organization* 121; Jessica L Weeks, 'Autocratic Audience Costs: Regime Type and Signaling Resolve' (2008) 62 *International Organization* 35.

¹⁰⁵Branislav L Slantchev, 'Feigning Weakness' (2010) 64 *International Organization* 357.

¹⁰⁶David M Kreps and Robert Wilson, 'Reputation and imperfect information' (1982) 27 *Journal of Economic Theory* 253; Paul Milgrom and John Roberts, 'Predation, reputation, and entry deterrence' (1982) 27 *Journal of Economic Theory* 280.

¹⁰⁷Robert M Townsend, 'Optimal contracts and competitive markets with costly state verification' (1979) 21 *Journal of Economic Theory* 265.

¹⁰⁸Ben Buchanan, *The Cybersecurity Dilemma: Hacking, Trust and Fear Between Nations* (Oxford University Press 2017).

¹⁰⁹Egloff (n 39); Egloff and Dunn Cavelti (n 39).

¹¹⁰Steven Shavell, 'The judgment proof problem' (1986) 6 *International Review of Law and Economics* 45.

analogue, and an earlier draft wrongly denied it: Shavell’s own comparison of liability with regulation of safety is precisely an account of enforcement migrating to a different instrument when liability fails against a judgment-proof party.¹¹¹ What differs internationally is the identity of the substitute. Domestically the substitute is a public regulator, which is a specialist body acting under a mandate and subject to review. Internationally the substitute is the injured party itself, acting unilaterally, under no mandate and subject to no review. The migration is the same phenomenon; the institution it migrates to is materially worse.

Proportionality is not so tightly binding in practice. True, and §17.2 now says so directly: where (3) fails, states use retorsion. The model predicts that behaviour rather than being embarrassed by it.

Condition (4) is mutual assured damage, not deterrence. Agreed, and (4) now says so. On fragility: the arrangement is only as durable as the exposure structure, which §19.1 shows is measurable and which moves.

T’s incentives are never derived. Correct. The retaliation decision of T — its own cost, its posterior after a disruption, its willingness to execute a threat it would prefer to withdraw — is not modelled, and the claim that a punishment is “credible” is therefore an assumption rather than a result. A full treatment would need to specify T’s problem and solve it; what is offered here is a feasibility condition, and its limits are those of a feasibility condition.

Part V — What Follows

21. The design problem

A proposal follows from the analysis, and it is constrained by the finding that produced it: it may not assume that attribution becomes cheap, because the whole argument is that it does not. A reform that works only if proof becomes easier is not a reform; it is a restatement of the problem.

The proposal therefore takes the cost as given and moves an *evidential* burden on one issue, leaving the legal burden and the standard of proof exactly where they are. That is a much weaker device than the one an earlier draft proposed, and the reason for the retreat is set out at §22.2: a presumption of attribution arising from market position would be squarely contrary to the passage in *Corfu Channel* which holds that exclusive control neither involves prima facie responsibility nor shifts the burden of proof. What survives is an inference of fact, available to a tribunal, on a conjunction of contingent facts rather than a standing attribute.

22. Proposal two: an effects-based evidential mechanism for economic coercion

22.0 What already exists, and why this is an extension rather than an invention

An earlier draft presented the mechanism below as new. It is not, and the omission is serious enough to correct prominently rather than in a footnote.

The European Union has operated an effects-based determination mechanism for economic coercion since December 2023. Regulation (EU) 2023/2675 defines economic coercion as the application or threatened application by a third country of a measure affecting trade or investment in order to prevent or obtain the cessation, modification or adoption of a particular act by the Union or a Member State, thereby interfering in its legitimate sovereign choices.¹¹² Determination is made case by case by the Commission and the Council against stated criteria, and where coercion is determined and engagement fails, the Union may adopt response measures including tariffs and restrictions.

A draft of this paper made two claims about that instrument which are the opposite of what it says. Both are corrected here in the body rather than in a footnote, because the corrected reading is not a concession — it is considerably better evidence for the analysis in Part IV than the version it replaces.

¹¹¹Steven Shavell, ‘Liability for Harm versus Regulation of Safety’ (1984) 13 *The Journal of Legal Studies* 357.

¹¹²Regulation (EU) 2023/2675 of the European Parliament and of the Council of 22 November 2023 on the protection of the Union and its Member States from economic coercion by third countries [2023] OJ L2023/2675, in force 27 December 2023.

It does require attribution, expressly and by reference to the customary rules. The draft asserted that the Regulation dispenses with attribution in favour of a finding about purpose. It does not. A “third-country measure” is defined as “any action or omission attributable to a third country under international law”.¹¹³ The recitals invoke the customary law of state responsibility as reflected in the ILC Articles, reproduce the art 8 formula of conduct carried out on the instructions of or under the direction or control of the state, and require that action under the Regulation be consistent with international law.¹¹⁴ The attribution requirement this paper spends Part II establishing to be unsatisfiable is therefore written into the trigger of the instrument built to address the conduct that makes it unsatisfiable.

That is the finding, and it is a stronger one than the error it replaces. Faced with conduct designed to defeat attribution, the most capable trading bloc in the world legislated an attribution requirement — not from oversight, but because a measure taken against a state that had not acted would itself be unlawful, and no legislature can draft that problem away. The Regulation relocates the *finding* from a tribunal to the Commission and Council, which is a candid acknowledgement that no tribunal could make it on the available evidence. It does not relocate the *requirement*, because the requirement is not the Union’s to move.

Its response measures are channel-agnostic, and that is the point. Two successive drafts of this paper got this wrong in opposite directions — the first calling the response measures retorsion, the second calling them countermeasures. Neither is right, and the correct reading is more useful than either.

The conditions for adopting response measures are set out in art 8(1) and contain **no wrongfulness predicate at all**: engagement must have failed, the measures must be necessary, and they must be in the Union interest.¹¹⁵ The instrument’s operative trigger is economic coercion as defined in art 2(1), not an internationally wrongful act. Article 8(4) then provides that “insofar as” the third-country measure constitutes an internationally wrongful act, response measures “may” consist of measures amounting to the non-performance of international obligations. That is a conditional permission to escalate onto the countermeasures channel, not a characterisation of the measures as countermeasures. Recital 13 says so in terms, contemplating measures “not only” consistent with the Union’s international obligations “but also” non-performing ones where wrongfulness obtains. Annex I is drafted to match: each listed measure is qualified as one “which may amount, as necessary, to the non-performance of applicable international obligations”, so none of them is inherently treaty-inconsistent.¹¹⁶

The instrument is therefore the two-channel structure of Part IV §17 enacted in legislation.

It acts on a determination of coercion, using measures which are lawful unfriendly acts by default; and it reserves the right to escalate to non-performance only in the subset of cases where the wrongfulness predicate — which requires the attribution art 3(1) imports — can actually be made out. That is precisely the architecture §7 derives: a lower channel always available and an upper channel gated by proof. The Union has not legislated around the problem this paper describes. It has legislated a mechanism whose more powerful half is conditioned on exactly the finding that Part II shows cannot ordinarily be made, and whose default half is the unadjudicated channel. A recent treatment sets out the difficulties this creates.¹¹⁷

That the author read this instrument wrongly twice, in opposite directions, before reading it correctly is recorded in the Appendix. It is the clearest instance in the paper’s own production of the thing the paper is about: the text was available throughout, freely, in the author’s working language; what was expensive was not obtaining it but attending to it.

Its declared purpose is deterrence. The Regulation states the objective as deterring coercion or obtaining its cessation, with response measures as a last resort. Whether an instrument of this design deters is exactly what condition (4) asks, and it is unanswered.

What follows is therefore an extension, and a narrower one than the draft claimed. It substitutes a **mea-**

¹¹³Regulation (EU) 2023/2675 of the European Parliament and of the Council of 22 November 2023 on the protection of the Union and its Member States from economic coercion by third countries [2023] OJ L2023/2675, art 3(1).

¹¹⁴ibid recitals 5 and 16, and art 1(3).

¹¹⁵ibid art 8(1).

¹¹⁶ibid art 8(4), recital 13 and Annex I.

¹¹⁷Miles Jackson and Federica Paddeu, ‘To deter and counteract economic coercion: the puzzle of countermeasures in the European Union’s Anti-Coercion Instrument’ (2025) 28 *Journal of International Economic Law* 567.

surable trigger for a political determination, removing one discretion and creating a different vulnerability considered at §22.3. It operates on the **evidential question within existing adjudicative frameworks** rather than establishing a parallel political procedure, and so is available to states without the Union’s market power. And it is **reciprocal by construction**: a trigger computed from public data on a published method applies to whoever satisfies it. What it does not do, and what the Union’s experience shows cannot be done by drafting, is remove the attribution requirement itself.

22.1 What it must do

Part IV established that where attribution cost is high, no proportionate punishment deters, and that the surviving channel does not run through law. A legal proposal must therefore be modest about what it can achieve. It cannot restore deterrence. What it can do is stop the evidential burden falling, as it now does, entirely on the party least able to discharge it.

22.2 The mechanism

Where (a) a measurable interruption occurs in a traded input, (b) the leading supplier of that input holds a demonstrable grip on world exports above a stated threshold, and (c) the interruption is not explained by a publicly available and independently verifiable non-discriminatory cause, then the tribunal may **draw an inference of fact** that the interruption is attributable to the state exercising jurisdiction over the supplier, and an evidential burden of displacing that inference falls on that state.

A note on what changed and why. An earlier draft framed this as a *rebuttable presumption of attribution*. That framing does not survive contact with the law. Attribution under art 2(a) of the ILC Articles means attribution under the secondary rules — an organ, a delegation of governmental authority, instructions or direction or control, or subsequent adoption. World-export share is none of those. A presumption that a state is *responsible* on the strength of the aggregate export share of firms located in its territory supplies no connecting factor at all, and *Bosnian Genocide* declined a connecting factor considerably stronger than that one. The device proposed here is therefore evidential and not attributive: it does not alter what must be established, only what may be inferred from what has been established.

So framed, it has a judicial pedigree the presumption version lacked — though the pedigree is qualified, and the qualification belongs before the support rather than after it. In *Corfu Channel* the Court held that a state denied access to the territory in which the facts occurred should be allowed “a more liberal recourse to inferences of fact and circumstantial evidence”, where the inference rests on “a series of facts linked together and leading logically to a single conclusion”.¹¹⁸ **The same passage states expressly that exclusive territorial control “neither involves prima facie responsibility nor shifts the burden of proof”.** A mechanism that shifted a burden on the strength of a standing fact about a state’s position would be squarely contrary to it, and this one is drafted not to.

The distinction is that limbs (a) to (c) are not a standing fact. Exclusive control, or a dominant supply position, is a permanent attribute; what is proposed here is an inference from a *conjunction* of contingent ones — a measured interruption, a measured world-supply position, and the absence of any published and verifiable ordinary explanation for the first. That is the “series of facts linked together” the Court contemplated, and whether they lead logically to a single conclusion is a question for the tribunal in the case before it. They will not always. That is a property of an inference and a defect in a presumption, and it is the reason for the change.

A further limit from the same authority, and it is not a soft one. The requirement that circumstantial inference “leave no room for reasonable doubt” is the general standard *Corfu Channel* applied to an ordinary question of fact, and it is not confined to charges of exceptional gravity — only the stricter “fully conclusive” formulation in *Bosnian Genocide* is.¹¹⁹ Part II §8.3A sets out the distinction. The mechanism proposed here must therefore be capable of supporting an inference which admits of no reasonable alternative, and that is why limb (c) is drafted as it is: an interruption in a chokepoint held by a dominant supplier is consistent

¹¹⁸ *Corfu Channel* (n 58) 18.

¹¹⁹ *Bosnian Genocide* (n 48) [209].

with several explanations, and it is the *published absence* of any verifiable ordinary explanation, produced by a party who could supply one at negligible cost if it existed, that closes them off. Whether it closes them off in a given case is for the tribunal. What the mechanism supplies is the conjunction that makes the question answerable.

Six features are deliberate.

It attaches to capability and effect, not to intent. Both limbs (a) and (b) are measurable from public data, as Part III demonstrates across the whole of world trade. Neither requires establishing what anyone intended, which is what cannot be done.

It shifts an evidential burden only. The legal burden of establishing responsibility remains with the claimant throughout and the standard of proof is unaltered. What limbs (a) to (c) do is establish a *prima facie* case sufficient to require an answer, which is the ordinary operation of an inference of fact and not a departure from art 2 of the ILC Articles.

It is answerable, and cheaply, by a state acting in good faith. A state facing the inference displaces it by producing the commercial or regulatory record: the licensing decision, the customs data, the sanitary finding and its evidential basis. That is a low burden for a state acting *bona fide* and a high one for a state that is not, which is the correct asymmetry. Shifting an evidential burden to the party with access to the information is a settled response of legal systems to structurally obstructed proof, and the economic analysis of evidence supports it on those grounds.¹²⁰

It does not create a new prohibition. The mechanism operates within existing frameworks: it addresses the evidential route to a predicate that already exists, and could be given effect in the dispute-settlement provisions of a trade agreement without altering any substantive obligation. Part II showed the doctrines to be substantively reasonable and evidentially unusable; the proposal repairs the second without disturbing the first.

It must be given effect through a forum, not asserted unilaterally. This is a constraint on implementation and it is not optional. A member of the World Trade Organization may not determine for itself that a violation has occurred except through recourse to dispute settlement.¹²¹ A trigger computed unilaterally from public data, generating a public finding that another state is responsible, would be such a determination, and the objection is the standing one to unilateral trade instruments generally. The mechanism proposed here is therefore an evidential rule *for a tribunal*, and its unilateral application is expressly outside what is proposed.

The threshold is empirical and public. Part III supplies the distribution against which any threshold must be set. A grip threshold of 0.5 would engage 6.4 per cent of production-input cells; a threshold set at the ninety-ninth percentile of observed grip, 0.73, engages one per cent of them by construction. These are not hypothetical calibrations. They are computed from the reference year and can be recomputed annually by anyone.

22.3 The objections, including the fatal one

“This is protectionism with a security label.” This is the serious objection and it must be met, not deflected. A mechanism attaching to market position gives every state an instrument for harassing efficient foreign suppliers, and the history of the security exception is a history of exactly that abuse.¹²²

Three features limit it, and none eliminates it — and the third is weaker than it looks. First, the inference requires a *measurable interruption*: absent one, market position alone triggers nothing, so an efficient supplier that keeps supplying is untouched. Second, an answer by production of the regulatory record is available and cheap to a state acting in good faith, so the cost of a bad-faith invocation falls on the invoking state in

¹²⁰Richard A Posner, ‘An Economic Approach to the Law of Evidence’ (1999) 51 *Stanford Law Review* 1477; Bruce L Hay and Kathryn E Spier, ‘Burdens of Proof in Civil Litigation: An Economic Perspective’ (1997) 26 *The Journal of Legal Studies* 413. For the contrary view see Louis Kaplow, ‘Burden of Proof’ (2012) 121 *Yale Law Journal* 738, discussed at §22.3.

¹²¹Understanding on Rules and Procedures Governing the Settlement of Disputes, art 23.1 and art 23.2(a).

¹²²Pinchis-Paulsen (n 54).

the form of a displaced claim. Third, the grip threshold is computed from public data on a stated method, so the trigger cannot be *drafted* to catch a particular supplier without catching all similarly situated ones.

That third point must be put more honestly than it was in an earlier draft, because the paper's own evidence undercuts it. Part III §12.2 establishes that one state is the leading supplier in 73.3 per cent of severe production-input chokepoints. **A neutrally drafted trigger would therefore operate overwhelmingly against a single state.** Formal generality is not substantive generality when the underlying distribution is that concentrated, and no drafting device conceals it.

The honest position is that this is a cost of the proposal rather than a flaw in it, but it has to be argued rather than mitigated away. If the concentration is real — and it is measured, published and checkable — then a mechanism triggered by concentration falls where concentration falls, and the alternative is a mechanism that does not operate at all. What cannot be claimed is that such a mechanism is neutral in effect. It is neutral in form; its incidence is determined by the world; and a state receiving three-quarters of its operation will not regard that distinction as meaningful. Any serious attempt to adopt it must expect to be met on exactly that ground.

Those are mitigations and not answers, and the honest position is that the proposal trades one abuse risk for another. The question is which is worse: a regime in which coercion is unreachable because authorship cannot be proved, or one in which the evidential burden shifts and is sometimes shifted opportunistically. Part IV's threshold result is what makes that a real choice rather than a rhetorical one — the current position is not a functioning regime with an evidential difficulty, it is a regime whose operative condition is systematically unmet.

“The rebuttal condition destroys the trigger, and the paper's own case study proves it.” This is the objection that comes closest to being fatal, and it is generated by the paper's own evidence rather than by an opponent's. Limb (c) provides that no inference arises where the interruption is explained by a publicly available and independently verifiable non-discriminatory cause. But Part III §11 records that the measures against Lithuania were expressed in sanitary and phytosanitary terms, and that the stated reasons were not the operative ones. A state that has read this proposal will publish a pretext. Worse, the mechanism *creates an incentive* to publish one, where previously the coercing state had no reason to generate any record at all.

Three things must be said, and only the first two are answers.

The word doing the work in limb (c) is “verifiable”, and it is doing more than it looks. A sanitary assertion and a sanitary finding are different objects. What limb (c) requires is not a stated reason but a published evidential basis: the sampling protocol, the test results, the laboratory, the risk assessment. That is not costless to fabricate, it is falsifiable once published, and it exposes the state that produces it to a second and better-understood contest — whether the measure is based on scientific principles and a risk assessment, which the sanitary agreement already requires and which was precisely what was pleaded in DS610.¹²³ The mechanism therefore does not defeat the pretext. It *converts* the pretext from an unfalsifiable assertion into a falsifiable claim, and moves the dispute from a question no tribunal can answer — what did you intend — to a question tribunals answer routinely, which is whether a stated scientific basis exists and supports the measure taken.

That conversion is the whole logic of this paper applied to itself. The argument throughout has been that institutions fail where the cost of proof exceeds the cost of the act. The response is not to prove the unprovable but to *raise the cost of the concealment relative to the cost of acting honestly*. A state with a genuine phytosanitary concern already holds the record and produces it at negligible cost. A state without one must manufacture it, at a cost that rises with the scrutiny it invites and with the number of episodes across which it must remain consistent. The incentive to generate a record is not a defect of the mechanism; it is its operative effect. A coercing state that is forced onto the record has given up the property that made deniable coercion attractive in the first place.

But the residual is real and is not answered. A well-resourced state, generating a competent and internally consistent scientific record, will defeat the mechanism. Nothing here prevents that, and it should

¹²³Agreement on the Application of Sanitary and Phytosanitary Measures, arts 2.2, 5.1 and 5.2; and see the provisions invoked in DS610 (n 30).

not be claimed that anything does. What the mechanism achieves against such a state is a rise in the cost and the delay of each episode and a body of published assertions that can be tested against each other over time. That is a smaller achievement than deterrence and it is the honest measure of what an evidential rule can do.

“You have not priced the false positives, and your own model says that is the central parameter.” Correct, and the omission should be stated rather than covered. Part IV §19 concedes that a commitment to respond without proof buys deterrence “at the price of punishing the innocent”, and the size of that price is q , the rate of innocent disruption. Nothing in the trade data estimates q : the instrument measures structural exposure, and limb (a) requires an interruption, which is an event the data does not contain. Estimating it would require an event dataset of supply interruptions at tariff-line resolution with independently established causes, which does not exist and which this paper does not supply. Until it does, any threshold set under §22.2 is a judgement and not a calibration, and the mechanism should be adopted with a conservative threshold and a corroboration requirement for cells near the boundary. A proposal whose central design parameter is unestimated should say so in the section where it is proposed, and this is that section.

“Its incidence would offend most-favoured-nation treatment.” A mechanism whose operation falls overwhelmingly on one state, and which generates measures in consequence, raises questions under the most-favoured-nation and non-discrimination obligations independently of the political objection addressed above. The answer is that the mechanism as reframed generates no measures at all — it is an evidential rule applied by a tribunal to whoever is before it, and any measure that follows is the ordinary consequence of an adverse finding. That answer holds only so long as the mechanism is not implemented unilaterally, which is why §22.2 makes that a condition rather than a preference.

“Presumptions of this kind offend the principle that he who asserts must prove.” As reframed the mechanism is an inference of fact and does not disturb the legal burden, so the objection loses most of its force. To the extent an evidential burden shifts, the departure is justified on the ground that the burden should sit with the party who can discharge it at least cost, and the state holding the regulatory record can discharge it at trivially lower cost than the state attempting to infer authorship from outside.¹²⁴ It should be recorded that the leading economic treatment is *sceptical* of that rationale: Kaplow argues that conventional justifications for proof burdens, including the least-cost-provider rationale, do not survive welfare analysis, and that the optimal burden turns instead on the ratio of harmful to benign acts in the region of the threshold.¹²⁵ On that analysis the right question is the one the preceding objection concedes cannot yet be answered — what proportion of triggered cases are genuine coercion — which is a further reason for a conservative threshold and against treating the least-cost rationale as settling anything.

“States will not agree to it.” Probably not, at least not multilaterally, and the paper does not pretend otherwise. The mechanism is more likely to appear first in bilateral or plurilateral instruments among parties with symmetrical exposure, and Part IV §19.1 shows that to be a *small* group: only two of the eighteen interpretable mutual pairs are within a factor of two on magnitude. That is an argument for expecting slow and partial adoption rather than none, and the value of specifying the mechanism now is that it exists when the political conditions arrive.

“The measurement is too coarse to bear the weight.” Correct, and Part III §10 sets out the limits without softening them: HS six-digit cannot see sub-heading dependencies, customs origin is not true origin, and evidenced capacity is not potential capacity. The mitigations differ by limit, and an earlier draft offered one mitigation for all three. Only the evidenced-capacity limit produces false *positives*, and a conservative threshold answers it. Resolution and customs origin run the other way: Part III §10 states that both cause the measure to under-read, and names rare earths as a case where measured grip on the compounds heading is far below the concentration of the refining that matters. Those produce false *negatives*, and a conservative threshold makes them worse. The mechanism will therefore under-fire, and it will under-fire hardest on mineral inputs — which are the goods most likely to be coerced through. That is a serious limitation and it is stated rather than mitigated: under-firing leaves the status quo, which is the lesser of the two policy

¹²⁴Posner (n 120); Hay and Spier (n 120).

¹²⁵Kaplow (n 120).

errors, but it means the instrument is weakest exactly where it would be most wanted. The answer is better resolution on the supply side rather than a looser threshold, because a looser threshold would buy the missing coverage by importing the false positives the first limit already threatens.

23. The English courts have already been here

Neither proposal invents a technique. Both apply reasoning the English courts have developed while confronting, in other contexts, the precise problem this paper describes: a decision that must be made when the material required to make it properly cannot be obtained.

Where the state must show its reasons. *Bank Mellat v HM Treasury (No 2)* concerned a direction restricting a bank's access to the United Kingdom financial sector.¹²⁶ The Supreme Court's treatment of proportionality required the Treasury to demonstrate that the measure was rationally connected to its objective and no more intrusive than necessary — that is, to produce its reasoning rather than assert its conclusion. The parallel with Proposal Two is close, though not exact: *Bank Mellat* is a domestic court applying proportionality to its own executive under a statutory scheme, whereas Proposal Two operates on a foreign sovereign in an international forum, and the difference in the court's authority over the respondent is not trivial. What transfers is the technique rather than the jurisdiction. The mechanism proposed at §22 does not ask a tribunal to divine a state's motive; it asks the state to produce the record on which it acted, which is what *Bank Mellat* required domestically of the executive and what the ordinary law of evidence requires of any party holding material the other cannot reach.

Where the standard of proof is doing the work. *Youssef v Secretary of State for Foreign and Commonwealth Affairs* addressed the basis on which a person may be listed, and the extent to which a decision-maker must satisfy itself of the underlying facts.¹²⁷ The question is the same one Part IV's — formalises: a regime conditioned on proof must specify how much proof, and the specification determines whether the regime operates at all. *Shagang Shipping Co Ltd v HNA Group Co Ltd* addresses the converse and is the more instructive authority.¹²⁸ The Supreme Court's treatment of an allegation that could not be established to the requisite standard confirms that the common law's answer to unprovable assertion is to disregard it — which is correct as between litigants, and is precisely why the burden's allocation matters so much when the party bearing it structurally cannot discharge it.

Where the court declines the question altogether. *Belhaj v Straw* concerned the limits of the foreign act of state doctrine and the circumstances in which an English court will adjudicate upon the conduct of a foreign sovereign.¹²⁹ The decision permits adjudication further than had been assumed, but the underlying tension is unresolved and it bears directly on Proposal Two: a mechanism that invites a finding against a foreign state invites the objection that the forum is not competent to make it. The answer available on the analysis here is that the inference arises from limbs (a) to (c) — a measured interruption, a measured world-supply position, and the published absence of a verifiable ordinary explanation — all established from public data, none of which requires the court to characterise a foreign sovereign's purpose.

Where a duty is imposed rather than implied. *Barton v Morris* is a caution against the proposal rather than support for it.¹³⁰ The Supreme Court declined to imply a term to fill a gap the parties' bargain had left, on the basis that the absence of provision was itself a feature of the bargain. The lesson for the verification duty proposed in the companion paper is that a verification duty on a vendor cannot be reached by implication into a licence agreement whose terms disclaim exactly that responsibility. It must be imposed, which is why the proposal is framed as a duty in law and not as a matter of construction.

Where loss is allocated long after the act. *URS Corporation Ltd v BDW Trading Ltd* concerned the recovery of losses arising from defects discovered many years after the work.¹³¹ The temporal structure is the one both proposals must survive: a fabricated citation may not be discovered until an appeal, and a

¹²⁶ *Bank Mellat v HM Treasury (No 2)* [2013] UKSC 39, [2014] AC 700 [20] (Lord Sumption), [74] (Lord Reed).

¹²⁷ *Youssef v Secretary of State for Foreign and Commonwealth Affairs* [2016] UKSC 3, [2016] AC 1457 [55]–[59].

¹²⁸ *Shagang Shipping Co Ltd v HNA Group Co Ltd* [2020] UKSC 34, [2020] 1 WLR 3549 [90]–[100].

¹²⁹ *Belhaj v Straw* [2017] UKSC 3, [2017] AC 964 [11], [123]–[124].

¹³⁰ *Barton v Morris* [2023] UKSC 3, [2023] AC 684 [22]–[28].

¹³¹ *URS Corporation Ltd v BDW Trading Ltd* [2025] UKSC 21.

dependency created by a sourcing decision may not be exploited for a decade. A liability rule that expires before the harm materialises allocates nothing, and any implementation of the verification duty proposed in the companion paper would need a limitation period running from discoverability rather than from the assertion.

Taken together these authorities establish something modest but useful. The English courts have repeatedly had to decide what follows when the evidence that would resolve a question is unavailable, and their answers have been consistent in structure: identify who is best placed to produce what is missing, require them to produce it, and treat failure to do so as significant. That is the whole of what this paper proposes, applied to a problem the common law has not yet been asked to confront directly.

24. Enforcement, and why a rule nobody enforces is worse than none

Both proposals impose duties, and a duty is only as good as the probability of its being enforced. The enforcement literature is unambiguous about what follows when that probability is low, and its conclusions are unflattering to the instinct that a rule is better than nothing.

Becker's framework and Stigler's refinement establish the arithmetic: what governs behaviour is the expected sanction, and a rule with negligible enforcement probability changes nothing except the vocabulary in which non-compliance is described.¹³² The institutional literature adds that the *form* of enforcement matters as much as its intensity. Milgrom, North and Weingast's account of the Law Merchant shows how a private institution can sustain contracting without a state, provided it maintains a *reliable record* of who did what.¹³³ Dixit's model of trade expansion under weak enforcement reaches the same conclusion from the theoretical side: the constraint on the extent of the market is the reach of whatever mechanism verifies conduct.¹³⁴

That is the identical structure this paper has traced twice, and its recurrence in a body of work on medieval trade is worth pausing on. A market grows to the extent that conduct can be verified and recorded; where the record fails, exchange retreats to parties who can monitor each other directly. Both proposals here are, in that light, attempts to maintain a record adequate to the extent of the market — the citation register in the verification duty proposed in the companion paper, the public measurement in Proposal Two.

The political-economy literature supplies the caution. Shleifer and Vishny's analysis of corruption shows how enforcement discretion becomes a rent, and Olson's account of stationary banditry explains why an enforcing authority's incentives are not automatically aligned with the rule it enforces.¹³⁵ Applied to Proposal Two, the risk is concrete and was conceded at §22.3: a mechanism triggered by market position hands the invoking state an instrument, and states with instruments use them. The mitigation offered — that the trigger is computed from public data on a published method — is a constraint on discretion of exactly the kind these literatures suggest is necessary and, on their evidence, insufficient on its own.

Two further observations bear on feasibility.

Formalism is a feature here, not a vice. Weisbach's defence of formalism in the tax context argues that bright-line rules are preferable where the alternative is a standard that sophisticated parties will plan around.¹³⁶ Both proposals are deliberately formal — resolve the citation, compute the grip — because the parties subject to them are sophisticated and a standard would be planned around within a season.

Measurement of legal texts is now feasible enough to monitor compliance. The empirical legal literature has demonstrated that large bodies of legal text can be analysed computationally at scale: Livermore and colleagues on judicial writing, Ash and Hansen's survey of text algorithms in economics, and

¹³²Becker (n 97); George J Stigler, 'The Optimum Enforcement of Laws' (1970) 78 *Journal of Political Economy* 526.

¹³³Paul R Milgrom, Douglass C North and Barry R Weingast, 'The role of institutions in the revival of trade: the Law Merchant, private judges, and the Champagne fairs' (1990) 2 *Economics & Politics* 1.

¹³⁴Avinash Dixit, 'Trade Expansion and Contract Enforcement' (2003) 111 *Journal of Political Economy* 1293.

¹³⁵Andrei Shleifer and Robert W Vishny, 'Corruption' (1993) 108 *Quarterly Journal of Economics* 599; Mancur Olson, 'Dictatorship, Democracy, and Development' (1993) 87 *American Political Science Review* 567.

¹³⁶David A Weisbach, 'Formalism in the Tax Law' (1999) 66 *The University of Chicago Law Review* 860.

Bybee and colleagues’ extraction of macroeconomic signal from business news all establish the capability.¹³⁷ Whether a jurisdiction’s filings show a rising share of unresolved citations is, on that evidence, a measurable quantity — which matters, because a duty whose breach cannot be observed in aggregate cannot be shown to be working, and an unmonitored duty decays into an incantation.

Part VI — Conclusion

25. Reflective analysis

This paper began from an asymmetry: the cost of applying economic pressure has fallen and the cost of proving who applied it has not. Four conclusions follow and they are of unequal strength.

The doctrinal survey is secure and its result is narrow. Six routes exist; five fail at the same point, which is proof rather than substance; the sixth reasons from effect to authorship and fails on standing and forum instead. The doctrines are, in each case, substantively reasonable. What has changed is not the law’s content but the cost of satisfying it, and a rule whose satisfaction has become prohibitively expensive is not a rule that has been weakened. It is a rule that has been repealed without anyone voting on it.

The migration result is the paper’s strongest theoretical claim and it is a feasibility argument rather than a model. Where the attainable probability of attribution is low, no proportionate countermeasure deters, and both premises are defensible. What that establishes is narrower than an earlier version of this argument asserted: it closes the countermeasure channel, not retaliation. Untargeted retorsion remains available and requires no attribution, so retaliation continues in a form that has shed its predicate and its forum. It is not bounded by nothing — the retaliator’s own cost and the risk of escalation bound it — but those bounds are matters of self-interest rather than of law, and they are neither adjudicable nor stated in advance. The threshold is also a property of a single dyad and a single episode; with heterogeneity in the aggressor’s benefit the measure of deterred episodes falls smoothly, and an earlier formulation conflated the two.

The empirical work is the most secure element and the most limited. The measurement is exhaustive at the resolution available, reproducible by anyone, and produces a finding that was not obvious: that severe supply chokepoints in production inputs are rare, carrying under one per cent of input trade by value, while three-quarters of those that exist have the same leading supplier. But it cannot see below six-digit resolution, it reads customs origin rather than true origin, and evidenced capacity is not potential capacity. Every figure in Part III is a lower bound and should be read as one.

The reciprocity result is negative, and it removes a consolation rather than adding a warning. If deterrence survives below the threshold at all, it survives through mutual exposure. That reciprocity is not there: reciprocal severe dependency occurs at exactly the rate the degree distribution of the trade network predicts, and where it occurs the two sides’ exposures differ by a median factor of fifteen. The direction of that answer would not have surprised Hirschman. What is new is that it has been measured rather than asserted, across the whole of world trade, on a published instrument — and that reaching it required discarding two positive results, each of which looked like a finding until the null was specified properly.

The most uncomfortable conclusion is the one about what remains, and it deserves restating without softening. A state facing deniable economic coercion that invests in doctrine and litigation is spending on the channel that has closed, while the channel that remains open is the one general international law leaves without proportionality discipline or terminating rule. The prediction is therefore not that such states will be defenceless. It is that the observable response to unattributable economic pressure will be an escalating exchange of measures that neither party characterises as a legal claim, because characterising it as one would require proving something. That is close to what is presently observed, which is either confirmation of the analysis or the reason the analysis was available to be made; the paper cannot distinguish those.

¹³⁷Michael A Livermore, Allen B Riddell and Daniel N Rockmore, ‘The Supreme Court and the Judicial Genre’ (2017) 59 *Arizona Law Review* 837; Elliott Ash and Stephen Hansen, ‘Text Algorithms in Economics’ (2023) 15 *Annual Review of Economics* 659; Leland Bybee and others, ‘Business News and Business Cycles’ (2024) 79 *The Journal of Finance* 3105.

Reducing exposure remains the instrument that works, but it should be labelled correctly. It removes the aggressor's motive rather than raising its expected cost, which is *prevention*, not deterrence. And the reciprocity result means the usual caution against decoupling — that severing interdependence destroys a deterrent — is weaker than it is usually stated: in most dyads the deterrent is not there to destroy, and a state relying on mutual exposure to restrain a supplier should first check whether the exposure runs in the direction it needs.

26. What would change the analysis

A durable shift in the standard of proof for state responsibility. If tribunals moved toward inference from effect and capability — which is what §22.2 proposes — the threshold would be relaxed, because the attainable probability of attribution is a function of the standard applied as well as of the evidence available. Nothing in current practice suggests this is imminent, though *Corfu Channel* supplies more room than the survey in Part II initially allowed.

A change in the reciprocity structure. Part IV reports that reciprocal severe dependency sits at its degree-preserving expectation, so the mutual-exposure condition is rarely satisfied. That is a fact about 2024 and not a law of nature, and it could move in either direction. If reciprocity rose above its expectation — which is the falsification test named at §4.3 — the condition would become generally available and the analysis would be materially less bleak. If instead the largest holder of chokepoints reduced its own exposure while retaining its supply positions, the small number of dyads in which the condition does hold would shrink further. Both are measurable annually on public data, and the published instrument tracks them.

Wider adoption of the interim appellate route. The security-exception jurisprudence rests on one adopted panel report and four appealed into a vacancy, which is the institutional half of the problem this paper describes. Part II §6B.1 records that an arbitration under DSU art 25 produced a binding award against a major economy and an assertion of compliance in response. What that route cannot do is bind a Member which has not joined it, and the Member whose measures generated the unadopted reports has not.

27. Final remarks

The argument can be put in a sentence. **Law allocates consequence on proof of authorship, and a technology has arrived that decouples the cost of acting from the cost of proving who acted.**

Where that decoupling occurs, the failures do not appear as failures. An unattributable act looks exactly like no act at all. A market with no price for coercion risk looks like a market that has judged the risk to be small. A doctrine that cannot be satisfied looks like a doctrine that is not being breached. In each case the institution reports success while doing nothing, and it will go on reporting success until something forces the question.

That is the reason for insisting throughout on measurement and on the honest statement of what a measurement cannot see. The instrument published with this paper does not solve the problem; a trade-concentration score is a modest thing. What it does is make a specific class of silence audible: a dependency that has no substitute, a reciprocity that is not there, a complaint withdrawn without a finding. None of those is a remedy. Each is a prompt to look, in a domain where the default is not to look because looking has become expensive.

The mechanism proposed in Part V is correspondingly modest, and it does not restore the world in which authorship was provable. That world is not coming back. What can be done is to stop the cost of proof falling by default on whoever happens to be standing closest to the harm — the state whose supply stopped, the firm whose consignment was held, the tribunal reading a sanitary finding it cannot test. That allocation is not a policy. It is what a legal system arrives at when it has not noticed that the question has changed.

It has changed. The institutions have not.

The datasets, derivation code and verification apparatus described in this paper are published in full and run offline. Every figure reported is reproducible from the stated sources by the stated method.

Bibliography

Books

- Buchanan B, *The Cybersecurity Dilemma* (Oxford University Press 2017)
Dolzer R and Schreuer C, *Principles of International Investment Law* (Oxford University Press 2012)
Hirschman AO, *National Power and the Structure of Foreign Trade* (University of California Press 1945)
Hufbauer GC and others, *Economic Sanctions Reconsidered* (Peterson Institute for International Economics 2007)
Keohane RO and Nye JS, *Power and Interdependence: World Politics in Transition* (Little, Brown 1977)
Mulder N, *The Economic Weapon: The Rise of Sanctions as a Tool of Modern War* (Yale University Press 2022)

Contributions to edited collections

- Barberis N, ‘Psychology-Based Models of Asset Prices and Trading Volume’ in *Handbook of Behavioral Economics: Applications and Foundations 1* (Elsevier 2018) 79
Kosseff J, ‘Retorsion as a Response to Ongoing Malign Cyber Operations’ in *12th International Conference on Cyber Conflict: 20/20 Vision - The Next Decade* (NATO CCDCOE Publications 2020) 9
Roscini M, ‘Evidentiary Issues in International Disputes Related to State Responsibility for Cyber Operations’ in *Cyber War* (Oxford University Press 2015) 215
Ruys T, ‘Sanctions, Retorsions and Countermeasures: Concepts and International Legal Framework’ in *Research Handbook on UN Sanctions and International Law* (Edward Elgar 2017) 19

Journal articles

- Abreu D, Pearce D and Stacchetti E, ‘Toward a Theory of Discounted Repeated Games with Imperfect Monitoring’ (1990) 58 *Econometrica* 1041
Acemoglu D and others, ‘The Network Origins of Aggregate Fluctuations’ (2012) 80 *Econometrica* 1977
Acemoglu D, Ozdaglar A and Tahbaz-Salehi A, ‘Systemic Risk and Stability in Financial Networks’ (2015) 105 *American Economic Review* 564
Alstadsæter A, Johannesen N and Zucman G, ‘Tax Evasion and Inequality’ (2019) 109 *American Economic Review* 2073
Amiti M, Redding SJ and Weinstein DE, ‘The Impact of the 2018 Tariffs on Prices and Welfare’ (2019) 33 *Journal of Economic Perspectives* 187
Ash E and Hansen S, ‘Text Algorithms in Economics’ (2023) 15 *Annual Review of Economics* 659
Baker SR, Bloom N and Davis SJ, ‘Measuring Economic Policy Uncertainty’ (2016) 131 *Quarterly Journal of Economics* 1593
Baldwin DA, ‘The Sanctions Debate and the Logic of Choice’ (2000) 24 *International Security* 80
Baldwin R, Freeman R and Theodorakopoulos A, ‘Hidden Exposure: Measuring US Supply Chain Reliance’ (2023) 2023 *Brookings Papers on Economic Activity* 79
Bapat NA and Clifton Morgan T, ‘Multilateral Versus Unilateral Sanctions Reconsidered: A Test Using New Data’ (2009) 53 *International Studies Quarterly* 1075
Barrot JN and Sauvagnat J, ‘Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks’ (2016) 131 *Quarterly Journal of Economics* 1543
Bartels L, ‘Withdrawing Provisional Application of Treaties: Has the EU Made a Mistake?’ (2012) 1 *Cambridge Journal of International and Comparative Law* 112
Becker GS, ‘Crime and Punishment: An Economic Approach’ (1968) 76 *Journal of Political Economy* 169
Bianchi F, Ludvigson SC and Ma S, ‘Belief Distortions and Macroeconomic Fluctuations’ (2022) 112 *American Economic Review* 2269
Biener C, Eling M and Wirfs JH, ‘Insurability of Cyber Risk: An Empirical Analysis’ (2015) 40 *Geneva*

Papers on Risk and Insurance - Issues and Practice 131

Bloom N, 'The Impact of Uncertainty Shocks' (2009) 77 *Econometrica* 623

Bonadio B and others, 'Global supply chains in the pandemic' (2021) 133 *Journal of International Economics* 103534

Borghard ED and Lonergan SW, 'The Logic of Coercion in Cyberspace' (2017) 26 *Security Studies* 452

Bown CP, 'The 2018 US-China trade conflict after forty years of special protection' (2019) 12 *China Economic Journal* 109

Bybee L and others, 'Business News and Business Cycles' (2024) 79 *Journal of Finance* 3105

Caldara D and Iacoviello M, 'Measuring Geopolitical Risk' (2022) 112 *American Economic Review* 1194

Caldara D and others, 'The economic effects of trade policy uncertainty' (2020) 109 *Journal of Monetary Economics* 38

Carson A, 'Facing Off and Saving Face: Covert Intervention and Escalation Management in the Korean War' (2015) 70 *International Organization* 103

Carvalho VM and others, 'Supply Chain Disruptions: Evidence from the Great East Japan Earthquake' (2021) 136 *Quarterly Journal of Economics* 1255

Cassese A, 'The Nicaragua and Tadić Tests Revisited in Light of the ICJ Judgment on Genocide in Bosnia' (2007) 18 *European Journal of International Law* 649

Cha VD, 'Collective Resilience: Deterring China's Weaponization of Economic Interdependence' (2023) 48 *International Security* 91

Cormac R and Aldrich RJ, 'Grey is the new black: covert action and implausible deniability' (2018) 94 *International Affairs* 477

Damgaard J, Elkjaer T and Johannesen N, 'What is real and what is not in the global FDI network?' (2024) 140 *Journal of International Money and Finance* 102971

Dixit A, 'Trade Expansion and Contract Enforcement' (2003) 111 *Journal of Political Economy* 1293

Draca M and others, 'On Target? Sanctions and the Economic Interests of Elite Policymakers in Iran' (2022) 133 *Economic Journal* 159

Drezner DW, 'The Hidden Hand of Economic Coercion' (2003) 57 *International Organization* 643

Easley D and others, 'Liquidity, Information, and Infrequently Traded Stocks' (1996) 51 *Journal of Finance* 1405

Efrony D and Shany Y, 'A Rule Book on the Shelf? Tallinn Manual 2.0 on Cyberoperations and Subsequent State Practice' (2018) 112 *American Journal of International Law* 583

Egloff FJ, 'Public attribution of cyber intrusions' (2020) 6 *Journal of Cybersecurity*

Egloff FJ and Dunn Cavelty M, 'Attribution and Knowledge Creation Assemblages in Cybersecurity Politics' (2021) 7 *Journal of Cybersecurity*

Eling M and Wirfs J, 'What are the actual costs of cyber risk events?' (2019) 272 *European Journal of Operational Research* 1109

Elliott M, Golub B and Jackson MO, 'Financial Networks and Contagion' (2014) 104 *American Economic Review* 3115

Elliott M, Golub B and Leduc MV, 'Supply Network Formation and Fragility' (2022) 112 *American Economic Review* 2701

Fajgelbaum PD and others, 'The Return to Protectionism' (2020) 135 *Quarterly Journal of Economics* 1

Farrell H and Newman AL, 'Weaponized Interdependence: How Global Economic Networks Shape State Coercion' (2019) 44 *International Security* 42

Fearon JD, 'Rationalist explanations for war' (1995) 49 *International Organization* 379

Felbermayr G and others, 'The global sanctions data base' (2020) 129 *European Economic Review* 103561

Ferguson VA, Waldron S and Lim DJ, 'Market adjustments to import sanctions: lessons from Chinese restrictions on Australian trade, 2020–21' (2023) 30 *Review of International Political Economy* 1255

Findley MG, Nielson DL and Sharman J, 'Causes of Noncompliance with International Law: A Field Experiment on Anonymous Incorporation' (2015) 59 *American Journal of Political Science* 146

Gai P and Kapadia S, 'Contagion in financial networks' (2010) 466 *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2401

Garlaschelli D and Loffredo MI, 'Patterns of Link Reciprocity in Directed Networks' (2004) 93 *Physical Review Letters* 268701

Gartzke E, 'The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth' (2013) 38 *International*

tional Security 41

Gaulier G and Zignago S, 'BACI: International Trade Database at the Product-Level (the 1994-2007 Version)' (2010) CEPII Working Paper No 2010-23

Gaustad G and others, 'Circular economy strategies for mitigating critical material supply issues' (2018) 135 Resources, Conservation and Recycling 24

Gilboa I and Schmeidler D, 'Maxmin expected utility with non-unique prior' (1989) 18 Journal of Mathematical Economics 141

Glasserman P and Young HP, 'How likely is contagion in financial networks?' (2015) 50 Journal of Banking & Finance 383

Glosten LR and Milgrom PR, 'Bid, ask and transaction prices in a specialist market with heterogeneously informed traders' (1985) 14 Journal of Financial Economics 71

Green EJ and Porter RH, 'Noncooperative Collusion under Imperfect Price Information' (1984) 52 Econometrica 87

Green JA, 'Fluctuating Evidentiary Standards for Self-Defence in the International Court of Justice' (2009) 58 International and Comparative Law Quarterly 163

Gulley AL, Nassar NT and Xun S, 'China, the United States, and competition for resources that enable emerging technologies' (2018) 115 Proceedings of the National Academy of Sciences 4111

Gutmann J, Neuenkirch M and Neumeier F, 'Sanctioned to Death? The Impact of Economic Sanctions on Life Expectancy and its Gender Gap' (2021) 57 Journal of Development Studies 139

Hansen LP and Sargent TJ, 'Robust Control and Model Uncertainty' (2001) 91 American Economic Review 60

Hassan TA and others, 'Firm-Level Political Risk: Measurement and Effects' (2019) 134 Quarterly Journal of Economics 2135

Hassan TA and others, 'Sources and Transmission of Country Risk' (2024) 91 Review of Economic Studies 2307

Hay BL and Spier KE, 'Burdens of Proof in Civil Litigation: An Economic Perspective' (1997) 26 Journal of Legal Studies 413

Jackson M and Paddeu F, 'To deter and counteract economic coercion: the puzzle of countermeasures in the European Union's Anti-Coercion Instrument' (2025) 28 Journal of International Economic Law 567

Jamnejad M and Wood M, 'The Principle of Non-intervention' (2009) 22 Leiden Journal of International Law 345

Jervis R, 'Cooperation under the Security Dilemma' (1978) 30 World Politics 167

Johannessen N and Zucman G, 'The End of Bank Secrecy? An Evaluation of the G20 Tax Haven Crackdown' (2014) 6 American Economic Journal: Economic Policy 65

Kaplow L, 'Burden of Proof' (2012) 121 Yale Law Journal 738

Kirshner J, 'The microfoundations of economic sanctions' (1997) 6 Security Studies 32

Kreps DM and Wilson R, 'Reputation and imperfect information' (1982) 27 Journal of Economic Theory 253

Kyle AS, 'Continuous Auctions and Insider Trading' (1985) 53 Econometrica 1315

Lee MJ and others, 'Uncovering hidden dependency in weighted networks via information entropy' (2021) 3 Physical Review Research 043136

Lim DJ and Ferguson VA, 'Informal economic sanctions: the political economy of Chinese coercion during the THAAD dispute' (2022) 29 Review of International Political Economy 1525

Lindsay JR, 'Tipping the scales: the attribution problem and the feasibility of deterrence against cyberattack' (2015) Journal of Cybersecurity tyv003

Livermore MA, Riddell AB and Rockmore DN, 'The Supreme Court and the Judicial Genre' (2017) 59 Arizona Law Review 837

Mancheri NA and others, 'Effect of Chinese policies on rare earth supply chain resilience' (2019) 142 Resources, Conservation and Recycling 101

Milanovic M, 'State Responsibility for Genocide' (2006) 17 European Journal of International Law 553

Milgrom P and Roberts J, 'Predation, reputation, and entry deterrence' (1982) 27 Journal of Economic Theory 280

Milgrom PR, North DC and Weingast BR, 'The role of institutions in the revival of trade: the Law Merchant, private judges, and the Champagne fairs' (1990) 2 Economics & Politics 1

Morgan TC, Bapat N and Krustev V, ‘The Threat and Imposition of Economic Sanctions, 1971–2000’ (2009) 26 *Conflict Management and Peace Science* 92

Neuenkirch M and Neumeier F, ‘The impact of UN and US economic sanctions on GDP growth’ (2015) 40 *European Journal of Political Economy* 110

Norrlof C, ‘Dollar hegemony: A power analysis’ (2014) 21 *Review of International Political Economy* 1042

Nye JS, ‘Deterrence and Dissuasion in Cyberspace’ (2017) 41 *International Security* 44

Olson M, ‘Dictatorship, Democracy, and Development’ (1993) 87 *American Political Science Review* 567

Pape RA, ‘Why Economic Sanctions Do Not Work’ (1997) 22 *International Security* 90

Pástor L and Veronesi P, ‘Uncertainty about Government Policy and Stock Prices’ (2012) 67 *Journal of Finance* 1219

Pinchis-Paulsen M, ‘Trade Multilateralism and U.S. National Security: The Making of the GATT Security Exceptions’ (2020) 41 *Michigan Journal of International Law* 109

Posner RA, ‘An Economic Approach to the Law of Evidence’ (1999) 51 *Stanford Law Review* 1477

Poznansky M, ‘Feigning Compliance: Covert Action and International Law’ (2019) 63 *International Studies Quarterly* 72

Rid T and Buchanan B, ‘Attributing Cyber Attacks’ (2015) 38 *Journal of Strategic Studies* 4

Riddle M and others, ‘Global critical materials markets: An agent-based modeling approach’ (2015) 45 *Resources Policy* 307

Sartori AE, ‘The Might of the Pen: A Reputational Theory of Communication in International Disputes’ (2002) 56 *International Organization* 121

Shavell S, ‘Liability for Harm versus Regulation of Safety’ (1984) 13 *Journal of Legal Studies* 357

Shavell S, ‘The judgment proof problem’ (1986) 6 *International Review of Law and Economics* 45

Shleifer A and Vishny RW, ‘Corruption’ (1993) 108 *Quarterly Journal of Economics* 599

Slantchev BL, ‘Feigning Weakness’ (2010) 64 *International Organization* 357

Sonnenfeld J and others, ‘Business Retreats and Sanctions Are Crippling the Russian Economy’ (2022) SSRN Electronic Journal

Squartini T, Fagiolo G and Garlaschelli D, ‘Randomizing world trade. I. A binary network analysis’ (2011) 84 *Physical Review E* 046117

Squartini T, Fagiolo G and Garlaschelli D, ‘Randomizing world trade. II. A weighted network analysis’ (2011) 84 *Physical Review E* 046118

Stigler GJ, ‘The Optimum Enforcement of Laws’ (1970) 78 *Journal of Political Economy* 526

Talmon S, ‘The Responsibility of Outside Powers for Acts of Secessionist Entities’ (2009) 58 *International and Comparative Law Quarterly* 493

Townsend RM, ‘Optimal contracts and competitive markets with costly state verification’ (1979) 21 *Journal of Economic Theory* 265

Tsagourias N, ‘Cyber attacks, self-defence and the problem of attribution’ (2012) 17 *Journal of Conflict and Security Law* 229

Voon T, ‘The Security Exception In WTO Law: Entering a New Era’ (2019) 113 *AJIL Unbound* 45

Wagner RH, ‘Economic interdependence, bargaining power, and political influence’ (1988) 42 *International Organization* 461

Weeks JL, ‘Autocratic Audience Costs: Regime Type and Signaling Resolve’ (2008) 62 *International Organization* 35

Weisbach DA, ‘Formalism in the Tax Law’ (1999) 66 *University of Chicago Law Review* 860

Zucman G, ‘The Missing Wealth of Nations: Are Europe and the U.S. net Debtors or net Creditors?’ (2013) 128 *Quarterly Journal of Economics* 1321

Reports, working papers and official documents

European Commission, ‘Strategic dependencies and capacities’ (2021) SWD(2021) 352 final Commission Staff Working Document

McKinsey Global Institute, ‘The complication of concentration in global trade’ (2023) McKinsey Global Institute

Sheffrin A, ‘Predicting Market Power Using the Residual Supply Index’ (2002) FERC Market Monitoring Workshop

Appendix — Verification apparatus

Every authority in this paper was checked against a primary or authoritative record before citation, by the following method. It is set out in full because the paper’s argument makes the author’s own practice material, and because the most instructive failures it records are the author’s own.

Channel 1 — Crossref. Each candidate reference was resolved against the Crossref REST API and accepted only where the returned title, first author and year agreed with the claim. Where a working paper and a journal article shared a title, candidates were ranked to prefer the version of record. Entries initially accepted against a preprint were re-resolved to the published version.

Channel 2 — Find Case Law. English authority was checked against a citation index of 75,837 judgments spanning 2001 to 2026 across thirteen court identifiers, harvested under a computational-analysis permission granted by The National Archives.

Channel 3 — manual. The judgments of the International Court of Justice, the reports and awards of WTO panels and of arbitrators under DSU art 25, the awards of investment tribunals, and the instruments of the European Union are outside the reach of both automated channels. Each was checked against the official text, provision by provision, and the evidence recorded. This is the channel that carries the paper’s load, and it is the one with no register behind it.

Continuous validation. Every citation is re-checked against the registers on each revision for title, year, volume, journal and first page. The build refuses to produce a manuscript in which a footnote is referenced but undefined, defined but unreferenced, defined in two parts at once, pointed at by a cross-reference that resolves to nothing, or carrying a marker used more than once. It also refuses one in which a “Part X §N” reference names the wrong Part, or in which this Appendix discusses an authority the reader cannot find in the paper — both checks added after the failures at classes five and six below.

What the apparatus cannot do, established by six classes of failure

The apparatus answers one question: does this authority exist, and is it what it is called. Everything below survived it. They are set out because they measure the limit of mechanical verification more precisely than any description of the method could, and because a paper making this argument that reported only its successes would be worth less than one that reported none of them.

Class one — the source was read, twice, and misread both times in opposite directions. This is the most instructive failure in the paper and it concerns the instrument closest to its subject. Regulation (EU) 2023/2675 was first described as dispensing with attribution and as deploying retorsion. Examination established that art 3(1) imports customary attribution expressly, so the first half was wrong. The correction then over-shot: relying on art 8(4), the paper described the response measures as countermeasures and recast the whole instrument as a counter-example to Part IV, and wrote a qualification into §17.2 on the strength of it. That was also wrong. Article 8(1) contains no wrongfulness predicate; art 8(4) is a conditional permission to escalate rather than a characterisation; and every measure in Annex I is drafted as one “which may amount, as necessary” to non-performance. The instrument is channel-agnostic, which is what §17 predicts, and both misreadings had to be removed before that could be seen.

Nothing about this failure was a resource problem. The Regulation is short, free, in the author’s first language, and was open on the screen each time. What was expensive was not obtaining the text but attending to the difference between a permission and a characterisation across four provisions and a recital. That is the verification cost of this paper’s own argument, incurred twice, on the one instrument it could least afford to get wrong.

Class two — correctly cited, wrongly used. Two load-bearing authorities were cited for propositions they do not support. *Manchester Building Society* was used to *generate* a duty; it is a scope-of-duty principle limiting recoverable loss once a duty exists, and it has never generated one. *Steel v NRAM* was cited for the

ratio of *Playboy Club* — that the representee was not the person to whom responsibility had been assumed — when in *NRAM* the representee *was* the addressee and the duty failed for unreasonable reliance instead. Both citations were perfectly formed and pointed at the works named, so no register could have caught either.

A third instance is more serious because it concerns the authority on which this paper’s own proposal rests. *Corfu Channel* was quoted for the passage permitting “a more liberal recourse to inferences of fact and circumstantial evidence” where the respondent holds the material, and the same passage’s express holding — that exclusive territorial control “neither involves prima facie responsibility nor shifts the burden of proof” — was omitted. The omitted sentence is the one that defeats a presumption and permits an inference, and the mechanism at §22.2 was reframed from the first to the second once it had been read.

Class three — the law moved and the manuscript did not. The worked case at §6B was drafted while the dispute was pending and reported it as such. The European Union suspended the panel twice, each time within days of the twelve-month lapse under DSU art 12.12, and withdrew the complaint on 1 December 2025 with no report and no finding. An earlier draft carried an italicised note advising the reader to check the procedural position at the date of use. That is not an acceptable substitute for checking it, in a paper whose method note stakes its credibility on having checked everything — and the completed history turned out to support the argument more strongly than the pending one had.

Class four — the statistic that confirms whatever it is pointed at. Section 19.1 was drafted three times. The first version reported that twenty-one of the twenty-three states China depends on are dependent on China in return, and called it confirmation of a mutual-hostage structure. Against the base rate — 142 of 152 states holding any severe dependency are dependent on China — coincidence predicts 21.5. The second version excluded China and reported reciprocity at 4.4 times the rate expected under a uniform random graph. Against a degree-preserving null, which respects the hub structure the data plainly has, the observed 38 reciprocated dependencies sit against 37.5 expected. Both positive findings were artefacts of a null model that did not represent the data, and each was discovered only because the previous one had been. The published result is negative.

No apparatus catches this. The arithmetic was correct at every stage; what was wrong was the comparison, and a register has nothing to say about a comparison.

Class five — the check passed and the artefact was wrong. Two footnote markers were each used twice in the assembled manuscript. That is valid markdown, and every check was satisfied: each note was defined, none orphaned, no number defined twice, every cross-reference resolved. But the renderer emits a *fresh* note for each occurrence of a marker, so the rendered document carried two more notes than the manuscript, and every cross-reference above the first duplicate landed one or two notes short — naming one authority and pointing at another. The manuscript was correct throughout, which is why nothing caught it: the apparatus was validating the source and the reader was holding the output.

Class six — the reference that resolved, to the wrong thing. This material was written as one paper and later divided into two. The division renumbered the Parts and the sections, and sixteen references of the form “Part IV §17.2” survived pointing at sections which existed but were not the ones meant. Every one resolved, so the citation checker passed them, and the footnote checker does not look at section references at all. They were found by a check written afterwards, which validates the *pair* rather than the target and now refuses the build. The same operation carried this Appendix over from the companion paper unchanged, so that for one revision it discussed eighteen authorities a reader of this paper could not find anywhere in it — a defect now caught by requiring every name here to appear in the body.

What this establishes. A register can answer whether an authority exists and whether it bears the identifier given. It cannot answer whether the authority supports the proposition it is cited for, whether the passage quoted is the operative one, whether the procedural position has moved since the draft, whether a statistical comparison means anything, whether the document the reader holds says what the document the author checked says, or whether a cross-reference points where it claims to. Those are the questions on which every failure recorded here turned, and not one of them is a register question. That is the irreducible verification cost this paper is about, measured on the paper itself. The apparatus reduced the cost of checking

by a large factor; it did not remove it, and what it left behind was not the easy part.

The datasets, the derivation code for Part III, the registers and the verification tooling are published in full.