

The Two Margins of Economic Coercion: Leverage and Partial Compliance under Imposed Sanctions

Abstract

When a state imposes trade sanctions, does economic leverage buy compliance? Outcomes are coded objective by objective and then pooled into success or failure, yet targets concede nothing, part, or all. I study that choice in 193 imposed sanctions by single-state senders, 1949 to 2019, matched to pre-episode trade exposure. Exposure shows no detectable association with whether targets concede. It predicts what they concede. Across the interquartile range of the target's stake relative to the sender's, partial compliance rises about thirteen points and full compliance falls about nine, while outright refusal stays flat. The relative structure of exposure carries that information where the bilateral levels do not. A model that separates the decision to accommodate at all from the decision of how far to go accounts for the composition of compliance. A threat-stage selection account of the same pattern finds no support in data graded at both stages.

JEL Classification: F13, F51, C35, D74

Keywords: Economic coercion, sanctions, partial compliance, selection, trade dependence

1 Introduction

When a state imposes trade sanctions to extract a policy change, does economic leverage buy compliance? The literature answers with a binary: sanctions succeed or fail. A sanctioned government can concede nothing, concede part of what is demanded, or concede everything. This paper asks on which of those margins leverage operates. In a sample of imposed trade sanctions spanning seven decades, bilateral trade exposure has no statistically detectable association with whether targets concede at all, and it predicts the composition of what they concede. Where the sender’s relative advantage is high, imposed sanctions end in partial compliance. Where it is low, full compliance is more common.

Three 2025 episodes illustrate the margins. None enters the main sample. In January 2025 the United States threatened steep tariffs on Colombia over deportation flights. Colombia complied in full within a day and the tariffs never took effect. Mexico and Canada faced the same instrument weeks later, the tariffs were imposed in March 2025, and both governments delivered concessions in installments over the following year without ever settling the demand in full. The one full complier never entered the imposed pool, and the imposed cases produced partial compliance.

The evidence comes from release 4 of the Global Sanctions Data Base, GSDB ([Felbermayr et al., 2020](#); [Yalcin et al., 2025](#)). The main sample is 193 resolved episodes with a trade-sanction instrument, a single-state sender, and a single-state target, beginning between 1949 and 2019, with outcomes coded none, partial, or full from GSDB’s per-objective success grades. Leverage is measured from pre-episode trade data in several theory-consistent ways: the target’s import dependence on the sender, which is the first-order channel in [Liu and Yang \(2025\)](#), a gravity welfare exposure measure equal to bilateral trade over target GDP scaled by the trade elasticity, outside-option and concentration measures, and two asymmetry ratios that compare the target’s exposure to the sender’s. All measures are three-year pre-episode averages.

No bilateral leverage measure has a statistically detectable association with whether the target concedes anything. Across all four core measures the first-stage estimates are small relative to their sampling variation. The data do not separate the concession rate of high-leverage senders from that of low-leverage senders among imposed trade sanctions. The one measure that moves that margin at all is dependence with third-country substitution priced in, and it moves it weakly.

Higher export-share asymmetry comes with a higher probability of partial compliance and a lower probability of full compliance, while the probability of no concession stays flat. Moving from the lower to the upper quartile of the asymmetry ratio, the probability that a

target concedes part of what is demanded rises from 27 to 40 percent and the probability that it concedes everything falls from 42 to 33 percent. Most of that 13-point rise comes out of full compliance, which gives up 9 points, and the remaining 4 come off outright refusal, a movement the data cannot distinguish from zero. The rise in partial compliance reaches conventional levels unadjusted and the decline in full compliance does not. In episodes with a single stated objective, where the two grading rules coincide, the estimates are larger on both margins and reach conventional levels on both: the same quartile move raises the probability of partial compliance by 22 percentage points and lowers the probability of full compliance by 16.

The result is specific to relative exposure. The target’s import dependence on the sender, the measure a first-order theory of trade pressure prefers, shows no statistically detectable association on either margin in the pooled sample, and the same holds for the other bilateral level measures: sender dependence, export dependence, bilateral concentration, outside options, and partner concentration. That null is a pooled-sample statement about bilateral shares. One subsample cell puts a level measure on the full margin. The asymmetry ratio is size-free by construction. It treats a target that sends 9 percent of its exports to a sender who returns 1 percent the same as a pair at 0.9 percent against 0.1 percent. The negative margin therefore comes from the relative structure of the bilateral relationship. Measuring dependence with substitution through the trade network priced in returns the rotation at the same size with intervals about half as wide, and makes the sender’s own exposure informative: a standard deviation of it lowers the probability of partial compliance by 15 percentage points, where the sender’s bilateral export share shows nothing. Entered together in one model, the data reject the joint null on the relative direction and do not reject it on the scale direction, and a contrast of the same size pointing elsewhere finds nothing. The four stakes entered without restriction do not reject either, so it is the model’s restriction that locates the signal. What these facts leave open is why economic leverage should govern how far a target concedes without affecting whether it concedes at all.

One answer was available from the start, and this paper tests it instead of adopting it. The sanctions literature holds that imposed sanctions are a selected sample. Threats resolve the easy cases before imposition. [Drezner \(2003\)](#) argued that coercion succeeds at the threat stage, where it is never observed as a sanction. [Lacy and Niou \(2004\)](#) formalized the logic. [Whang et al. \(2013\)](#) estimated it structurally and found that threats work through expected cost, the coercive channel. [Miller \(2014\)](#) documented it for nonproliferation. In the TIES data, which unlike GSDB record threats, 45 percent of coded coercion successes occur at the threat stage, before any sanction is imposed ([Morgan et al., 2014](#)). If leverage makes threats effective, then leverage clears prospective full compliers out of the imposed pool, and

the survivors sanctioned at high leverage are the resolute. A negative association between leverage and full compliance among imposed cases can be produced by this mechanism, even when leverage raises compliance everywhere.

Imposed-only data cannot separate that reading from the other one. Under strategic partial compliance, a target facing relative pressure concedes part of the demand and withholds the rest. Under threat-stage selection, prospective full compliers settle before imposition and never enter the imposed pool. The benchmark of Section F of the supplementary appendix generates both margins and produces the same observable pattern under either, so nothing in GSDB tells them apart. Grades on both sides of the imposition threshold can, and they do not support it. In TIES, leverage leaves threat-stage full compliance flat and raises it after imposition, so the imposed pool is not drained of full compliers.

These are associations among sanctions that reached imposition, and nothing in the design isolates exogenous variation in leverage.

What survives that test divides the decision in two. Whether a target accommodates at all runs through political admissibility: a government that cannot be seen to yield under duress refuses whatever is asked and whatever the economic cost, and that decision does not trade off against pressure. It is not the only route to no concession, since a politically willing target whose concession cost is high enough also concedes nothing, and reproducing the flat extensive margin requires that second route to be small, which the data bound and do not establish. Conditional on entering, economic leverage governs the extent of accommodation, and it does so through two instruments the sender chooses, how hard to squeeze and how much to ask. Relative exposure emerges from that structure instead of being imposed on it, because pressure is produced with the target's stake and paid for with the sender's. That political component is bounded from below and not identified in these data. The composition margin does not require identifying it, since the ratio of full to any compliance cancels it.

The effectiveness literature descends from [Hufbauer et al. \(2007\)](#), whose case codings put the overall success rate near one third, and from the skeptical rejoinder of [Pape \(1997\)](#), and its large-sample form runs on the coding projects that followed ([Bapat et al., 2013](#); [Morgan et al., 2014](#); [Weber and Schneider, 2022](#)). Within it the trade-linkage results are mixed. [McLean and Whang \(2010\)](#) find that target dependence on the sender raises success, with the largest effect when the sender is the target's largest trading partner, [Peterson \(2020\)](#) reports that aggregate dependence performs poorly once network position is accounted for, and [Akoto et al. \(2020\)](#) show that the composition of trade moves acquiescence to threats where its volume does not. Bilateral exposure delivers unstable predictions, which the extensive-margin null here confirms and the full-compliance margin qualifies. The selection branch of that literature is wider: [Whang \(2010\)](#) estimates initiation and outcome jointly, [Walentek](#)

et al. (2021) simulate threat-stage dynamics, and Cilizoglu and Bapat (2020) model targets that sanctions-proof in anticipation. This paper moves from the binary success margin to the composition of compliance among imposed cases, where the selection channel bears on the full-compliance margin more than on the any-concession margin.

A second literature treats economic coercion as an equilibrium phenomenon in world trade. Clayton et al. (2026) develop a framework in which a hegemon extracts concessions up to a participation constraint that binds exactly, so compliance in their equilibrium is complete by construction, Clayton et al. (2024) study the fragmentation dynamics that coercion sets off, and Clayton et al. (2025) measure geoeconomic pressure. Liu and Yang (2025) model international power as arising from asymmetric gains from trade and place the import channel first in order, Kleinman et al. (2024) show in a quantitative spatial model that bilateral trade shares are poor sufficient statistics for dependence once network substitution is accounted for, and Itskhoki and Mukhin (2022) and Bianchi and Sosa-Padilla (2025) price the macroeconomic side of sanctions. These models are tested, when tested at all, on engagement, pressure, or trade flows, and none of them delivers or confronts a partial-compliance margin. A third and smaller literature does take partial outcomes as its object: Mangini (2024) models the trilemma among concession size, enforcement, and duration, with convex concession costs of the kind used below, and Mangini (2026) studies trade conditionality in the Generalized System of Preferences, where compliance is graded. The leverage question is older (Hirschman, 1945; Baldwin, 1985; Drezner, 1999; Farrell and Newman, 2019).

Section 2 describes the sample, the outcome coding, and the exposure measures. Section 3 establishes the three facts and takes the selection account to threat-stage data. Section 4 separates political admissibility from economic bargaining and states what the data do and do not identify. Section 5 concludes.

2 Data, Outcomes, and Exposure

2.1 Sample and outcome construction

The episode data are the Global Sanctions Data Base, release 4 (Felbermayr et al., 2020; Yalcin et al., 2025), which records 1,547 sanction cases with senders, targets, instruments, objectives, and outcome grades. I keep cases that meet three conditions: a trade-sanction instrument is present, the sender is a single state mappable to an ISO3 code, and the target is likewise a single state. Cases sent by organizations or coalitions, the European Union, the United Nations, and multi-country sender strings, are excluded, which is what unilateral

means throughout this paper. Cases whose only stated objective is “other” are dropped. The filter yields 379 episodes. All episodes beginning in 2020 or later are then removed, since they are mechanically unresolved at coding time, leaving a pre-2020 universe of 266 episodes beginning between 1949 and 2019. Section A of the supplementary appendix documents the build, the censoring rule that takes the universe to the main sample, and the trade-data coverage that takes the main sample to the estimation samples.

GSDB codes success as a comma-separated string with one grade per objective, drawn from full success, partial success, negotiated settlement, failure, and ongoing. Because the grades are recorded objective by objective, a multi-objective episode can carry several grades at once. The outcome variable reads “concede everything” strictly: an episode is full only when every coded objective is a full success, none when no objective drew a concession or a settlement, and partial in every other case. Negotiated settlements are partial, following GSDB’s own semantics, in which a settlement extracts some but not all of the stated demand. The partial category bundles concessions that partially met a single demand with episodes in which some objectives drew concessions while others failed outright, and both are partial extractions of the posted demand set. In the main sample the coding gives 60 episodes with no concession, 66 partial, and 67 full (Table 1, Panel C), so any concession occurs in 133 of 193 episodes and the identity any equals partial plus full holds by construction.

Alternative codings serve as robustness checks. Section B of the supplementary appendix reports the alternative dominance-ladder coding, which grades a multi-objective episode by its best objective, along with its construction and the five-category distribution of the pre-2020 universe in Panel B of Table 1. The other two place negotiated settlements in a category of their own in a four-category outcome and drop them from the sample. Graded imposed-stage outcomes are standard in the coding manuals: TIES distinguishes partial from complete acquiescence after imposition (Morgan et al., 2014), EUSANCT carries graded outcomes (Weber and Schneider, 2022), and GSDB grades success objective by objective. Empirical practice has pooled the grades into a success dummy at the first step, usually on sample-size grounds.

The 73 ongoing episodes carry a sentinel end year in the raw data, and treating them as failures would code censoring as outcome. The main sample excludes them, leaving 193 resolved episodes. The exclusion is a censoring rule, and its complement is checked: the robustness results re-estimate everything with ongoing episodes included as failures to date. The same logic already removed the post-2020 block at the universe step, since an episode that has not ended cannot have its outcome graded. Durations for ongoing episodes are lower bounds throughout, and no analysis in this paper uses duration as an outcome. The universe is the 266 pre-2020 episodes, the main sample is the 193 that survive the censoring

rule, and the estimation sample for a given measure is the subset of the main sample on which that measure is non-missing.

The baseline specification carries four covariates, and the robustness rows vary the control set as described where they appear. A demand-type proxy k grades the political cost of the sender’s objective on an ordinal scale, from policy and terrorism-related demands at the bottom, through war-prevention and human-rights or democracy demands, to territorial demands and destabilization at the top, taking the maximum over an episode’s objectives. Its five levels appear in the universe with frequencies 81, 47, 76, 30, and 32 (Table 1, Panel D). An indicator marks the United States as sender, which covers 109 of the 266 universe episodes. A Cold War indicator marks episodes beginning in or before 1991, which covers 119. The fourth control is the log of one plus the number of prior GSDB episodes against the same target, a sanction-history control with a scripted construction rule from the full data base. One alternative control set uses a fourteen-country target list, covering 98 universe episodes, states such as Cuba, Iran, and North Korea that appear repeatedly as sanction targets. That list has no documented construction rule, so it serves as a control and as a sample split, never as a measured regime characteristic. Episodes in which trade restrictions are the only coercive instrument, with no financial, arms, military, travel, or other measure attached, number 108 in the universe and 92 in the main sample, and form the instrument-purity subsample.

Table 1 assembles the construction, the taxonomy, and the covariate counts in one place.

GSDB records imposed sanctions only. There is no threat stage in the data, so every estimate conditions on the coercion attempt having escalated to imposition. Imposition is itself an equilibrium outcome, since senders choose targets, demands, and whether to follow through. Section F of the supplementary appendix shows what the imposed-only conditioning does to observable outcome shares under an explicit selection mechanism.

The three categories record two decisions, and the paper keeps them apart throughout. The first is whether the target concedes anything, which contrasts no concession against any concession. The second is how far it goes conditional on conceding, which contrasts partial against full among the episodes that conceded something. Section 4 gives the second a purely economic determinant and the first a mixed one, political admissibility together with an economic route that has to be small for the extensive margin to be flat, and the empirical sections are organised so the two are never read off a single coefficient.

Table 1: Sample construction and outcome taxonomy

	Episodes
<i>Panel A: construction</i>	
GSDB V4 cases	1,547
Trade instrument, unilateral single-state sender and target, non-other objective	379
Began 2020 or later (mechanically unresolved, dropped)	113
Pre-2020 universe	266
Ongoing at coding (end-year sentinel, censored, excluded)	73
Main sample	193
Any-concession episodes (stage-two denominator)	133
Trade-only subsample (no non-trade instrument), universe	108
Trade-only subsample, main sample	92
<i>Panel B: outcomes, five-category dominance ladder (pre-2020 universe)</i>	
None	60
Negotiated settlement	22
Partial	34
Full	77
Ongoing (censored)	73
<i>Panel C: outcomes, baseline three-category coding (main sample)</i>	
None	60
Partial (includes negotiated settlement)	66
Full (every coded objective a full success)	67
<i>Panel D: covariates (pre-2020 universe)</i>	
United States sender	109
Cold War era (begin ≤ 1991)	119
Fourteen-country target list	98
Demand-type proxy $k \in \{1, 2, 3, 5, 9\}$	81/47/76/30/32

Notes: Episode counts computed in a single scripted pass from the raw GSDB V4 case file (Felbermayr et al., 2020; release 4, Yalcin et al., 2025), with every count asserted against an externally maintained case ledger at build time. Episodes begin between 1949 and 2019. Panel C applies the baseline coding, in which an episode is full only when every coded objective is a full success and negotiated settlements are partial, following the GSDB semantics that a settlement extracts some but not all of the stated demand. Any concession equals partial plus full by construction ($66 + 67 = 133$). The alternative dominance-ladder coding of Appendix ??, which grades an episode by its best objective grade, gives main-sample counts of 60 none, 56 partial, and 77 full (Table ??) and generates the five-category distribution of Panel B. Ongoing episodes carry a sentinel end year and are censored, so the main sample excludes them.

2.2 Trade exposure measures

All measures derive from a composite bilateral flow built from CEPII Gravity V202211. For each origin-destination-year, the flow is the first non-missing of five sources in a fixed priority order: BACI reconciled flows, destination-reported Comtrade, origin-reported Comtrade, destination-reported IMF Direction of Trade, and origin-reported IMF Direction of Trade. Reconciled data are preferred to importer reports and importer reports to exporter reports, and the IMF series ranks last but is indispensable before the 1960s, when it is the only source

available, so early episodes keep their measurement windows. Country totals of imports and exports aggregate the composite flow, GDP is the CEPII current-dollar series, and a missing bilateral flow within a needed dyad-year is set to zero only when both endpoints report positive totals that year, so genuine non-reporting is never coerced to zero trade. Each measure is computed for an episode beginning in year b over the window of the three preceding years, as the mean of the annual ratios computable within the window. Pre-episode windows limit the mechanical feedback from the sanction itself to the trade data, though not the deeper simultaneity of sanctions and trade relationships. Section C of the supplementary appendix varies the measurement window, from the single year before an episode begins to a five-year window and a buffered window that omits the two years closest to onset.

The theory-preferred measure is the target’s import dependence on the sender,

$$m_1 = \frac{F(s \rightarrow t)}{M(t)}, \quad (1)$$

the share of the target’s imports supplied by the sender. In [Liu and Yang \(2025\)](#) the gains a country derives from purchases it cannot easily replace are the first-order source of power, which makes m_1 the natural regressor. Its mirror images are the sender’s import dependence on the target m_2 , the target’s export dependence on the sender m_3 , and the sender’s export dependence on the target m_4 . Bilateral concentration m_5 measures the sender’s share of the target’s total trade, outside options $m_6 = 1 - m_5$ measure the share with third countries, and a Herfindahl of the target’s partner shares proxies the thinness of its outside market. A gravity welfare exposure measure follows the first-order approximation of welfare loss from severed bilateral trade,

$$\delta = \frac{F(s \rightarrow t) + F(t \rightarrow s)}{\text{GDP}(t) \cdot \theta}, \quad (2)$$

with trade elasticity $\theta = 4$ at baseline. Because θ enters as a pure rescaling of the regressor, its calibrated value moves coefficients mechanically and leaves inference unchanged. Finally, two asymmetry ratios compare exposures across the dyad:

$$A = \frac{m_3}{m_3 + m_4}, \quad A_M = \frac{m_1}{m_1 + m_2}. \quad (3)$$

The export-share ratio A compares the target’s export dependence to the sender’s. The import-share ratio A_M applies the same construction to the import channel. Both run from zero to one, with values near one indicating that the target’s relative exposure dwarfs the sender’s. Section C of the supplementary appendix reports the baseline specification with the log odds of each ratio in place of the bounded share.

Table 2 lists definitions, coverage, and distributions in the main sample.

Table 2: Leverage measures: definitions, coverage, and distribution

Measure	Definition	N	Mean	SD	p25	p50	p75
m_1 target import dependence	$F(s \rightarrow t)/M(t)$: sender share of target imports	181	0.104	0.158	0.006	0.042	0.129
m_2 sender import dependence	$F(t \rightarrow s)/M(s)$: target share of sender imports	180	0.018	0.080	0.000	0.001	0.013
m_3 target export dependence	$F(t \rightarrow s)/X(t)$: sender share of target exports	180	0.104	0.158	0.001	0.037	0.147
m_4 sender export dependence	$F(s \rightarrow t)/X(s)$: target share of sender exports	181	0.019	0.083	0.000	0.002	0.011
m_5 bilateral concentration	$[F(s \rightarrow t) + F(t \rightarrow s)]/[M(t) + X(t)]$: bilateral concentration	180	0.109	0.156	0.005	0.039	0.163
m_6 outside options	$1 - m_5$: target trade with third countries	180	0.891	0.156	0.837	0.961	0.995
m_6^H partner-share HHI	Herfindahl of target partner shares	180	0.140	0.115	0.073	0.101	0.166
δ gravity exposure ($\theta = 4$)	$[F(s \rightarrow t) + F(t \rightarrow s)]/[\text{GDP}(t)^\theta]$: welfare exposure	163	0.014	0.021	0.000	0.004	0.017
A export-share asymmetry	$m_3/(m_3 + m_4)$: export-share asymmetry ratio	166	0.736	0.331	0.541	0.911	0.985
A_M import-share asymmetry	$m_1/(m_1 + m_2)$: import-share asymmetry ratio	166	0.851	0.251	0.826	0.967	0.996

Notes: $F(o \rightarrow d)$ is the composite bilateral flow from CEPII Gravity V202211 (BACI, then Comtrade destination- and origin-reported, then IMF DOTS, in that fixed priority order). $M(c)$ and $X(c)$ are total imports and exports and GDP is the CEPII current-dollar series. Each measure is the mean of annual ratios over the three years before the episode begins. N is the count of non-missing values in the main sample of 193. Summary statistics are computed on the main sample.

A ratio is size-free: A treats a target sending nine percent of its exports to a sender who returns one percent identically to a pair at 0.9 percent against 0.1 percent, though the economic stakes differ by an order of magnitude. A ratio cannot proxy the level of damage a sanction can inflict, and it is not a measure of network position: Kleinman et al. (2024) show that bilateral shares, however arranged, are poor substitutes for network-based dependence measures, so A is read here as a dyadic index, never a network statistic. Section 3.3 rebuilds the same ratio from a general-equilibrium exposure measure that carries third-country substitution inside it. Section 4.3 derives the ratio from the coercion technology of the model: under the export-market-access reading, A is the share transform of the target-to-sender stake ratio. In the main sample A has mean 0.736 and median 0.911 (Table 2), so the typical imposed trade sanction already features a sender with a large relative advantage. Section C of the supplementary appendix plots the ratio against the level for the main sample. It also re-estimates the baseline specification on samples that drop the episodes with the smallest total stake $m_3 + m_4$ (Table ??). The two stakes that form A enter jointly there in place of their ratio (Table ??).

2.3 Specification and inference

The baseline specification estimates a multinomial logit for the three-category outcome (none, partial, full). Each leverage measure enters separately, with demand-type, US-sender, Cold War, and prior-sanction controls. Standard errors are clustered by target. An episode is graded full only when every stated objective is met in full. The strict outcome coding, the

sanction-history control, and target-clustered inference define the baseline specification for the whole paper.

Formally, the outcome model is

$$\Pr(y_i = j) = \frac{\exp(\alpha_j + \beta_j x_i + \mathbf{z}_i' \gamma_j)}{\sum_l \exp(\alpha_l + \beta_l x_i + \mathbf{z}_i' \gamma_l)}, \quad j \in \{\text{none, partial, full}\}, \quad (4)$$

with x_i the leverage measure, $\mathbf{z}_i$ the controls, and the none category as base. The reported average marginal effects aggregate $\partial \Pr(y_i = j) / \partial x_i$ over the estimation sample, so the three effects for a measure sum to zero by construction. The sequential decomposition estimates a logit for any concession on the main sample and a logit for full compliance on the concession subsample, which reproduces the two-stage structure the binary literature works with.

Section B of the supplementary appendix reports the same estimates under the dominance-ladder coding, with the target-list control and robust standard errors. Estimation samples per measure reflect trade-data coverage (Table 2). Coefficients appear with standard errors and exact p -values, without significance stars, and every small cell is flagged.

3 The Empirical Facts

3.1 The extensive and intensive margins

Figure 1 puts the two decisions side by side.




Figure 1: The two margins against relative exposure. Panel (a) plots the probability that a target makes any concession. Panel (b) gives the probability of full compliance among the episodes that conceded something. Points are means within eight equal-count bins of A with 95 percent intervals from a cluster bootstrap over targets. The line is an *uncontrolled* logit in A alone, with a band from the same resamples: it describes the raw relationship and is not a picture of the controlled estimates of Table 3, whose stage-one coefficient is positive and whose stage-two coefficient is -1.418 . Binning on A instead of $\ln \lambda$ retains all 166 episodes, since the log-odds transform is undefined at $A = 0$ for three of them and at $A = 1$ for one. The fitted slope is -0.024 on the left and -1.439 on the right.

The two decisions of Section 2.1 are estimated as they are defined, in sequential form: a first-stage logit for whether the target concedes anything and a second-stage logit for full compliance among the episodes that conceded something. This is the representation the model of Section 4 speaks to, and it is not a robustness variant of the three-category estimator. Table 3 reports both stages for the four core measures under the baseline specification, with the single-objective subsample as a fifth row.

Table 3: Sequential decomposition under the baseline specification: any concession, then full given any

Measure	Stage 1: any concession			Stage 2: full any		
	Coef. (SE)	p	N	Coef. (SE)	p	N
m_1 target import dependence	1.306 (1.492)	0.382	181	0.969 (1.547)	0.531	126
δ gravity exposure ($\theta = 4$)	-3.629 (11.850)	0.759	163	-4.633 (12.588)	0.713	115
A export-share asymmetry	0.400 (0.751)	0.595	166	-1.418 (0.761)	0.063	118
A_M import-share asymmetry	0.580 (0.706)	0.411	166	-2.259 (0.945)	0.017	118
A , single-objective episodes	0.585 (0.831)	0.481	134	-3.156 (1.138)	0.006	97

Notes: Logit coefficients on the leverage measure under the baseline specification: strict outcome coding, the same controls as Table 4 with the sanction-history control, and standard errors clustered by target. Stage 1 is any concession versus none on the main sample. Stage 2 is full versus partial on the any-concession subsample. The strict recode moves mass only from full to partial, so the stage-1 outcome equals its dominance-coding counterpart by construction. Stage-2 cluster counts: 59 targets (m_1), 53 (δ), 52 (A and A_M), 45 (single-objective row). Negotiated settlements count as partial concessions and ongoing episodes leave the denominator.

Every stage-one coefficient in the table has a p -value above 0.3, for ratios and levels alike. The threat-stage selection literature expects that null: by the time a sanction is imposed, the cases leverage could have resolved cheaply have already settled off sample. Threat-stage data do not bear that expectation out. With bilateral exposure the extensive margin is undetectable throughout, and the one stage-one estimate in the paper that departs from that pattern comes from a different measurement system and is reported in Section 3.3.

The second-stage estimates are negative for both asymmetry ratios. The A coefficient among concession episodes is -1.418 (0.761, $p = 0.063$, $N = 118$, 52 target clusters), the A_M coefficient is -2.259 (0.945, $p = 0.017$), and on the single-objective subsample the A coefficient is -3.156 (1.138, $p = 0.006$, $N = 97$). The stage-two estimates agree in sign and pattern with the three-category estimates reported next.

The same two margins can be read jointly, from a single three-category model that imposes no ordering on the outcomes. It gives in levels what the sequential coefficients give in odds, and the two agree.

Table 4 reports average marginal effects from the baseline multinomial logit for the four core leverage measures, with a fifth row for the single-objective subsample.

Table 4: Leverage and the composition of outcomes: baseline multinomial logit average marginal effects

Measure	$P(\text{none})$		$P(\text{partial})$		$P(\text{full})$		N
	AME (SE)	p	AME (SE)	p	AME (SE)	p	
m_1 target import dependence	-0.246 (0.267)	0.357	0.006 (0.297)	0.983	0.240 (0.255)	0.347	181
δ gravity exposure ($\theta = 4$)	0.655 (2.059)	0.751	0.425 (2.126)	0.842	-1.079 (1.964)	0.583	163
A export-share asymmetry	-0.087 (0.139)	0.530	0.279 (0.136)	0.040	-0.192 (0.125)	0.125	166
A_M import-share asymmetry	-0.153 (0.123)	0.213	0.440 (0.153)	0.004	-0.287 (0.161)	0.074	166
A , single-objective episodes	-0.121 (0.138)	0.380	0.453 (0.126)	0.000	-0.331 (0.155)	0.033	134

Notes: Average marginal effects of each leverage measure on the probability of each outcome under the baseline specification: the strict outcome coding, in which full compliance requires every coded objective to be a full success, with demand-type indicators, a United States sender indicator, a Cold War indicator, and the log prior-episode sanction-history control, estimated one measure at a time on the main sample with standard errors clustered by target and the clustered variance propagated to the margins. Cluster counts: 74 targets (m_1), 65 (δ), 68 (A and A_M), 60 (single-objective row). The single-objective row restricts to episodes with one coded objective, where the strict and dominance codings coincide by construction. All estimates are associations among imposed trade sanctions. Table ?? reports the same quantities under the alternative dominance-ladder coding.

Higher export-share asymmetry is associated with no movement in the probability of outright failure (-0.087 , $p = 0.530$), a higher probability of partial compliance ($+0.279$, $p = 0.040$), and a lower probability of full compliance that does not reach conventional significance on its own (-0.192 , $p = 0.125$). The two success-margin effects are close in size and opposite in sign, which is the rotation. The import-share ratio A_M reproduces the shape with more precision: $+0.440$ ($p = 0.004$) on partial and -0.287 ($p = 0.074$) on full. The single-objective estimates, from episodes where the strict and dominance codings coincide by construction, are larger on both margins and reach conventional levels on both: $+0.453$ ($p < 0.001$) and -0.331 ($p = 0.033$). The standard error on the full margin rises from 0.125 to 0.155 on the smaller sample, and the estimate clears conventional levels because the point estimate moves further than the standard error.

Figure ?? in Section G of the supplementary appendix traces the same estimates as predicted probabilities across the observed support of A .

Table 5 puts the magnitudes on observed support, comparing predicted probabilities at each ratio's own 25th and 75th percentiles.

Table 5: Finite probability contrasts at the 25th and 75th percentiles, baseline specification

Measure	Measure		$P(\text{full})$ (SE)		$\Delta P(\text{full})$		$P(\text{partial})$ (SE)		$\Delta P(\text{partial})$		N
	p25	p75	at p25	at p75	(SE)	p	at p25	at p75	(SE)	p	
A export-share asymmetry	0.541	0.985	0.417 (0.050)	0.329 (0.053)	-0.088 (0.057)	0.123	0.273 (0.042)	0.401 (0.056)	0.128 (0.063)	0.043	166
A_M import-share asymmetry	0.826	0.996	0.393 (0.042)	0.342 (0.051)	-0.051 (0.028)	0.067	0.308 (0.034)	0.388 (0.051)	0.079 (0.029)	0.006	166

Notes: Predicted probabilities from the baseline multinomial logit of Table 4 for the two asymmetry ratios, evaluated at the estimation-sample 25th and 75th percentiles of each measure with other covariates at observed values, and the finite contrasts between them, with standard errors from the target-clustered variance. Contrasts stay on the observed support of each measure. On the single-objective subsample the same contrasts for A (percentiles 0.521 and 0.985) are -0.162 (SE 0.076, $p = 0.034$) on $P(\text{full})$ and $+0.219$ (0.063, $p = 0.000$) on $P(\text{partial})$, $N = 134$.

For A , the interquartile move (from 0.541 to 0.985) raises the predicted probability of partial compliance from 0.273 to 0.401, a contrast of $+0.128$ (standard error 0.063, $p = 0.043$), while full compliance falls from 0.417 to 0.329, a contrast of -0.088 (0.057, $p = 0.123$). Probability moves from full to partial compliance in the same direction, the two movements leaving a residual of $+0.040$ on any concession, which is itself not distinguishable from zero. The A_M contrasts are smaller and signed the same way, and they are more precise ($+0.079$, $p = 0.006$ on partial, -0.051 , $p = 0.067$ on full), in part because A_M is compressed near one in this sample, so its interquartile range is short. On the single-objective subsample the A contrasts reach $+0.219$ ($p < 0.001$) and -0.162 ($p = 0.034$).

Cutting the estimation sample into terciles of A gives the observed shares directly. In the bottom tercile they are 0.268 none, 0.214 partial, and 0.518 full on 56 episodes. In the middle tercile they are 0.345, 0.418, and 0.236 on 55 episodes, and in the top tercile 0.255, 0.364, and 0.382 on 55. Full compliance falls between the bottom and the middle tercile and recovers partially in the top tercile, so the rotation appears through the interior of the support of A and is not generated by the most asymmetric episodes.

Neither the tercile indicators entered against the linear term ($p = 0.429$) nor a three-knot restricted cubic spline ($p = 0.395$) rejects linearity in A , so the functional form is not determined at this sample size. And the terciles are of unequal width, because A is compressed near one: the bottom bin covers about three quarters of the observed range of the ratio and the top bin under three percent of it. The same compression puts the spline's upper two knots at 0.911 and 0.998, close enough together that the nonlinear term is poorly identified.

The result is specific to relative exposure. For target import dependence m_1 , the measure a first-order theory of trade pressure prefers, the three average marginal effects are -0.246 , $+0.006$, and $+0.240$, with p -values of 0.357, 0.983, and 0.347, and the gravity exposure measure δ shows no statistically detectable association either (-1.079 , $p = 0.583$ on full).

The measure sweep in Section C of the supplementary appendix extends the null to sender dependence, export dependence, bilateral concentration, outside options, and partner-share concentration. The null on the level measures is a statement about bilateral dependence levels in the pooled main sample. It has one exception in the subsample grid: in the post-Cold-War era the average marginal effect of m_1 on full compliance is -1.913 ($p = 0.022$, $N = 72$), the only bilateral level-measure cell in the subsample grid that moves either margin (Section C of the supplementary appendix). It has a second exception once the measurement system changes: Section 3.3 reports that the sender’s exposure at the level does move the partial margin when substitution through the trade network is priced into the measure. Coefficient scales differ across measures because the regressors are defined on different supports, so Figure ?? in Section G of the supplementary appendix reports the change in predicted probability across each measure’s own interquartile range.

The ratio measures the relative structure of the dyad, the quantity on which behaviour depends when each side weighs its own stake against the other’s. Section 4.3 derives that quantity inside the model, with the two stakes proxied by m_3 and m_4 and $A = m_3/(m_3 + m_4)$ their share transform.

Everything here is an association among imposed trade sanctions. Senders choose targets, demands, instruments, and whether to impose, bilateral trade structure is itself an equilibrium object, and nothing in the design isolates exogenous variation in leverage. The partial-compliance margin reaches the five percent level in all four columns of Table ??, which vary the episode-level grading rule between the strict and dominance-ladder codings and then change the sample to single-objective episodes and the unit of analysis to the objective and under both standard-error schemes reported for the multinomial logit, robust and clustered by target, and a cluster bootstrap on the baseline puts the same margin at $p = 0.075$ against the analytic 0.040. The score bootstrap does not extend to the multinomial logit, so it has never been applied to this margin. The full-compliance margin does not reach conventional significance on its own in the baseline specification ($p = 0.125$) and sits at the ten percent level in the sequential stage. Neither margin survives the multiplicity adjustment of Section 3.3. The strongest evidence on both margins sits in the single-objective episodes, 157 of the 193 in the main sample, of which 134 enter the estimation with A non-missing.

The pattern is consistent with two mechanisms and the data cannot separate them. Under strategic partial compliance, a target facing overwhelming relative exposure concedes part of the demand and withholds the rest. Under threat-stage selection, prospective full compliers leave the imposed pool before imposition, and more so where leverage is high, so the imposed cases observed at high leverage are disproportionately the resolute. Both mechanisms can generate a flat extensive margin and a negative full-compliance margin among imposed

cases, and the discriminating observation, what happened to the cases that never reached imposition, is outside GSDB. The benchmark of Section F of the supplementary appendix nests both mechanisms and shows that they are observationally equivalent in imposed-only data. They are not equivalent once threats are recorded.

3.2 Relative exposure versus exposure levels




Figure 2: The relative direction and the scale direction, stage two. In logs the target's and the sender's stakes split into their difference, which is the relative direction, and their sum, which is the scale direction. The two are an orthogonal reparameterisation of the same four regressors, so neither is favoured by the functional form. Points are stage-two coefficients on each contrast with 95 percent intervals clustered by target. The relative contrasts are -0.263 and -0.365 with standard errors near 0.15 and 0.19 . The scale contrasts are $+0.373$ and -0.499 with standard errors near 0.62 , four times wider.

The model distinguishes relative exposure from absolute exposure, and that distinction is a restriction on one fitted model, and not something to be read off nine separate regressions. Figure 2 puts the two directions side by side. In logs the restriction is exact. Entering the four log stakes freely, the target and sender coefficients split into a ratio direction, their difference, and a scale direction, their sum, and the two directions are an orthogonal reparameterisation of the same four regressors, so neither is favoured by the functional form. The data reject the joint null on the ratio direction, at $\chi^2(4) = 10.21$ with $p = 0.037$ in the multinomial logit and $\chi^2(2) = 6.04$ with $p = 0.049$ at stage two, and they do not reject it on the scale direction, at $p = 0.528$ and $p = 0.444$. The four stakes entered without any restriction do not reject at all ($\chi^2(8) = 12.04$, $p = 0.149$), so it is the model's own restriction that locates

the signal, by spending degrees of freedom on one direction instead of spreading them over four. A contrast of the same size pointing elsewhere finds nothing: imports against exports, in place of target against sender, returns $\chi^2(4) = 1.54$ with $p = 0.820$ and $\chi^2(2) = 0.64$ with $p = 0.725$.

The same contrast appears in levels, where the two asymmetry ratios entered together with the four measures they are built from reject at $\chi^2(4) = 16.57$ ($p = 0.002$) and $\chi^2(2) = 8.94$ ($p = 0.012$) while the level block does not ($p = 0.162$ and $p = 0.414$). That reading holds only for levels entered linearly. With a quadratic in the four levels the level block rejects in the multinomial ($\chi^2(16) = 57.27$, $p < 0.001$) and the ratio block falls to $\chi^2(4) = 8.70$ ($p = 0.069$), while stage two keeps its shape ($p = 0.021$ on the ratios and $p = 0.129$ on the levels). The log parameterisation carries the weight here for that reason: its two directions are symmetric by construction and the levels parameterisation is not.

None of this searches across measures, so the family adjustment above is not its correction and it says nothing about whether any single measure survives that adjustment. Both results stand as answers to different questions. Three limits belong with the reading. A non-rejection on the scale direction is not evidence that scale is absent. The ratio tests are joint over the export and import channels and separate neither from the other. And in the floored log domain, which admits eighteen episodes with a zero leg that the strict domain excludes, the per-channel scale tests move to $p = 0.075$ and $p = 0.081$ while the joint one stays at $p = 0.193$.

3.3 Stability, multiplicity, and measurement

Table 6 collects the design choices that can move the full-compliance estimate: the grading unit, the clustering scheme, the censoring rule, and the multiplicity adjustment across leverage measures.

Table 6: Robustness and inference for the export-share asymmetry ratio A

Specification	Stage 2: full any			AME on $P(\text{full})$		
	Coef. (SE)	p	N	AME (SE)	p	N
Baseline specification	-1.418 (0.761)	0.063	118	-0.192 (0.125)	0.125	166
Single-objective episodes	-3.156 (1.138)	0.006	97	-0.331 (0.155)	0.033	134
Objective-level estimates	-2.085 (1.215)	0.086	137	-0.225 (0.134)	0.093	203
Dyad clustering, target list, dominance ladder	-2.683 (1.056)	0.011	118	-0.355 (0.132)	0.007	166
Ongoing episodes coded as failures	-1.522 (0.811)	0.061	118	-0.148 (0.085)	0.083	236
Multiplicity adjustment, nine distinct measures	-1.418	[0.500]	118	-0.192	[0.998]	166

Notes: Each row reports the stage-two logit coefficient on A for full compliance given any concession and the multinomial-logit average marginal effect of A on $P(\text{full})$. Row 1 is the baseline specification of Table 4: strict outcome coding, demand-type, US-sender, Cold War, and $\ln(1 + \text{prior episodes})$ controls, standard errors clustered by target (52 target clusters in stage two, 68 in the average-marginal-effect sample). Row 2 holds the baseline specification and keeps the episodes with one coded objective, where the strict and dominance codings coincide by construction (45 and 60 target clusters). Row 3 estimates the strict grades at the level of the individual objective, with standard errors clustered by episode (118 and 166 episode clusters). Row 4 clusters by dyad (86 dyad clusters in stage two and 111 in the average-marginal-effect sample). Row 5 returns the 73 censored ongoing episodes as failures to date, which enlarges the average-marginal-effect sample to 236 and leaves the stage-two subsample unchanged. Row 6 repeats the baseline estimates of row 1 and reports in brackets their Holm-adjusted p -values over the family of nine distinct leverage measures, dropping outside options because it is one minus bilateral concentration (Table ??). Romano–Wolf stepdown over the same family gives 0.518 on the stage-two coefficient and 0.688 on the full-compliance effect. Rows 3 to 5 carry the fourteen-country target-list control and row 4 the dominance-ladder coding, the settings of Appendix ?. Bracketed entries are p -values that do not come from the row’s own estimated variance matrix. The alternative outcome codings, the alternative controls, the score bootstrap, and the measure sweep are in Appendices ?? and ??.

The estimate is negative in every row. Grading the outcome without aggregation, estimating on the individual objective in place of the episode, returning the 73 censored ongoing episodes to the sample as failures, and clustering by dyad all hold the sign, and the first of these gives the largest sequential estimate, so aggregation across objectives is not what produces the result. The dyad row differs from the baseline on three dimensions at once and is not a clustering comparison. Multiplicity is where the inference weakens. Nine leverage measures are estimated on the baseline, dropping outside options because it is one minus bilateral concentration, and treating A as one of them moves every margin past conventional levels: on the partial margin the raw p -value of 0.040 becomes 0.318 under Holm and 0.418 under a Romano–Wolf stepdown with 1,000 resamples, on the full margin 0.125 becomes 0.998 and 0.688, and the sequential coefficient moves from 0.063 to 0.500 and 0.518. A cluster bootstrap on the unadjusted partial margin returns 0.075 against the analytic 0.040, so

bootstrap inference is less favourable than the analytic clustered approximation before any multiplicity question arises. The import-share ratio A_M is the only measure to clear Holm on any margin, at 0.036 on partial compliance, and it does not clear the stepdown (0.177). The rotation is a pattern in the point estimates, and once the measures tried are accounted for it is not a pattern that survives a family-wise correction.

Removing countries one at a time leaves the point estimates where they are. Across 169 refits that drop each target and then each sender in turn, the stage-two coefficient and both average marginal effects retain their baseline sign in all 265 estimates, and standard errors fall in none of them, so this is evidence on stability and none on precision. Figure ?? in Section G of the supplementary appendix plots the refits against the baseline.

Section B of the supplementary appendix reports the dominance-ladder coding, the alternative treatments of negotiated settlements, and the ordered-logit diagnostics. Section C of the supplementary appendix reports the subsample splits by era, sender, and instrument, the alternative control sets, the score bootstrap, the specifications that add GDP asymmetry and absolute exposure, and the full measure sweep.

Network-adjusted dependence. A bilateral share prices no substitution. When a sanctioned target can redirect sales to third markets, the share of its exports that goes to the sender overstates what severing the relationship costs it, and [Kleinman et al. \(2024\)](#) show that shares are poor sufficient statistics for dependence once general-equilibrium substitution is accounted for. Their framework delivers a real income exposure matrix whose entry U_{ni} is the elasticity of country n 's real income per capita to country i 's productivity, so third-country substitution sits inside the measure. Applying the construction of A to the exposure pair of a sanction dyad gives a network-adjusted ratio built on the same episodes, the same three-year pre-episode window, and the same trade data as A . The version used here, A^{rup} , takes the exposure pair from a bilateral-rupture counterfactual in which the pair's trade costs are set prohibitive in both directions, which is the object closest to a sanction and which makes both components non-negative by construction. The replication package for that framework could not be obtained, so the exposure matrix is rebuilt from the published equations, and every estimate that rests on it comes from that reconstruction, not from a rerun of published code. Section C.3 of the supplementary appendix documents the construction, its verification, and the sensitivity of the estimates.

The two ratios correlate at 0.650, so the adjustment is not a relabelling of A . Table 7 compares them on the intersection of their estimation samples, 144 episodes carrying 41, 45 and 58 outcomes of each kind, with both measures re-estimated there so that no comparison confounds the measure with the sample.

Table 7: Network-adjusted dependence and the composition of outcomes, common sample

Specification	$\Delta P(\text{partial})$			$\Delta P(\text{full})$			N
	Estimate (SE)	p	95 pct. interval	Estimate (SE)	p	95 pct. interval	
<i>Panel A: interquartile change, one measure at a time</i>							
Bilateral ratio A	0.113 (0.068)	0.098	[-0.021, 0.246]	-0.086 (0.059)	0.141	[-0.201, 0.029]	144
Network-adjusted ratio A^{up}	0.117 (0.038)	0.002	[0.042, 0.193]	-0.064 (0.025)	0.010	[-0.112, -0.015]	144
<i>Panel B: both entered jointly, effect of one standard deviation</i>							
Network-adjusted ratio A^{up}	0.150 (0.054)	0.005	[0.045, 0.256]	-0.058 (0.055)	0.295	[-0.165, 0.050]	144
Bilateral ratio A	0.012 (0.046)	0.788	[-0.077, 0.102]	-0.047 (0.065)	0.475	[-0.174, 0.081]	144

Notes: The network-adjusted ratio applies the construction of $A = m_3/(m_3 + m_4)$ to a general-equilibrium exposure pair built in the framework of Kleinman et al. (2024), in which U_{ni} is the elasticity of country n 's real income per capita to country i 's productivity, so third-country substitution is priced inside the measure. A^{rup} uses the bilateral-rupture counterfactual, in which the pair's trade costs are set prohibitive both ways, so its two components are non-negative by construction. It is built over the same three-year pre-episode window and the same trade data as A , and the exposure matrix is reconstructed from the published equations, because the replication package could not be obtained. The common sample is the intersection of the two estimation samples, with both measures re-estimated on it, so the two rows of Panel A are the same 144 episodes by construction, in 56 target clusters, with outcome counts 41 none, 45 partial and 58 full, all above the forty-event floor this paper applies. Panel A reports the change in the predicted probability across each measure's own 25th to 75th percentile, 0.567 to 0.982 for A and 0.819 to 0.992 for A^{rup} , under the baseline specification of Table 4 with the measure swapped and nothing else changed. The two point estimates are close on both margins and the intervals are not: the network interval is 56 percent as wide on partial compliance and 42 percent as wide on full compliance, and it excludes zero on both margins where the bilateral interval contains it. Panel B enters both ratios in the same multinomial logit, each standardised, and reports average marginal effects with intervals from the same clustered variance. The two ratios correlate at 0.650 on the main sample. Sequential stages for A^{rup} : stage one 1.360 (0.809), $p = 0.093$, stage two -2.741 (0.929), $p = 0.003$ on 103 episodes. The linearised variant $A^{\text{net}} = U_{ts}/(U_{ts} + U_{st})$ is defined on 119 of the 193 main-sample episodes, correlates at 0.710 with A , and has outcome counts 33, 37 and 49, two of them below the floor, so it is reported in Table ?? and not here. Entered alone and standardised, the sender's network exposure U_{st} moves partial compliance by -0.154 (0.048), $p = 0.001$, where the raw sender export share returns -0.191 (0.162), $p = 0.238$ (Table ??). Standard errors are clustered by target throughout. All estimates are associations among imposed trade sanctions.

The two measures agree on the size of the rotation and disagree on how well it is determined. Across its own quartiles the paper's ratio raises the predicted probability of partial compliance by +0.113 and lowers the probability of full compliance by -0.086, and the network-adjusted ratio gives +0.117 and -0.064. Neither bilateral interval excludes zero ($p = 0.098$ and $p = 0.141$) and both network intervals do ($p = 0.002$ and $p = 0.010$). The network interval is 56 percent as wide as the bilateral one on partial compliance and 42 percent as wide on full compliance. Entered jointly and standardised, the network-adjusted ratio carries an effect on partial compliance of +0.150 ($p = 0.005$) and the bilateral ratio

one of +0.012 ($p = 0.788$), so conditional on the network-adjusted measure the bilateral index contains little additional predictive information. The joint specification establishes that and nothing further. The network measure is a reconstruction defined on 144 of the 193 main-sample episodes, and it does not rank the two measures. The adjustment buys precision and not magnitude: the rotation the paper reports for bilateral shares is the same rotation, measured with less noise, once dependence is allowed to substitute through third countries.

The measure carries limitations. The linearised version of the ratio, $A^{\text{net}} = U_{ts}/(U_{ts}+U_{st})$, is estimable on 119 of the 193 main-sample episodes, where its none and partial categories hold 33 and 37 episodes and breach the forty-event floor this paper applies elsewhere, which is why the rupture variant carries the main text. A further 38 episodes fall outside A^{net} because their network exposure is negative, which is what the framework returns for pairs whose third-market competition dominates the direct trade channel, and on those pairs the rotation is neither confirmed nor refuted. Section C.3 of the supplementary appendix shows that what changes the inference is the measure and not the subsample, and that no specification clears the floor and the sign rule at once.

The extensive margin is also less flat under this measure than under the bilateral shares. The stage-one coefficient on A^{rup} is +1.360 ($p = 0.093$), short of conventional significance and inside the ten percent level, where every stage-one coefficient in Table 3 sits above $p = 0.3$. The extensive margin is undetectable throughout the bilateral evidence, so this is weaker evidence that economically adjusted dependence may also move it, and it is a secondary indication. Section C.3 of the supplementary appendix reports all of these.


The extension also amends the measure-specificity finding of Section 3.1. Bilateral dependence levels carry no detectable association with either margin, and that survives for the target’s own network exposure, since U_{ts} entered alone moves nothing ($p = 0.960$ on partial compliance and $p = 0.669$ on full). It does not survive for the sender’s. Entered alone and standardised, U_{st} lowers the probability of partial compliance by 0.154 ($p = 0.001$) and raises the probability of full compliance by 0.072 ($p = 0.081$), where the raw sender export share returns -0.191 ($p = 0.238$) on the partial margin. On the extensive margin the two estimators disagree: the multinomial puts the effect on no concession at +0.082 ($p = 0.017$) and a sequential logit of any concession on the same measure and sample puts it at -0.053 ($p = 0.186$), so this sample does not measure that margin closely enough to read a sign off it. What the network adjustment delivers is the composition result. Dependence levels measured as bilateral shares are uninformative, and once substitution through the trade network is incorporated the sender’s exposure carries information that its bilateral share does not.

3.4 Threat-stage evidence


How outcomes are composed at each stage of resolution is a description the GSDB sample cannot provide, since GSDB records only imposed cases. The TIES database ([Morgan et al., 2014](#)) codes final outcomes by stage for 1,412 sanction cases from 1945 to 2005, including threatened-only cases. TIES is a distinct dataset from the GSDB episodes used everywhere else in this paper, with its own universe, coding rules, and coverage window, and the tabulations below are uncontrolled shares of resolved cases. They describe the composition of outcomes on each side of the imposition threshold. They carry no information about what sorts cases across it. Leverage enters only in the test that closes the subsection.

Complete acquiescence occurs in 30.9 percent of threat-resolved cases and 22.0 percent of imposed-resolved cases, and among successes the complete-acquiescence share falls from 55.7 percent at the threat stage to 38.5 percent after imposition, with negotiated settlement absorbing the difference (12.7 against 24.0 percent). Overall, 45.5 percent of all coded coercion successes occur before any sanction is imposed, and under the strict acquiescence-only definition the threat stage accounts for the majority, 52.6 percent. Restricting to trade instruments lowers the threat-stage share (37.5 percent relaxed, 43.2 percent strict) and preserves the composition shift: complete acquiescence falls from 25.2 to 20.8 percent of resolved cases and from 45.2 to 35.7 percent of successes across the imposition threshold.

The any-success rate is flat across stages, 55.4 percent among threat-resolved and 57.0 percent among imposed-resolved cases, so the composition of success shifts across the imposition threshold, away from complete acquiescence and toward negotiated settlement. That composition shift has the shape of the two-margin pattern, in an independent dataset that observes the stage GSDB cannot. It is a description of stage composition and not evidence on the leverage mechanism. Stage of resolution is itself an equilibrium outcome, the shares are uncontrolled, and leverage plays no part in them, so these tabulations neither support nor adjudicate the selection account.



(a) Does leverage keep threats from being imposed?



(b) Full compliance, before and after imposition

Figure 3: The selection account against threat-stage data. Panel (a) reports the imposition margin, which the benchmark signs in neither direction and which is shown for completeness: leverage raises the probability that a threat reaches imposition. Panel (b) carries the comparison. What would evidence the mechanism is a wedge, leverage raising full compliance before imposition while lowering it after, so the amber intervals would be positive and the navy ones negative. The amber intervals cover zero and the navy ones are positive, so the wedge does not appear. Coefficients with 95 percent intervals clustered by target, on the unrestricted TIES universe with frozen controls.

The account can be tested directly, and Figure 3 is that test. If leverage sorts the compliant cases out before imposition, it should raise full compliance at the threat stage and drain the imposed-stage cell. Across the 816 TIES episodes for which the network ratio is computable, leverage leaves threat-stage full compliance flat, at $p = 0.497$ and $p = 0.924$ on

the two network measures and $p = 0.435$ on the bilateral ratio, and raises it after imposition, so none of the twelve frozen specifications produces the wedge that would evidence it. What the design fixed in advance is that the wedge, a positive threat-stage gradient against a negative imposed-stage one, is what would support the mechanism, and that no test of the difference between the two margins is available, since they sit on nested samples with different base rates. The benchmark itself signs the imposed margin only under the moderate selection intensity of Proposition ??(c). Part (b) leaves the sign to pressure net of selection and part (d) makes every imposed share invariant. A flat threat-stage estimate and a positive imposed-stage one are a configuration in which pressure dominates selection, so what the exercise establishes is absence of support for the selection reading and not its refutation. The network measure buys no precision on this universe, where its standard error per standard deviation runs between one and one and a fifth times the bilateral one and the two measures correlate at 0.906 against 0.650 in the estimation sample. Leverage does raise the probability that a threat reaches imposition at all, at $p = 0.093$ and $p = 0.032$, a margin the benchmark signs in neither direction.

4 Political Entry and Economic Bargaining

The three facts of Section 3 ask for a model in which one thing decides whether a target accommodates and another decides how far it goes. This section builds the smallest such structure. A target carries two independent types: a political admissibility level, which settles whether accommodation can be contemplated at all, and a concession cost, which settles how much is delivered once it can. The sender chooses how hard to press and how much to demand, and buys pressure with a technology that uses the target's stake and is paid for out of its own. Divisible demands and convex concession costs make the generic outcome partial. All proofs are in Section I of the supplementary appendix.

The model the paper began with is a different one, in which a threat stage clears prospective full compliers out of the sample before it is formed. Section 3.4 takes that account to data that grade outcomes on both sides of the imposition threshold and does not confirm it. It is kept as a benchmark in Section F of the supplementary appendix, since it states the hypothesis those tests are aimed at, and it is not the model used here.

4.1 Political admissibility

If leverage does not clear prospective full compliers out of the pool before imposition, then whatever separates the decision to concede at all from the decision of how much to concede

has to operate among the imposed cases themselves.

Let the target carry a second and independent type: a political admissibility level a_i , with accommodation of any size politically available only when $a_i \geq \underline{a}$. Write $A_i = \mathbf{1}\{a_i \geq \underline{a}\}$ and $\pi_A = \Pr(A_i = 1)$. Admissibility is a disposition to accommodate at all, not a judgement about the particular demand on the table: a government whose survival, constitutional order or governing coalition rules out being seen to yield under duress refuses whatever is asked and whatever the economic cost. Preferences over a are lexicographic in that sense, so admissibility does not trade off against pressure and $\underline{a}$ does not depend on the size of the demand. That is a restriction, and it is the one that keeps π_A out of the sender's problem: a sender that could soften a target's willingness to engage by asking for less would face a different optimisation, and the identification results below would not go through unchanged. The economic side is unchanged. A target with $A_i = 1$ solves the problem of Proposition 2, conceding $\theta^*(k) = \min\{x/(k\bar{\theta}), \bar{\theta}\}$, where x is now the effective pressure the sender chooses, no longer a primitive.

The sender chooses effort e and the demand $\bar{\theta}$. Effort buys pressure through a coercion technology $x = m_T g(e)$ and costs the sender $m_S c(e)$ in its own trade. When g and c share a curvature the cost of delivering pressure x is x/λ with $\lambda = m_T/m_S$: relative exposure emerges from the technology of coercion instead of being written into a cost function. Section 4.2 derives the two cutoffs that follow from the target's problem, $R = x/\bar{\theta}^2$ for full compliance and $k_A = x/(\tau\bar{\theta})$ for the reporting threshold τ . Taking them as given, the observed categories follow.

Proposition 1 (Outcome probabilities under political admissibility). *With a and k independent, $\underline{a}$ invariant to exposure, and $0 < \tau \leq \bar{\theta}$ so that full compliance clears the reporting threshold,*

$$\Pr(\text{any}) = \pi_A F_k(k_A), \quad \Pr(\text{full}) = \pi_A F_k(R), \quad \Pr(\text{partial}) = \pi_A [F_k(k_A) - F_k(R)].$$

Conceding nothing arises in two ways,

$$\Pr(\text{none}) = (1 - \pi_A) + \pi_A [1 - F_k(k_A)],$$

a political route and an economic one. If the reporting threshold is small enough that k_A sits far in the upper tail of F_k , the density term $f_k(k_A)k_A$ is negligible and, for elasticities of the order the sender's problem generates, $\partial \Pr(\text{any})/\partial \ln \lambda \approx 0$, so the extensive margin is flat without any cancellation of elasticities. That condition is not free. It requires the economic route to carry little of the observed no-concession mass, which by the identity above means

π_A close to $\Pr(\text{any})$. Section 4.5 asks how much of that the data will bear.

4.2 Economic bargaining

Take a target for which accommodation is politically available, $A_i = 1$. It faces a sanction of effective pressure x and a demand of size $\bar{\theta}$, both chosen by the sender in Section 4.3, and it picks how much of the demand to meet. Concessions are divisible and irreversible, and the cost of conceding $\theta \in [0, \bar{\theta}]$ is

$$C(\theta) = \frac{1}{2} k \theta^2,$$

where the type $k > 0$ is the target's private marginal cost of delivering a concession, drawn from a distribution F_k with density f_k on $(0, \infty)$. It is the cost of implementing and absorbing a policy change that is already politically available, and it is distinct from the admissibility type a of Section 4.1, which decides whether any accommodation can be contemplated. Keeping the two apart is what lets the intensive margin depend on economics alone while the extensive one depends on both. Convexity captures the standard logic that the first items conceded are the cheapest (Mangini, 2024).

Two symbols change meaning here. In this section θ is the concession level and $\bar{\theta}$ the demand, not the trade elasticity of equation (2), and k is the target's concession cost, not the ordinal demand-type proxy of Section 2.1.

The sanction imposes a loss that compliance relieves pro rata: with a concession θ in place the target bears

$$x \left(1 - \frac{\theta}{\bar{\theta}}\right),$$

so a target that meets the demand in full escapes the pressure and one that concedes nothing bears all of it. The target solves $\min_{\theta \in [0, \bar{\theta}]} \frac{1}{2} k \theta^2 + x \left(1 - \theta/\bar{\theta}\right)$.

Proposition 2 (Interior partial compliance). *The target's optimal concession is unique and equals*

$$\theta^*(k) = \min \left\{ \frac{x}{k \bar{\theta}}, \bar{\theta} \right\}.$$

Compliance is full if and only if $k \leq R \equiv x/\bar{\theta}^2$. Above that margin the concession is interior, strictly decreasing in k , strictly increasing in pressure, and strictly decreasing in the size of the demand, with $\theta^ \rightarrow 0$ as $k \rightarrow \infty$.*

Divisibility and convex costs deliver the partial-compliance margin: pressure equates the marginal cost of conceding to the marginal relief, and only types below R find full compliance worthwhile. A reporting threshold τ , with $0 < \tau \leq \bar{\theta}$ so that a full concession always clears it, maps the continuous margin into the observed categories. A concession is recorded at all

when $\theta^* \geq \tau$, which happens for

$$k \leq k_A \equiv \frac{x}{\tau \bar{\theta}}.$$

The two cutoffs are the whole economic content of the model. They respond differently to the demand, R through $\bar{\theta}^2$ and k_A through $\bar{\theta}$. A larger demand shrinks both sets, the full-compliance set twice as fast in logs, and it is that gap and not any invariance of the any-concession margin that lets a more ambitious demand rotate full compliance into partial.

4.3 The sender's problem

Effective pressure and demand ambition are chosen, not given. The sender picks an effort e , which buys pressure through a coercion technology $x = m_T g(e)$ and costs it $m_S c(e)$ in its own trade. When the technology and the cost share a curvature, $g(e) = e^\zeta$ and $c(e) = e^\zeta$, inverting the first gives $e = (x/m_T)^{1/\zeta}$ and the cost of delivering pressure x is

$$m_S \left(\frac{x}{m_T} \right) = \frac{x}{\lambda}, \quad \lambda = \frac{m_T}{m_S}.$$

Nothing about relative exposure is assumed here. It falls out of the technology, because pressure is produced with the target's stake and paid for with the sender's, and it is the only way the two stakes enter.

The sender also picks the demand $\bar{\theta}$, at a cost $K(\bar{\theta})$ that is increasing and convex, and values what it obtains at w per unit of concession together with a prize V for complete capitulation. Writing $\Psi(R) = F_k(R) + R \int_R^\infty k^{-1} dF_k(k)$ for the expected concession per unit of demand, so that $\mathbb{E}[\theta^*] = \bar{\theta} \Psi(R)$ and $\Psi'(R) = \int_R^\infty k^{-1} dF_k(k) > 0$, the problem is

$$\max_{x, \bar{\theta}} \pi_A \left[w \bar{\theta} \Psi(R) + V F_k(R) \right] - \frac{x}{\lambda} - K(\bar{\theta}), \quad R = \frac{x}{\bar{\theta}^2}.$$

Proposition 3 (The sender's choices). *An interior optimum satisfies*

$$\pi_A \left[\frac{w \Psi'(R)}{\bar{\theta}} + \frac{V f_k(R)}{\bar{\theta}^2} \right] = \frac{1}{\lambda}, \quad \pi_A \left[w \Psi(R) - 2wR \Psi'(R) - \frac{2V f_k(R)R}{\bar{\theta}} \right] = K'(\bar{\theta}).$$

Both conditions involve the two stakes only through λ , so x^ , $\bar{\theta}^*$ and every outcome share are homogeneous of degree zero in (m_T, m_S) .*

The prize is what turns the demand response in the direction the data require. Set $V = 0$ and take a quadratic ambition cost $K(\bar{\theta}) = \frac{\kappa}{2} \bar{\theta}^2$. The first condition gives $\bar{\theta} = \pi_A w \lambda \Psi'(R)$,

and substituting it into the second collapses the pair to one equation in the cutoff,

$$\Phi(R) \equiv \frac{\Psi(R) - 2R\Psi'(R)}{\Psi'(R)} = \lambda\kappa, \quad \frac{dR^*}{d\lambda} = \frac{\kappa}{\Phi'(R^*)}, \quad \Phi'(R) = -1 + \frac{\Psi(R)f_k(R)}{R\Psi'(R)^2}.$$

The cutoff does respond to leverage, so it is not pinned. What fails without the prize is the sign. Across the type families and admissible equilibria swept in Section E of the supplementary appendix, Φ' came out positive in every cell in which this equation has a solution, so R^* rises with λ and full compliance becomes *more* likely as leverage grows, the opposite of the estimated rotation. Adding the prize reversed that sign throughout the same sweep. Neither statement is proved here: both are properties of the parameter families examined, not general comparative statics.

Homogeneity is not a modelling convenience. Writing $g(e) = e^\zeta$ and $c(e) = e^\xi$, the cost of delivering pressure x is $m_S(x/m_T)^{\xi/\zeta}$, and this is a function of λ and x alone only when $\xi = \zeta$. Exact invariance to the common scale of the relationship therefore *forces* the cost of effective pressure to be linear in pressure with slope $1/\lambda$, whatever the curvature of the underlying technology. The relative-versus-scale distinction that Section 3.2 finds in the data is the observable counterpart of that restriction, and the model would have to be given a different coercion technology to lose it.

4.4 Identification and comparative statics




Figure 4: The architecture. A target that cannot politically afford to accommodate refuses whatever the economic cost, and that decision does not respond to leverage. Among targets that can, relative leverage and the concession-cost type determine how far the accommodation goes, including the branch in which the cost is high enough that nothing is conceded even though accommodation was politically available. The two observed margins follow. The extensive one is $\pi_A F_k(k_A)$: when the reporting threshold puts k_A far into the tail, $F_k(k_A)$ approaches one and the margin approaches π_A , which the quantitative exercise of Section 4.5 requires to be well below one. Leverage moves both cutoffs. What makes its effect on the extensive margin small is the little probability mass sitting at k_A , and not any absence of leverage from that margin.

Figure 4 sets out the two routes and the cutoffs that separate them. Three claims follow, and they do not stand on the same footing.

Proposition 4 (Identification of the intensive composition). *At every λ at which $\Pr(\text{any} \mid \lambda) > 0$, which is every λ when $\pi_A > 0$ and $F_k(k_A(\lambda)) > 0$,*

$$\frac{\Pr(\text{full} \mid \lambda)}{\Pr(\text{any} \mid \lambda)} = \frac{F_k(R(\lambda))}{F_k(k_A(\lambda))}.$$

The composition of compliance conditional on conceding is identified without identifying π_A and without restricting how π_A varies with leverage.

That is the strongest identification statement in the paper and it carries no exclusion

restriction. The next one does.

Proposition 5 (Comparative statics under the political-entry exclusion). *Let $\pi_A > 0$, let F_k be strictly increasing with F_k and its density f_k positive at the two cutoffs, $F_k(R(\lambda)) > 0$, $F_k(k_A(\lambda)) > 0$, $f_k(R(\lambda)) > 0$ and $f_k(k_A(\lambda)) > 0$ over the range of λ considered, so that both probabilities are strictly positive and their logarithms are defined, and suppose $\partial \ln \pi_A / \partial \ln \lambda = 0$. Writing ε_R , ε_{k_A} , $\varepsilon_{\bar{\theta}}$ and ε_x for the leverage elasticities of the two cutoffs, the demand and the pressure,*

$$\frac{\partial \ln \Pr(\text{full})}{\partial \ln \lambda} = \frac{Rf_k(R)}{F_k(R)} \varepsilon_R, \quad \frac{\partial \ln \Pr(\text{any})}{\partial \ln \lambda} = \frac{k_A f_k(k_A)}{F_k(k_A)} \varepsilon_{k_A},$$

with both factors strictly positive, so the signs of ε_R and ε_{k_A} are those of the observed slopes for every admissible π_A and every F_k meeting the density condition above. Since $k_A/R = \bar{\theta}/\tau$ and $x = R\bar{\theta}^2$,

$$\varepsilon_{\bar{\theta}} = \varepsilon_{k_A} - \varepsilon_R, \quad \varepsilon_x = 2\varepsilon_{k_A} - \varepsilon_R,$$

and therefore

$$\frac{1}{2}\varepsilon_x < \varepsilon_{\bar{\theta}} \iff \varepsilon_R < 0, \quad \varepsilon_{\bar{\theta}} < \varepsilon_x \iff \varepsilon_{k_A} > 0.$$

The last line is worth stating plainly. The elasticity band $\frac{1}{2}\varepsilon_x < \varepsilon_{\bar{\theta}} < \varepsilon_x$, under which a more exposed target faces both more pressure and a demand that grows fast enough to put full compliance further out of reach, is exactly equivalent to the pair $\{\varepsilon_R < 0, \varepsilon_{k_A} > 0\}$, and carries no content beyond those two signs. What that pair represents is a full-compliance cutoff that falls with leverage together with an any-concession cutoff that rises with it. It is not a statement that the extensive margin is flat: with the factor $k_A f_k(k_A)/F_k(k_A)$ strictly positive, $\varepsilon_{k_A} > 0$ makes $\partial \ln \Pr(\text{any})/\partial \ln \lambda$ strictly positive. The observed margin is close to flat only when that factor is small, which is the density condition of Section 4.1 and the restriction Section 4.5 examines.

Without the exclusion the second proposition fails, since $\partial \ln \Pr(\text{full})/\partial \ln \lambda = \partial \ln \pi_A/\partial \ln \lambda + [Rf_k(R)/F_k(R)]\varepsilon_R$ and the sign of ε_R cannot be recovered from the full-compliance margin alone. The model assumes that political admissibility does not respond to bilateral leverage. The data are consistent with that assumption and do not establish it.

The political-entry probability itself is not point identified. Since $F_k \leq 1$, $\pi_A \geq \Pr(\text{any} \mid \lambda)$ at every λ , so π_A lies in $[L, 1]$ with $L = \sup_{\lambda} \Pr(\text{any} \mid \lambda)$, and L is unknown. A simultaneous confidence procedure over the leverage support puts $\pi_A \geq 0.667$ at the 95 percent level without imposing a flat extensive margin, and $\pi_A \geq 0.647$ under the additional restriction that the extensive margin is constant. These are lower confidence bounds for L and not estimates of it. Section D of the supplementary appendix gives the construction.

Separating π_A from $F_k(k_A)$ would need a political shifter that moves whether a target accommodates at all without moving how much it concedes once it does. These data do not contain one. The seven-country sublist inherited with the target list, Cuba, Iran, Libya, North Korea, Somalia, Sudan and Syria, carries no documented source, and thirteen of its twenty-four episodes in the estimation sample are Iran. The objective-type block has documented provenance, is inert on the intensive margin, and does not move the extensive margin once leverage is controlled, which is unsurprising given that objective type explains nineteen percent of the variation in $\ln \lambda$. Nothing here is estimated off a constructed proxy, and the political component is carried as latent.

4.5 Quantitative discipline

The results above are about signs. The next discipline to impose is quantitative, and it is demanding: when can the same model also account for the *size* of the two movements, and for the level of the extensive margin as well as its flatness?

Calibrating the type distribution to the observed split and letting relative leverage run across its interquartile range, the model generates movements of the estimated size in a region of the parameter space and not at a point, across three type families, ambition costs spanning two orders of magnitude, and both cost curvatures. But that region depends on how no concession is divided between politics and economics. Reproducing the observed level of the extensive margin as well as its movement pins $F_k(k_A) = \Pr(\text{any})/\pi_A$, and under that restriction the number of matching parameterisations falls from 54 when π_A is near $\Pr(\text{any})$, where the cutoff is capped in the far tail and the count depends on where it is capped, to 6 at $\pi_A = 0.75$ and to none at 0.85 and above.

The reason is worth stating, because it comes from the model and not from the grid. Leverage moves the economic any-concession cutoff by a good deal, with an elasticity between 0.70 and 0.91 throughout the sweep, so the extensive margin can only be flat if the density of types at that cutoff is negligible. That in turn requires $F_k(k_A)$ close to one, which by Proposition 1 means almost all of no concession is political.

Nothing here is offered as validation of the model, and the paper does not rest on it. The intensive result stands on Propositions 4 and 5, which do not use the exercise at all. What the exercise adds is a restriction and a testable implication: the architecture accounts quantitatively for both margins only when political inadmissibility carries nearly all of non-accommodation, so a direct measure of political admissibility would put the model at risk in a way the present data cannot. Section E of the supplementary appendix gives the design, the sweep, and the sensitivity.

5 Conclusion

This paper separates two decisions that empirical practice pools into one, whether a sanctioned government concedes and how far it goes once it does, and estimates them on 193 imposed trade sanctions by single-state senders between 1949 and 2019. Trade exposure predicts the second, and on the first it shows no detectable association.

I find no association between bilateral trade exposure and whether a target concedes anything. Where it does show up is in what the target concedes. Moving from the lower to the upper quartile of the target's stake relative to the sender's, the probability that a target concedes part of what is demanded rises from 27 to 40 percent and the probability that it concedes everything falls from 42 to 33 percent. Most of that 13-point rise comes out of full compliance, which gives up 9 points, and the remaining 4 come off outright refusal, a movement the data cannot distinguish from zero. Among episodes with a single stated objective, where the two grading rules coincide, the same move raises the probability of partial compliance by 22 percentage points and lowers the probability of full compliance by 16. A binary success regression adds the two movements and returns nothing, which is what a literature working with a success dummy has been estimating.

I find that information in the relative structure of exposure and not in the bilateral levels beneath it. Entered together, the two stakes carry the pattern through their difference, and their sum is not distinguishable from zero. The evidence is stronger along the relative direction than along the scale direction, and the interval on the sum still admits scale effects large enough to matter. Target import dependence, the measure a first-order theory of trade pressure prefers, moves neither margin in the pooled sample, and one post-Cold-War subsample cell is the only exception anywhere in the level measures. A ratio is size-free, so it cannot stand for how much damage a sanction can inflict. What it describes is bargaining position, and bargaining position sorts what targets concede and not whether they concede. Measuring dependence with third-country substitution priced in returns the same rotation with intervals about half as wide, and makes the sender's own exposure informative: a standard deviation of it lowers the probability of partial compliance by 15 percentage points, where the sender's bilateral export share shows nothing.

I find no support for the account that would explain both margins by selection. If the threat stage takes the prospective full compliers first and leaves imposition to the resolute, leverage should raise full compliance before imposition and drain it after. On the threat-stage data, leverage leaves full compliance before imposition undetectably different from zero and raises it after imposition. The wedge that would evidence selection does not appear, which leaves the account unsupported and not refuted.

The model I offer divides the decision instead of selecting on it. Whether a government accommodates runs through political admissibility: one that cannot be seen to yield under duress refuses whatever is asked, and that refusal does not respond to how much trade is at stake. How far it goes, once accommodation is available, is settled between the pressure the sender applies and the ambition of what it demands. Relative exposure governs the second because pressure is produced with the target's stake and paid for with the sender's. The composition of compliance is identified without identifying the political margin, which these data put at 67 percent of targets or more and leave latent above that.

These are associations among sanctions that reached imposition, and nothing in the design isolates exogenous variation in leverage.

Three questions follow. Whether political admissibility can be measured is the first: the architecture accounts for both margins together only when inadmissibility carries nearly all of non-accommodation, so a direct measure would put it at risk in a way these data cannot. Whether the composition result replicates outside GSDB is the second: TIES separates partial from complete acquiescence and EUSANCT carries graded outcomes, and the rotation has been estimated on one coding of one database. What the sender chooses is the third: the model has it picking pressure and ambition jointly, and these data code the demand by type and not by size, so the ambition margin is untested.

References

- Akoto, W., Peterson, T. M., and Thies, C. G. (2020). Trade composition and acquiescence to sanction threats. *Political Research Quarterly*, 73(3):526–539.
- Baldwin, D. A. (1985). *Economic Statecraft*. Princeton University Press, Princeton.
- Bapat, N. A., Heinrich, T., Kobayashi, Y., and Morgan, T. C. (2013). Determinants of sanctions effectiveness: Sensitivity analysis using new data. *International Interactions*, 39(1):79–98.
- Bianchi, J. and Sosa-Padilla, C. (2025). International sanctions and dollar dominance. *The Economic Journal*. Advance article, doi:10.1093/ej/ueaf052.
- Cilizoglu, M. and Bapat, N. A. (2020). Economic coercion and the problem of sanctions-proofing. *Conflict Management and Peace Science*, 37(4):385–408.
- Clayton, C., Coppola, A., Maggiori, M., and Schreger, J. (2025). Geoeconomic pressure. Working Paper 34020, National Bureau of Economic Research.

- Clayton, C., Maggiori, M., and Schreger, J. (2024). A theory of economic coercion and fragmentation. Working Paper 33309, National Bureau of Economic Research.
- Clayton, C., Maggiori, M., and Schreger, J. (2026). A framework for geoeconomics. *Econometrica*, 94(1):105–136.
- Drezner, D. W. (1999). *The Sanctions Paradox: Economic Statecraft and International Relations*. Cambridge University Press, Cambridge.
- Drezner, D. W. (2003). The hidden hand of economic coercion. *International Organization*, 57(3):643–659.
- Farrell, H. and Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security*, 44(1):42–79.
- Felbermayr, G., Kirilakha, A., Syropoulos, C., Yalcin, E., and Yotov, Y. V. (2020). The global sanctions data base. *European Economic Review*, 129:103561.
- Hirschman, A. O. (1945). *National Power and the Structure of Foreign Trade*. University of California Press, Berkeley.
- Hufbauer, G. C., Schott, J. J., Elliott, K. A., and Oegg, B. (2007). *Economic Sanctions Reconsidered*. Peterson Institute for International Economics, Washington, DC, 3rd edition.
- Itskhoki, O. and Mukhin, D. (2022). Sanctions and the exchange rate. Working Paper 30009, National Bureau of Economic Research.
- Kleinman, B., Liu, E., and Redding, S. J. (2024). International friends and enemies. *American Economic Journal: Macroeconomics*, 16(4):350–385.
- Lacy, D. and Niou, E. M. S. (2004). A theory of economic sanctions and issue linkage: The roles of preferences, information, and threats. *The Journal of Politics*, 66(1):25–42.
- Liu, E. and Yang, D. Y. (2025). International power. Working Paper 34006, National Bureau of Economic Research.
- Mangini, M.-D. (2024). The economic coercion trilemma. *Journal of Conflict Resolution*, 68(6):1226–1251.
- Mangini, M.-D. (2026). How effective is trade conditionality? Economic coercion in the Generalized System of Preferences. *The Review of International Organizations*, 21(1):143–162.

- McLean, E. V. and Whang, T. (2010). Friends or foes? Major trading partners and the success of economic sanctions. *International Studies Quarterly*, 54(2):427–447.
- Miller, N. L. (2014). The secret success of nonproliferation sanctions. *International Organization*, 68(4):913–944.
- Morgan, T. C., Bapat, N., and Kobayashi, Y. (2014). Threat and imposition of economic sanctions 1945–2005: Updating the TIES dataset. *Conflict Management and Peace Science*, 31(5):541–558.
- Pape, R. A. (1997). Why economic sanctions do not work. *International Security*, 22(2):90–136.
- Peterson, T. M. (2020). Reconsidering economic leverage and vulnerability: Trade ties, sanction threats, and the success of economic coercion. *Conflict Management and Peace Science*, 37(4):409–429.
- Walentek, D., Broere, J., Cinelli, M., Dekker, M. M., and Haslbeck, J. M. B. (2021). Success of economic sanctions threats: Coercion, information and commitment. *International Interactions*, 47(3):417–448.
- Weber, P. M. and Schneider, G. (2022). Post-Cold War sanctioning by the EU, the UN, and the US: Introducing the EUSANCT dataset. *Conflict Management and Peace Science*, 39(1):97–114.
- Whang, T. (2010). Structural estimation of economic sanctions: From initiation to outcomes. *Journal of Peace Research*, 47(5):561–573.
- Whang, T., McLean, E. V., and Kuberski, D. W. (2013). Coercion, information, and the success of sanction threats. *American Journal of Political Science*, 57(1):65–81.
- Yalcin, E., Felbermayr, G., Kariem, H., Kirilakha, A., Kwon, O., Syropoulos, C., and Yotov, Y. V. (2025). The global sanctions data base—Release 4: The heterogeneous effects of the sanctions on Russia. *The World Economy*, 48(9):2003–2017.

Funding. This research received no specific grant from any funding agency, commercial or not-for-profit sector.

Competing interests. The author declares none.

Use of AI tools. Claude (Anthropic) was used for editorial revision of the manuscript, for cross-checking each reported quantity against the table or estimation log that produces it,

for numerically auditing the propositions and their hypotheses, and for writing the scripts that build the exhibits from stored estimation output. It did not carry out the estimation, and no estimate, standard error, sample definition or specification was produced or altered by it. The author is accountable for the accuracy, integrity and originality of the paper.