Supplementary Appendix

This appendix accompanies “The Two Margins of Economic Coercion: Leverage and Partial Compliance under Imposed Sanctions.” Section, table and figure numbers referred to without a letter are those of the manuscript.

A Sample Construction and Attrition

The dataset is built in a single scripted pass from the raw GSDB release-4 case file, and every count reported in Table 1 is asserted programmatically at build time against an externally maintained case ledger, so the construction fails loudly if a count moves. The raw file has one row per case identifier, and the episode collapse is the identity map. The filter proceeds in a fixed order: trade instrument present, sender a single state mappable to ISO3 under a documented name map, target likewise, and objective string not exactly “other.” Organization senders, including the European Union and the United Nations, and multi-country sender strings are excluded by the single-state rule. Cases listing “other” among several objectives are retained.

Outcome grades are parsed from GSDB’s per-objective success string by the dominance ladder set out in Appendix B. One case in the extract mixes a failed grade with a negotiated settlement and is coded negotiated by the ladder. No case mixes an ongoing grade with a success grade, an assertion checked at build time, so the ladder’s ordering between ongoing and the success grades never binds in this sample. One Soviet-successor episode carries an ongoing grade with a pre-1992 end year and is excluded from the main sample by the censoring rule, a documented boundary case. Episodes sharing a dyad and start year remain separate observations when GSDB assigns separate case identifiers, which is an inference caveat carried forward, and robust standard errors do not adjust for it.

Attrition from the universe runs in two steps. The censoring rule removes the 73 ongoing episodes and takes the 266-episode universe to the 193-episode main sample, and trade-data coverage then takes the main sample to the measure-specific estimation samples (Table 2). The composite flow ladder exists to make that attrition small and documented: reconciled BACI flows are used where available, importer-reported and exporter-reported Comtrade fill the middle decades, and the IMF Direction of Trade series, ranked last on quality, is indispensable for the earliest episodes, which would otherwise lose their entire measurement window. Within a needed dyad-year, a missing bilateral flow is zero-filled only when both countries report non-missing totals that year, so the zeros are recorded silences of active traders and not gaps in coverage. The gravity exposure measure requires GDP and has the thinnest coverage of the four core measures.

B Alternative Outcome Coding and Estimators

The dominance ladder collapses the GSDB per-objective success string to an episode-level category in a fixed order: an episode containing a full-success grade is full, else one containing a partial-success grade is partial, else one containing a negotiated settlement is negotiated, else one containing an ongoing grade is ongoing, else it is none. In the pre-2020 universe this yields 60 episodes with no concession, 22 negotiated settlements, 34 partial, 77 full, and 73 ongoing, the five-category distribution of Table 1, Panel B. Collapsed to three categories on the main sample it gives 60 none, 56 partial, and 77 full. Because the ladder grades a multi-objective episode by its best outcome for the sender, an episode that met one demand in full and failed another counts as full compliance. The strict coding of Section 2.1 withholds the full grade unless every coded objective is a full success, and 10 of the 77 dominance-full episodes in the main sample reclassify under that rule. The recode moves mass only from full to partial, so the any-concession stage is identical under the two codings by construction.

Table 1 carries the dominance-ladder coding through the same measures, sequential stages, and interquartile contrasts as the baseline tables, with the fourteen-country target-list control in place of the sanction history and robust standard errors. That combination, the dominance-ladder coding with the target-list control and robust standard errors, is the comparison specification for the appendix tables that vary one design choice at a time.

Table 1: Alternative outcome coding: dominance ladder with robust standard errors

Panel A: multinomial logit average marginal effects

	P(none)		P(partial)		P(full)		
Measure	AME (SE)	p	AME (SE)	p	AME (SE)	p	N
m_1 target import dependence	-0.132 (0.203)	0.515	0.170 (0.249)	0.494	-0.038 (0.288)	0.895	181
δ gravity exposure ($\theta = 4$)	1.808 (1.613)	0.262	1.726 (1.654)	0.297	-3.534 (1.842)	0.055	163
A export-share asymmetry	-0.063 (0.119)	0.598	0.417 (0.138)	0.003	-0.355 (0.134)	0.008	166
A_M import-share asymmetry	-0.130 (0.128)	0.308	0.390 (0.147)	0.008	-0.260 (0.155)	0.093	166

Panel B: sequential decomposition

	Stage 1: any concession			Stage 2: full | any		
Measure	Coef. (SE)	p	N	Coef. (SE)	p	N
m_1 target import dependence	0.727 (1.040)	0.484	181	-0.663 (1.541)	0.667	126
δ gravity exposure ($\theta = 4$)	-8.817 (8.959)	0.325	163	-20.463 (11.062)	0.064	115
A export-share asymmetry	0.225 (0.635)	0.723	166	-2.683 (1.087)	0.014	118
A_M import-share asymmetry	0.522 (0.700)	0.456	166	-2.214 (0.879)	0.012	118

Panel C: finite p25–p75 probability contrasts

	Measure		ΔP(full)		ΔP(partial)		
Measure	p25	p75	(SE)	p	(SE)	p	N
m_1 target import dependence	0.006	0.129	-0.005 (0.035)	0.899	0.021 (0.030)	0.489	181
δ gravity exposure ($\theta = 4$)	0.000	0.017	-0.061 (0.033)	0.063	0.030 (0.027)	0.280	163
A export-share asymmetry	0.541	0.985	-0.160 (0.061)	0.009	0.191 (0.065)	0.003	166
A_M import-share asymmetry	0.826	0.996	-0.046 (0.026)	0.082	0.070 (0.028)	0.011	166

Notes: The alternative outcome coding: the dominance ladder, which grades a multi-objective episode by its best objective grade, with demand-type indicators, a United States sender indicator, a Cold War indicator, and the fourteen-country target-list indicator, estimated one measure at a time on the main sample with robust standard errors. Panel A reports multinomial-logit average marginal effects, Panel B the sequential logits, and Panel C the interquartile contrasts in predicted probabilities on each measure's observed support. Under this coding A's interquartile move takes P(full) from 0.492 to 0.332 and P(partial) from 0.205 to 0.396. Table 4 varies this specification row by row, and Table 2 places the two codings side by side. All estimates are associations among imposed trade sanctions.

For A the full-compliance average marginal effect is -0.355 ($p = 0.008$) under the dominance ladder and -0.219 ($p = 0.082$) under the strict coding, with the controls and robust standard errors held fixed, and the stage-two coefficients are -2.683 ($p = 0.014$) and -1.522 ($p = 0.061$). The partial-compliance effect is at the five percent level under both ($+0.417$, $p = 0.003$ and $+0.277$, $p = 0.048$). The interquartile contrasts under the dominance ladder are -0.160 and $+0.191$. The ten mixed-grade episodes carry the difference between the two columns, so the full margin is a ten percent result under the coding that requires every objective to be met.

Where negotiated settlements are filed and whether ongoing episodes are censored are the other two coding choices, and Table 4 carries both alternatives through the full specification. Treating negotiated settlements as their own category in a four-way multinomial logit leaves the full-compliance average marginal effect at -0.375 ($p = 0.004$), because the full category itself is untouched and only the composition of the middle moves. Excluding settlements entirely gives a stage-two coefficient of -4.217 ($p < 0.001$, $N = 102$): with the boundary cases removed, the remaining partial-versus-full contrast is sharper. Recoding the 73 censored ongoing episodes as failures to date enlarges the sample to 236 and yields an average marginal effect of -0.234 ($p = 0.013$). The two-margin pattern is a property of the outcome structure, and no single coding decision manufactures it.

The single-objective episodes, where the two grading rules coincide by construction, number 157 of the 193 in the main sample, and the regressions on them use 134 episodes with A non-missing (97 in stage two). There the stage-two coefficient is -3.296 ($p = 0.002$, $N = 97$) and the full-compliance effect is -0.397 ($p = 0.006$, $N = 134$) under this appendix's inference, and the target-clustered counterparts appear in Tables 4 and 3. Regressions at the level of the individual objective grade, clustered by episode, reproduce the rotation ($+0.413$, $p = 0.014$ on partial and -0.225, $p = 0.093$ on full). The gap between the two codings is a fact about the ten reclassified multi-objective episodes and not about the core of the sample. The strict recode was run for the other core measures too, holding this appendix's controls and robust standard errors fixed. The import-share ratio A_M holds five percent on the sequential margin there (-2.237, $p = 0.010$, $N = 118$), while the gravity measure sits at a stage-two p-value of 0.224 ($N = 115$). Both cells are reported in the note to Table 2.

Table 2: Outcome coding and the full margin

	Dominance-ladder coding			Strict coding			Single-objective episodes			Objective level		
	Coef. (SE)	p	N	Coef. (SE)	p	N	Coef. (SE)	p	N	Coef. (SE)	p	N
Stage 1: any concession	0.225 (0.635)	0.723	166	0.225 (0.635)	0.723	166	0.248 (0.811)	0.760	134	0.851 (0.618)	0.169	203
Stage 2: full | any	-2.683 (1.087)	0.014	118	-1.522 (0.811)	0.061	118	-3.296 (1.061)	0.002	97	-2.085 (1.215)	0.086	137
AME on P(partial)	0.417 (0.138)	0.003	166	0.277 (0.140)	0.048	166	0.464 (0.132)	0.000	134	0.413 (0.168)	0.014	203
AME on P(full)	-0.355 (0.134)	0.008	166	-0.219 (0.126)	0.082	166	-0.397 (0.146)	0.006	134	-0.225 (0.134)	0.093	203

Notes: All cells use the export-share asymmetry ratio A with the reference-specification controls (demand-type indicators, US sender, Cold War, and the target-list indicator) and robust standard errors. The dominance-ladder coding grades a multi-objective episode by its best objective grade. The strict coding grades an episode full only when every coded objective is a full success, and every mixed-grade episode with a concession is partial. In the main sample 10 of the 77 dominance-ladder-full episodes reclassify to partial (strict counts: 60 none, 66 partial, 67 full). The recode moves mass only from full to partial, so the any-concession stage is identical across the first two columns by construction. The single-objective column restricts the dominance-ladder coding to episodes with one coded objective (157 of 193 main-sample episodes), where the dominance-ladder and strict codings coincide by construction. The objective-level column estimates on one row per coded objective grade with leverage and controls at the episode level and standard errors clustered by episode (203 grade rows in 166 episodes, stage two 137 rows in 118 episodes). Negotiated grades count as partial throughout. The strict-coding column is estimated for the other measures as well: its stage-two estimate is -2.237 (0.868) at $p = 0.010$, $N = 118$ for the import-share ratio A_M and -13.555 (11.147) at $p = 0.224$, $N = 115$ for the gravity exposure measure δ.

If outcomes were ordered and leverage shifted them proportionally, a single ordered-logit slope would summarize the two margins. It does not, and the way it fails is itself evidence. Table 3 collects the diagnostics, estimated under the dominance-ladder coding.

Table 3: Ordered logit and proportional-odds diagnostics

	Ordered logit		LR test of	Generalized ordered logit (autofit)		
Measure	Coef. (SE)	p	proportionality	Constraint released	Any-cut coef.	Full-cut coef.
m_1 target import dependence	0.116 (0.962)	0.904	8.69 (df 8), $p = 0.369$	none ($p = 0.574$)	–	–
δ gravity exposure ($\theta = 4$)	-12.552 (6.857)	0.067	8.32 (df 8), $p = 0.403$	cold-war control ($p = 0.041$)	–	–
A export-share asymmetry	-0.824 (0.611)	0.177	16.26 (df 8), $p = 0.039$	A itself ($p = 0.002$)	-0.041 (0.590), $p = 0.944$	-1.251 (0.564), $p = 0.027$
A_M import-share asymmetry	-0.375 (0.804)	0.641	11.65 (df 8), $p = 0.167$	A_M itself, cold-war control ($p = 0.002$)	0.963 (0.805), $p = 0.231$	-1.313 (0.746), $p = 0.078$

Notes: Estimated under the alternative outcome coding of Appendix B: dominance-ladder coding, target-list control, and robust standard errors. The LR test compares the ordered logit with the multinomial logit on the identical sample under plain maximum likelihood. The generalized ordered logit (gologit2, autofit at the .05 level) reports which variable's parallel-lines constraint is released, with the release-test p-value, and the resulting cut-specific leverage coefficients where the leverage measure itself is released. The unconstrained A_M model produced ten in-sample negative predicted probabilities, so the multinomial logit is the reference model and the generalized fit a diagnostic.

The pooled ordered-logit slope for A is -0.824 ($p = 0.177$), an average of a flat any-margin and a negative full-margin that understates both. The likelihood-ratio test against the multinomial model rejects proportionality for A ($p = 0.039$) and for no other measure. The generalized ordered logit is sharper: fitting cut-specific slopes variable by variable, it releases the parallel-lines constraint for exactly one variable in the A model, A itself ($p = 0.002$), and the released coefficients are -0.041 ($p = 0.944$) at the any-concession cut and -1.251 ($p = 0.027$) at the full-compliance cut. The A_M model repeats the pattern (-1.313 at the full cut, $p = 0.078$, with a null any cut), while for m_1 nothing violates proportionality because nothing moves. The two-margin shape is a statement the data make about the leverage measure specifically, and only in the ratio family. The unconstrained A_M fit produced ten negative in-sample predicted probabilities, a known fragility of the estimator at this sample size, so the multinomial logit is the reference model and the generalized fit a diagnostic.

C Robustness, Additional Measures, and Inference

Table 4 varies every design choice that can be varied without changing the question: the three outcome codings, the censoring rule, era, sender identity, the target list, influential cases, and the instrument-purity subsample, for the ratio measure A (Panel A) and the level measure m_1 (Panel B). The rows that do not vary the coding hold the comparison specification of Appendix B fixed, and the two coding rows depart from it by construction.

Table 4: Robustness results for the two margins: asymmetry ratio A and import dependence m_1

	Stage 1: any			Stage 2: full | any			AME on P(full)		
	Coef. (SE)	p	N	Coef. (SE)	p	N	AME (SE)	p	N
Panel A: export-share asymmetry A									
Dominance-ladder coding, main sample	0.225 (0.635)	0.723	166	-2.683 (1.087)	0.014	118	-0.355 (0.134)	0.008	166
Coding: negotiated settlement as own category	0.225 (0.635)	0.723	166	-2.683 (1.087)	0.014	118	-0.375 (0.129)	0.004	166
Coding: negotiated settlements excluded	-0.088 (0.687)	0.898	150	-4.217 (1.139)	0.000	102	-0.417 (0.132)	0.002	150
Censoring: ongoing included as failure	0.189 (0.476)	0.691	236	-2.683 (1.087)	0.014	118	-0.234 (0.094)	0.013	236
Era: Cold War (begin ≤ 1991)	0.925 (0.838)	0.270	99	-0.799 (1.594)	0.616	65	-0.023 (0.171)	0.892	99
Era: post-Cold-War	-0.489 (1.213)	0.687	67	-3.972 (1.284)	0.002	49	-0.649 (0.157)	0.000	67
Sender: United States	0.201 (1.126)	0.858	67	-0.257 (2.691)	0.924	46	-0.025 (0.367)	0.946	69
Sender: non-US	0.575 (0.733)	0.432	97	-3.052 (0.983)	0.002	70	-0.364 (0.122)	0.003	97
Drop the fourteen target-list countries	0.340 (0.697)	0.626	121	-1.654 (1.139)	0.146	88	-0.205 (0.167)	0.221	121
Trade-only (no other instrument) subsample	0.082 (0.846)	0.922	80	-1.700 (1.115)	0.127	50	-0.225 (0.158)	0.153	80
Drop top-5 influential cases (stage 2 only)	–	–	–	-4.598 (1.037)	0.000	113	–	–	–
Panel B: target import dependence m_1									
Dominance-ladder coding, main sample	0.727 (1.040)	0.484	181	-0.663 (1.541)	0.667	126	-0.038 (0.288)	0.895	181
Coding: negotiated settlement as own category	0.727 (1.040)	0.484	181	-0.663 (1.541)	0.667	126	-0.038 (0.293)	0.897	181
Coding: negotiated settlements excluded	0.394 (1.056)	0.709	159	-0.785 (1.384)	0.571	104	-0.031 (0.277)	0.909	159
Censoring: ongoing included as failure	1.317 (0.979)	0.179	252	-0.663 (1.541)	0.667	126	0.059 (0.213)	0.782	252
Era: Cold War (begin ≤ 1991)	1.530 (1.258)	0.224	109	1.153 (1.386)	0.406	70	0.230 (0.210)	0.272	109
Era: post-Cold-War	-1.296 (2.645)	0.624	72	-9.125 (5.424)	0.093	51	-1.913 (0.834)	0.022	72
Sender: United States	1.015 (1.828)	0.579	68	2.663 (2.704)	0.325	46	0.349 (0.343)	0.309	70
Sender: non-US	0.948 (1.187)	0.425	111	-1.591 (2.493)	0.523	78	-0.154 (0.440)	0.726	111
Drop the fourteen target-list countries	0.153 (1.112)	0.891	134	-0.535 (1.457)	0.714	95	-0.053 (0.299)	0.858	134
Trade-only (no other instrument) subsample	1.200 (1.448)	0.407	84	-0.845 (2.762)	0.760	52	0.077 (0.416)	0.854	84
Drop top-5 influential cases (stage 2 only)	–	–	–	-1.797 (1.465)	0.220	121	–	–	–

Notes: Cell entries as in Tables 4 and 3. The rows that do not vary the outcome coding hold the dominance-ladder specification of Appendix B fixed: dominance-ladder coding, target-list control, and robust standard errors. Controls are demand-type indicators, US sender, Cold War, and the target-list indicator, with the split variable dropped from the controls within its own split. The two coding rows depart from that specification by construction. Negotiated settlement as its own category is a four-category multinomial logit, so its stages 1 and 2 repeat the reference-specification row while its AME comes from the four-category model, and the excluded-settlement row drops those episodes from the sample. In the ongoing-as-failure row stage 2 is mechanically unchanged. The influential-case row drops the five largest-dbeta episodes from the stage-two logit: for A, cases 46 (USA–ROU 1957), 89 (IDN–MYS 1963), 145 (ESP–GIB 1969), 224 (USA–BRA 1978), 271 (GBR–ARG 1982); for m_1, cases 89 (IDN–MYS 1963), 224 (USA–BRA 1978), 228 (USA–IND 1978), 231 (CHN–VNM 1978), 255 (USA–LBY 1981). No cell falls below $N = 40$. The smallest cells are US-sender stage 2 ($N = 46$), post-Cold-War stage 2 ($N = 49$), and trade-only stage 2 ($N = 50$).

The stage-two coefficient and the full-compliance average marginal effect stay negative under the three codings that vary the treatment of negotiated settlements, and the margin is largest when negotiated settlements are excluded entirely (-4.217, $p < 0.001$, $N = 102$), which says the result is a fact about the partial-versus-full boundary and not an artifact of where settlements are filed. Recoding the 73 censored ongoing episodes as failures gives an average marginal effect of -0.234 ($p = 0.013$) on the enlarged sample of 236, so the censoring rule is not doing the work. With the five most influential stage-two cases dropped the coefficient is -4.598. The margin is present among non-US senders (-3.052, $p = 0.002$, $N = 70$) and in the post-Cold-War era (-3.972, $p = 0.002$, $N = 49$), where the average marginal effect reaches -0.649.

The estimate attenuates without reversing sign in the complementary halves of those splits. The Cold War subsample returns -0.799 ($p = 0.616$, $N = 65$). The US-sender subsample returns -0.257 ($p = 0.924$) on a stage-two cell of 46, too small to carry much.

Dropping the fourteen target-list countries attenuates the coefficient to -1.654 ($p = 0.146$, $N = 88$), and restricting to episodes with no non-trade instrument gives -1.700 ($p = 0.127$) on a stage-two cell of 50. Each split cuts the sample roughly in half, so lost significance is partly lost power. The splits are informative about composition, one split at a time. The estimate is larger in magnitude than the pooled -2.683 in the post-Cold-War half (-3.972) and larger again among non-US senders (-3.052), and each of those two cells still contains episodes from the other side of the other split. Their intersection holds 33 stage-two episodes, below the forty-observation floor the cells of this table respect, so the two splits are reported one at a time and no joint cell is estimated. Part of the pooled magnitude comes from the fourteen heavily sanctioned targets, and part comes from episodes where trade restrictions traveled with other instruments.

Panel B runs the identical grid for m_1 and returns coefficients indistinguishable from zero almost everywhere, which extends the measure specificity of Section 3.1 through every subsample. One cell is an exception. In the post-Cold-War era even m_1 turns negative on the full margin (average marginal effect -1.913, $p = 0.022$, $N = 72$), suggesting that in the recent period the pattern begins to reach the level measures. The pooled m_1 rows remain null, and I do not build on this cell.

Table 5 turns to the controls and the standard errors. The fourteen-country target list has no documented construction rule. Dropping the target-list control has little effect on the estimate: the stage-two coefficient is -2.642 ($p = 0.012$) and the average marginal effect -0.356 ($p = 0.007$). Replacing it with sanction-history controls built from the full GSDB, the count of prior episodes against the same target, its logarithm, or a repeat-target indicator, leaves the coefficient on A between -2.454 and -2.684 across variants. The history controls return one side finding: prior sanction exposure predicts full compliance positively (stage-two coefficient 0.079, $p = 0.036$), so repeat targets concede fully more often, conditional on conceding at all.

The robust standard errors of this appendix treat episodes as independent, and repeated targets and repeated dyads argue otherwise. Clustering by target moves the stage-two p-value to 0.017 (52 clusters) and the full-compliance average marginal effect to $p = 0.030$ (68 clusters), and clustering by dyad gives 0.011 and 0.007. A score bootstrap with the null imposed puts the stage-two estimate of this table, which carries the dominance-ladder coding and the target-list control, at $p = 0.015$ under target clustering and at $p = 0.069$ under two-way sender-and-target clustering, the weakest cell in the inference block. The same bootstrap run on the baseline specification, the strict coding with the sanction-history control and target-clustered standard errors, puts the stage-two estimate of Table 3 at $p = 0.052$. The Holm row of the same table adjusts over the family of ten leverage measures in the measure sweep, taking the stage-two p-value from 0.014 to 0.122 and the average-marginal-effect p-value from 0.008 to 0.084, with Romano–Wolf stepdown over 1,000 resamples at 0.226 for the sequential estimate. The outside-options measure is one minus the concentration measure and therefore the same test twice, and the nine distinct members give Holm values of 0.108 and 0.075. The bootstrap does not extend to the multinomial logit, so for the average marginal effect the clustered rows report alternative one-way clustering specifications.

Table 5: Inference and controls for the export-share asymmetry ratio A

	Stage 2: full | any			AME on P(full)		
Specification	Coef. (SE)	p	N	AME (SE)	p	N
Dominance-ladder specification: robust standard errors	-2.683 (1.087)	0.014	118	-0.355 (0.134)	0.008	166
Drop the fourteen-country target list	-2.642 (1.053)	0.012	118	-0.356 (0.132)	0.007	166
Replace list with prior-episode count	-2.454 (1.056)	0.020	118	-0.318 (0.126)	0.011	166
Replace list with $\ln(1 + \text{prior episodes})$	-2.520 (1.064)	0.018	118	-0.322 (0.125)	0.010	166
Replace list with repeat-target indicator	-2.684 (1.066)	0.012	118	-0.344 (0.133)	0.009	166
Cluster by target	-2.683 (1.127)	0.017	118	-0.355 (0.164)	0.030	166
Cluster by dyad	-2.683 (1.056)	0.011	118	-0.355 (0.132)	0.007	166
Score bootstrap, cluster by target	-2.683	[0.015]	118	–	–	–
Score bootstrap, two-way (sender and target)	-2.683	[0.069]	118	–	–	–
Add ln GDP ratio (target/sender)	-3.206 (1.832)	0.080	105	-0.536 (0.216)	0.013	146
Add m_1 absolute exposure	-2.727 (1.118)	0.015	118	-0.361 (0.141)	0.010	166
Add ln GDP ratio and m_1 jointly	-2.963 (1.809)	0.101	105	-0.492 (0.217)	0.024	146
Scale test: coefficient on $M = m_3 + m_4$ (dominance ladder)	0.874 (1.163)	0.452	118	0.252 (0.191)	0.187	166
Scale test: coefficient on $M = m_3 + m_4$ (baseline)	1.600 (1.195)	0.181	118	0.344 (0.183)	0.060	166
Holm-adjusted p, family of ten measures	-2.683	[0.122]	118	-0.355	[0.084]	166

Notes: Both column blocks report estimates for A: the stage-two logit of full compliance given any concession and the multinomial-logit average marginal effect on P(full). Except in the scale-test row labelled baseline, the outcome coding is the dominance ladder and the controls are those of Table 1 with robust standard errors, and each row label names the one element it changes. The replacement rows swap the fourteen-country target list for sanction-history controls built from the full GSDB. The prior-episode count itself enters the stage-two logit at 0.079 (SE 0.038, $p = 0.036$). Cluster counts: 52 targets and 86 dyads in the stage-two sample, 68 targets and 111 dyads in the AME sample. Bracketed entries are p-values that do not come from the row's own estimated variance matrix. The score-bootstrap rows impose the null (bootstrap-t, Rademacher weights, 9,999 replications, seed 20260809), and the bootstrap does not extend to the multinomial logit, so those AME cells are empty and the clustered rows above report alternative one-way clustering specifications. The Holm row adjusts the first row's p-values over the family of ten leverage measures in the measure sweep (Table 6). Romano–Wolf stepdown over the same family (1,000 resamples, seed 20260809) gives 0.226 for the stage-two estimate, and the nine-member family that drops the mechanical mirror of the concentration measure gives Holm 0.108 and 0.075. The scale-test rows report the coefficient on the absolute scale $M = m_3 + m_4$ when it enters jointly with A: the dominance-ladder row keeps this table's specification, and the baseline row uses the strict coding with the $\ln(1 + \text{prior episodes})$ control and target-clustered standard errors. A's own stage-two coefficient in those two rows is -2.801 ($p = 0.010$) and -1.615 ($p = 0.039$). With $\ln M$ in place of M the stage-two scale coefficient is 0.150 (SE 0.131, $p = 0.252$) under the dominance ladder and 0.198 (SE 0.104, $p = 0.057$) under the baseline specification.

A correlates at -0.695 with the log GDP ratio on the estimation sample, so a natural worry is that the asymmetry ratio is country size in disguise. Adding the log GDP ratio leaves the GDP term at -0.063 ($p = 0.697$) while the coefficient on A grows in absolute value to -3.206, and the stage-two p-value rises to 0.080 through standard-error inflation under collinearity and the loss of 13 observations and not through any movement of the point estimate toward zero. The full-compliance average marginal effect stays at five percent (-0.536, $p = 0.013$). Absolute exposure behaves the same way: adding m_1 leaves the coefficient at -2.727 ($p = 0.015$), and entering the log GDP ratio and m_1 together gives -2.963 ($p = 0.101$) with the average marginal effect at -0.492 ($p = 0.024$), on the 105 stage-two and 146 multinomial episodes that carry both controls. Table 5 reports the three rows. The scale test, which enters $M = m_3 + m_4$ jointly with A, is a test of the invariance that Proposition 3 delivers, and not of a prediction the model makes about scale: the political-entry model is homogeneous of degree zero in the two stakes, so it implies no effect of common scale at all. That test appears in the last rows of the same table, and Appendix C.2 reads the same restriction off the two stakes entered in logs.

Table 6 reports the full measure sweep under the dominance-ladder coding.

Table 6: All leverage measures under the dominance-ladder coding

	Stage 1: any		Stage 2: full | any		AME on P(full)		
Measure	Coef. (SE)	p	Coef. (SE)	p	AME (SE)	p	N
m_1 target import dependence	0.727 (1.040)	0.484	-0.663 (1.541)	0.667	-0.038 (0.288)	0.895	181
m_2 sender import dependence	2.019 (2.091)	0.334	6.759 (7.352)	0.358	1.205 (0.928)	0.194	180
m_3 target export dependence	1.334 (1.107)	0.228	-0.322 (1.291)	0.803	0.099 (0.232)	0.668	180
m_4 sender export dependence	0.891 (1.553)	0.566	14.423 (9.637)	0.134	1.523 (1.029)	0.139	181
m_5 bilateral concentration	0.754 (1.052)	0.473	-0.665 (1.507)	0.659	-0.031 (0.286)	0.913	180
m_6 outside options	-0.754 (1.052)	0.473	0.665 (1.507)	0.659	0.031 (0.286)	0.913	180
m_6^H partner-share HHI	2.534 (1.953)	0.194	-0.895 (2.055)	0.663	0.119 (0.404)	0.768	180
δ gravity exposure ($\theta = 4$)	-8.817 (8.959)	0.325	-20.463 (11.062)	0.064	-3.534 (1.842)	0.055	163
A export-share asymmetry	0.225 (0.635)	0.723	-2.683 (1.087)	0.014	-0.355 (0.134)	0.008	166
A_M import-share asymmetry	0.522 (0.700)	0.456	-2.214 (0.879)	0.012	-0.260 (0.155)	0.093	166

Notes: Cell entries as in Tables 4 and 3, one measure at a time, under the alternative outcome coding of Appendix B. N refers to the multinomial logit sample. The gravity exposure measure is reported at $\theta = 4$; rescaling θ rescales the regressor, so its coefficients scale by θ and its p-values are unchanged by construction.

No bilateral share moves either margin at conventional levels in the pooled sample. Sender-side dependence, export-side dependence, bilateral concentration, outside options, and partner concentration all return coefficients indistinguishable from zero on both stages, so the measure specificity of the ratio family is a property of the share menu as a whole and not of one regressor. The gravity exposure measure is the one level measure that comes close under this coding, at -20.463 ($p = 0.064$) in stage two and -3.534 ($p = 0.055$) on the full-compliance average marginal effect, where the baseline specification of Table 4 leaves it at $p = 0.583$. It is a single measure at the ten percent level under one of the two codings, and I do not build on it. The elasticity scaling of δ is inert by construction: rescaling θ rescales the regressor, so coefficients scale by θ and every p-value is unchanged, and nothing in the δ evidence depends on the calibrated elasticity.

The two families of measures carry different information in this sample, so the choice between them matters. Figure 1 plots the ratio against the level for the main sample.

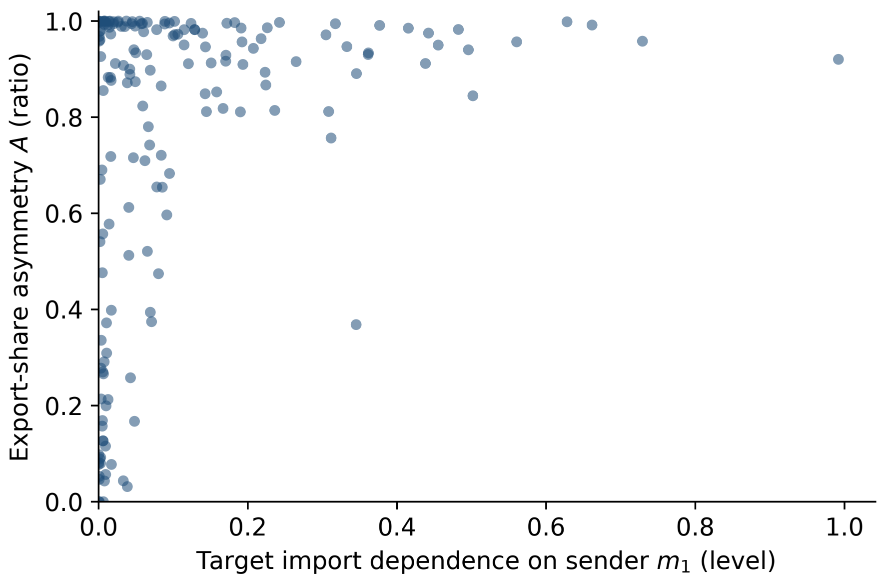

Figure 1: Ratio versus level: export-share asymmetry A against target import dependence m_1, main sample. Each point is one episode. The ratio is size-free, so high asymmetry coexists with negligible exposure levels. Episodes with near-identical target import dependence sit at different asymmetry ratios.

The two asymmetry ratios have an unbounded counterpart in their log odds. Table 7 reports the baseline specification with $R = \ln(m_3/m_4)$ and $R_M = \ln(m_1/m_2)$ in place of A and A_M, on the episodes where both legs of each ratio are strictly positive. In both channels the average marginal effect on partial compliance is positive, the average marginal effect on full compliance is negative, and the stage-two coefficient is negative, with the full-compliance average marginal effect the least precise of the three.

Table 7: Exposure ratios in log-odds form

	AME on P(partial)		AME on P(full)		Stage 2: full | any		N	
Measure	AME (SE)	p	AME (SE)	p	Coef. (SE)	p	Mlogit	Stage 2
$R = \ln(m_3/m_4)$, export channel	0.035 (0.013)	0.010	-0.014 (0.014)	0.310	-0.144 (0.076)	0.059	162	116
$R_M = \ln(m_1/m_2)$, import channel	0.022 (0.012)	0.067	-0.013 (0.016)	0.402	-0.125 (0.078)	0.107	162	116

Notes: The baseline specification of Table 4 with the log odds of each exposure ratio as the regressor: $R = \ln(m_3/m_4)$, which is the logit of A, and $R_M = \ln(m_1/m_2)$, which is the logit of A_M. Strict outcome coding, demand-type indicators, a United States sender indicator, a Cold War indicator, and the log prior-episode sanction-history control, one measure at a time on the main sample, with standard errors clustered by target and the clustered variance propagated to the margins. The first two column pairs are multinomial-logit average marginal effects on the probability of partial and of full compliance, and the third is the coefficient on the measure in the stage-two logit of full compliance given any concession. The log odds are defined where both legs of a ratio are strictly positive, so an episode whose share sits exactly at zero or exactly at one leaves the sample: 162 of the 166 main-sample episodes with A defined carry R (4 dropped), and 162 of the 166 with A_M defined carry R_M (4 dropped). Effects are per log point. Over the estimation-sample interquartile range of R (0.313 to 4.192, a span of 3.880 log points) the predicted probability of partial compliance moves by $+0.137$ (SE 0.054, $p = 0.011$) and the predicted probability of full compliance by -0.056 (0.055, $p = 0.311$). Over the interquartile range of R_M (1.560 to 5.422, a span of 3.861) the same contrasts are $+0.088$ (0.050, $p = 0.076$) and -0.053 (0.064, $p = 0.404$). Cluster counts: 67 targets in the multinomial models and 52 in stage two. All estimates are associations among imposed trade sanctions.

The exposure measures are built from trade recorded before an episode begins, which leaves open whether trade in the years closest to onset had already responded to the dispute that produced the sanction. Table 8 rebuilds A, A_M, and m_1 over four pre-episode measurement windows and estimates the baseline specification on each: the single year before onset, the three-year window used throughout the paper, a five-year window, and a buffered window running from $t-5$ to $t-3$ that drops the two years most likely to carry a response to the developing dispute. The set of years changes and the estimation sample changes with it, since Gravity coverage differs across windows: the multinomial sample runs from 156 to 168 episodes and the stage-two sample from 110 to 119. The exercise is joint stability to the window and to the coverage it induces, and not to the window alone. The three-year rows reproduce Tables 4 and 3.

Table 8: Predetermined exposure: pre-episode measurement windows

	AME on P(partial)		AME on P(full)		Stage 2: full | any		N	
Window	AME (SE)	p	AME (SE)	p	Coef. (SE)	p	Mlogit	Stage 2
Panel A: A export-share asymmetry								
$t-1$	0.126 (0.135)	0.353	-0.079 (0.129)	0.538	-0.587 (0.735)	0.424	163	116
$t-3$ to $t-1$ (baseline)	0.279 (0.136)	0.040	-0.192 (0.125)	0.125	-1.418 (0.761)	0.063	166	118
$t-5$ to $t-1$	0.485 (0.112)	0.000	-0.331 (0.134)	0.013	-2.891 (0.824)	0.000	168	119
$t-5$ to $t-3$ (buffered)	0.462 (0.124)	0.000	-0.327 (0.138)	0.018	-2.691 (0.864)	0.002	156	110
Panel B: A_M import-share asymmetry								
$t-1$	0.441 (0.169)	0.009	-0.267 (0.169)	0.115	-2.199 (1.030)	0.033	163	116
$t-3$ to $t-1$ (baseline)	0.440 (0.153)	0.004	-0.287 (0.161)	0.074	-2.259 (0.945)	0.017	166	118
$t-5$ to $t-1$	0.449 (0.147)	0.002	-0.333 (0.148)	0.024	-2.441 (0.885)	0.006	168	119
$t-5$ to $t-3$ (buffered)	0.295 (0.194)	0.128	-0.205 (0.159)	0.197	-1.545 (1.036)	0.136	156	110
Panel C: m_1 target import dependence								
$t-1$	-0.073 (0.278)	0.794	0.257 (0.239)	0.281	1.369 (1.538)	0.373	180	126
$t-3$ to $t-1$ (baseline)	0.006 (0.297)	0.983	0.240 (0.255)	0.347	0.969 (1.547)	0.531	181	126
$t-5$ to $t-1$	0.072 (0.278)	0.797	0.176 (0.266)	0.507	0.550 (1.501)	0.714	181	126
$t-5$ to $t-3$ (buffered)	0.047 (0.288)	0.870	-0.103 (0.276)	0.710	-0.384 (1.641)	0.815	163	115

Notes: The baseline specification of Table 4 with each measure rebuilt over four pre-episode measurement windows. For an episode beginning in year b the windows are the single year $b-1$, the three years $b-3$ to $b-1$ used throughout the paper, the five years $b-5$ to $b-1$, and a buffered window covering $b-5$ to $b-3$ that omits the two years closest to onset. Only the set of years changes. The composite flow ladder, the country-year aggregation, the zero-fill rule, and the clip to Gravity coverage from 1948 to 2020 are those of Table 2, and the $t-3$ to $t-1$ rows reproduce Table 4 and Table 3. Windows reaching further back lose episodes whose early years fall outside Gravity coverage: the buffered window leaves 8 of the 193 main-sample episodes with no window year in range. Main-sample episodes with A defined across the four windows number 163, 166, 168, 156, with stage-two cells of 116, 118, 119, 110, and the counts for A_M are identical. For m_1 they are 180, 181, 181, 163. Strict outcome coding, demand-type indicators, a United States sender indicator, a Cold War indicator, and the log prior-episode sanction-history control, one measure at a time on the main sample, with standard errors clustered by target and the clustered variance propagated to the margins. The first two column pairs are multinomial-logit average marginal effects on the probability of partial and of full compliance, and the third is the coefficient on the measure in the stage-two logit of full compliance given any concession. Cluster counts at the baseline window are 68 targets in the multinomial models and 52 in stage two, and they fall to 63 and 48 at the buffered window. Varying the window changes when exposure is measured and nothing else. All estimates are associations among imposed trade sanctions.

The rotation keeps its sign at every window for both ratios. For A it is larger and more significant at the five-year and buffered windows than at the three-year window, with average marginal effects of $+0.485$ and $+0.462$ on partial compliance against -0.331 and -0.327 on full compliance, so moving the measurement further from onset does not attenuate the pattern on the samples each window leaves. Two cells lose precision. At the single-year window A gives $+0.126$ ($p = 0.353$) on partial compliance and -0.079 ($p = 0.538$) on full, and at the buffered window A_M gives $+0.295$ ($p = 0.128$) and -0.205 ($p = 0.197$), with both signs unchanged. Target import dependence returns estimates indistinguishable from zero at all four windows, so the measure specificity of Section 3.1 is not a property of the three-year window.

C.1 Episodes with negligible total stakes

A ratio carries no information about the size of the relationship it describes, so the association reported for A could in principle come from dyads where neither side has much at stake. Table 9 re-estimates the baseline specification on samples that drop the episodes with the smallest total export-side stake $M = m_3 + m_4$. The cut points are the tenth, twenty-fifth and fiftieth percentiles of M on the full estimation sample of 166 episodes, computed once and applied to every model. The two terms of M are normalised by different totals, the target's exports in m_3 and the sender's in m_4, so their sum is an index of the two exposures and not a share of either country's trade or of the two combined. In those index units the tenth percentile is 0.0010 and the twenty-fifth is 0.0102.

The estimate steepens as the small-stake episodes leave. The stage-two coefficient on A is -1.418 ($p = 0.063$, $N = 118$) on the full sample, -2.059 ($p = 0.016$, $N = 106$) once the bottom tenth of M is dropped, and -2.718 ($p = 0.033$, $N = 86$) once the bottom quarter is dropped. The multinomial average marginal effect on full compliance follows the same path, at -0.192, -0.227, and -0.308. The spread of A narrows as the trimming proceeds, and netting that out leaves the movement intact: at a common step of one estimation-sample standard deviation the same three coefficients are -0.469, -0.633, and -0.733. Dropping the bottom half of M gives -6.657 ($p = 0.012$) on 57 stage-two episodes, which is 30 full compliers against 27. Every outcome category on that rung falls below the forty-event floor fixed for this exercise before estimation, so the rung is reported and not interpreted. The same floor flags the multinomial and stage-one models at the bottom-quarter rung, each on the 39 episodes with no concession that remain there, while stage two at that rung clears the floor at 45 against 41. The four samples are nested, so the table is one sample being trimmed and not four independent tests.

Interacting A with M returns the expected negative sign and no precision. The stage-two coefficient on the interaction is -6.954 (standard error 9.106, $p = 0.445$) with M in levels and -0.204 (0.209, $p = 0.330$) with $\ln M$. What the interacted model does deliver is the same ordering in its marginal effects: the effect of A on full compliance given any concession runs -0.309, -0.372, and -0.518 across the twenty-fifth, fiftieth and seventy-fifth percentiles of M, against -0.316 ($p = 0.053$) in the uninteracted model on the same sample. Stage one stays null on every rung, which is the pattern Table 3 already reports for A.

Table 9: Total stake and the export-share asymmetry ratio

	AME on P(partial)		AME on P(full)		Stage 2: full | any		N	
Sample	AME (SE)	p	AME (SE)	p	Estimate (SE)	p	Mlogit	Stage 2
Panel A: samples trimmed by total stake $M = m_3 + m_4$								
Full sample	0.279 (0.136)	0.040	-0.192 (0.125)	0.125	-1.418 (0.761)	0.063	166	118
Drop the bottom tenth of M ($M < 0.0010$)	0.368 (0.134)	0.006	-0.227 (0.137)	0.098	-2.059 (0.852)	0.016	150	106
Drop the bottom quarter of M ($M < 0.0102$)	0.447 (0.146)	0.002	-0.308 (0.178)	0.084	-2.718 (1.272)	0.033	125[a]	86
Drop the bottom half of M ($M < 0.0569$)	0.765 (0.205)	0.000	-0.676 (0.267)	0.011	-6.657 (2.636)	0.012	83[a]	57[a]
Panel B: A interacted with M, average marginal effects of A								
M in levels, A at p_{25} of M	0.256 (0.153)	0.095	-0.207 (0.130)	0.110	-0.309 (0.168)	0.067	166	118
M in levels, A at p_{50} of M	0.313 (0.145)	0.032	-0.237 (0.133)	0.074	-0.372 (0.167)	0.026	166	118
M in levels, A at p_{75} of M	0.449 (0.205)	0.029	-0.315 (0.167)	0.060	-0.518 (0.260)	0.046	166	118
$\ln M$, A at p_{25} of M	0.336 (0.135)	0.013	-0.288 (0.126)	0.022	-0.469 (0.150)	0.002	166	118
$\ln M$, A at p_{50} of M	0.393 (0.127)	0.002	-0.325 (0.136)	0.017	-0.524 (0.158)	0.001	166	118
$\ln M$, A at p_{75} of M	0.428 (0.141)	0.002	-0.348 (0.152)	0.022	-0.554 (0.171)	0.001	166	118

Notes: The baseline specification of Table 4 for the export-share asymmetry ratio A, re-estimated on samples defined by the total export-side stake $M = m_3 + m_4$: strict outcome coding, the controls of Table 4, standard errors clustered by target, and the clustered variance propagated to the margins. The first two column pairs are multinomial-logit average marginal effects on the probability of partial and of full compliance, and the third is the coefficient on A in the stage-two logit of full compliance given any concession. In Panel B every entry is an average marginal effect of A, including the stage-two column, because the coefficient on A in an interacted model is its effect at $M = 0$, which is off the support of the data. Panel A cut points are the percentiles of M on the full estimation sample of 166 episodes, computed once and applied to all three models. M adds the target's export dependence on the sender to the sender's export dependence on the target, two shares normalised by different totals, so it indexes the two exposures jointly and is not itself a share of any one quantity. Its tenth percentile is 0.0010, its twenty-fifth 0.0102 and its median 0.0569. A floor on the sum places no floor under either leg: of the 125 episodes above the twenty-fifth-percentile cut, 40 carry both legs at or above 1 percent of their own country's exports, and the median episode there has the target's leg at 9.33 percent against the sender's at 0.63 percent. Outcome counts, none against partial against full, run 48 / 55 / 63, 44 / 50 / 56, 39 / 41 / 45 and 26 / 27 / 30 down the four rungs, and the stage-two cells hold 63 against 55, 56 against 50, 45 against 41 and 30 against 27 in the same order. Superscript a marks a cell whose smallest outcome category falls below the 40-event floor fixed for this exercise before estimation. With the bottom quarter dropped, the flagged cells are the multinomial at 39 and stage one at 39. With the bottom half dropped, they are the multinomial at 26, stage one at 26 and stage two at 27. Those cells are reported and not interpreted. The spread of A narrows as small-stake episodes leave the sample, so the stage-two coefficients are also reported at a common step of one estimation-sample standard deviation of A: -0.469, -0.633, -0.733 and -1.480 down the four rungs. Stage one is null throughout: 0.400 (0.751) at $p = 0.595$, 0.663 (0.860) at $p = 0.441$, 0.555 (0.797) at $p = 0.486$ and -0.049 (1.387) at $p = 0.972$. Panel B evaluates A at the twenty-fifth, fiftieth and seventy-fifth percentiles of M on the estimation sample, which are 0.0102, 0.0569 and 0.1837. The stage-two coefficient on the interaction is -6.954 (9.106) at $p = 0.445$ with M in levels and -0.204 (0.209) at $p = 0.330$ with $\ln M$, and the multinomial interaction does not reject jointly ($\chi^2(2) = 0.84$, $p = 0.658$ in levels and $\chi^2(2) = 0.79$, $p = 0.674$ in logs). The stage-two average marginal effect of A without the interaction is -0.316 (0.164) at $p = 0.053$ on the full sample. Cluster counts: 68 targets in the multinomial models and 52 in stage two on the full sample, falling to 57 and 44 once the bottom quarter of M is dropped. All estimates are associations among imposed trade sanctions.

The exercise answers a measurement question and does not identify a threshold. At the smallest total stakes the two flow shares that form A are themselves measured against near-zero denominators, so the ratio is noisiest where its coefficient is weakest, and classical attenuation would generate this ladder on its own. So would a level of trade below which leverage is not worth exercising. The design cannot separate the two, and the interaction is too imprecise to arbitrate between them. What the table settles is the size-free objection: the association is not carried by relationships in which the two exposures are jointly negligible, and it holds once the index of joint exposure is bounded away from zero. It settles nothing about either country taken alone. A floor on the sum leaves each term free, and one sizeable term can carry a cut point on its own. Of the 125 episodes above the twenty-fifth-percentile cut, 40 have both terms at or above one percent of their own country's exports, and the median episode there pairs a target term of 9.33 percent with a sender term of 0.63 percent (Table 9). The trimmed samples are relationships that matter to the target, and in most of them the sender's own exposure stays small.

C.2 The two stakes entered jointly

One notational warning for this appendix and the log-odds table that precedes it. The symbol R here is the empirical log ratio $\ln(m_3/m_4)$ and has nothing to do with the full-compliance cutoff $R = x/\bar{\theta}^2$ of Section 4.2, and the total stake $M = m_3 + m_4$ has nothing to do with the function Φ of Section 4.3.

The measure sweep enters each exposure measure on its own, so the two stakes that form A have not appeared in the same equation. Table 10 enters them together in logs at stage two, on the episodes where both stakes are strictly positive. That domain is the one Table 7 already uses and leaves 116 stage-two episodes in 52 target clusters.

The sender's stake carries the opposite sign to the target's in both pairs and reaches significance in neither. The sender-stake coefficient is $+0.144$ (standard error 0.100, $p = 0.147$) in the export pair and $+0.112$ (0.078, $p = 0.153$) in the import pair, against target-stake coefficients of -0.143 (0.076) and -0.255 (0.147). A variant that keeps the ten stage-two episodes carrying a zero stake by placing them at half the smallest positive value returns a significant sender coefficient in both pairs. That floor is arbitrary, and the variant is not used here.

With both stakes in logs the index splits into a relative term and a scale term, so the restriction that only relative exposure matters is that the two coefficients sum to zero. The sum is $+0.001$ (0.092, $p = 0.992$) in the export pair and -0.143 (0.124, $p = 0.249$) in the import pair. Neither rejects, so the unrestricted two-stake specification is consistent with the relative-exposure restriction.

The interval on the sum is what that consistency is worth. Doubling both stakes leaves A unchanged and moves the log odds of full compliance by $\ln 2$ times the sum, so the export-pair interval of $[-0.179, 0.181]$ admits anything from a twelve percent fall in the odds of full compliance to a thirteen percent rise. Scale effects large enough to matter substantively sit inside that interval, so the non-rejection is weak evidence. On the export pair the relative term does not move when the scale term is freed, at -0.144 against -0.144 under the restriction, while on the import pair the same non-rejection coexists with the relative term moving from -0.125 under the restriction to -0.183 freed, a change of about 47 percent. Collinearity is not what costs the precision, since the two log stakes correlate at 0.328 and 0.413 and no variance inflation factor in the level specification exceeds 1.21. At stage one the sender-stake coefficient is negative in both pairs and far from significance, so the positive sign is a stage-two pattern.

Table 10: The two stakes of each ratio entered jointly, stage two

Term	Coef. (SE)	p	95 percent interval	N
Panel A: export pair, $A = m_3/(m_3 + m_4)$				
$\ln m_3$, target stake	-0.143 (0.076)	0.060	[-0.293, 0.006]	116
$\ln m_4$, sender stake	0.144 (0.100)	0.147	[-0.051, 0.340]	116
Scale term $b_T + b_S$	0.001 (0.092)	0.992	[-0.179, 0.181]	116
Relative term $(b_T - b_S)/2$	-0.144 (0.076)	0.058	[-0.293, 0.005]	116
Relative term, restriction imposed	-0.144 (0.076)	0.059	[-0.293, 0.005]	116
Panel B: import pair, $A_M = m_1/(m_1 + m_2)$				
$\ln m_1$, target stake	-0.255 (0.147)	0.082	[-0.542, 0.033]	116
$\ln m_2$, sender stake	0.112 (0.078)	0.153	[-0.042, 0.266]	116
Scale term $b_T + b_S$	-0.143 (0.124)	0.249	[-0.385, 0.100]	116
Relative term $(b_T - b_S)/2$	-0.183 (0.100)	0.067	[-0.379, 0.013]	116
Relative term, restriction imposed	-0.125 (0.078)	0.107	[-0.277, 0.027]	116

Notes: The stage-two logit of full compliance given any concession, with the two exposure stakes that form each asymmetry ratio entered together in logs in place of the ratio itself: strict outcome coding, the controls of Table 4, and standard errors clustered by target. b_T is the coefficient on the target's log stake and b_S the coefficient on the sender's log stake. Both stakes in logs give $b_T \ln m_3 + b_S \ln m_4 = \frac{1}{2}(b_T - b_S)\ln(m_3/m_4) + \frac{1}{2}(b_T + b_S)\ln(m_3 m_4)$, so the sum $b_T + b_S$ is the effect of scaling both stakes by a common factor and the relative-exposure restriction is $b_T + b_S = 0$. The last row of each panel is the restricted fit, which is the log-odds cell of Table 7. Intervals are 95 percent, from the same clustered variance matrix as the coefficient. The estimation domain is both stakes strictly positive, which is the domain Table 7 already uses: of the 193 main-sample episodes, 162 carry both stakes strictly positive and 116 of those reach stage two, in 52 target clusters, with 63 full compliers against 53. The scale term is 0.001 (SE 0.092) at $p = 0.992$ for the export pair and -0.143 (SE 0.124) at $p = 0.249$ for the import pair, and neither rejects. Doubling both stakes moves the log odds of full compliance by $\ln 2$ times the scale term, so the intervals above admit odds ratios from 0.883 to 1.134 for the export pair and 0.766 to 1.072 for the import pair. On that same doubling scale the smallest effect the test could detect at 5 percent with 80 percent power is an odds ratio of 1.196 for the export pair and 1.272 for the import pair. The sender-stake coefficient is 0.144 (SE 0.100) at $p = 0.147$ in the export pair and 0.112 (SE 0.078) at $p = 0.153$ in the import pair. At stage one it is -0.107 at $p = 0.179$ in the export pair and -0.044 at $p = 0.527$ in the import pair. A variant that keeps the 10 stage-two episodes carrying a zero stake by placing them at half the smallest positive value returns a sender-stake coefficient of 0.175 (0.084) at $p = 0.036$ for the export pair and 0.173 (0.072) at $p = 0.017$ for the import pair. That floor is arbitrary and those estimates are not used here. On the stage-two estimation sample the two log stakes correlate at 0.328 for the export pair and 0.413 for the import pair. In levels, on the multinomial estimation sample, the correlation is 0.097 for the export pair and 0.030 for the import pair, and no variance inflation factor exceeds 1.21 in the level specification or 1.71 in the log specification. Freeing the scale term moves the relative term by 0.01 percent in the export pair and 46.61 percent in the import pair. All estimates are associations among imposed trade sanctions.

C.3 Network-adjusted dependence

On the 118 episodes where the paper's own ratio is also defined, that ratio rotates too, at -0.098 ($p = 0.067$) on full compliance against -0.088 ($p = 0.123$) when A is estimated on all 166 episodes where it is defined, so what changes the inference is the measure and not the subsample. Retaining the 38 episodes with negative network exposure leaves neither the signed ratio nor the difference form $U_{ts} - U_{st}$ at conventional significance on the 157 episodes that carry a computable exposure matrix.

Excluding those episodes and clearing the forty-event floor are requirements this sample cannot meet at once. The variant carried in the main text keeps them, because the rupture components are non-negative by construction, and dropping every episode the sign rule of A^{net} discards takes the common sample of Table 7 from 144 episodes to 118, whose none and partial categories hold 33 and 37 and fall below the floor. No specification satisfies both criteria, so one of them has to give, and the main text keeps the floor.

The exposure matrix behind Section 3.3 is the real income exposure of Kleinman et al. (2024), whose entry U_{ni} is the elasticity of country n's real income per capita to country i's productivity. It is assembled year by year from the matrix of expenditure shares, the matrix of income shares, and a trade elasticity, with the baseline at $\theta = 4$, the value this paper already uses for the gravity measure δ. The deposited replication package for that framework returned an authorization failure behind an automated-traffic challenge on every access route tried on the build date, so none of the authors' code was used. The equations were read from the published working paper and reimplemented, and everything that rests on the rebuilt matrix comes from that reconstruction and not from a rerun of published code.

Exposure is computed in each year of the window running from three years before an episode to one year before it and then averaged, which is the rule the paper applies to its own measures, and the trade flows are the same composite ladder documented in Appendix A. Domestic absorption is approximated by GDP minus exports, where the framework uses gross output minus exports, which overstates openness relative to their construction and is a documented deviation. Countries that fail the year's admissibility conditions are folded into a single rest-of-world node whose income share has a median of 0.017 and a maximum of 0.078, so the truncation is economically small.

Four checks bound the reconstruction error. The closed form of the exposure matrix reproduces a direct solve of the underlying linear system to 1.1×10^{-16}. The residual left by unbalanced trade is constant across countries to 6.3×10^{-16}, so it is a numeraire shift that leaves the exposures unchanged. The raw export-share ratio rebuilt from the same matrices correlates at 0.958 with the paper's own A, so a difference between the network ratio and A is the network adjustment and not a difference in underlying trade data. And running the estimation file on the paper's own A reproduces the primary results to 3.3×10^{-8} in the coefficients, so the specification is the paper's, unchanged.

Coverage is the binding constraint. Of the 193 main-sample episodes, 157 have a computable exposure matrix, and 36 do not, because a sender or a target sits outside the admissible country set of the relevant year. The rupture ratio A^{rup} is defined on 144 of those and the linearised ratio A^{net} on 119, because the sign rule that A^{net} requires discards 38 episodes whose exposure is negative on one side or the other, 30 on the target's side and 30 on the sender's. The 13 episodes that carry a matrix but no rupture ratio record no bilateral trade in the measurement window, so the rupture loss is zero on both sides and the ratio has no denominator. The 119 episodes where A^{net} exists are selected on the outcome, with full compliance in 41 percent of them against 24 percent of the rest, which is why every estimate is reported beside the paper's own ratio on the same episodes. The two differ by one: A^{net} is computable on 119 episodes and the bilateral ratio on 118 of them.

Table 11 carries the measurement variants, the trade elasticity, and the episodes the sign rule discards.

Table 11: The network-adjusted ratio: measurement variants, elasticity, and the negative-exposure episodes

Measure	ΔP(partial) Estimate (SE)	p	ΔP(full) Estimate (SE)	p	N	Stage two Coef. (SE)	p	N
Panel A: measurement								
Bilateral ratio A, main sample	0.128 (0.063)	0.043	-0.088 (0.057)	0.123	166	-1.418 (0.761)	0.063	118
Bilateral ratio A, on the A^{net} episodes	0.112 (0.063)	0.077	-0.098 (0.054)	0.067	118	-1.822 (1.231)	0.139	85
Network-adjusted ratio A^{net}, $\theta = 4$	0.137 (0.050)	0.006	-0.104 (0.037)	0.004	119	-3.085 (1.299)	0.018	86
Network-adjusted ratio, exact-hat algebra	0.134 (0.050)	0.007	-0.103 (0.037)	0.005	118	-3.083 (1.355)	0.023	85
Rupture ratio A^{rup}	0.117 (0.038)	0.002	-0.064 (0.025)	0.010	144	-2.741 (0.929)	0.003	103
Panel B: trade elasticity								
$\theta = 2$	0.108 (0.053)	0.042	-0.082 (0.040)	0.038	120	-2.661 (1.543)	0.085	87
$\theta = 3$	0.132 (0.047)	0.005	-0.099 (0.034)	0.004	119	-3.104 (1.285)	0.016	86
$\theta = 4$ (baseline)	0.137 (0.050)	0.006	-0.104 (0.037)	0.004	119	-3.085 (1.299)	0.018	86
$\theta = 6$	0.136 (0.052)	0.008	-0.106 (0.038)	0.006	118	-3.070 (1.363)	0.024	85
$\theta = 8$	0.137 (0.053)	0.009	-0.108 (0.039)	0.006	118	-3.046 (1.371)	0.026	85
Panel C: negative-exposure episodes retained								
Signed ratio, all episodes with a matrix	0.025 (0.029)	0.377	-0.030 (0.019)	0.105	157	-1.364 (1.032)	0.186	111

Notes: Each row is the baseline specification of Table 4 with the dependence measure swapped and nothing else changed. The first four columns give the change in the predicted probability across the measure's own 25th to 75th percentile, with N the multinomial estimation sample, and the last three give the stage-two logit coefficient for full compliance given any concession on the concession subsample, with its own N. A^{net} is built from the linearised real income exposure matrix and A^{rup} from a bilateral-rupture counterfactual in which the pair's trade costs are set prohibitive both ways, whose two components are non-negative by construction, so it is defined on 144 episodes against 119 for A^{net} and correlates at 0.650 with A. The exact-hat row replaces the linearisation with the framework's exact-hat algebra at a ten percent productivity shock, and the two measures correlate at 0.99998. Across the elasticity grid the correlation of the network ratio with A stays between 0.706 and 0.715. The signed ratio keeps the 38 episodes whose network exposure is negative, where third-market competition dominates the direct trade channel, so its denominator may cross zero and it runs from -1.932 to 3.331 on the main sample. Dropping instead every episode that the sign rule of A^{net} discards takes the common sample of Table 7 from 144 episodes to 118, with outcome counts 33, 37 and 48, two of them below the forty-event floor. On that sample A^{rup} gives 0.118 (0.036), $p = 0.001$ on partial compliance and -0.078 (0.022), $p = 0.000$ on full, against 0.112 (0.063), $p = 0.077$ and -0.098 (0.054), $p = 0.067$ for A. No sample both excludes those episodes and clears the floor, so the floor rule binds and Table 7 is estimated on the larger sample. The difference form $U_{ts} - U_{st}$ on the same 157 episodes gives average marginal effects per standard deviation of 0.008 (0.040), $p = 0.833$ on partial compliance and -0.020 (0.041), $p = 0.616$ on full compliance. Standard errors are clustered by target. All estimates are associations among imposed trade sanctions.

The linearised ratio gives a larger interquartile rotation than the rupture ratio, at $+0.137$ ($p = 0.006$) on partial compliance and -0.104 ($p = 0.004$) on full, with a stage-two coefficient of -3.085 ($p = 0.018$), and it sits on a sample whose none and partial categories hold 33 and 37 episodes. Replacing the linearisation with the framework's exact-hat algebra changes nothing, since the two measures correlate at 0.99998 and no estimate moves by more than 0.003. Nothing depends on the trade elasticity either: across a grid from 2 to 8 the interquartile change on partial compliance runs from $+0.108$ to $+0.137$ and on full compliance from -0.082 to -0.108, and the correlation of the network ratio with A stays between 0.706 and 0.715. Keeping the negative-exposure episodes costs the result its precision. The signed ratio, whose denominator may cross zero and which runs from -1.93

to 3.33, gives -0.030 ($p = 0.105$) on full compliance, and the difference form $U_{ts} - U_{st}$ gives -0.020 ($p = 0.616$) per standard deviation on the same 157 episodes. Neither rejects, and neither refutes the rotation.

The requirements collide on those episodes. Excluding every episode the sign rule discards takes the common sample of Table 7 from 144 episodes to 118, whose none and partial categories hold 33 and 37 and therefore breach the forty-event floor. On that smaller sample the rupture ratio gives $+0.118$ ($p = 0.001$) on partial compliance and -0.078 ($p < 0.001$) on full, against $+0.112$ ($p = 0.077$) and -0.098 ($p = 0.067$) for the paper's ratio, so the comparison does not turn on those episodes. No sample both excludes them and clears the floor, and the floor rule is the one that binds in the main text.

Table 12 turns to the exposure levels and to the two measures entered jointly.

Table 12: Exposure levels: raw bilateral shares against network exposure

	P(none)		P(partial)		P(full)		
Measure	AME (SE)	p	AME (SE)	p	AME (SE)	p	N
Panel A: raw bilateral shares, one at a time							
m_3 target export dependence, main sample	-0.054 (0.035)	0.124	0.011 (0.037)	0.774	0.043 (0.033)	0.189	180
m_3 target export dependence, A^{net} episodes	-0.022 (0.049)	0.661	0.039 (0.041)	0.336	-0.017 (0.048)	0.717	119
m_4 sender export dependence, main sample	0.059 (0.074)	0.429	-0.191 (0.162)	0.238	0.133 (0.094)	0.160	181
Panel B: network exposure levels, one at a time							
U_{ts} target to sender	0.015 (0.045)	0.734	0.002 (0.041)	0.960	-0.017 (0.040)	0.669	157
U_{st} sender to target	0.082 (0.034)	0.017	-0.154 (0.048)	0.001	0.072 (0.042)	0.081	157
D_t^{rup} rupture loss, target	0.009 (0.044)	0.837	-0.001 (0.041)	0.990	-0.009 (0.041)	0.832	157
D_s^{rup} rupture loss, sender	0.086 (0.033)	0.010	-0.157 (0.043)	0.000	0.072 (0.043)	0.099	157
Panel C: two measures entered jointly							
Network-adjusted ratio A^{net}	-0.037 (0.059)	0.525	0.144 (0.076)	0.057	-0.107 (0.078)	0.172	118
with bilateral ratio A	0.011 (0.069)	0.878	0.018 (0.068)	0.796	-0.028 (0.083)	0.735	118
Rupture ratio A^{rup}	-0.093 (0.051)	0.069	0.150 (0.054)	0.005	-0.058 (0.055)	0.295	144
with bilateral ratio A	0.034 (0.069)	0.620	0.012 (0.046)	0.788	-0.047 (0.065)	0.475	144
U_{ts} network exposure	0.117 (0.078)	0.132	-0.034 (0.077)	0.657	-0.083 (0.068)	0.223	156
with m_3 export dependence	-0.152 (0.097)	0.119	0.061 (0.084)	0.466	0.091 (0.068)	0.182	156
Panel D: difference form							
$U_{ts} - U_{st}$	0.012 (0.044)	0.787	0.008 (0.040)	0.833	-0.020 (0.041)	0.616	157

Notes: Average marginal effects per standard deviation of the measure, under the baseline specification of Table 4 with the dependence measure swapped and nothing else changed, standard errors clustered by target. Panels A, B and D enter one measure at a time. Panel C enters two standardised measures in the same multinomial logit and reports both, so each pair occupies two rows. U_{ts} is the elasticity of the target's real income per capita to the sender's productivity and U_{st} its mirror image, both from the reconstructed exposure matrix, and D^{rup} is the corresponding loss under a bilateral rupture counterfactual. At the level of the individual exposure the network and raw measures are close: U_{ts} correlates at 0.803 with m_3 and at 0.947 with the gravity exposure measure δ, and U_{st} at 0.821 with m_4. The network adjustment does its work on the ratio. All estimates are associations among imposed trade sanctions.

The paper's null on dependence levels survives on the target's side of the relationship and fails on the sender's. Neither the target's network exposure U_{ts} nor its rupture counterpart moves any outcome, which is what the raw export share m_3 also does. The sender's network exposure does move outcomes, lowering partial compliance by 0.154 per standard deviation ($p = 0.001$) and raising the probability of no concession by 0.082 ($p = 0.017$), with the rupture counterpart giving -0.157 and $+0.086$, where the raw sender share m_4 returns -0.191 at

$p = 0.238$. The partial and full signs are the ones the ratio implies, since a larger sender stake lowers A. The no-concession effect is the one the two estimators of Section 3.3 disagree on. What the network adjustment adds at the level is on the sender's side, and the claim that dependence levels are uninformative holds for bilateral shares and not for exposure in general. Entered jointly, the network 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 adds little predictive information, which is all the joint specification establishes. The same pairing of the target's network exposure with its raw export share leaves neither term significant.

D Bounding the Political-Entry Probability

Proposition 1 puts $\Pr(\text{any} \mid \lambda) = \pi_A F_k(k_A(\lambda))$, and $F_k \leq 1$ gives $\pi_A \geq \Pr(\text{any} \mid \lambda)$ at every λ. The population object is $\pi_A \in [L, 1]$ with $L = \sup_\lambda \Pr(\text{any} \mid \lambda)$, and L is unknown.

The largest of a set of binned proportions is not an estimator of L. Choosing the biggest of eight cells of about twenty episodes each is a selection problem, and in this sample it inflates the bound by more than fifteen percentage points. The construction used instead is a simultaneous lower band. Let $\text{Lo}(\cdot)$ satisfy $\Pr\{\Pr(\text{any} \mid \lambda) \geq \text{Lo}(\lambda) \text{ for all } \lambda\} \geq 0.95$. Then $\sup_\lambda \text{Lo}(\lambda) \leq L \leq \pi_A$ with at least that coverage, so $\sup_\lambda \text{Lo}(\lambda)$ is a valid one-sided lower confidence bound for π_A.

The band comes from a cluster bootstrap over targets, matching the clustering used throughout the paper, with 1,000 resamples multiplied to 4,000 for the quantiles and a sup-t critical value so the coverage is joint across bins. That critical value is 2.33 against 1.645 pointwise.

The support matters. The log-odds transform $\lambda = A/(1 - A)$ is undefined at $A = 0$, which holds for three episodes, and at $A = 1$, which holds for one. Binning on $\ln \lambda$ drops those four. Binning on A retains all 166 and estimates the same supremum, since A is a monotone transform of λ. The choice is not innocuous. On the reduced support the bound is 0.749, and on the full support it is 0.667. Four episodes are enough to change which bin attains the maximum, and because the estimand is a supremum that difference passes straight into the bound. The full-support figure is the one reported in Section 4.4 and the reduced-support value is superseded.

	primary	under flatness	raw maximum
$\ln \lambda$ support, $N = 162$	0.749	0.652	0.905
full support of A, $N = 166$	**0.667**	0.647	0.857

The primary column imposes nothing on the shape of $\Pr(\text{any} \mid \lambda)$. The second imposes a constant extensive margin, which a slope test does not reject ($p = 0.446$) and does not establish, so it is reported as the restricted reading and not the headline.

No value of π_A is treated as an estimate. By Proposition 5 the signs of the economic comparative statics do not depend on where in the admissible range the truth lies. The magnitudes do, and by a wide margin: Appendix E shows that the model reproduces the observed movements only when π_A sits close to $\Pr(\text{any})$.

E The Quantitative Region

This appendix asks whether the model of Section 4 can move the outcome shares by as much as the estimates do. The answer has two layers, and they should not be run together.

The first is the intensive margin, and two claims about it have to be kept apart. Proposition 4 needs no restriction on politics at all: the ratio $\Pr(\text{full})/\Pr(\text{any}) = F_k(R)/F_k(k_A)$ cancels π_A whatever its value and even if it moves with leverage. Proposition 5 is what licenses reading the signs of ε_R and ε_{k_A} off the unconditional slopes, and it does require the political-entry exclusion. Nothing below touches either.

The second is quantitative, and it does depend on the latent decomposition of no concession. Reproducing the observed *level* of the extensive margin as well as its movement pins $F_k(k_A) = \Pr(\text{any})/\pi_A$, and the exercise below shows that the model has quantitative capacity only when π_A sits near $\Pr(\text{any})$. No parameterisation is put forward as preferred, and none is tuned to reproduce the estimated movements.

Design. The calibration targets are the outcome shares of the same 166 episodes that generate the quartiles and contrasts of A, since those quartiles exist only where A is observed: no concession 0.289, partial 0.331, full 0.380. The political-entry probability is normalised to $1 - \Pr(\text{none}) = 0.711$, and the economic full-compliance cutoff is placed where the split then requires it, $F_k(R) = 0.380/\pi_A = 0.534$. The reporting threshold is set so that the economic any-concession cutoff sits at the 99.9th percentile of the type distribution, which is what makes $f_k(k_A)k_A$ negligible. The sender's problem of Proposition 3 is then solved at the empirical quartiles of relative leverage, $\lambda = 1.198$ and $\lambda = 63.269$, and the movement over that range is a prediction.

Normalising π_A is not an estimate of it. Section 4.4 bounds π_A from below and does not identify it, and the equality $\Pr(\text{none}) = (1 - \pi_A) + \pi_A[1 - F_k(k_A)]$ means the normalisation assigns the whole extensive margin to the political component up to the 0.001 left by the reporting threshold. The signs required by the observed slopes do not depend on that assignment, by Proposition 5. The sender's own comparative statics are a different matter: π_A enters both first-order conditions of Proposition 3. The magnitudes below are conditional on it, and the sensitivity at the end of this appendix shows how much they move across admissible values.

What is swept. The type law and its parameters over log-logistic, Weibull, gamma and lognormal families. The prize V runs over $\{0.25, 1, 4\}$ and the ambition cost $K(\bar{\theta}) = \kappa\bar{\theta}^b$ over sixteen values of κ from 0.001 to 0.5 in log steps and $b \in \{1.5, 2.5\}$. That is 4,992 cells.

Result. Six hundred and sixteen cells put the equilibrium cutoff where the observed split requires it. Of those, **38 reproduce the signs and the magnitudes** at a cutoff placed at the 99.9th percentile, and 54 at the 99.95th, taking a match to mean that both the partial and the full movements fall within a factor of two of the estimated +0.13 and −0.09 and that the no-concession movement stays within 0.05. The matching cells span thirteen of the sixteen values of κ, from 0.001 to 0.144, three type families, and both cost curvatures.

Sensitivity to the normalisation. Varying π_A while holding the observed moments fixed means imposing both

$$F_k(k_A) = \frac{\Pr(\text{any})}{\pi_A} \qquad \text{and} \qquad F_k(R) = \frac{\Pr(\text{full})}{\pi_A},$$

so a larger political-entry probability forces the economic any-concession cutoff down out of the tail. Repeating the sweep under that restriction gives

π_A	$F_k(k_A)$	$F_k(R)$	cells placed	cells matching
0.711	1.000	0.534	616	54
0.75	0.948	0.506	653	6
0.85	0.836	0.446	671	0
1.00	0.711	0.380	524	0

The counts trace the quantitative frontier of the mechanism. They are not a verdict on the model, which rests on Propositions 4 and 5, but a statement of how far the political decomposition can be pushed before the architecture stops accounting for both margins at once. At $\pi_A = \Pr(\text{any})$ the restriction sets $F_k(k_A) = 1$, which a distribution on $(0, \infty)$ cannot attain, so the cutoff is capped in the far tail and the count depends on where: 38 cells at the 99.9th percentile and 54 at the 99.95th.

The decomposition says why the region collapses. Under the political-entry exclusion,

$$\frac{\partial \ln \Pr(\text{any})}{\partial \ln \lambda} = \underbrace{\frac{k_A f_k(k_A)}{F_k(k_A)}}_{\text{mass at the extensive cutoff}} \varepsilon_{k_A},$$

and the sweep puts ε_{k_A} between $+0.70$ and $+0.91$ at every π_A tried. Leverage moves the economic any-concession cutoff by a good deal. The flatness therefore has to come from the other factor, and it does: among matching cells the density term $k_A f_k(k_A)$ runs from 0.0006 to 0.0076 at $\pi_A = 0.711$ and from 0.064 to 0.139 at $\pi_A = 0.75$, twenty times larger.

That restricts the political side of the model. Since

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

a negligible density at the cutoff means $F_k(k_A) \approx 1$ and hence $\Pr(\text{none}) \approx 1 - \pi_A$: nearly all of non-concession has to come from political inadmissibility. The restriction is easier to read as a split of the observed 28.9 percent of episodes with no concession into a political part $1 - \pi_A$ and an economic part $\pi_A[1 - F_k(k_A)]$.

π_A	political share of no concession	economic share	economic as a fraction of all none
0.711	0.289	0.000	0.00
0.75	0.250	0.039	0.14
0.85	0.150	0.139	0.48
1.00	0.000	0.289	1.00

At $\pi_A = 0.85$ about sixteen percent of politically willing targets would concede nothing for economic reasons, which is close to half of all no-concession episodes. That is where the matching region disappears. The cutoff has left the tail, and because leverage moves it the extensive margin moves with it. At $\pi_A = 1$ every no-concession episode would be economic and no parameterisation delivers a flat Pr(any) alongside a declining Pr(full | any).

The reading is therefore narrower than a bare non-identification result. The model can reproduce both margins together when most non-accommodation reflects political inadmissibility. It does not provide a quantitative account of the flat extensive margin when a substantial share of no concession is attributed to economic resistance. That is a partial falsification of the high-π_A versions of this architecture, and it is informative in its own right.

Two cautions on how far this can be pushed. The value 0.711 is not the lower endpoint of the identified set: Section 4.4 reports a lower *confidence* bound of 0.667, the population endpoint $L = \sup_\lambda \Pr(\text{any} \mid \lambda)$ is unknown, and it may lie above 0.711. And $\pi_A = \Pr(\text{any})$ implies $F_k(k_A) = 1$, which a distribution with unbounded support reaches only in the limit, so the first row of the table is a limiting case approached through the tail and not a point the model attains.

Two things follow, and they are worth keeping apart. The intensive result does not depend on any of this: by Proposition 4 the ratio of full to any compliance cancels π_A whatever its value, and by Proposition 5 the signs of the comparative statics are those of the observed slopes for every admissible π_A. What does depend on it is the tail explanation for the flat extensive margin, which requires political admissibility to account for nearly all of non-concession. If a substantial share of no-concession episodes were economic instead of political, the model would need a different reason for the extensive margin to be flat, and this exercise does not supply one.

One further limit. The band of matching κ reaches the lower bound of the grid, so it may extend further down. The matches are spread over thirteen values with the highest density in the interior instead of piled at the boundary, which is not the pattern of a corner solution, but the bound is a feature of the grid and is stated as such.

F The Threat-Stage Benchmark

This appendix collects the model the paper began with and no longer uses as its account of the data. In it the damage $L(m)$ and the demand $\bar{\theta}$ are primitives and not choices, with L differentiable and strictly increasing in leverage, $L'(m) > 0$, which is what makes the cutoff of Proposition 1 move with leverage at all, a target that resists is sanctioned with a credibility determined in equilibrium, and settlement at the threat stage removes prospective full compliers before the imposed sample is formed. That channel is what Section 3.4 tests and does not confirm. The benchmark is kept because it states the hypothesis that test is aimed at, and because its scale-free case is a useful reference point. It is not the model of Section 4, and where the two disagree the political-entry model is the one the paper defends.

Lemma 1 (Threat-stage binarization)**.** *For $\pi < 1$ the threat stage is binary in equilibrium: a target either settles by conceding $\bar{\theta}$ in full, or concedes nothing and awaits the imposition* *lottery, and interior threat-stage concessions are strictly dominated. The argument prices the no-sanction state, which carries probability $1-\pi$. At $\pi = 1$ imposition is certain, an early concession below θ_I^* is topped up to θ_I^* at the same total cost, and interior concessions are payoff-equivalent and not dominated, so the settle-or-resist representation holds there weakly.*

The settle-or-resist comparison then pins down a cutoff in the normalized type u. Settling costs $k\bar{\theta}^2/2$ for sure. Resisting costs, in expectation, π times the sunk damage plus the optimized loss under pressure. Writing $\pi^* \equiv 1/(1+2\sigma)$:

Proposition 1 (Settlement cutoff and threat-stage selection)**.** *There is a unique cutoff $\hat{u}(\pi,\sigma)$ such that the target settles if and only if $u \le \hat{u}$, where*

$$\hat{u}(\pi,\sigma) = \begin{cases} \dfrac{2\pi\sigma}{1-\pi}, & \pi < \pi^*, \\ \pi(1+\sigma) + \sqrt{\pi^2(1+\sigma)^2 - \pi}, & \pi \ge \pi^*, \end{cases}$$

continuous at π^* with $\hat{u}(\pi^*,\sigma) = 1$. The cutoff in k is $\hat{k} = \hat{u}\, r$, strictly increasing in leverage, in credibility π, and in the sunk share σ, and $\hat{u} \to 0$ as $\pi \to 0$, and as $\sigma \to 0$ at any fixed $\pi < 1$. At $\pi = 1$ the upper branch applies for every σ and $\hat{u} \to 1$ instead. Settlement is full compliance, so every settler exits before imposition. Consequently:

- *(i) the imposed sample is the truncation $k > \hat{k}$, whose type distribution first-order stochastically dominates the population distribution (imposed targets are selected toward resolute types),*
- *(ii) if $\pi \ge \pi^*$, every type that would comply in full under imposition settles at the threat stage, and full compliance is* never *observed among imposed cases,*
- *(iii) if $\pi < \pi^*$, the imposed full compliers are exactly the types $k \in (\hat{k}, r]$: targets that gambled on non-imposition and lost.*

Threats work only because imposition destroys something that later compliance cannot recover: as $\sigma \to 0$ the settle set vanishes, since a target loses nothing by waiting out the threat and conceding under pressure. Observed full compliance under imposed sanctions is then a symptom of *imperfect credibility.* Once credibility reaches π^* the model reproduces the classic result that would-be successes are resolved before imposition (Eaton and Engers, 1992; Drezner, 2003; Lacy and Niou, 2004), and the imposed sample contains partial resisters only. Below that threshold the imposed sample still contains full compliers, the types that gambled on non-imposition and lost.

Credibility is an equilibrium object. Given the resisting pool, the sender imposes when its cost draw is below the expected concession value, and the pool in turn depends on the credibility targets assign to the threat.

Lemma 2 (Equilibrium credibility)**.** *The map $\pi \mapsto H\big(w\,\bar{\theta}\,\mathbb{E}[\min\{1, 1/u\} \mid u > \hat{u}(\pi,\sigma)]\big)$ is decreasing and continuous, so an equilibrium π exists and is unique. Suppose in addition that the fixed point is interior, that H is differentiable with $H' > 0$ at the equilibrium argument, and that Θ is differentiable. If the elasticity $kf_k'(k)/f_k(k)$ is strictly decreasing (as for log-normal, gamma, Weibull, and log-logistic types), higher leverage shifts the resisting pool's* *normalized types down in the likelihood-ratio order, raises the value of imposing, and so raises equilibrium credibility: $\pi'(m) > 0$. Continuity and monotonicity on their own give a weakly nondecreasing $\pi(m)$ and no derivative at all, since a credibility distribution that is flat at the equilibrium argument leaves π unmoved however the value of imposing shifts. If instead F_k is Pareto and both cutoffs stay above its lower support point $k_{\min}$ over the leverage range considered, the resisting pool is scale free, the value of imposing does not vary with m, and $\pi'(m) = 0$.*

Because every resister faces the same imposition probability, conditioning on imposition reweights nothing beyond the truncation at $\hat{k}$. For a compliance margin $z \in (0, \bar{\theta}]$, the observed share of imposed cases conceding at least z is

$$Q_z \;=\; \Pr\bigl(\theta_I^* \geq z \mid \text{imposed}\bigr) \;=\; \frac{F_k(c_z L) - F_k(bL)}{1 - F_k(bL)}, \qquad c_z = \frac{1}{z\,\bar{\theta}}, \quad b = \frac{\hat{u}}{\bar{\theta}^2}, \tag{1}$$

with the full-compliance margin at $z = \bar{\theta}$ (so $c_{\bar{\theta}} L = r$) and the any-concession margin at $z = \tau$. The expression is a probability only where the conditioning event is nonempty, that is where $c_z L \geq bL$. Where $c_z L < bL$ every imposed type concedes less than z and $Q_z = 0$. On the branch $\pi < \pi^*$ used below, $\hat{u} < 1$ gives $bL < r = c_{\bar{\theta}} L \leq c_z L$ for every $z \leq \bar{\theta}$, so the restriction binds only on the branch $\pi \geq \pi^*$, where Proposition 1(ii) already puts observed full compliance at zero. Let $g(x) \equiv x f_k(x)/\bigl(1 - F_k(x)\bigr)$ denote the survival elasticity of the type distribution. Differentiating (1) along m gives the exact decomposition

$$\frac{d\,\ln\bigl(1 - Q_z\bigr)}{dm} \;=\; \underbrace{g(bL)\,\frac{\hat{u}'(m)}{\hat{u}}}_{\text{selection}} \;-\; \underbrace{\frac{L'(m)}{L}\Bigl[g(c_z L) - g(bL)\Bigr]}_{\text{pressure}}. \tag{2}$$

Proposition 2 (Two margins of leverage)**.** *Suppose $0 < \pi < \pi^*$, $\pi'(m) \geq 0$ and $0 < \tau < \bar{\theta}$.*

(a) Unconditionally*, leverage raises compliance: the settlement share $F_k(\hat{k})$, expected total concessions $\mathbb{E}[\theta]$, and the probability of any concession are all nondecreasing in m, and strictly increasing wherever $f_k(\hat{k}) > 0$. Positive credibility is what the strict conclusion needs. At $\pi = 0$ the cutoff is $\hat{k} = 0$ at every m, nothing settles, no resister is ever sanctioned, and all three quantities are identically zero.*

(b) Among imposed cases*, the sign of dQ_z/dm is the sign of the pressure term net of the selection term in (2). Selection pushes every margin down, because higher credibility clears more would-be compliers out at the threat stage.*

(c) If g is strictly increasing and the selection intensity is moderate, in the sense that

$$g(r) - g(bL) \;<\; \frac{L}{L'}\, g(bL)\, \frac{\hat{u}'}{\hat{u}} \;<\; g(c_\tau L) - g(bL),$$

then higher leverage lowers *the observed full-compliance share and* raises *the observed any-concession share among imposed cases,*

$$\frac{dQ_{\bar{\theta}}}{dm} < 0 < \frac{dQ_\tau}{dm}.$$

The interval is nonempty because $\tau < \bar{\theta}$ gives $c_\tau > c_{\bar{\theta}}$ and g is increasing. At the boundary $\tau = \bar{\theta}$ the two endpoints coincide and the interval is empty, which is why the hypotheses exclude it.

(d) Scale-free benchmark. *If F_k is Pareto with lower support point $k_{\min}$, if π is constant in m, and if $\hat{k} = \hat{u}r \geq k_{\min}$ over the leverage range considered, so that every argument at which the survival function is evaluated lies inside the Pareto support, then every imposed-sample share Q_z is invariant to leverage: $Q_z = 1-(\hat{u}z/\bar{\theta})^\alpha$, with no dependence on L at all. The support condition is not cosmetic. Once $\hat{k}$ falls below $k_{\min}$ the survival function is one at that argument, the powers of L no longer cancel, and scale invariance fails: with $k_{\min} = 1$, $\alpha = 1$, $\sigma = 1/2$, $\pi = 1/3$ and $\bar{\theta} = 1$, the observed full-compliance share is zero at $L = 0.1$ and one half at $L = 2$.*

G Additional Figures

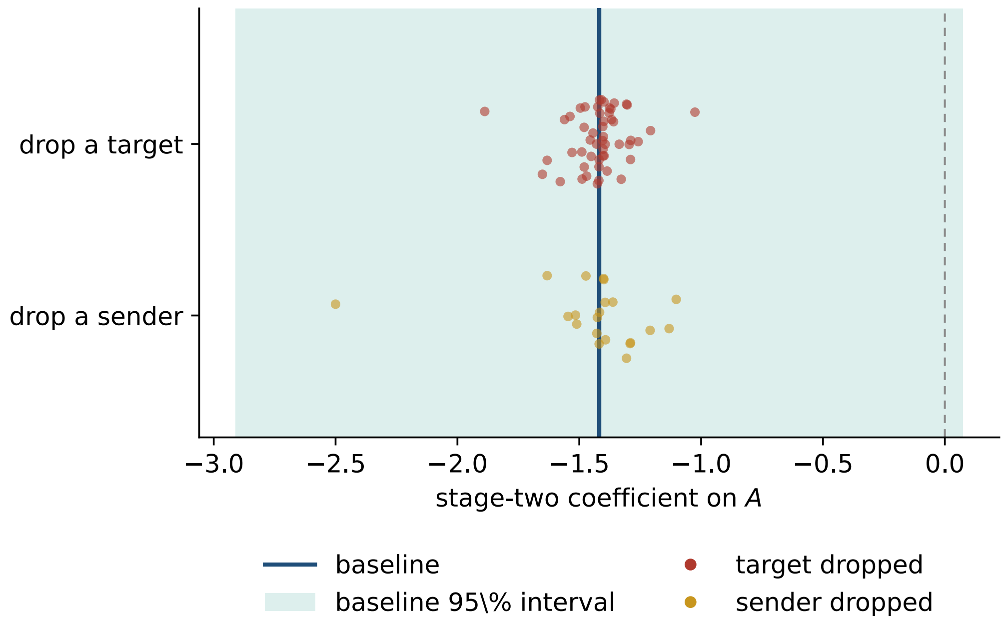

(a) Stage-two coefficient

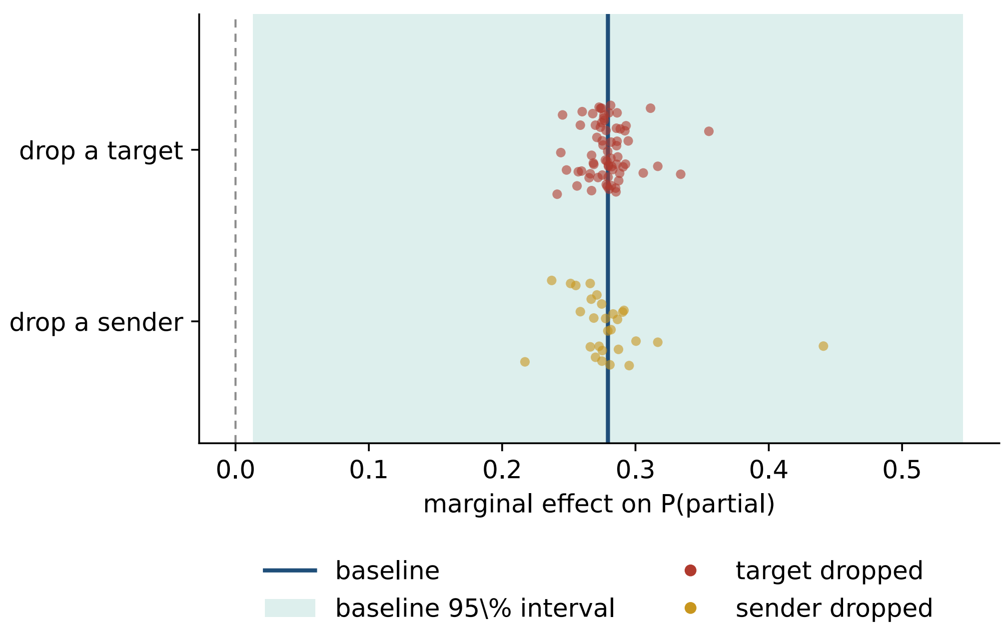

(b) Effect on partial compliance

Figure 2: Leave-one-out estimates against the baseline. Each point is a refit that drops one target or one sender. The vertical line is the baseline and the shaded strip its 95 percent interval. Every refit keeps the baseline sign, 73 of them for the stage-two coefficient and 96 for the partial-compliance effect, and all sit well inside the baseline interval. That interval still covers zero in panel (a), where the stage-two coefficient carries $p = 0.063$, and does not in panel (b), where the partial-compliance effect carries $p = 0.040$. That is the whole content of the exercise: the pattern is not the work of particular episodes, and dropping them buys no precision.

These figures support the main text without carrying its argument. The leave-one-out refits document stability and not a separate result, and the interquartile contrasts report by measure what the body reports for the asymmetry ratio alone.

G.1 Predicted outcome probabilities from the three-category model

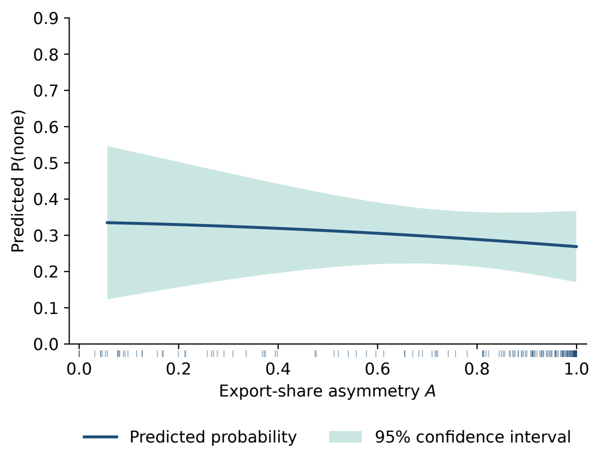

(a) P(none)

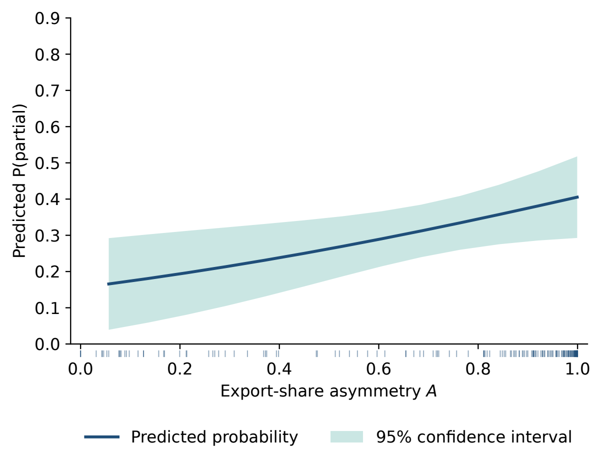

(b) P(partial)

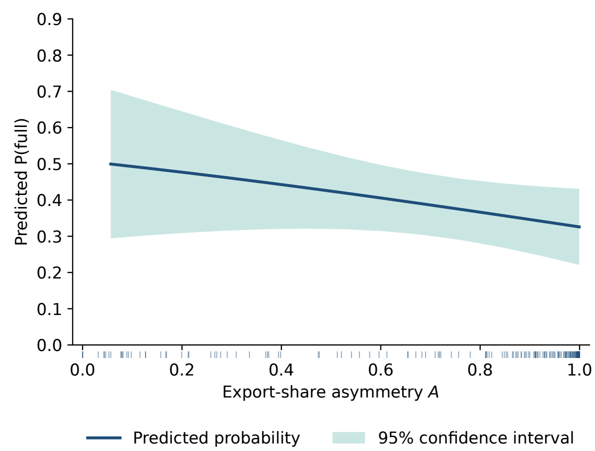

(c) P(full)

Figure 3: Predicted outcome probabilities across the observed support of the asymmetry ratio A. Predictions from the baseline multinomial logit of Table 4, which applies the strict outcome coding, the sanction-history control, and standard errors clustered by target, evaluated at thirteen evenly spaced points between the estimation-sample 5th and 95th percentiles of A, other covariates at observed values, with 95 percent confidence bands. The probability of no concession is flat across the range of relative exposure, the probability of partial compliance rises, and the probability of full compliance falls.

G.2 Interquartile contrasts by leverage measure

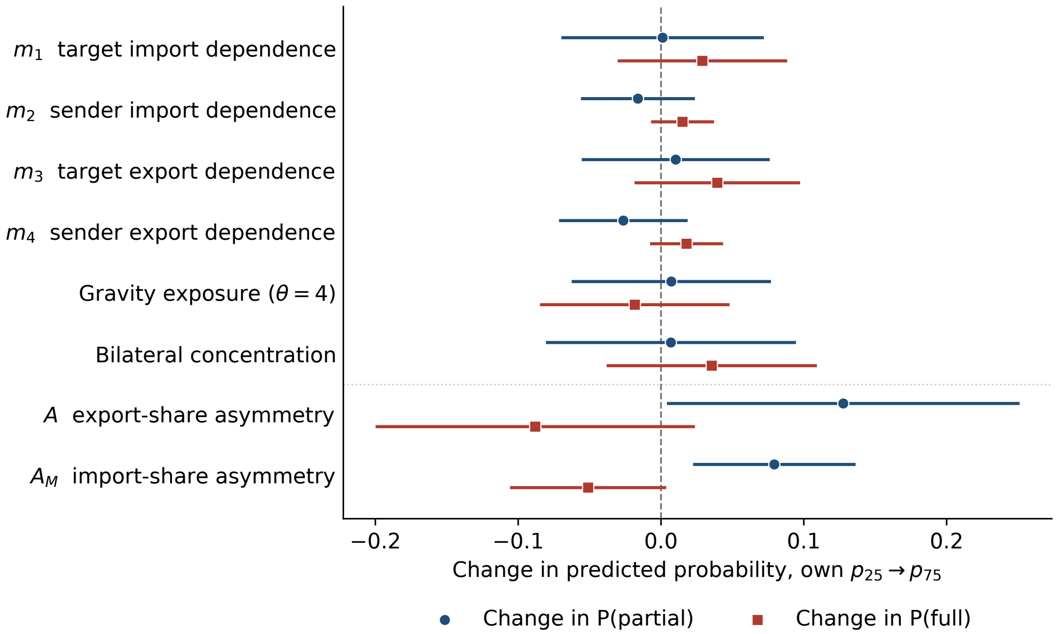

Figure 4: Interquartile contrasts by leverage measure. Change in the predicted probability of partial compliance and of full compliance for a move from each measure's own 25th to its own 75th percentile, with 95 percent confidence intervals, under the baseline specification. Interquartile probability changes provide a common metric across measures that live on different supports. The bilateral level measures show no detectable association on either margin in the pooled sample. The two asymmetry ratios produce the rotation, positive on partial compliance and negative on full compliance. Eight of the nine distinct measures are drawn, ten rows if the mechanical complement is retained: outside options is one minus bilateral concentration, so the two carry the same contrast with opposite sign and only concentration appears.

H Threat-Stage Tabulations

Table 13 reports the TIES tabulations summarized in Section 3.4, by stage of resolution and for the trade-instrument subset.

Table 13: TIES: outcomes by stage of resolution, 1945–2005

	Resolved at threat stage	Resolved after imposition
Panel A: outcome rates within each subsample		
Cases with a coded final outcome	473	551
Any success (acquiescence or negotiated)	$262/473 = 55.4\%$	$314/551 = 57.0\%$
Acquiescence only (strict)	$202/473 = 42.7\%$	$182/551 = 33.0\%$
Complete acquiescence	$146/473 = 30.9\%$	$121/551 = 22.0\%$
Partial acquiescence	$56/473 = 11.8\%$	$61/551 = 11.1\%$
Negotiated settlement	$60/473 = 12.7\%$	$132/551 = 24.0\%$
Complete acquiescence as share of successes	$146/262 = 55.7\%$	$121/314 = 38.5\%$
Panel B: share of all coercion successes occurring at the threat stage		
All cases, relaxed definition	$262/576 = 45.5\%$	
All cases, strict acquiescence only	$202/384 = 52.6\%$	
Trade-instrument cases, relaxed	$126/336 = 37.5\%$	
Trade-instrument cases, strict	$86/199 = 43.2\%$	
US primary sender, relaxed	$138/289 = 47.8\%$	
Panel C: trade-instrument cases, complete acquiescence		
Share of resolved cases	$57/226 = 25.2\%$	$75/360 = 20.8\%$
Share of successes	$57/126 = 45.2\%$	$75/210 = 35.7\%$

Notes: Computed from the public TIES v4 case file (Morgan et al., 2014; 1,412 cases). Stage of resolution follows the final-outcome code (codes 1–5 threat stage, 6–10 after imposition). Success is target acquiescence (partial or complete) plus negotiated settlement. Sender capitulation and stalemate are failures. The 262/314 split of the 576 relaxed-definition successes reproduces the published count exactly. Trade instrument means embargo, partial embargo, import restriction, or export restriction, threatened or imposed. TIES is a separate dataset from the GSDB episodes used everywhere else in this paper, with different coverage and coding rules, and these tabulations are raw shares with no controls.

TIES codes final outcomes by stage for 1,412 sanction cases from 1945 to 2005, including cases that were threatened and never imposed. It 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 entries are uncontrolled shares and not estimates.

I Proofs

The proofs of the political-entry model come first. Those of the threat-stage benchmark of Appendix F follow, and are unchanged.

Proof of Proposition 1. By Proposition 2 a target with $A_i = 1$ concedes $\theta^*(k) = \min\{x/(k\bar{\theta}), \bar{\theta}\}$, so it complies fully exactly when $k \le R = x/\bar{\theta}^2$ and concedes at least the reporting threshold exactly when $x/(k\bar{\theta}) \ge \tau$, that is when $k \le k_A = x/(\tau\bar{\theta})$. A target with $A_i = 0$ concedes nothing whatever k is. The admissibility threshold $\underline{a}$ does not depend on exposure, so $\pi_A = \Pr(a \ge \underline{a})$ is a constant with respect to $(x, \bar{\theta})$. Independence of a and k then gives $\Pr(\text{full}) = \Pr(A_i = 1)\Pr(k \le R) = \pi_A F_k(R)$ and $\Pr(\text{any}) = \pi_A F_k(k_A)$, and partial compliance is the difference because the two events are nested, $R \le k_A$ whenever $\tau \le \bar{\theta}$. $\square$

Proof of Proposition 2. The objective $\frac{1}{2}k\theta^2 + x(1-\theta/\bar{\theta})$ is strictly convex in θ with derivative $k\theta - x/\bar{\theta}$, so the unconstrained minimiser is $\theta = x/(k\bar{\theta})$ and it is unique. It exceeds the upper bound $\bar{\theta}$ exactly when $k < x/\bar{\theta}^2 = R$, in which case the constraint binds and $\theta^* = \bar{\theta}$. At $k = R$ the two coincide. On the interior branch $\partial\theta^*/\partial k = -x/(k^2\bar{\theta}) < 0$, $\partial\theta^*/\partial x = 1/(k\bar{\theta}) > 0$ and $\partial\theta^*/\partial\bar{\theta} = -x/(k\bar{\theta}^2) < 0$, and $\theta^* \to 0$ as $k \to \infty$. ☐

Proof of Proposition 3. Integrating the concession over types,

$$\mathbb{E}[\theta^*] = \int_0^R \bar{\theta}\, dF_k(k) + \int_R^\infty \frac{x}{k\bar{\theta}}\, dF_k(k) = \bar{\theta}\Big[F_k(R) + R\int_R^\infty \frac{dF_k(k)}{k}\Big] = \bar{\theta}\,\Psi(R),$$

using $x/\bar{\theta} = R\bar{\theta}$. Differentiating Ψ and cancelling the two terms in $f_k(R)$,

$$\Psi'(R) = f_k(R) + \int_R^\infty \frac{dF_k(k)}{k} - R\,\frac{f_k(R)}{R} = \int_R^\infty \frac{dF_k(k)}{k} > 0.$$

With $R = x/\bar{\theta}^2$ we have $\partial R/\partial x = 1/\bar{\theta}^2$ and $\partial R/\partial\bar{\theta} = -2R/\bar{\theta}$, so

$$\Pi_x = \pi_A\Big[\frac{w\Psi'(R)}{\bar{\theta}} + \frac{Vf_k(R)}{\bar{\theta}^2}\Big] - \frac{1}{\lambda}, \qquad \Pi_{\bar{\theta}} = \pi_A\Big[w\Psi(R) - 2wR\Psi'(R) - \frac{2Vf_k(R)R}{\bar{\theta}}\Big] - K'(\bar{\theta}),$$

and setting both to zero gives the stated conditions. Only two objects in the problem involve the exposure pair: the cost of pressure, which is x/λ by the construction of Section 4.3, and nothing else. Hence Π depends on (m_T, m_S) only through λ, and replacing (m_T, m_S) by (am_T, am_S) for $a > 0$ leaves the objective, its maximisers x^* and $\bar{\theta}^*$, the cutoffs R and k_A and every share in Proposition 1 unchanged. ☐

Proof of Proposition 4. Divide the second expression of Proposition 1 by the first. Both are positive by hypothesis, so the quotient is defined, and the factor π_A appears in both and cancels at each λ separately, so the ratio is $F_k(R(\lambda))/F_k(k_A(\lambda))$ whether or not π_A is constant in λ. No restriction on $\pi_A(\cdot)$ is used. ☐

Proof of Proposition 5. Take logarithms of the two shares in Proposition 1 and differentiate in $\ln\lambda$. Under $\partial\ln\pi_A/\partial\ln\lambda = 0$ the constant drops out and

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

and both factors are strictly positive by the density assumption. Strict monotonicity of F_k alone would not do: a density may vanish at isolated points while the distribution still increases strictly, and where $f_k(R) = 0$ the observed slope is zero whatever ε_R is, so the sign is not recovered. The parametric families of Appendix E all have densities positive at their cutoffs. The signs of ε_R and ε_{k_A} therefore equal those of the observed slopes for any admissible π_A and any such F_k. From $k_A/R = \bar{\theta}/\tau$, taking logarithms and differentiating gives $\varepsilon_{\bar{\theta}} = \varepsilon_{k_A} - \varepsilon_R$. From $x = R\bar{\theta}^2$ it gives $\varepsilon_x = \varepsilon_R + 2\varepsilon_{\bar{\theta}} = 2\varepsilon_{k_A} - \varepsilon_R$. Hence

$$\tfrac{1}{2}\varepsilon_x < \varepsilon_{\bar{\theta}} \iff \varepsilon_{k_A} - \tfrac{1}{2}\varepsilon_R < \varepsilon_{k_A} - \varepsilon_R \iff \varepsilon_R < 0,$$

and $\varepsilon_{\bar{\theta}} < \varepsilon_x \iff \varepsilon_{k_A} - \varepsilon_R < 2\varepsilon_{k_A} - \varepsilon_R \iff \varepsilon_{k_A} > 0$. Without the exclusion the first display carries an extra additive term $\partial\ln\pi_A/\partial\ln\lambda$ and the sign of ε_R is not recoverable from Pr(full) alone. ☐

The benchmark's target problem. The benchmark of Appendix F uses the same target optimisation as Proposition 2 with the primitive damage L in place of the chosen pressure x, so no separate proof is needed: the first-order condition $k\theta = L/\bar{\theta}$ gives the interior solution $L/(k\bar{\theta})$, which reaches the corner $\bar{\theta}$ exactly at $k = L/\bar{\theta}^2 = r$, and every statement below is read with that substitution.

Proof of Lemma 1. Consider a resister conceding $t \in (0, \bar{\theta})$ early. If $t \le \theta_I^*(k)$ the sanction-state outcome is unchanged (the target tops up to θ_I^* and total cost depends only on the total level), while in the no-sanction state, reached with probability $1 - \pi > 0$, the early concession is a pure loss $kt^2/2$. Expected cost is strictly increasing in t on this range, and for $t > \theta_I^*(k)$ it is strictly increasing as well (the derivative is $kt - \pi L/\bar{\theta} > 0$ there). Finally, conceding $\bar{\theta}$ *as a resister* costs $\pi\sigma L$ more than settling, because settlement alone avoids the sunk damage. So the optimum is at $t = 0$ or settlement at $\bar{\theta}$. □

Proof of Proposition 1. For $u \le 1$ the resist cost is $\pi(\sigma L + k\bar{\theta}^2/2)$ and the comparison reduces to $u \le 2\pi\sigma/(1-\pi)$. For $u > 1$ it is $\pi(\sigma L + L - L^2/(2k\bar{\theta}^2))$ and the comparison reduces to $u^2 - 2\pi(1+\sigma)u + \pi \le 0$, which holds between the two roots, the larger being the expression above. Direct algebra shows $2\pi\sigma/(1-\pi) \ge 1$ exactly when $\pi \ge \pi^*$ and that both branches equal one at π^*. The upper branch needs two cases, since its discriminant $\pi^2(1+\sigma)^2 - \pi$ is negative for $\pi < (1+\sigma)^{-2}$ and $(1+\sigma)^{-2} < \pi^*$. For $\pi < (1+\sigma)^{-2}$ the quadratic has no real root and, opening upward, is positive everywhere, so no type with $u > 1$ settles. For $(1+\sigma)^{-2} \le \pi < \pi^*$ the roots are real and both lie below one, because the larger root increases in π and equals one at π^*, so again no type with $u > 1$ settles. Either way the settle set is the interval $(0, \hat{u}]$ with $\hat{u}$ as stated. Monotonicity follows from $\partial_\pi\big(2\pi\sigma/(1-\pi)\big) = 2\sigma/(1-\pi)^2 > 0$ and the corresponding derivative on the upper branch, which is positive on $[\pi^*, 1]$. Truncation from below implies survival functions ordered as $\big(1 - F_k(k)\big)/\big(1 - F_k(\hat{k})\big) \ge 1 - F_k(k)$ for $k \ge \hat{k}$, which is (i). Parts (ii) and (iii) restate $\hat{u} \ge 1$ and $\hat{u} < 1$ using Proposition 2. □

Proof of Lemma 2. Write $\Theta(\hat{u}, r) = \mathbb{E}[\min\{1, 1/u\} \mid u > \hat{u}]$ for the value of the resisting pool, so that the best-response map is $\Phi(\pi) = H\big(w\bar{\theta}\,\Theta(\hat{u}(\pi,\sigma), r)\big)$ with $r = L(m)/\bar{\theta}^2$. Truncating away the lowest normalized types removes the largest values of $\min\{1, 1/u\}$, so Θ is decreasing in $\hat{u}$, and $\hat{u}$ is continuous and increasing in π (Proposition 1). With H continuous, Φ is decreasing and continuous and maps $[0, 1]$ into itself, so $\pi - \Phi(\pi)$ is strictly increasing, nonpositive at $\pi = 0$ and nonnegative at $\pi = 1$, and the fixed point exists and is unique.

For the sign in leverage, the normalized type $u = k/r$ has density $r f_k(ru)$, so for $r_2 > r_1$ the likelihood ratio $r_2 f_k(r_2 u)/\big(r_1 f_k(r_1 u)\big)$ has log derivative $\big[\eta(r_2 u) - \eta(r_1 u)\big]/u$ in u, where $\eta(k) = k f_k'(k)/f_k(k)$. If η is strictly decreasing this log derivative is strictly negative, so the normalized types at r_2 are smaller than those at r_1 in the likelihood-ratio order, and that ordering is inherited by the distributions truncated to the resisting pool $\{u > \hat{u}\}$. Since $u \mapsto \min\{1, 1/u\}$ is decreasing, $\Theta_r > 0$. At an interior fixed point the implicit-function theorem gives $\pi'(m) = \Phi_m/(1 - \Phi_\pi)$ with $1 - \Phi_\pi \ge 1$, because $\Phi_\pi = h \cdot w\bar{\theta}\,\Theta_{\hat{u}}\hat{u}_\pi \le 0$, and with $\Phi_m = h \cdot w\bar{\theta}\,\Theta_r\, L'(m)/\bar{\theta}^2 > 0$, so $\pi'(m) > 0$. If F_k is Pareto then η is constant, the likelihood ratio does not vary with u, and the pool-value computation at the end of this appendix gives Θ free of r on both branches, so $\Phi_m = 0$ and $\pi'(m) = 0$. □

Proof of Proposition 2. (a) With $\pi > 0$ and L strictly increasing the cutoff $\hat{k} = \hat{u}(\pi(m), \sigma)L(m)/\bar{\theta}^2$ is nondecreasing in m and strictly increasing wherever $\hat{u} > 0$, and each type crossing it moves from expected concession $\pi\theta_I^*(\hat{k}) < \bar{\theta}$ to $\bar{\theta}$, which raises the settlement share strictly wherever $f_k(\hat{k}) > 0$ and weakly otherwise. Among remaining resisters θ_I^* rises pointwise in L and the imposition probability is nondecreasing. (b) is (2), whose selection term is nonnegative for every z whenever $\hat{u}' \geq 0$, with $\hat{u}' = \hat{u}'(\pi)\pi'(m) \geq 0$ by Proposition 1 and the hypothesis $\pi'(m) \geq 0$, which Lemma 2 delivers under its differentiability conditions and not from continuity alone. (c) applies the sign rule at $z = \bar{\theta}$ and $z = \tau$. (d) substitutes $1 - F_k(x) = (k_{\min}/x)^\alpha$ into (1), which is legitimate only at arguments above $k_{\min}$: under the support condition the powers of L cancel between numerator and denominator, and without it the substitution is invalid at the lower argument and the cancellation fails. □

Pareto pool value used in Lemma 2. For $1-F_k(k) = (k_{\min}/k)^\alpha$ with cutoffs above $k_{\min}$, direct integration gives $\Theta = 1 - \hat{u}^\alpha/(\alpha+1)$ for $\hat{u} \leq 1$ and $\Theta = \alpha/\big((\alpha+1)\hat{u}\big)$ for $\hat{u} > 1$, continuous at $\hat{u} = 1$, decreasing in $\hat{u}$, and free of r on both branches, which is the scale-freeness the Pareto case of Proposition 2(d) rests on. □

References

Drezner, D. W. (2003). The hidden hand of economic coercion. *International Organization*, 57(3):643–659.

Eaton, J. and Engers, M. (1992). Sanctions. *Journal of Political Economy*, 100(5):899–928.

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.

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.