One-Sided Violence in Two-Sided Wars

Yuri M. Zhukov

Georgetown University

May 1, 2026

Abstract

When and where do two-sided armed conflicts become one-sided? In contested territories, one combatant often sustains coercive violence persistently and unilaterally while its rival restrains. This asymmetry can manifest as wartime terrorism, repression, or other patterns of one-sided violence during conflict. We argue that two conditions jointly produce such behavior: targeting selectivity and outside options. A combatant sustains unilateral escalation when its violence is precise enough to extract civilian compliance, or when it can survive without civilian cooperation (e.g., through external recruits, lootable resources, or expeditionary forces). We develop this argument theoretically and derive a scalar measure of coercive asymmetry — a combatant's behavioral commitment to sustained unilateral punishment — that researchers can estimate from any dyadic conflict panel. We validate our claims with cross-national data spanning dozens of modern armed conflicts and with declassified Soviet secret police records on counterinsurgency in western Ukraine (1944–1950).

Armed conflicts are not always as two-sided as their labels suggest. In wars where both sides field organized forces, one combatant often escalates persistently and unilaterally while its rival restrains. The phenomenon takes many forms: wartime state repression (Schubiger, 2021), terrorism in civil war (Stanton, 2016), indiscriminate counterinsurgency (Lyall, 2009), mass atrocities against civilians (Eck and Hultman, 2007). This asymmetry is among the most common structural features of armed conflict, with direct bearing on how wars unfold (Leites and Wolff, Jr., 1970; Mack, 1975; Arreguin-Toft, 2001), who pays their costs (Davenport et al., 2019; Anderton and Brauer, 2021), and the line between two-sided warfare and one-sided terror (Besley and Persson, 2009; Rohner et al., 2025).

The scope of our study extends to any setting where two organized armed actors simultaneously claim authority over the same civilian population and neither holds a monopoly on the legitimate use of force. This category encompasses insurgencies, civil wars, military occupations, and expeditionary conflicts, and excludes settings where one side has effectively consolidated control.[1] We ask why, when both sides compete over the same population, one finds sustained unilateral violence worthwhile and the other does not.

Existing literatures illuminate pieces of the puzzle, but no existing theory directly explains the dyadic asymmetry. The rationalist bargaining tradition can generate inferences about coercive persistence (e.g., Slantchev, 2003; Powell, 2012), but treats the continuation of fighting — not the allocation of coercive effort between two rivals — as its core explanandum.[2] Research on the dynamics of civil conflict (e.g., Kalyvas, 2006; Salehyan et

[1]This context distinguishes our subject from the large literature on state repression (e.g., Davenport, 1995; Ritter and Conrad, 2016; Bueno de Mesquita et al., 2024), which models a coercive relationship between a state and its (usually unarmed) domestic opposition, with no organized rival making a simultaneous claim on the same population.

[2]Asymmetric escalation can be an equilibrium outcome in bargaining models when one actor possesses private information about its resolve or capabilities (Fearon, 1995), when a shifting power distribution creates commitment problems (Powell, 2012), or when one side's private beliefs about relative strength take longer to update (Slantchev, 2003; Wolford et al., 2011). However, predictions about asymmetric persistence are byproducts of these models, not primary objects of explanation. The civilian population, furthermore, is not an optimizing actor in their equilibrium logic.

al., 2014; Blair, 2023) has identified key mechanisms that bear on the puzzle (e.g., information, outside options), but tends to analyze each combatant's violence in isolation rather than modeling the strategic interaction that produces the asymmetry.[3] A third strand of research, in political economy (e.g., Besley and Persson, 2009, 2011), has shown that the same structural factors can drive one- and two-sided violence, but explains cross-national selection into violence regimes rather than intra-conflict asymmetries between combatants.[4]

We argue that two conditions jointly produce persistent coercive asymmetry: targeting selectivity and outside options. A combatant sustains unilateral punishment when its violence is precise enough to extract civilian compliance, or when it can afford to forgo that compliance entirely. How a combatant delivers punishment determines whether one-sided violence succeeds. When a combatant strikes selectively, targeting enemy collaborators and sparing bystanders, civilians can protect themselves by cooperating with the punishing side. When a combatant strikes indiscriminately, civilians absorb harm regardless of their loyalties, and cooperation offers no protection. Indiscriminate violence can still degrade the opponent's support base, and combatants with robust outside options (e.g., external recruitment, lootable resources, expeditionary forces) can afford the backlash that indiscriminate punishment provokes. A combatant turns to one-sided violence because it

[3] The information-economy literature shows that a combatant's ability to sustain selective coercion depends on distinguishing combatants from noncombatants (Kalyvas, 2006) through ethnic ties (Lyall, 2010), service provision (Berman et al., 2011), organizational capacity (Hoover Green, 2016), or technological surveillance (Gohdes, 2020). The outside-options literature shows that when a combatant decouples its survival from local civilian cooperation — through access to lootable or external resources (Wood, 2010; Salehyan et al., 2014), or border access (Blair, 2023) — it severs the disciplinary feedback loop that would otherwise constrain its violence. These mechanisms identify the conditions under which a single actor can sustain one-sided force. Yet neither literature models what happens when two rivals, facing the same civilian population, simultaneously differ in their informational capacity or outside options.

[4] Besley and Persson (2009, 2011) model peace, (one-sided) repression, and (two-sided) civil war as ordered equilibrium outcomes of a common rent-seeking game. Rohner et al. (2025) extend this logic empirically. These models illuminate the common roots of one- and two-sided violence, but abstract from three features central to our puzzle: civilian agency, within-conflict variation, and the dyadic structure in which two combatants simultaneously allocate coercive effort. As Anderton and Brauer (2021) observe, formal models of strategic interaction between perpetrators and civilian targets remain rare.

expects escalation to produce submission, or because it does not care if it sparks defiance.

We formalize this logic as a stochastic game between two armed rivals competing for a civilian population. From this model, we derive a scalar measure of *coercive asymmetry*, a combatant's behavioral commitment to sustained one-sided violence. We show how to estimate this measure from any dyadic conflict panel data. Our measure goes beyond existing empirical efforts by capturing the dynamic, dyadic nature of coercive interaction.

Our empirical analysis proceeds in two stages. First, we estimate coercive asymmetry across a cross-national sample of armed conflicts, spanning civil wars, counterinsurgencies, occupations, and expeditionary wars. We show that our theoretical predictions hold across conflict types and are not artifacts of any single case. Second, we validate these patterns with micro-level data from a single, well-documented conflict — the Soviet campaign against the Ukrainian Insurgent Army (1944–1950) — where archival completeness reduces the reporting and misclassification concerns that affect cross-national event data.

We make three contributions to the literature on armed conflict. First, we derive the conditions under which two-sided conflicts produce persistent one-sided violence, by either armed actor. Second, we introduce a scalar measure of coercive asymmetry that researchers can estimate from standard dyadic conflict data — identifying which conflicts are relatively symmetric, which are asymmetric, and in whose direction. Third, we show that coercive asymmetry shifts predictably within conflicts over space and time, in ways that country-year approaches cannot capture. Both empirical contributions rest on a two-stage design that pairs cross-national breadth with archival validation.

1 Theory of coercive asymmetry

Our theory applies to armed conflicts in which two organized combatants compete for the loyalty of a civilian population that neither side fully controls. The combatants may be states or non-state actors; the conflict may be an insurgency, civil war, military occupation, or colonial counterinsurgency. The framework requires contested sovereignty, where neither combatant holds a monopoly on legitimacy or violence in the contested area, and civilians face real costs from defying either side. We refer to the two combatants as Side 1 and Side 2, where Side 1 is the incumbent (the armed actor exercising greater de facto administrative control at the outset) and Side 2 is the challenger. We summarize our theoretical model's logic here, and present its full structure and proofs in Appendix A1.

The conflict involves three sets of actors: a civilian population and two combatants (Sides 1 and 2). Each combatant's campaign depends to some degree on civilian cooperation (e.g., intelligence, shelter, labor, material support) although the degree of dependence varies. To secure that cooperation, each combatant chooses whether to restrain or escalate punishment.[5] The expected mechanism is coercive. By raising the cost of supporting the opponent, a combatant hopes to redirect civilian cooperation toward itself. Punishment reaches civilians through two channels. One is selective: the combatant identifies and targets cooperators of the opposing side, making support for that side more costly. The other is indiscriminate: violence inflicts collateral damage on nearby civilians regardless of their allegiances, raising the cost of being present rather than the cost of taking sides.[6] Indiscriminate violence also hurts some enemy collaborators, albeit by chance, at high cost to neutral civilians, and without the deterrent effect of selective violence.

How heavily a combatant depends on local civilian support shapes its incentives to escalate. Combatants with robust outside options (e.g., external recruitment, lootable resource revenue, materiel sourced outside the contested locality) can absorb the cooperation loss that indiscriminate punishment tends to provoke.[7] Combatants with few outside options

[5]We denote punishment levels L (low) and H (high). Each combatant $k \in \{1, 2\}$ selects $\rho_k \in \{L, H\}$.

[6]Formally, combatant k hunts identified enemy cooperators at rate $\lambda_k \rho_k$, where $\lambda_k \in (0, 1)$ is selectivity and ρ_k is punishment level; collateral damage accrues at rate $(1 - \lambda_k)\rho_k$.

[7]We capture outside options with the parameter $c_k \geq 0$.

face stronger incentives to keep the population on side.

Civilians, in turn, are not a homogeneous mass. Each civilian occupies a distinct exposure profile: some face greater risk from selective targeting, because their location, occupation, or social network marks them as potential collaborators. Others face greater risk from indiscriminate fire, perhaps because they live near infrastructure combatants routinely strike or lack the means to evacuate.[8] Combatants know the population's general composition but not which specific civilians belong to which exposure category.

How civilians respond to punishment depends on how precisely that punishment discriminates.[9] When Side 1's violence is precise enough to punish Side 2's collaborators, while sparing bystanders and loyalists, the relative costs of cooperating with Side 2 rise, and civilians protect themselves by cooperating with Side 1: submission. When Side 1's violence is too blunt to spare bystanders or loyalists, cooperation with the punishing side no longer offers protection, and civilians instead cooperate with Side 2: defiance.

Combatants' choices to escalate or restrain hinge on whether their targeting technologies are precise enough to make escalation pay.[10] Each combatant has a selectivity threshold: a minimum share of violence that must befall enemy collaborators and not bystanders, in order to deter cooperation with the enemy. A combatant whose selectivity meets or

[8]We index civilian j's exposure type by $\theta_j \in (0,1)$: low θ_j denotes greater vulnerability to selective targeting, high θ_j to collateral damage. Combatants observe distribution $F(\theta \mid x_i)$, where x_i are observable locality characteristics, but cannot assign types to individuals before punishing (see Assumption 1).

[9]Civilian j chooses $e_j^* \in [0,1]$, the share of cooperation directed toward Side 1, to maximize

$$U_j(e_j) = -\big[\theta_j(1-\lambda_1)\rho_1 + (1-\theta_j)\lambda_2\rho_2\big]e_j - \big[\theta_j(1-\lambda_2)\rho_2 + (1-\theta_j)\lambda_1\rho_1\big](1-e_j) - \frac{\gamma}{2}\big(e_j - \tfrac{1}{2}\big)^2, \quad (1)$$

where $\gamma > 0$ governs the curvature of civilian utility (equivalently, the cost of deviating from perfect neutrality $e_j = 1/2$). The first-order condition gives $\partial e_j^*/\partial\rho_1 = (\lambda_1 - \theta_j)/\gamma$, which is positive (submission) when $\lambda_1 > \theta_j$ and negative (defiance) when $\lambda_1 < \theta_j$ (see Proposition 1).

[10]Combatant k chooses $\rho_k \in \{L, H\}$ to maximize

$$U_k(\rho_k, \rho_{-k}) = b \cdot \mathbb{E}[e_k^*(\rho_k, \rho_{-k}, \theta_j) \mid x_i] + c_k(1-\lambda_k)\rho_k - \kappa_k(\rho_k), \quad (2)$$

where $b > 0$ is the value of a unit of civilian cooperation, e_k^* is the share of cooperation that k receives, $c_k \geq 0$ is k's outside option, and punishment cost $\kappa_k(\cdot)$ satisfies $\kappa_k(L) = 0$ and $\kappa_k(H) > 0$. Taking the exceeds this threshold will always prefer to escalate, no matter what the opponent does. A less selective combatant will always prefer to restrain, for the same reason.[11]

When both sides fall short of their selectivity thresholds, mutual restraint follows. When both clear them, mutual escalation results. The more revealing cases are the asymmetric ones, where one side escalates and the other restrains: one-sided violence. These asymmetries arise not because either side miscalculates, but because they face the same population with different technologies and reach different conclusions about whether escalation pays.

Armed conflicts are not static. Front lines move, populations change, surveillance capacity evolves. As these conditions shift, so does each side's calculus about whether escalation pays. At any point in time and space, the two combatants' behavior falls into one of four patterns: mutual restraint, mutual escalation, or unilateral escalation by Side 1 or 2.[12] Which of these patterns will emerge in the future depends on which one currently prevails: the chances of mutual restraint tomorrow may be different for localities already experiencing mutual restraint as opposed to one-sided violence or mutual escalation.[13]

Each combatant's behavior reduces to two probabilities: the chance it escalates when

difference $U_k(H,\cdot) - U_k(L,\cdot)$ and signing it yields a threshold. k strictly prefers high punishment when

$$\lambda_k > \lambda_k^\dagger \;\equiv\; \bar{\theta}_i + \gamma \frac{\kappa_k(H) - c_k(1-\bar{\theta}_i)\Delta}{(b - c_k\gamma)\Delta}, \tag{3}$$

where $\bar{\theta}_i = \mathbb{E}[\theta_j \mid x_i]$ is the mean civilian type in locality i, and $\Delta > 0$ is the punishment increment (see Theorem 1). The threshold falls as c_k rises: outside options expand the region where escalation dominates.

[11] Aggregating the civilian response over $F(\theta \mid x_i)$ shows that whether escalation attracts or repels cooperation depends only on λ_k relative to $\bar{\theta}_i$. The outside-options term likewise involves only k's own parameters. Each side's calculus is therefore independent of the opponent's choice, making the game dominance-solvable with a unique Nash equilibrium (see Theorem 1).

[12] We model these dynamics as a stochastic game, where combatants re-optimize punishment as the strategic environment shifts, rather than executing a fixed long-run plan. Within any period, each combatant's action is independent of what the opponent does (Theorem 1). But across periods, the opponent's most recent action serves as an observable signal of how conditions have shifted. Combatants condition on this signal to update their assessment of the current environment, not to react strategically within a period. This time-varying structure gives the model empirical traction: the sequence of observed punishment choices carries recoverable information about which state prevails and why.

[13] The observed sequence of joint punishment profiles forms a first-order Markov chain, collecting tran-

the opponent restrained last period, and the chance it restrains when the opponent escalated last period. Together, these determine how durable each form of one-sided violence is. We define *coercive asymmetry* ω as how much more persistent one-sided violence by Side 1 is than one-sided violence by Side 2.[14] When ω is positive, Side 1 escalates persistently while Side 2 restrains. When ω is negative, Side 2 escalates while Side 1 restrains. When ω is zero, neither form of one-sided violence dominates. Positive ω does not imply that Side 2 uses no force, only that it is less likely to initiate or sustain unilateral escalation.

So where and when should we expect one-sided violence, and by which side? Our theory yields several testable predictions about how ω varies across localities and conflicts.[15] Where Side 1's violence is more selective relative to Side 2's, escalation generates a larger cooperation premium for Side 1, and Side 1 becomes more likely to use one-sided violence

sition probabilities in a 4×4 matrix $\mathbf{M} = [m_{a'a}]$, where $m_{a'a} = P(a' \mid a)$ and $\mathcal{A} = \{LL, LH, HL, HH\}$:

$$\mathbf{M} = \begin{array}{c} \\ LL \\ LH \\ HL \\ HH \end{array} \!\!\begin{array}{c} \begin{array}{cccc} LL & LH & HL & HH \end{array} \\ \begin{pmatrix} m_{LL,LL} & m_{LL,LH} & m_{LL,HL} & m_{LL,HH} \\ m_{LH,LL} & m_{LH,LH} & m_{LH,HL} & m_{LH,HH} \\ m_{HL,LL} & m_{HL,LH} & m_{HL,HL} & m_{HL,HH} \\ m_{HH,LL} & m_{HH,LH} & m_{HH,HL} & m_{HH,HH} \end{pmatrix} \end{array} \tag{4}$$

We reduce $\mathbf{M}$'s 12 free parameters (one normalizing constraint per row) to four by restricting each combatant to condition its current action on the opponent's *previous* action. This assumes military decisions reflect adaptation to observed enemy behavior (Smith et al., 2000). For each side k, define $p_k \equiv P(\rho_{kt} = L \mid \rho_{-k,t-1} = L)$ and $q_k \equiv P(\rho_{kt} = L \mid \rho_{-k,t-1} = H)$. Under this restriction, combatants' current actions are conditionally independent and all $m_{a'a}$ become products of p_k and q_k (Proposition 2):

$$\mathbf{M} = \begin{array}{c} \\ LL \\ LH \\ HL \\ HH \end{array} \!\!\begin{array}{c} \begin{array}{cccc} LL & LH & HL & HH \end{array} \\ \begin{pmatrix} p_1p_2 & q_1p_2 & p_1q_2 & q_1q_2 \\ p_1(1-p_2) & q_1(1-p_2) & p_1(1-q_2) & q_1(1-q_2) \\ (1-p_1)p_2 & (1-q_1)p_2 & (1-p_1)q_2 & (1-q_1)q_2 \\ (1-p_1)(1-p_2) & (1-q_1)(1-p_2) & (1-p_1)(1-q_2) & (1-q_1)(1-q_2) \end{pmatrix} \end{array} \tag{5}$$

[14]Formally, $\omega_1 \equiv m_{HL,HL} - m_{LH,LH} = (1-p_1)q_2 - q_1(1-p_2)$ (Proposition 2). ω_2 reverses the roles of the two sides. For brevity, we drop the subscript and write $\omega \equiv \omega_1$ unless we state otherwise. Positive ω requires $p_1 \approx 0$ (Side 1 escalates when Side 2 restrains) and $q_2 \approx 1$ (Side 2 restrains under pressure).

[15]These follow from Corollary 1 and comparative statics on $\lambda_k^\dagger$.

(i.e., ω rises). Where Side 1 has better outside options, it escalates in hopes of degrading the opponent's base through attrition, even when doing so alienates the population; again, ω rises. Where escalating is costlier for Side 1 (e.g., due to international pressure, logistical constraints, political opposition), Side 1 tilts toward restraint and ω falls. The effect of civilian vulnerability on coercive asymmetry depends on selectivity: higher vulnerability pushes ω upward if Side 1 is more selective, and downward if it is less selective. Estimating ω and how it varies with observable proxies for these parameters is our main empirical test.

2 Estimation

We empirically estimate coercive asymmetry ω in two steps. First, we model the joint sequence of combatant decisions using dyadic conflict data. Second, we recover predicted transition probabilities from that model, and assemble them into an estimate $\widehat{\omega}$.

We observe the two combatants' decisions to escalate or restrain as pairs of binary outcomes across localities and time periods.[16] We assume that these decisions are potentially correlated, in that unobserved local and time-specific shocks affecting both sides (e.g., unit rotations, agricultural cycles, population flows) might simultaneously push both combatants toward escalation or restraint.[17] If we ignore this dependence and estimate

[16]We index combatants $k = 1, 2$, localities i, and periods t. We define $Y_{k,it} = 1$ if combatant k escalates and $Y_{k,it} = 0$ otherwise. Observed outcomes emerge from latent indices via $Y_{k,it} = \mathbf{1}[Y^*_{k,it} \geq 0]$. A zero does not mean no violence occurred. Event datasets record only the most visible acts, which survive archival declassification or publication by media outlets. Violence that available measurement instruments cannot detect goes unrecorded. We treat this observation mechanism as the boundary between restraint and escalation: low punishment is empirically invisible, high punishment is not.

[17]We allow for this by modeling the two latent indices as jointly normal:

$$\begin{bmatrix} Y^*_{1,it} \\ Y^*_{2,it} \end{bmatrix} \sim N_2\left(\begin{bmatrix} \mu_{1,it} \\ \mu_{2,it} \end{bmatrix}, \begin{bmatrix} 1 & r \\ r & 1 \end{bmatrix} \right),$$

where r is a scalar cross-equation correlation parameter. We test whether the dependence is statistically negligible with likelihood ratio tests, and retain the joint specification where it is not.

each combatant's behavior in isolation, we could misstate the uncertainty around $\widehat{\omega}$.

In our empirical model, each combatant's propensity to escalate depends on three inputs: the opponent's most recent action, its own most recent action, and an interaction between the two.[18] The opponent's past action raises or lowers the baseline propensity to escalate directly. It also conditions how strongly a combatant's own past behavior carries forward: facing an escalating opponent shifts the weight each side places on its own prior action.[19] We estimate this system by bivariate probit, controlling for pre-conflict local attributes, with corrections for unobserved locality-specific and period-specific confounders.[20]

We recover $\widehat{\omega}$ from the fitted model's predicted transition probabilities. For each combatant, we compute two sets of predicted probabilities: the probability of escalation after the opponent restrained last period, and the probability of restraint after the opponent escalated.[21] Our estimate $\widehat{\omega}$ is the difference between the probability that Side 1 sustains the former configuration (unilateral escalation) and the probability that Side 1 sustains the latter (unilateral restraint).[22] We obtain confidence intervals by parametric bootstrap.[23]

[18]The estimating equations take the form, suppressing locality subscript i,

$$\begin{bmatrix}\mu_{1,t}\\ \mu_{2,t}\end{bmatrix} = \begin{bmatrix} y_{2,t-1}\zeta_1 + y_{1,t-1}\big(\alpha_1 + y_{2,t-1}\phi_1\big) + \mathbf{x}_{1,t}\beta_1 \\ y_{1,t-1}\zeta_2 + y_{2,t-1}\big(\alpha_2 + y_{1,t-1}\phi_2\big) + \mathbf{x}_{2,t}\beta_2 \end{bmatrix}, \tag{6}$$

where $y_{k,t-1}$ is combatant k's lagged outcome, $\mathbf{x}_{k,t}$ are vectors of covariates including pre-conflict locality attributes, regional, year, and month fixed effects; α_k, β_k, ζ_k, ϕ_k are regression coefficients.

[19]In equation (6), the effect of $y_{1,t-1}$ on μ_1 equals α_1 when Side 2 restrained last period and $\alpha_1 + \phi_1$ when Side 2 escalated; ϕ_k captures this moderating shift.

[20]We augment the estimating equations with locality-level and period-level time-means of each side's lagged behavior to partial out time-invariant heterogeneity and aggregate period-specific shocks (Mundlak, 1978; Wooldridge, 2010). We use Mundlak-corrected specifications for general inference about $\widehat{\omega}$, and the base interactive specification for comparative statics tests on time-invariant covariates. We use likelihood ratio tests to select between the fully interactive model and more restricted variants, including non-strategic (opponent lags excluded, $\zeta_k = \phi_k = 0$) and additive (interaction excluded, $\phi_k = 0$).

[21]For combatant k: $(1-\hat{p}_k) \equiv \widehat{P}(Y_k = 1 \mid y_{-k,t-1} = 0)$, the predicted probability of escalation when the opponent restrained, and $\hat{q}_k \equiv \widehat{P}(Y_k = 0 \mid y_{-k,t-1} = 1)$, the probability of restraint when the opponent escalated. The diagonal entries of transition matrix $\widehat{\mathbf{M}}$ are $\widehat{m}_{HL,HL} = (1-\widehat{p}_1)\hat{q}_2$, $\widehat{m}_{LH,LH} = \widehat{q}_1(1-\hat{p}_2)$.

[22]Formally, $\widehat{\omega} \equiv \widehat{m}_{HL,HL} - \widehat{m}_{LH,LH} = (1-\hat{p}_1)\hat{q}_2 - \hat{q}_1(1-\hat{p}_2)$.

[23]Since $\widehat{\omega}$ combines predictions from both equations in (6), propagating uncertainty through the system requires joint simulation. We draw from the joint coefficient distribution and recompute $\widehat{\omega}$ at each draw.

To test the comparative statics predictions, we apply the same procedure under counterfactual covariate values, shifting each variable from its observed minimum to its maximum while holding all others at their median or modal values. We report the resulting contrast in $\widehat{\omega}$ with a 95% bootstrap confidence interval. A contrast whose sign matches the theoretical prediction confirms the comparative static. We apply this procedure identically in both empirical applications below, differing only in data sources and covariate specifications.

3 Application I: Cross-National Evidence

We estimate coercive asymmetry across a broad cross-national sample to assess whether the theoretical predictions hold beyond any single conflict type, region, or data source.

Our cross-national sample draws on two primary event datasets that jointly maximize the number of conflict-spell observations meeting our scope conditions (i.e., dyadic violence in contested territories): the Uppsala Conflict Data Program's Georeferenced Event Dataset (UCDP-GED; Sundberg and Melander, 2013) and the Armed Conflict Location and Event Data project (ACLED; Raleigh et al., 2010). We harmonize UCDP-GED and ACLED into consistent actor categories and spatial units through the xSub data integration platform (Zhukov et al., 2019), ensuring that the threshold between restraint and escalation carries the same operational meaning across cases. We restrict the sample to two-sided armed conflicts in which combatants directed violence against each other and civilians, and exclude periods of non-violent contention and residual post-conflict violence.[24] Within each country, we partition observations into discrete conflict spells (Davenport and Appel, 2022): contiguous intervals during which the primary belligerent dyad remained relatively stable.[25] This procedure yields 138 spells in 51 countries across Africa, Asia, Europe, the

[24] We identified and removed out-of-scope cases through manual review (Appendix A2).

[25] Spell boundaries correspond to discontinuities in dyadic composition, due to regime change, state collapse, the entry or exit of a dominant armed actor, or a foreign military intervention.

Middle East, and the Americas, ranging from one to six per country (see Appendix A2).[26]

Our spatial unit of observation is the second-order administrative division (to which we refer generally as "district") and our temporal unit is the calendar month. Throughout, we assign armed formations fighting on the side of incumbent government or occupation authorities to Side 1, and their armed opponents (including state and non-state actors) to Side 2.[27] The covariates — across all conflict spells — include road density, log distance to provincial capital, forest cover, elevation, number of local ethnic groups, petroleum presence, log population (1990 baseline), built-up land area, and fixed effects for province, year and month. For each conflict spell, we fit the same bivariate probit specification, recover each side's predicted behavioral tendencies under counterfactual opponent behavior, and assemble them into a coercive asymmetry score with bootstrapped confidence intervals.[28]

3.1 Patterns of coercive asymmetry across conflicts

Figure 1 reports coercive asymmetry estimates and 95% confidence intervals for 138 conflict spells across 51 countries.[29] A positive $\widehat{\omega}$ indicates that Side 1's coercion is self-reinforcing: the asymmetric state where incumbents escalate and the opposition restrains is more persistent than the reverse. A negative estimate indicates that the opposition sustains unilateral escalation more effectively than incumbents. Null estimates indicate coercive symmetry.

The distribution of $\widehat{\omega}$ spans a wide range, from strongly negative to strongly positive, with considerable heterogeneity within and across conflict types. 27 (20%) of the $\widehat{\omega}$ esti-

[26]Multi-spell cases are conflicts without a stable dyadic structure. We retain only spell-source combinations for which the variance of $\hat{\lambda}_k$ is positive on both sides.

[27]For each data source, country and conflict spell, the binary outcome is $Y_{1,it} = 1$ if sources report any Side 1 violence in district i and month t, and $Y_{2,it} = 1$ if sources report any violence by Side 2.

[28]For each source-spell combination, we estimate equation (6) with the two-way Mundlak correction. We recover transition probabilities $\hat{p}_k$, $\hat{q}_k$ by averaging predicted marginal probabilities over the observed covariate distribution under counterfactual opponent behavior, and compute $\widehat{\omega} = (1 - \hat{p}_1)\hat{q}_2 - \hat{q}_1(1 - \hat{p}_2)$ with 95% parametric bootstrap CI's drawn from the joint posterior of the bivariate probit coefficient vector.

[29]78% of likelihood ratio tests reject independent-probit in favor of joint estimation (Appendix A2).

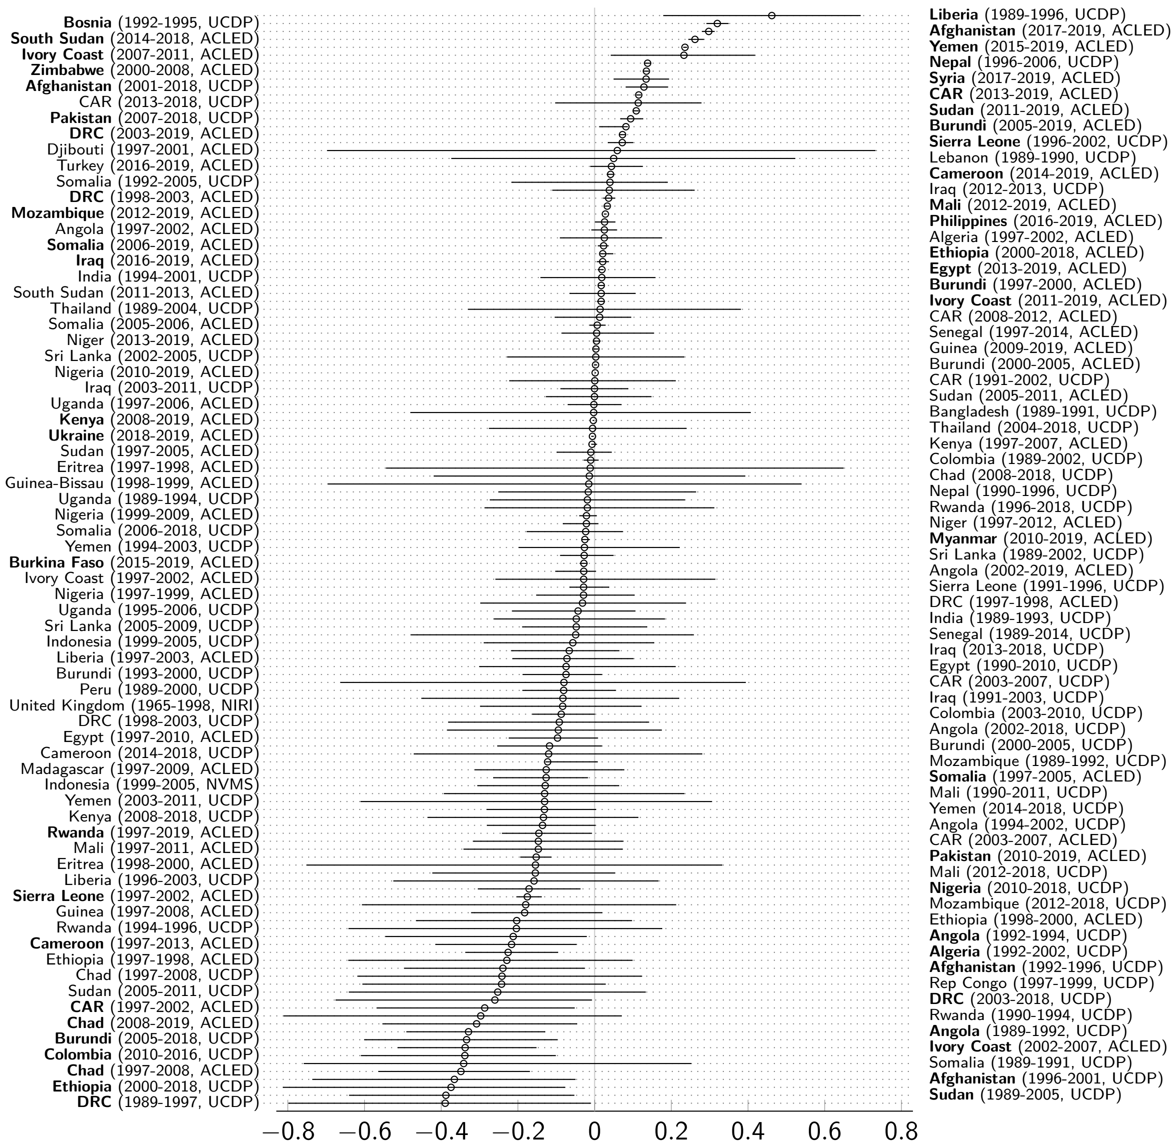

Figure 1: **Cross-national estimates of coercive asymmetry**. Point estimates and 95% bootstrap confidence intervals for $\hat{\omega}$. Bold labels indicate significance at 95% level. Of 138 estimates: 47 (34%) are positive, 91 (66%) are negative, including 27 (20%) positive and significant, 25 (18%) negative and significant.

mates are positive and significant at the 95% confidence level, and 25 (18%) are negative and significant.[30] The remaining 86 (62%) are indistinguishable from zero.

The estimates in Figure 1 reveal two distinct types of sign variation in $\widehat{\omega}$: shifts across conflict spells within the same country, and disagreements over the same spell across data sources. Within-country sign reversals often track documented shocks to the structural parameters that govern $\widehat{\omega}$, particularly shifts in relative selectivity or outside options of Sides 1 and 2 following events like external military intervention or the loss of material support. Afghanistan is a clear example: $\widehat{\omega}$ is negative in both pre-2001 spells (1992–1996, 1996–2001), reflecting the persistent coercive advantage of opposition forces over a fragmented central government. The estimates turn sharply positive following the U.S.-led intervention in October 2001, when external support enabled the new Afghan government (and coalition forces) to sustain operations even where civilian cooperation was scarce.

The Central African Republic (CAR) traces a similar arc. $\widehat{\omega}$ is significantly negative during the period of state fragmentation and overlapping rebel factions (1990–2002), drifts toward zero through the transitional period (2003–2007, 2008–2012), and turns significantly positive in the post-2013 spell. The latter shift coincides with the arrival of Russia's Wagner Group forces in 2017–2018, which substantially expanded government-side outside options, supplying external personnel, materiel, and logistical capacity that insulated the government from dependence on local civilian cooperation.

Sudan shows a parallel trajectory: negative and significant in the early civil war period (1989–2005), near-zero during the Comprehensive Peace Agreement era (2005–2011), and significantly positive in the subsequent post-secession Darfur/South Kordofan period. Cameroon follows the same pattern at a smaller scale: negative during the Boko Haram spillover period in the Far North, turning positive after 2014, as government and paramil-

[30] Negative $\widehat{\omega}$ estimates tend to carry wider confidence intervals because $\hat{p}_1$ draws on HL periods, which are rare when Side 1 rarely escalates unilaterally. The scarcity of observations in this conditioning cell inflates bootstrap variance for $\hat{p}_1$, and that uncertainty propagates into wider CIs for $\widehat{\omega} < 0$.

itary forces consolidated control over the Anglophone crisis.

A second source of sign variation is less common: disagreement between UCDP-GED and ACLED over the same spell. This pattern affects only five spells — Burundi post-2005, Democratic Republic of Congo post-2003, Ethiopia 2000–2018, Pakistan post-2007, Sierra Leone 1996–2002 — and likely reflects systematic differences in the two sources' collection protocols rather than genuine uncertainty about conflict dynamics.[31] The rarity of such disagreement (five out of 138) testifies to the broad convergence of results across datasets.

3.2 Targeting selectivity and outside options

Our theory predicts that more selective targeting increases coercive asymmetry: the more precisely a combatant identifies enemy collaborators, the higher the return to sustained punishment.[32] We test this prediction by computing $\widehat{\omega}$ at the bottom and top quintiles of the selectivity distribution within each source-spell sample. We operationalize selectivity as a district-level time average of the share of Side 1's violence directed at armed opponents rather than civilians, consistent with the theoretical role of selectivity as the share of punishment falling on identified enemy collaborators.[33]

Figure 2 reports the estimated change in coercive asymmetry as selectivity moves from the bottom to the top of each conflict's distribution, with 95% confidence intervals.[34] Most

[31]UCDP-GED restricts coverage to events involving organized armed actors, requires documentation of fatalities, and draws primarily on global newswires (Sundberg and Melander, 2013). ACLED casts a broader net, collecting events from local and regional media alongside wire services, and includes non-fatal violence (Raleigh et al., 2010). Because these differences in sourcing and inclusion criteria do not inflate event counts symmetrically across both sides of a dyad, they can shift the relative magnitudes of $\hat{p}_1$ and $\hat{q}_1$, and thus the sign of $\widehat{\omega}$, in ways that are conflict-specific and difficult to predict from first principles.

[32]Formally, Corollary 1(i) establishes $\partial\omega/\partial\lambda_1 > 0$: greater selectivity by Side 1 raises $\widehat{\omega}$.

[33]In xSub, we compute $\bar{\lambda}_{1i} = \frac{1}{T}\sum_t \frac{\texttt{DYAD\_A\_B}it}{\texttt{INITIATOR\_SIDEA}_{it}}$, where $\texttt{DYAD\_A\_B}_{it}$ counts Side 1-initiated events targeting Side 2 combatants in location i at t and $\texttt{INITIATOR\_SIDEA}_{it}$ counts all Side 1-initiated events in i, t. The local time-average is a more credible proxy for the underlying technology parameter $\lambda_k \in (0, 1)$ than short-run tactical choices, and avoids zero-denominator and endogeneity problems.

[34]Formally, the contrast is $\Delta\widehat{\omega} = \widehat{\omega}^{\text{hi}} - \widehat{\omega}^{\text{lo}}$.

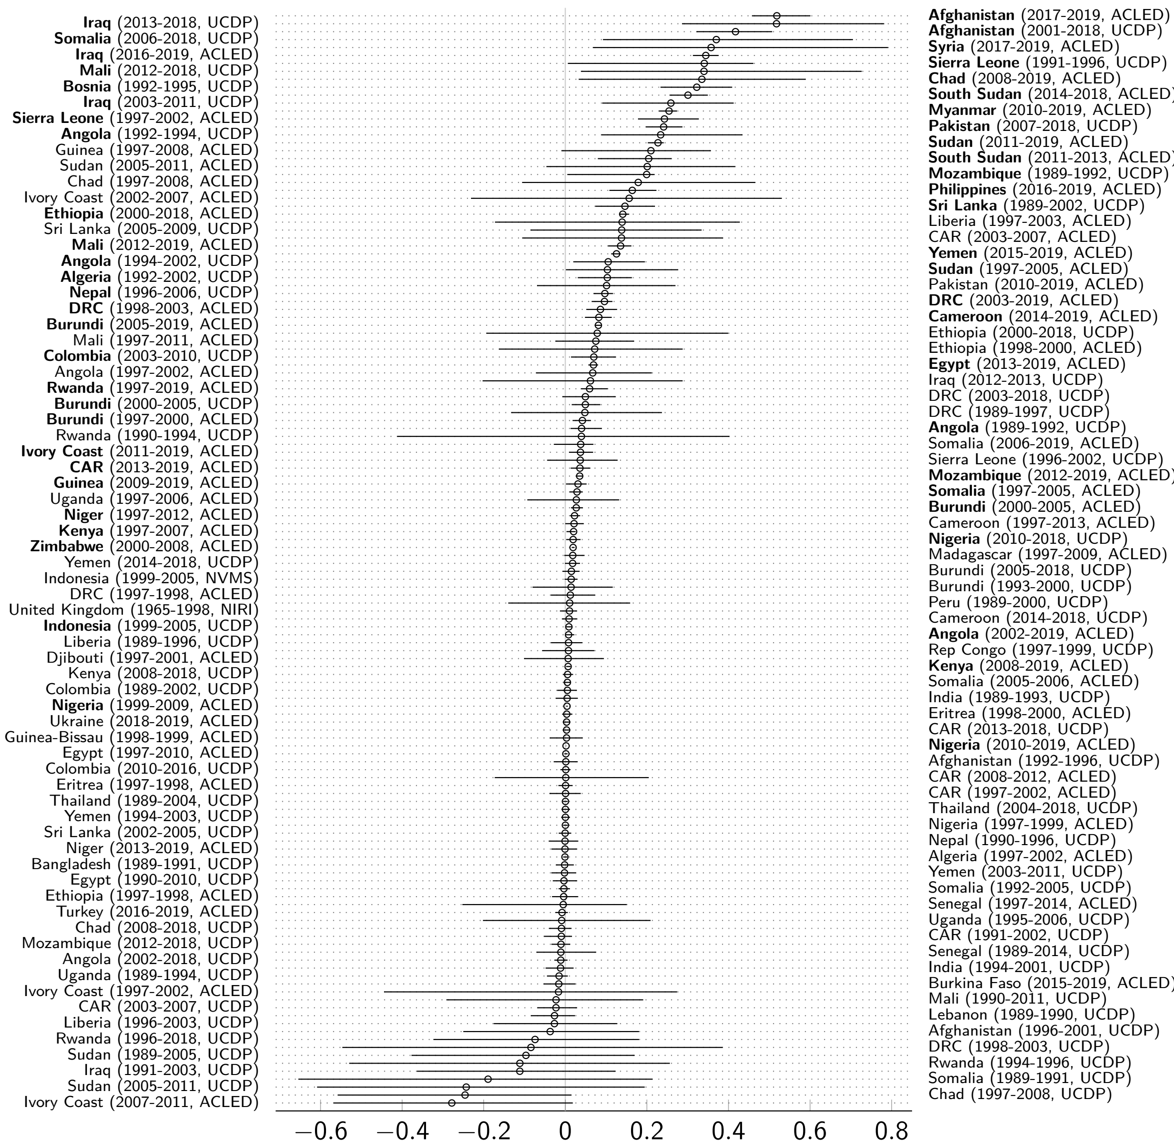

Figure 2: **Coercive asymmetry rises with selectivity of violence**. Point estimates and 95% bootstrap confidence intervals for $\Delta\hat{\omega}_1 = \hat{\omega}_1^{\text{hi}} - \hat{\omega}_1^{\text{lo}}$, where $\hat{\omega}_1^{\text{hi}}, \hat{\omega}_1^{\text{lo}}$ are average values of $\hat{\omega}_1$ for localities in the highest versus lowest quintiles of $\bar{\lambda}_1$. Bold labels indicate significance at 95% level. Of 138 estimates: 103 (75%) are positive, 35 (25%) are negative, including 53 (38%) positive and significant, 0 (0%) negative and significant.

cases (75%) show a positive contrast, consistent with our theoretical prediction. 53 (38%) of the estimated contrasts are positive and significant, and another 50 (37%) are positive but statistically indistinguishable from zero. No case shows a statistically significant negative contrast. Higher selectivity never reliably reduces government coercive asymmetry.

These results hold across all 138 source-spell samples, spanning five decades, four continents, and two data sources. They hold regardless of whether government coercive asymmetry is itself positive or negative.[35] Even where the opposition is generally more persistent, greater government selectivity narrows the gap and sometimes reverses it. In 16 cases where $\widehat{\omega} < 0$, the high-selectivity estimate turns positive, and 11 of these reversals are statistically significant (e.g., Iraq, Somalia, Myanmar).

Our theory also predicts that combatants with robust outside options (e.g., external sponsors, lootable resources, forces drawn from outside the locality) can sustain unilateral escalation even without precise targeting, because they do not depend on local civilian cooperation to keep fighting.[36] We test this prediction with two complementary approaches.

First, we ask how often $\widehat{\omega}$ is positive and significant even at low selectivity, a pattern that targeting precision alone cannot explain.[37] Of the 138 conflict spells, 16 (12%) show significantly positive coercive asymmetry even where government operations are mostly indiscriminate, consistent with outside options sustaining escalation independently of targeting capacity. Liberia's first civil war (1989–1996) fits this pattern: Charles Taylor's forces drew on timber and mineral revenues to sustain operations largely without rely-

[35]Confidence intervals for $\widehat{\omega}$ are substantially wider when the point estimate is negative (mean width 0.45) than when it is positive (mean width 0.18), and the same asymmetry carries over to the contrast estimates (0.24 vs. 0.12). This is not a statistical artifact. Negative $\widehat{\omega}$ arises when the opposition sustains unilateral escalation more reliably than the government. In this configuration, government coercion is predominantly indiscriminate and $\hat{p}_1$ is low and variable across space, producing a noisier estimate of the $HL \to HL$ transition and a wider bootstrap distribution for $\widehat{\omega}$.

[36]Formally, Corollary 1(iv) establishes $\partial \omega / \partial c_1 < 0$: greater outside options c_1 lower the selectivity threshold $\lambda_1^{\dagger}$, expanding the region in which Side 1 escalates. pushing $\widehat{\omega}$ toward negative values (Remark 1).

[37]Formally, we evaluate $\widehat{\omega}$ at the bottom quintile of the λ_1 distribution ($\widehat{\omega}^{\text{lo}}$) and identify cases where the 95% bootstrap CI excludes zero from below.

ing on civilian cooperation. Ivory Coast's post-election crisis (2007–2011) coincides with French *Licorne* force involvement that gave the government a coercive backstop. Yemen (2015–2019) reflects the Saudi-led coalition's access to external materiel and logistical support. Bosnia (1992–1995) and CAR (post-2013) similarly point to external reinforcement as the operative mechanism. In each case, access to resources or forces from outside the locality sustained coercive asymmetry where precision targeting could not.

Second, we ask whether local conditions that proxy for outside options and escalation costs shift coercive asymmetry in the directions our theory predicts.[38] For each spell, we shift each covariate from its observed minimum to maximum while holding all others at their medians, and report the resulting change in $\widehat{\omega}$. Table 1 summarizes the results.

Petroleum presence (i.e., a binary indicator of whether any petroleum deposit lies within the district, Lujala et al. 2007) shows the cleanest ratio of significantly positive to negative contrasts (6:1 across 73 spells), consistent with resource rents providing outside options that can sustain government escalation without requiring precise targeting.

Proximity to built-up areas and provincial capitals produce similar patterns. Such settings offer incumbents superior infrastructure, easier force projection, and established administrative presence, reducing the logistical costs of sustained escalation and providing resource flows that reduce government dependence on civilian support.

Ethnic co-habitation (i.e., the number of unique ethnic settlement zones overlapping each district, Weidmann et al. 2010) yields perhaps the most consistent results: 68% of spells show a positive contrast, with significantly positive cases outnumbering significantly negative ones by nearly 4:1. Where multiple ethnic groups' territories overlap, civilians face higher coordination costs for collective resistance, making sustained defiance harder to organize and reinforcing the government's coercive advantage.

[38] Formally, we estimate the contrast $\widehat{\omega}^{\mathrm{hi}} - \widehat{\omega}^{\mathrm{lo}}$ as each covariate moves from its observed minimum to maximum, holding all others at their medians, and ask whether the sign matches Corollary 1. Outside options correspond to c_k; escalation costs correspond to $\kappa_k(H)$.

Table 1: **Cross-national comparative statics.** Each cell reports the number of source-spell samples in which $\hat{\omega}$ increases or decreases when a covariate moves from its observed minimum to maximum, holding all other covariates at their medians. Sample sizes vary across covariates due to exclusion of cases with no within-spell variation or missing data on the covariate in question. 95% bootstrap confidence intervals determine significance.

			All contrasts		Significant at 95%	
Covariate	Spells	Countries	Positive	Negative	Positive	Negative
Number of ethnic groups	109	41	74 (68%)	35 (32%)	23 (21%)	6 (6%)
Built-up areas	111	43	66 (59%)	45 (41%)	30 (27%)	9 (8%)
Log population (1990)	111	43	64 (58%)	47 (42%)	16 (14%)	8 (7%)
Petroleum presence	73	29	42 (58%)	31 (42%)	12 (16%)	2 (3%)
Road density	111	43	60 (54%)	51 (46%)	10 (9%)	12 (11%)
Forest cover	111	43	56 (50%)	55 (50%)	14 (13%)	14 (13%)
Mean elevation	111	43	55 (50%)	56 (50%)	15 (14%)	17 (15%)
Log distance adm. center	111	43	46 (41%)	65 (59%)	9 (8%)	19 (17%)

4 Application II: Soviet Union vs. Ukrainian insurgents

The cross-national results establish that coercive asymmetry tracks selectivity and outside options across a wide range of conflicts. We now ask if the same patterns hold under more demanding measurement conditions. A post-World War II conflict in western Ukraine is well-suited for this test: Soviet security forces and the Ukrainian Insurgent Army (UPA) competed for the loyalty of the same civilian population across territories the USSR recently annexed and did not fully control; archival records document both sides' violence more completely than media-based event data, with less censoring of low-intensity incidents.

We draw on a published dataset on the Soviet-UPA conflict, which integrates incident reports from declassified secret police, Communist Party of Ukraine, and insurgent documents (Zhukov, 2015; Rozenas et al., 2017).[39] The dataset covers conflict events across Ukraine from 1943 to 1955, with locations, dates, casualties, and tactics for each incident. We aggregate these to the rayon-month level and restrict the sample to Ukraine's eight newly-annexed western oblasts during 1944–1950, where the UPA concentrated its activity and violently contested Soviet administrative control ($N = 20{,}076$ rayon-months).[40]

The binary outcome for Side 1 (Soviets) equals one if any act of Soviet violence (e.g., killing, arrest, household deportation) occurred in a rayon-month, and zero otherwise. The outcome for Side 2 (UPA) is one if an insurgent-initiated violent incident occurred. As in the cross-national application, we proxy selectivity as the rayon-level share of violence the Soviets directed at insurgents rather than civilians, averaged over the sample period.[41]

We include combatant-specific covariates reflecting the different structural constraints each side faced: measures of administrative reach and state capacity for the Soviets, and measures of mobility and concealment for the UPA.[42] Both sets include fixed effects for region (oblast), year, and month, and a spatial lag of each side's own lagged outcome.

4.1 Patterns of coercive asymmetry in Soviet Ukraine

Empirical estimates of coercive asymmetry are consistently positive and significant across all specifications, indicating that one-sided Soviet violence was more self-sustaining than

[39]Data available at doi.org/10.1177/0022002713520590 (under "Supplementary Material").

[40]Soviet rayons were second-tier administrative units roughly comparable to U.S. counties. Soviet oblasts were first-tier administrative units similar to U.S. states.

[41]We calculate $\bar{\lambda}_{1i} = \frac{1}{T}\sum_t^T \frac{\text{insurgent-targeted incidents}_{it}}{\text{all Soviet incidents}_{it}}$, where the numerator counts only the operations the Soviets recorded as targeting combatants in i, t (i.e., no family deportations, no civilian casualties).

[42]For the Soviets, $\mathbf{x}_{1,t}$ includes railway access, number of village soviets (proxying administrative penetration depth and state capacity), electrification, natural resources, agricultural land share, urbanization, rayon area, and pre-1944 Soviet partisan control. For the UPA, $\mathbf{x}_{2,t}$ includes railway access, forest cover (proxying cover and concealment), electrification, natural resources, agricultural land share, urbanization, number of border crossings to neighboring rayons, and rayon area. Village soviets and pre-1944 partisan control capture the depth of Soviet administrative and intelligence infrastructure, respectively, which shaped NKVD/MVD capacity but had no direct analogue for the UPA. Forest cover and local crossings governed insurgent mobility and concealment in ways that do not extend to state security forces.

one-sided violence by the UPA. In the best-fitting model, $\widehat{\omega} = 0.2$ (95% CI: [0.16, 0.26]).[43]

Figure 3 shows how estimates of Soviet coercive asymmetry $\widehat{\omega}$ vary across quintiles of selectivity (i.e., share of violence directed only at insurgents). Soviet coercive asymmetry increases with selectivity, from $\widehat{\omega} = 0.141$ in the lowest quintile to $\widehat{\omega} = 0.251$ in the highest. This aligns with our expectation that selectivity raises the net benefit of escalation.[44]

Figure 3 also speaks to whether Soviet outside options alone, rather than targeting precision, could explain the coercive asymmetry. According to their own records, roughly three quarters of Soviet violence in the median western Ukrainian locality fell on civilians rather than combatants, making this an exceptionally indiscriminate counterinsurgency campaign. If outside options fully explain Soviet escalation, coercive asymmetry should be uniformly positive across localities, irrespective of how indiscriminate the violence was.[45] The evidence is partially consistent with that prediction. Soviet coercive asymmetry is positive and significant even in the lowest-selectivity quintile, but it is not uniform. Even if the Soviets' material advantages enabled them to sustain unilateral escalation where targeting was completely indiscriminate, the positive gradient in Figure 3 suggests that one-sided violence was more durable where Soviet targeting was more precise.

[43]We use the two-way Mundlak specification for all $\widehat{\omega}$ inference. Likelihood ratio tests confirm significant cross-equation correlation, validating joint estimation over independent equations (Appendix A2). $\widehat{\omega}$ estimates from independent probit equations are all smaller, with confidence intervals spanning zero. The interactive specification outperforms both the additive and non-strategic models, and the Mundlak correction improves fit further by absorbing rayon-level heterogeneity and period-specific shocks (Table A4).

[44]This follows from Corollary 1.i.

[45]Formally, this predicts $\widehat{\omega}_1 > 0$ even where $\lambda_1 \approx 0$, because the outside option c_1 alone sustains escalation independently of targeting precision.

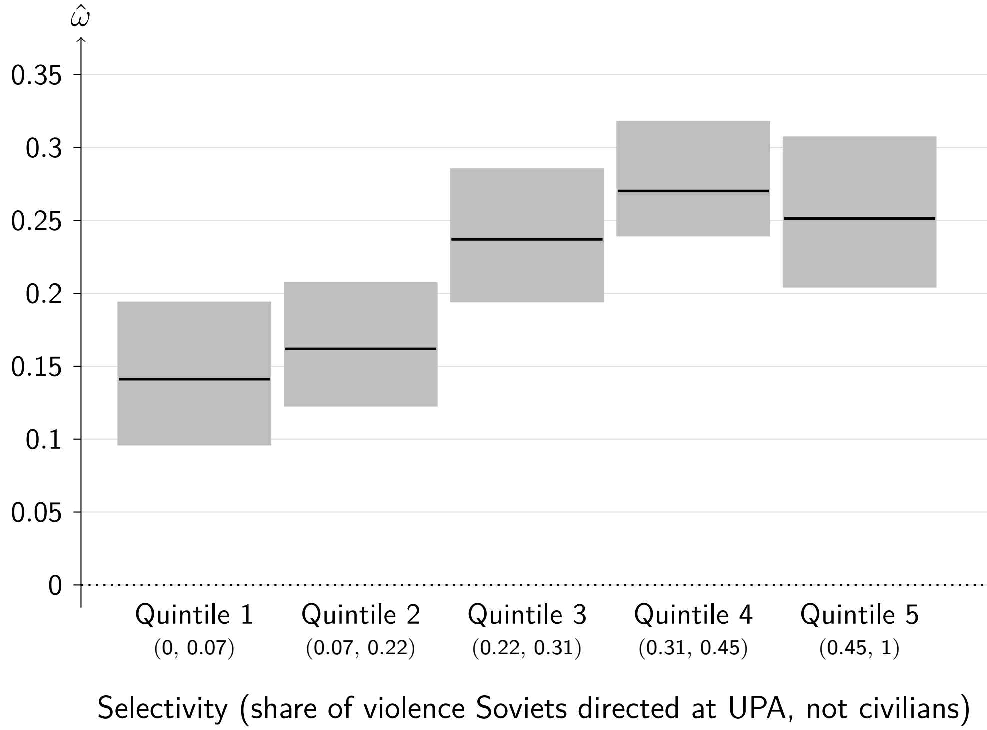

Figure 3: **Coercive asymmetry rises with Soviet selectivity**. Black lines are $\widehat{\omega}$ point estimates, grey bars are 95% bootstrap CIs. Parentheses show selectivity range in quintile.

What structural conditions drove variation in Soviet one-sided violence? To answer this, Table 2 shifts each covariate from its observed minimum to maximum, holding all others at their medians, and reports how much coercive asymmetry changes.[46] Administrative penetration depth, which we measure by the density of village soviets, produces the largest estimated contrast, $\Delta\widehat{\omega} = 0.17$ (95% CI: 0.09, 0.26).[47] A deeper institutional footprint connected localities to broader state supply chains, personnel rotations, and organizational resources, reducing dependence on local civilian goodwill. Rayons with denser local administrative networks sustained escalation more effectively than those without them.

[46]Each contrast is $\widehat{\omega}_{\text{hi}} - \widehat{\omega}_{\text{lo}}$ from specification (6), which we estimated here without the Mundlak correction to preserve identification of the time-invariant covariates.

[47]Village soviets here proxy for the state's outside options (c_1).

Soviet coercive asymmetry was also higher in territories with a history of WWII-era Soviet partisan control, $\Delta\widehat{\omega} = 0.05$ (95% CI: 0.02, 0.08). The intelligence networks these operations left behind improved targeting precision, making Soviet coercion potentially more efficient.[48] Urban share produces a modest but similarly positive contrast, $\Delta\widehat{\omega} = 0.02$ (95% CI: 0.003, 0.04). Cities were both easier to surveil and more strategically valuable, giving security forces incentives to sustain escalation regardless of short-run outcomes.[49] Rail access, rayon area, electrification, and natural resources fall short of conventional significance, suggesting they capture underlying structural channels with considerable noise.

Table 2: **Comparative statics**. Change in $\widehat{\omega}$ when each covariate moves from its observed minimum to maximum, holding all others at medians. 95% bootstrap CIs in parentheses.

Covariate	Range	$\Delta\hat{\omega}$ (95% CI)
Village soviets	$7 \to 82$	0.166 (0.093, 0.263)
Partisan control (pre-1944)	$0 \to 1$	0.045 (0.021, 0.075)
Urban share	$0 \to 1$	0.021 (0.003, 0.040)
Area (1000 sq km)	$0 \to 1.8$	0.077 (-0.012, 0.138)
Railway present	$0 \to 1$	0.008 (-0.012, 0.022)
Electrification	$0 \to 1$	0.034 (-0.013, 0.085)
Natural resources	$0 \to 3$	0.010 (-0.028, 0.050)
Agricultural land	$0 \to 1$	-0.005 (-0.034, 0.017)

5 Threats to Inference

Our estimates clear the hardest inferential hurdle: the core patterns hold across a broad cross-national sample and survive replication in archival data where reporting and misclassification concerns are less acute. The most serious remaining concerns are more likely to attenuate our comparative static results than to generate them spuriously.

[48] Partisan control proxies for prior intelligence infrastructure (λ_1).

[49] Urban density raises targeting selectivity (λ_1), but also concentrates economic and administrative assets that increase outside options (c_1), making the two channels difficult to separate.

First, reporting bias is the most direct threat to inference. Violence against civilians may be more observable than violence against combatants, and larger events are more likely to enter both archival and media-based records than smaller ones. Either pattern can inflate the level estimate $\widehat{\omega}$ if it makes Side 1's escalation more visible than Side 2's.[50] The contrast estimates $\Delta\widehat{\omega}$ are less vulnerable: reporting asymmetries would need to vary systematically with selectivity to generate a spurious positive gradient. If anything, the most visible events tend to be large and indiscriminate, which would compress differences between high- and low-selectivity localities and attenuate $\Delta\widehat{\omega}$ toward zero.

Second, binary coding of escalation imposes a visibility threshold: some escalatory behavior goes unrecorded, and some low-level violence appears as escalation.[51] Such misclassification adds noise to the transition probabilities we use to construct $\widehat{\omega}$, pulling the level estimate toward zero. It also attenuates $\Delta\widehat{\omega}$ by shrinking contrasts across localities.

Third, selectivity measures may be endogenous to local conditions that simultaneously sustain one-sided violence. Administrative capacity can raise targeting precision and also make unilateral escalation more durable, which means the positive association between selectivity and $\widehat{\omega}$ may partly reflect common underlying capability rather than the causal effect of selectivity. This concern does not bias the level estimate $\hat{\omega}$, which we identify from observed transition patterns rather than from the selectivity proxy, but it can inflate $\Delta\widehat{\omega}$ if high-capacity localities sort into both higher selectivity and higher coercive asymmetry. We reduce this risk by proxying selectivity with time-averaged targeting shares rather than short-run tactical choices, and by conditioning on pre-conflict locality attributes.

Fourth, the cross-national analysis relies on imperfect proxies for outside options and escalation costs, and on event data that vary in sourcing rules, media access, and visibility

[50] In media-accessible settings, journalists may document government operations more thoroughly than insurgent ones, biasing $\hat{\omega}$ upward. In closed authoritarian settings, the reverse may hold.

[51] Our observation mechanism assumes that low punishment is empirically invisible, and high punishment is not. This is a necessary simplification, not a claim that zeroes represent perfect peace.

of violence. Noisy proxies are unlikely to bias $\widehat{\omega}$ systematically, but they attenuate comparative statics by weakening the empirical stand-ins for our theoretical parameters. Source disagreement adds variance in the same direction, although UCDP-GED and ACLED agree in almost all overlapping spells. The Soviet-UPA application serves as a partial check on both concerns: the same theoretical gradients appear under archival measurement conditions where proxy noise and reporting gaps are substantially lower.

6 Conclusion

Combatants escalate unilaterally when and where their violence is precise enough to extract civilian compliance, or when outside options make compliance dispensable. Our framework derives these conditions, measures their behavioral imprint from sequences of observed violence, and shows that the resulting index shifts predictably across space and time.

Several extensions follow naturally from this contribution. First, future work could extend our theoretical framework to allow strategic investment in selectivity and outside options (e.g., surveillance infrastructure, patronage networks, alliances) or relax the dyadic structure to model asymmetry in multi-party conflicts. Second, $\hat{\omega}$ provides a common metric for phenomena the literature has treated as distinct (e.g., wartime repression and terrorism, atrocities), revealing them as points on a single continuum rather than categorically different objects of study. Researchers can now classify conflicts by which side holds the structural coercive advantage and why, sharpening downstream work on conflict duration, civilian casualties, and settlement. Third, our comparative statics rest on observational variation; identifying exogenous shocks to selectivity and outside options (e.g., foreign materiel influx, surveillance technology rollout) would move the study of one-sided violence from structural estimation of theoretical parameters toward credible causal identification.

References

Anderton, Charles H. and Jurgen Brauer, "Mass Atrocities and Their Prevention," *Journal of Economic Literature*, 2021, *59* (4), 1240–1292.

Arreguin-Toft, Ivan, "How the weak win wars: A theory of asymmetric conflict," *International Security*, 2001, *26* (1), 93–128.

Berman, Eli, Michael Callen, Joseph H Felter, and Jacob N Shapiro, "Do working men rebel? Insurgency and unemployment in Afghanistan, Iraq, and the Philippines," *Journal of Conflict Resolution*, 2011, *55* (4), 496–528.

Besley, Timothy and Torsten Persson, "Repression or civil war?," *American Economic Review*, 2009, *99* (2), 292–97.

_ and _ , "The logic of political violence," *The Quarterly Journal of Economics*, 2011, *126* (3), 1411–1445.

Blair, Graeme A., "The Fortification Dilemma: Border Control and Rebel Violence," *American Journal of Political Science*, 2023.

Davenport, Christian, "Multi-dimensional threat perception and state repression: An inquiry into why states apply negative sanctions," *American Journal of Political Science*, 1995, *39* (3), 683–713.

_ and Benjamin J Appel, "Stopping state repression: An examination of spells," *Journal of Peace Research*, 2022, *59* (5), 633–647.

_ , Håvard Mokleiv Nygård, Hanne Fjelde, and David Armstrong, "The Consequences of Contention: Understanding the Aftereffects of Political Conflict and Violence," *Annual Review of Political Science*, 2019, *22*, 361–377.

de Mesquita, Bruce Bueno, David P. Myatt, Alastair Smith, and Scott A. Tyson, "The Punisher's Dilemma," *The Journal of Politics*, 2024, *86* (2), 395–411.

Eck, Kristine and Lisa Hultman, "One-Sided Violence Against Civilians in War Insights from New Fatality Data," *Journal of Peace Research*, 2007, *44* (2), 233–246.

Fearon, James D., "Rationalist Explanations of War," *International Organization*, 1995, *49* (3), 379–414.

Fink, A. M., "Equilibrium in a stochastic n-person game," *Journal of Science of the Hiroshima University, Series A-I*, 1964, *28* (1), 89–93.

Gohdes, Anita R, "Repression Technology: Internet Accessibility and State Violence," *American Journal of Political Science*, 2020, *64* (3), 488–503.

Green, Amelia Hoover, "The Commander's Dilemma: Creating and Controlling Armed Group Violence," *Journal of Peace Research*, 2016.

Kalyvas, Stathis N., *The Logic of Violence in Civil War*, Cambridge, 2006.

Leites, Nathan and Charles Wolff, Jr., *Rebellion and Authority*, Santa Monica, CA: RAND, 1970.

Lujala, P., Jan Ketil Rod, and Nadia Thieme, "Fighting over Oil: Introducing A New Dataset," *Conflict Management and Peace Science*, 2007, *24* (3).

Lyall, Jason, "Does indiscriminate violence incite insurgent attacks? Evidence from Chechnya," *Journal of Conflict Resolution*, 2009, *53* (3), 331–362.

_ , "Are Coethnics More Effective Counterinsurgents? Evidence from the Second Chechen War," *American Political Science Review*, 2010, *104* (1).

Mack, Andrew, "Why Big Nations Lose Small Wars: The Politics of Asymmetric Conflict," *World Politics*, 1975, *27*, 175–200.

Mertens, Jean-François and T. Parthasarathy, "Equilibria for discounted stochastic games," CORE Discussion Paper 1987050, Center for Operations Research and Econometrics (CORE), Université catholique de Louvain 1987.

Mundlak, Yair, "On the Pooling of Time Series and Cross Section Data," *Econometrica*, 1978, *46* (1), 69–85.

Powell, Robert, "Persistent Fighting and Shifting Power," *American Journal of Political Science*, 2012, *56* (3), 620–637.

Raleigh, Clionadh, Andrew Linke, Håvard Hegre, and Joakim Karlsen, "Introducing ACLED: Armed Conflict Location and Event Data," *Journal of Peace Research*, 2010, *47* (5), 651–660.

Ritter, Emily Hencken and Courtenay R Conrad, "Preventing and responding to dissent: The observational challenges of explaining strategic repression," *American Political Science Review*, 2016, *110* (1), 85–99.

Rohner, Dominic, Oliver Vanden Eynde, and Emma Verhille, "Variants of Violence: Classifying Conflict Types and Policies for Peace," *Economic Policy*, July 2025, *40* (123), 621–645.

Rozenas, Arturas, Sebastian Schutte, and Yuri Zhukov, "The Political Legacy of Violence: The Long-Term Impact of Stalin's Repression in Ukraine," *Journal of Politics*, 2017, *79* (4), 1147–1161.

Salehyan, Idean, David Siroky, and Reed M Wood, "External Rebel Sponsorship and Civilian Abuse: A Principal-Agent Analysis of Wartime Atrocities," *International Organization*, 2014, pp. 1–29.

Schubiger, Livia Isabella, "State violence and wartime civilian agency: Evidence from Peru," *The Journal of Politics*, 2021, *83* (4), 1383–1398.

Slantchev, Branislav L., "The Power to Hurt: Costly Conflict with Completely Informed States," *American Political Science Review*, 2003, *97* (1), 123–135.

Smith, Ron, Martin Sola, and Fabio Spagnolo, "The Prisoners Dilemma and Regime-Switching in the Greek-Turkish Arms Race," *Journal of Peace Research*, 2000, *37* (6), 737–750.

Stanton, Jessica, *Violence and Restraint in Civil War*, Cambridge: Cambridge University Press, 2016.

Sundberg, Ralph and Erik Melander, "Introducing the UCDP georeferenced event dataset," *Journal of Peace Research*, 2013, *50* (4), 523–532.

Weidmann, Nils B, Jan Ketil Rød, and Lars-Erik Cederman, "Representing ethnic groups in space: A new dataset," *Journal of Peace Research*, 2010, *47* (4), 491–499.

Wolford, Scott, Dan Reiter, and Clifford J Carrubba, "Information, commitment, and war," *Journal of Conflict Resolution*, 2011, *55* (4), 556–579.

Wood, Reed M., "Rebel Capability and Strategic Violence against Civilians," *Journal of Peace Research*, 2010, *47* (5), 601–614.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, MA: MIT Press, 2010.

Zhukov, Yuri M., "Population Resettlement in War: Theory and Evidence from Soviet Archives," *Journal of Conflict Resolution*, 2015, *59* (7), 1155–1185.

Zhukov, Yuri M, Christian Davenport, and Nadiya Kostyuk, "Introducing xSub: A New Portal for Cross-National Data on Sub-National Violence," *Journal of Peace Research*, 2019, *56* (4), 604–614.

A1 Theoretical Appendix

Table A1: **Model notation.**

Symbol	Description
i, t, j, k	Indices for locality, time period, civilian, and combatant ($k \in \mathcal{K}$).
$\mathcal{K}$	Combatant set $\mathcal{K} = \{1, 2\}$, where Side 1 is the incumbent armed actor and Side 2 is the challenging armed actor.
A_k	Action set for combatant k: $A_k = \{L, H\} = \{0, \Delta\}$.
L, H, Δ	Low and high punishment; $L = 0$, $H = \Delta > 0$.
θ_j	Civilian j's type (relative collateral vs. selective vulnerability).
$F(\theta\|x_i), f(\theta\|x_i)$	Type distribution and density in locality i, conditional on local observables x_i. Mean type in locality i is $\bar{\theta}_i = \mathbb{E}[\theta_j\|x_i]$.
ρ_k	Punishment level chosen by combatant k; $\rho_k \in A_k$.
e_j, e_j^*, e_k^*	Civilian j's cooperation with Side 1 ($e_j \in [0, 1]$); e_j^* is optimal value; e_k^* is share directed to k ($e_1^* = e_j^*$, $e_2^* = 1 - e_j^*$).
$\bar{e}_i$	Mean cooperation with Side 1 in locality i: $\bar{e}_i = \mathbb{E}[e_j^*\|x_i]$.
λ_k	Selectivity of k's violence.
$\lambda_k^\dagger$	Critical selectivity level at which H becomes optimal for k.
λ_i^*	Aggregate civilian threshold in locality i: $\lambda_i^* = \bar{\theta}_i$.
b	Value of a unit of civilian cooperation to each combatant.
$\kappa_k(\rho_k)$	Punishment cost for k; $\kappa_k(L) = 0$, $\kappa_k(H) > 0$.
c_k	Combatant k's outside options to local support; $c_k \geq 0$.
γ	Curvature of civilian utility (cost of deviating from $e_j = 1/2$).
Φ_k	Cooperation premium of combatant k: $\Phi_k = \frac{\Delta}{\gamma}(\lambda_k - \bar{\theta}_i)$.
$U_j(e_j), U_k(\rho_k, \rho_{-k})$	Civilian j's utility, combatant k's utility.
S, s_{it}	State space; state of locality i at time t: $s_{it} = (\lambda_{it}, x_{it})$.
g, ε_{it}	Selectivity transition: $\lambda_{i,t+1} = g(\lambda_{it}, \rho_{1,it}, \rho_{2,it}, x_{it}) + \varepsilon_{i,t+1}$.
$v_k(s, \rho_k, \rho_{-k})$	Stage payoff to combatant k in state s under profile (ρ_k, ρ_{-k}).
$\xi_k, V_k(s)$	Strategy and continuation value of combatant k at state s.
p_k, q_k	Reactive strategies: $p_k = P(L_t \mid L_{t-1}^{-k})$, $q_k = P(L_t \mid H_{t-1}^{-k})$.
δ	Common discount factor in the stochastic game.
$\mathcal{A} = A_1 \times A_2$	Set of observable punishment profiles: $\mathcal{A} = \{LL, LH, HL, HH\}$.
$M, m_{a'a}$	4×4 transition matrix over $\mathcal{A}$; $m_{a'a} = P(a' \mid a)$ for $a, a' \in \mathcal{A}$.
ω	Coercive resolve index: $\omega = m_{HL,HL} - m_{LH,LH}$.

Two combatants ($k \in \{1, 2\}$) simultaneously choose $\rho_k \in \{L, H\} = \{0, \Delta\}$ in each locality i and period t. Each civilian j draws type $\theta_j \sim F(\theta|x_i)$, with $\bar{\theta}_i \equiv \mathbb{E}[\theta_j|x_i]$; θ_j is private information, but combatants observe $F(\theta|x_i)$.

Assumption 1 (Civilian utility)**.** Civilian j chooses $e_j \in [0, 1]$ to maximize

$$U_j(e_j) = -\Big[\theta_j(1-\lambda_1)\rho_1 + (1-\theta_j)\lambda_2\rho_2\Big]e_j - \Big[\theta_j(1-\lambda_2)\rho_2 + (1-\theta_j)\lambda_1\rho_1\Big](1-e_j) - \frac{\gamma}{2}\Big(e_j - \tfrac{1}{2}\Big)^2. \tag{7}$$

Assumption 2 (Combatant utilities)**.** Combatant k chooses $\rho_k \in \{L, H\}$ to maximize

$$U_k(\rho_k, \rho_{-k}) = b \cdot \mathbb{E}[e_k^* \mid x_i] + c_k(1 - \lambda_k)\rho_k - \kappa_k(\rho_k), \tag{8}$$

where $b > 0$, $c_k \geq 0$, $\kappa_k(L) = 0$, $\kappa_k(H) > 0$. Note that c_k does not enter civilian utility: defiance and submission depend solely on λ_k and θ_j.

Proposition 1 (Civilian response)**.** *Fix locality i.*

(a) (Interior solution.) If $\gamma > 2|(\lambda_1 - \theta_j)\rho_1 - (\lambda_2 - \theta_j)\rho_2|$, j has a unique best response:

$$e_j^* = \frac{1}{2} + \frac{(\lambda_1 - \theta_j)\rho_1 - (\lambda_2 - \theta_j)\rho_2}{\gamma}. \tag{9}$$

(b) (Defiance and submission.)

$$\frac{\partial e_j^*}{\partial \rho_1} = \frac{\lambda_1 - \theta_j}{\gamma}, \qquad \frac{\partial e_j^*}{\partial \rho_2} = -\frac{\lambda_2 - \theta_j}{\gamma}. \tag{10}$$

Civilian j defies Side 1 iff $\lambda_1 < \theta_j$; submits iff $\lambda_1 > \theta_j$. Symmetrically for Side 2.

(c) (Aggregation.)

$$\bar{e}_i(\rho_1, \rho_2) = \frac{1}{2} + \frac{(\lambda_1 - \bar{\theta}_i)\rho_1 - (\lambda_2 - \bar{\theta}_i)\rho_2}{\gamma}. \tag{11}$$

Define $\lambda_i^* \equiv \bar{\theta}_i$. The mean response $\bar{e}_i$ rises with ρ_1 iff $\lambda_1 > \lambda_i^*$, falls iff $\lambda_1 < \lambda_i^*$.

Proof. Part (a). The FOC $\partial U_j/\partial e_j = 0$ gives $-[(\theta_j - \lambda_1)\rho_1 - (\theta_j - \lambda_2)\rho_2] - \gamma(e_j - \frac{1}{2}) = 0$, which solves to (9). Strict concavity ($\partial^2 U_j/\partial e_j^2 = -\gamma < 0$) confirms uniqueness; the condition on γ ensures $e_j^* \in (0, 1)$.

Part (b). Differentiate (9) directly to obtain (10).

Part (c). Take expectations of (9) over $\theta \sim F(\theta|x_i)$ to obtain (11); the sign of $\partial\bar{e}_i/\partial\rho_1 = (\lambda_1 - \bar{\theta}_i)/\gamma$ follows immediately. □

Define the *cooperation premium*:

$$\Phi_k \;\equiv\; \frac{\Delta}{\gamma}(\lambda_k - \bar{\theta}_i). \tag{12}$$

Mean cooperation with k under each profile is $\mathbb{E}[e_k^*|H,L] = \frac{1}{2} + \Phi_k$, $\mathbb{E}[e_k^*|L,H] = \frac{1}{2} - \Phi_{-k}$, $\mathbb{E}[e_k^*|H,H] = \frac{1}{2} + \Phi_k - \Phi_{-k}$, $\mathbb{E}[e_k^*|L,L] = \frac{1}{2}$. Since $\rho_k \in \{0, \Delta\}$, net gain from escalating is

$$U_k(H,\cdot) - U_k(L,\cdot) \;=\; b\,\Phi_k + c_k(1-\lambda_k)\Delta - \kappa_k(H), \tag{13}$$

independent of the opponent's action and strictly increasing in λ_k (since $c_k < b/\gamma$).

Theorem 1 (Equilibrium regimes)**.** *Fix locality i with $\bar{\theta}_i$, (λ_1, λ_2), $\kappa_k(H) > 0$, and $c_k \geq 0$ with $c_k < b/\gamma$ for each k. Define*

$$\lambda_k^\dagger \;\equiv\; \bar{\theta}_i + \frac{\gamma[\kappa_k(H) - c_k(1-\bar{\theta}_i)\Delta]}{(b - c_k\gamma)\Delta}. \tag{14}$$

The game is dominance-solvable with unique Nash equilibrium:

(i) (H,H) *if* $\lambda_1 > \lambda_1^\dagger$ *and* $\lambda_2 > \lambda_2^\dagger$ (mutual escalation)*;*

(ii) (L,L) *if* $\lambda_1 < \lambda_1^\dagger$ *and* $\lambda_2 < \lambda_2^\dagger$ (mutual restraint)*;*

(iii) (H,L) *if* $\lambda_1 > \lambda_1^\dagger$ *and* $\lambda_2 < \lambda_2^\dagger$ (asymmetric escalation by Side 1)*;*

(iv) (L,H) *if* $\lambda_1 < \lambda_1^\dagger$ *and* $\lambda_2 > \lambda_2^\dagger$ (asymmetric escalation by Side 2)*.*

When $\kappa_1(H) = \kappa_2(H)$ and $c_1 = c_2$, the thresholds coincide and only regimes (i)–(ii) arise.

Proof. Substituting $\Phi_k = \Delta(\lambda_k - \bar{\theta}_i)/\gamma$ into (13):

$$U_k(H,\cdot) - U_k(L,\cdot) = \Delta\lambda_k\left(\frac{b}{\gamma} - c_k\right) + \Delta\left(c_k - \frac{b\bar{\theta}_i}{\gamma}\right) - \kappa_k(H).$$

Since $c_k < b/\gamma$, the coefficient on λ_k is positive. Setting equal to zero and solving yields (14). H strictly dominates iff $\lambda_k > \lambda_k^\dagger$; L strictly dominates iff $\lambda_k < \lambda_k^\dagger$. Since (13) is independent of the opponent's action, each combatant has a strictly dominant action, making the game dominance-solvable. Crossing the two independent conditions yields regimes (i)–(iv); symmetry of thresholds under $\kappa_1(H) = \kappa_2(H)$, $c_1 = c_2$ rules out asymmetric profiles. □

The threshold $\lambda_k^\dagger$ is strictly decreasing in c_k: outside options expand the region in which escalation dominates. When $c_k > \kappa_k(H)/[\Delta(1-\bar{\theta}_i)]$, the threshold falls below $\bar{\theta}_i$, and k escalates despite a negative cooperation premium (Remark 1).

Remark 1 (Attrition sub-regime)**.** When $\lambda_k^\dagger < \bar{\theta}_i$, any k with $\lambda_k \in (\lambda_k^\dagger, \bar{\theta}_i)$ escalates despite $\Phi_k < 0$: the outside-option term $c_k(1-\lambda_k)\Delta$ covers both the cooperation loss and $\kappa_k(H)$. From the perspective of ω, attrition-driven and cooperation-driven escalation are observationally equivalent; separating them empirically requires variation in proxies for c_k.

Suppose each combatant k conditions its current action on the opponent's previous action.[52] Define

$$p_k \equiv P(\rho_{kt} = L \mid \rho_{-k,t-1} = L), \qquad q_k \equiv P(\rho_{kt} = L \mid \rho_{-k,t-1} = H). \tag{15}$$

Proposition 2 (Coercive resolve index)**.**

(a) (Transition matrix.) Under the independence restriction,

$$m_{HL,HL} = (1-p_1)q_2, \qquad m_{LH,LH} = q_1(1-p_2). \tag{16}$$

(b) (Coercive resolve index.)

$$\omega \;\equiv\; m_{HL,HL} - m_{LH,LH} \;=\; (1-p_1)q_2 - q_1(1-p_2). \tag{17}$$

(c) (Structural decomposition.) Under the MPE of the stochastic game,

$$\omega = \omega(\lambda_1, \lambda_2, c_1, c_2, \kappa_1(H), \kappa_2(H), \bar{\theta}_i, b, \gamma, g). \tag{18}$$

$\omega > 0$ iff Side 1 escalates ($\lambda_1 > \lambda_1^\dagger$) and Side 2 restrains ($\lambda_2 < \lambda_2^\dagger$); $\omega < 0$ iff Side 1 restrains and Side 2 escalates; $\omega = 0$ iff both occupy the same regime.

Proof. Part (a). Under independence, $m_{a'a} = P(a_1'|a_2) \cdot P(a_2'|a_1)$. For $HL \to HL$: Side 1 sees L and plays H with probability $1-p_1$; Side 2 sees H and plays L with probability q_2; so $m_{HL,HL} = (1-p_1)q_2$. $m_{LH,LH} = q_1(1-p_2)$ follows by symmetry.

Part (b). Subtract (16) directly.

Part (c). The MPE pins down $\xi_k : S \to A_k$, making (p_k, q_k) functions of structural primitives, establishing (18). For sign identification: when Side 1 escalates ($p_1 = q_1 = 0$) and Side 2 restrains ($p_2 = q_2 = 1$), $\omega = 1 \cdot 1 - 0 \cdot 0 = 1 > 0$. Reversing gives $\omega = 0 \cdot 0 - 1 \cdot 1 = -1 < 0$. When both occupy the same regime, $p_k = q_k$ and $\omega = (1-p_1)p_2 - p_1(1-p_2) = p_2 - p_1 = 0$ under symmetry. In the stochastic game, intermediate values arise as weighted averages over periods in each configuration. ∎

[52] At any fixed state $s_{it} = (\lambda_{it}, x_{it})$, Theorem 1 implies $p_k = q_k$. The inequality $p_k \neq q_k$ can hold in the time series because λ_{it} evolves stochastically: the opponent's lagged action proxies for the unobserved current state, and the reactive strategies encode equilibrium belief-updating, not strategic history-dependence.

Remark 2 (MPE existence)**.** The stochastic game admits a stationary MPE satisfying

$$V_k(s) = \max_{\rho_k}\left\{ v_k(s, \rho_k, \xi_{-k}(s)) + \delta \int V_k(s')\, P(ds' \mid s, \rho_k, \xi_{-k}(s)) \right\}.$$

Existence for finite state and action spaces follows from Fink (1964); for continuous state spaces from Mertens and Parthasarathy (1987) under norm-continuous transition probabilities, satisfied by g.

Remark 3 (Folk theorem)**.** In an infinitely repeated game against a fixed population, the folk theorem makes virtually any payoff vector sustainable as a subgame-perfect equilibrium, eliminating testable predictions. The stochastic game resolves this: time-varying s_{it} makes dominant actions pin down behavior uniquely at each state, and the observed transition dynamics of punishment profiles carry structural information about the prevailing regime.

Remark 4 (Limiting case $\gamma \to 0$)**.** As $\gamma \to 0$, (9) degenerates to a corner: $e_j^* \to 1$ if $(\lambda_1 - \theta_j)\rho_1 > (\lambda_2 - \theta_j)\rho_2$, and $e_j^* \to 0$ otherwise. The threshold (14) converges to $\bar{\theta}_i$ as $\gamma \to 0$, collapsing the critical selectivity to the mean civilian type. The qualitative predictions of Proposition 1 and Theorem 1 hold across the full range of γ.

Corollary 1 (Comparative statics on ω)**.**

(i) $\partial\omega/\partial\lambda_1 \geq 0,\ \ \partial\omega/\partial\lambda_2 \leq 0.$

(ii) $\partial\omega/\partial\bar{\theta}_i \gtreqless 0$ *as* $\lambda_1 \gtreqless \lambda_2.$

(iii) $\partial\omega/\partial\kappa_1(H) \leq 0,\ \ \partial\omega/\partial\kappa_2(H) \geq 0.$

(iv) $\partial\omega/\partial c_1 \geq 0,\ \ \partial\omega/\partial c_2 \leq 0.$

Proof. Parts (i)–(iii) follow from Proposition 2(c) and Theorem 1. For part (iv): $\partial\lambda_k^\dagger/\partial c_k < 0$ since both the numerator $\kappa_k(H) - c_k(1-\bar{\theta}_i)\Delta$ and denominator $b - c_k\gamma$ decrease in c_k, with net derivative negative under $c_k < b/\gamma$. A lower $\lambda_1^\dagger$ expands the region $\lambda_1 > \lambda_1^\dagger$, raises $m_{HL,HL}$ and ω. The symmetric argument for c_2 lowers $\lambda_2^\dagger$, raises $m_{LH,LH}$, reduces ω. □

A2 Empirical Appendix

This appendix provides supporting documentation for both empirical applications. Table A2 lists all conflict spells we considered for the cross-national sample, with superscripts indicating exclusion criteria. After manual review, we removed spells with no active conflict, spells where the conflict ended before the sample period begins, and borderline cases. Spells with no superscript enter the estimation sample. Table A3 summarizes model fit across these spells: each cell reports the percentage of spells in which a likelihood ratio test rejects the column model in favor of the row model at the 95% confidence level. Tables A4 and A5 present the full model comparison and $\hat{\omega}$ estimates for the Soviet-UPA application. Table A4 reports LRT statistics comparing each nested specification against its alternatives. Table A5 reports the corresponding $\hat{\omega}$ point estimates and 95% bootstrap confidence intervals alongside benchmarks from independent probit specifications.

Country	Period	Country	Period	Country	Period	Country	Period
Afghanistan	(1989–1992)	Cuba	(1990–2012)[a]	Jordan	(2016–2019)[a]	Philippines	(2014–2019)
	(1992–1996)	Djibouti	(1997–2001)	Kazakhstan	(1986–2019)[a]	Romania	(2018–2019)[a]
	(1996–2001)		(2001–2019)[a]	Kenya	(1989–2007)	Rwanda	(1990–1994)
	(2001–2019)	Dominican Republic	(1991–2017)[a]		(2008–2008)[b]		(1994–1994)
Angola	(1989–1992)	Algeria	(1990–1991)		(2008–2019)		(1994–1996)
	(1992–1992)		(1992–2002)	Kyrgyzstan	(2018–2019)[a]		(1996–2019)
	(1992–1994)		(2003–2019)[c]	Cambodia	(1989–1991)	Sudan	(1989–2005)
	(1994–2002)	Egypt	(1990–2010)		(1991–1997)		(2005–2011)
	(2002–2019)		(2011–2012)[a]		(1997–2019)[a]		(2011–2019)
Albania	(1996–1997)		(2013–2013)	Lebanon	(1989–1990)	Senegal	(1989–2014)
	(1998–2019)[a]		(2013–2019)		(1990–2005)[c]		(2014–2019)[a]
Azerbaijan	(1987–1991)	Eritrea	(1997–1998)		(2005–2006)	Sierra Leone	(1990–1996)
	(1992–1994)		(1998–2000)		(2006–2019)[c]		(1996–2002)
	(1994–2019)[c]		(2000–2018)[c]	Liberia	(1989–1996)		(2002–2019)[a]
Burundi	(1990–1993)	Spain	(1989–2010)		(1996–2003)	Somalia	(1989–1991)
	(1993–2000)		(2010–2019)[a]		(2003–2019)[b]		(1992–2005)
	(2000–2005)	Estonia	(1987–1992)[a]	Sri Lanka	(1989–2002)		(2005–2006)
	(2005–2019)	Ethiopia	(1989–1991)		(2002–2005)		(2006–2019)
Benin	(1991–2019)[a]		(1991–1998)		(2005–2009)	Serbia	(2018–2019)[a]
Burkina Faso	(1991–2014)[c]		(1998–2000)		(2009–2019)[b]	South Sudan	(2011–2013)
	(2015–2019)		(2000–2018)	Lithuania	(1987–2013)[a]		(2014–2018)
Bangladesh	(1970–1971)		(2018–2019)[b]	Latvia	(1986–1995)[a]		(2018–2019)[b]
	(1972–1991)	Gabon	(1990–2019)[a]	Morocco	(1997–2019)[a]	Eswatini	(1997–2019)[a]
	(1992–2019)[c]	United Kingdom	(1965–1998)	Madagascar	(1990–2009)	Syria	(2017–2019)
Bulgaria	(2012–2019)[a]		(1998–2017)[a]		(2009–2019)[a]	Chad	(1989–1990)
Bosnia	(1992–1995)	Ghana	(1991–2019)[a]	Mali	(1990–2011)		(1991–1996)
	(1995–2019)[b]	Guinea	(1991–2008)		(2012–2012)[b]		(1997–2008)
Belarus	(1987–2019)[a]		(2009–2019)		(2012–2019)		(2008–2019)
Botswana	(1990–2019)[a]	Gambia	(1997–2019)[a]	Myanmar	(2010–2019)	Togo	(1990–2019)[a]
CAR	(1990–2002)	Guinea-Bissau	(1995–1998)	Mozambique	(1989–1992)	Thailand	(1989–2004)
	(2003–2007)		(1998–1999)		(1992–2012)[b]		(2004–2019)
	(2008–2012)		(1999–2019)[a]		(2012–2019)	Tajikistan	(2018–2019)[a]
	(2013–2019)	Greece	(2018–2019)[a]	Mauritania	(1990–2007)	Tunisia	(1991–2010)[a]
China	(1989–2008)[c]	Guatemala	(1990–1996)		(2008–2019)[c]		(2011–2019)[c]
	(2009–2017)[c]		(1997–2017)[b]	Malawi	(1990–2019)[a]	Turkey	(2016–2019)
Ivory Coast	(1990–2002)	Croatia	(2018–2019)[a]	Namibia	(1997–2019)[a]	Tanzania	(1990–2019)[a]
	(2002–2007)	Haiti	(1990–1994)	Niger	(1989–2012)	Uganda	(1989–1994)
	(2007–2011)		(1994–2004)[c]		(2013–2019)		(1995–2006)
	(2011–2019)		(2004–2014)[c]	Nigeria	(1990–1999)		(2006–2019)[b]
Cameroon	(1990–2013)	Indonesia	(1989–1998)		(1999–2009)	Ukraine	(1987–2013)
	(2014–2019)		(1999–2005)		(2010–2019)		(2014–2014)[c]
DRC	(1989–1997)		(2005–2019)[b]	Nicaragua	(1989–1990)		(2014–2019)
	(1997–1998)	India	(1972–1984)		(1990–2018)[a]	Uzbekistan	(1987–2019)[a]
	(1998–2003)		(1985–1993)	Nepal	(1990–1996)	Vietnam	(2010–2019)[a]
	(2003–2019)		(1994–2001)		(1996–2006)	Yemen	(1994–1994)
Rep Congo	(1992–1997)		(2002–2019)[c]		(2006–2019)[a]		(1994–2003)
	(1997–1999)	Iran	(2016–2019)[a]	Pakistan	(1988–2001)		(2003–2011)
	(2000–2019)[b]	Iraq	(1989–1990)		(2001–2007)[c]		(2011–2014)
Colombia	(1989–2002)		(1990–1991)		(2007–2019)		(2014–2019)
	(2003–2010)		(1991–2003)	Peru	(1989–2000)	South Africa	(1990–1994)
	(2010–2016)		(2003–2011)		(2001–2017)[b]		(1994–2019)[b]
	(2016–2018)		(2012–2013)	Philippines	(1989–2001)	Zambia	(1990–2019)[a]
Costa Rica	(1993–2013)[a]		(2013–2019)		(2002–2014)	Zimbabwe	(1990–2000)[b]
							(2000–2008)
							(2008–2019)[a]

[a] Spell contains no active conflict (excluded). [b] Conflict ends before spell period (excluded). [c] Borderline case (excluded).

Table A2: **Conflict spells in the cross-national sample.** Superscripts mark exclusion criteria (see note). 138 spells with no superscript enter the estimation sample.

Table A3: **Cross-national bivariate probit model comparison**. LRT columns report percent of tests across 134 conflict spells where χ^2 statistic is significant at 95% level.

	Likelihood Ratio Test				
Model	v. 2-probit	v. indep.	v. add.	v. inter.	v. Mundlak
Independent	78.4	—	—	—	—
Additive	77.6	61.9	—	—	—
Interactive	75.4	61.2	23.9	—	—
+ Mundlak	73.9	89.6	86.6	88.8	—
+ Mundlak (2-way)	73.1	88.8	86.6	86.6	26.1

Table A4: **Bivariate probit model comparisons (Soviet-UPA)**. LRT columns report χ^2 statistic vs. named nested model. Significance: $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$.

	Likelihood Ratio Test				
Model	v. 2-probit	v. indep.	v. add.	v. inter.	v. Mundlak
Independent	361.1**	—	—	—	—
Additive	374.1**	207.7**	—	—	—
Interactive	371.9**	227.9**	20.2**	—	—
+ Mundlak	320.0**	1379.8**	1172.1**	1151.9**	—
+ Mundlak (2-way)	343.2**	1497.7**	1290.0**	1269.8**	117.9**

Table A5: **Coercive asymmetry estimates (Soviet-UPA)**, bivariate probit. $\widehat{\omega}$: coercive asymmetry. $\widehat{\omega}_{2p}$: independent probit benchmark. 95% CIs from parametric bootstrap.

	Coercive Asymmetry		Fit	
Model	$\widehat{\omega}$ (95% CI)	$\widehat{\omega}_{2p}$ (95% CI)	N	AIC
Independent	0.10 (0.08, 0.12)	0.02 (-0.04, 0.08)	20076	27329
Additive	0.17 (0.14, 0.21)	0.02 (-0.04, 0.07)	20076	27125
Interactive	0.18 (0.14, 0.23)	0.02 (-0.04, 0.07)	20076	27109
+ Mundlak	0.17 (0.13, 0.21)	0.02 (-0.04, 0.07)	20076	25961
+ Mundlak (2-way)	0.20 (0.16, 0.26)	0.02 (-0.04, 0.07)	20076	25847