

Snakes in the Backyard*

Ravi Ashok Satpute [†]

August 15, 2026

Job Market Paper
Latest Version ([link](#))

Abstract

Why do states empower domestic armed groups that may later challenge their power? I develop a signaling model in which a state with private information about its strength chooses between direct violence and delegation to a non-state violent proxy. Direct attacks are informative: success signals strength, and failure reveals weakness. Delegation to a proxy offers a type-independent success probability but endogenously creates a political claimant who may later contest power with private information about the cost of doing so. I characterize *Perfect Bayesian Equilibria* and show that in the key separating equilibrium, powerful states attack directly while weak states delegate, accepting blowback risk because direct action's type-dependent success probability makes it a materially inferior lottery relative to delegation with probabilistic blowback. The model provides a formal explanation for why proxy delegation is systematically associated with state weakness and endogenous political instability.

JEL Classification Codes: D72, D74, D82

Keywords: Proxy Warfare, Conflict, Endogenous Blowback

*I thank

[†]Indian Institute of Management Calcutta; Email: satputerafp21@email.iimcal.ac.in

1 Introduction

Why do states deliberately empower domestic violent groups that they eventually cannot fully control? From paramilitaries in Colombia and Mexico to Sudan’s Janjaweed, from Hezbollah in Lebanon to Russia’s Kadyrovtsy, governments across regime types tolerate, sometimes even cultivate, and deploy armed non-state groups that later extract concessions, become autonomous centers of coercion, or contest political power. The pattern is puzzling. Canonical theories of state formation emphasize the imperative to monopolize violence (Weber, 1946, Tilly, 1985, North et al., 2009, Myerson, 2008); states that tolerate rival coercive organizations may invite challenges to their authority and undermine the foundations of political order. Yet, the deliberate empowerment and usage of such violent non-state actors, what I term *proxy delegation*, is widespread and systematic.

The phenomenon spans historical periods and political systems. Autocracies deploy militias to enforce political control while maintaining distance from regimes: Syria’s Shabiha, Sudan’s Janjaweed, and Zimbabwe’s war veterans operated with state support but outside formal command structures (Ahram, 2011, Carey et al., 2015). Even some of the democracies have relied on violent groups to pursue objectives that formal security forces could not openly pursue: paramilitaries in Colombia and Northern Ireland, vigilantes in the Philippines and Mexico, operated in gray zones between state and non-state violence (Mazzei, 2009, Carey et al., 2013, 2015).¹ Acemoglu et al. (2013) document how the Colombian state’s electoral incentives and relationship with paramilitaries undermined its monopoly on violence. In most of these settings, the proxy relationship created risks that materialized over time: groups turned against patrons, demanded political concessions, representation, or leveraged their coercive capacity for autonomous ends.

The costs of such arrangements are substantial and often predictable. Proxies accumulate organizational capacity, territorial control, and popular legitimacy through the violence they conduct. Operational success increases popularity and smoothens the recruitment of new fighters, but, importantly, it also breeds ambition. Armed groups that prove effective at violence frequently discover that the same capabilities translate into political leverage; leverage that may be directed against the state itself. This dynamic, which I call *blowback*, represents a fundamental tension in proxy relationships: the very effectiveness that makes proxies useful also makes them dangerous for the state.

The literature offers three plausible explanations across several fields for state reliance on violent proxies. First, proxies are *militarily efficient* in a sense that they are cheaper, more specialized, possess local intelligence and ethnic networks, or are better suited to particular conflict environments than regular forces (Kalyvas, 2006, Hughes and Tripodi, 2009, Lyall, 2010). This explanation captures real advantages but cannot account for why

¹In the international context, the US relied on proxies in Nicaragua and Afghanistan in the 1980s.

states use and empower proxies when direct action is feasible, more likely to be effective, and especially when the proxy deployment creates long-run political costs that outweigh short-run military gains. A second explanation, *plausible deniability*, emphasizes that proxies allow states to pursue violent objectives with reduced liability by obscuring state involvement from international audiences concerned with human rights or sovereignty norms (Salehyan, 2010, Carey et al., 2015). By delegating violence, states may avoid accountability, escape sanctions, reputational costs, or intervention (Mitchell et al., 2014, Stanton, 2015, Carey et al., 2015). This account has merit for explaining the use of proxies in international conflicts or an alternate foreign policy, but it focuses almost exclusively on external audiences. It does not address the domestic political consequences of delegation, which are often more salient for regime survival than international opprobrium. Lastly, proxy reliance can be seen as evidence of state weakness or *capacity constraint*: governments delegate because they lack the capacity to project force directly and effectively (Besley and Persson, 2010, 2011). This explanation faces two difficulties. First, it risks circularity: if state weakness is inferred from the very behavior that it seeks to explain, the account has limited predictive content. The challenge is to specify an *independent* measure of state capacity and to show that it predicts delegation. Second, the explanation cannot account for why strong states also delegate, despite possessing overwhelming conventional advantages. Taken together, these explanations fail to offer a credible reconciliation of the puzzle: why would a rational state, aware of the blowback risk, deliberately create a potential political rival?

This paper provides a formal theory of endogenous rival creation through proxy delegation, rooted in a simple observation: violence modes differ in their informational content. Direct attacks have type-dependent success probabilities; delegated attacks do not. Delegated attacks decouple success from state capacity; their outcomes depend on the proxy’s effectiveness and yield no information about the state’s underlying type. I show that this informational asymmetry shapes the material attractiveness of each violence mode and endogenizes the severity of political blowback.

I develop a signaling model where an incumbent state has private information about its strength and chooses between direct attack and proxy delegation to capture resources from a target. Delegation requires transferring resources for sponsoring the attack and rewarding the proxy for captured resources. Direct action involves passive political risk: the outcome is publicly observed, and failure perfectly reveals the incumbent as weak through Bayesian updating. Following a failed direct attack, the target remains undefeated, and voters replace the incumbent with a baseline political alternative. Delegation utilizes a type-independent way to attack the target but involves active political risk. Following a successful operation, the proxy privately observes the cost of political activation, the organizational and reputational investments required to contest power. Proxy’s decision to challenge the state is itself endogenous. This creates uncertainty about blowback: the

state cannot perfectly predict whether delegation will trigger a political challenge. The outcome of this challenge depends on perceived state strength relative to the proxy's known capability.

The equilibrium analysis yields qualitatively distinct configurations. In the *separating equilibrium*, powerful states attack directly, and weak states delegate to the proxy. Under separation, violence mode perfectly reveals type: direct action signals strength, and delegation signals weakness. Voters correctly infer that direct attackers are powerful and that delegators are weak. The weak state prefers delegation due to a material lottery comparison generated by the signaling environment. Mimicking the strong state via direct attack is a dangerous gamble. Because strong states succeed with certainty, a direct attack failure leads to perfect revelation of weakness and certain displacement. The weak state concedes the informational battle and admits weakness by delegating to secure a higher probability of material victory against the target. It accepts informational exposure and the probabilistic proxy blowback to avoid the high probability of both informational and material failure. This mechanism, in which the act of delegation endogenously creates a political competitor whose threat severity depends on equilibrium beliefs and material outcomes, is new to the formal political economy literature. This equilibrium captures the core “snakes in the backyard” logic: a weak state rationally empowers a potential domestic rival to avoid exposing its own weakness. The state's short-run interest in surviving the current period politically conflicts with its long-run interest in maintaining a monopoly on political power. Delegation essentially trades off current political survival against future political competition.

Two pooling equilibria also exist. When the cost of direct action is sufficiently high relative to delegation, both types delegate. Because no information is revealed, no blowback arises in this equilibrium. When the cost of direct action is sufficiently low relative to delegation, both types attack directly, and delegation never occurs. The separating equilibrium arises for intermediate cost configurations, precisely when weak states are militarily disadvantaged enough that the proxy lottery dominates the direct-attack lottery despite the blowback risk. Interestingly, in the region of equilibrium multiplicity, where both separating as well as pooling-on-proxy equilibria exist, identical material fundamentals support both the separating and the blowback-free pooling equilibrium. Whether the state creates a political rival ultimately depends not on the conflict environment but on how domestic audiences coordinate their beliefs about proxy delegation. When delegation is considered as usual statecraft by the voters, proxy blowback may not arise. Semi-pooling equilibria, where the strong state plays a pure strategy and the weak state mixes, exist only at knife-edge parameter values and are non-generic.

Comparative statics provide important insights into proxy delegation. Higher proxy compensation reduces the probability of blowback by increasing the proxy's opportunity cost of political entry, as they forfeit these ongoing rents if they challenge the state.

This indicates a tradeoff between short-run resource extraction and long-run political survival. In addition, higher proxy inefficiency reduces the parameter region supporting the separating equilibrium. A highly ineffective proxy poses less of a political threat, but its military weakness removes the material advantage of delegation. This forces the weak state to use direct action. The relevant constraint on blowback is therefore the state's incentive to delegate in the first place and not just the proxy's ability to win elections.

These results have several important implications. First, proxy deployment should correlate with incumbent weakness, especially where revealing weakness is politically costly. Strong states have less to hide and face lower costs from revealing their type through direct action. This prediction distinguishes the theory from capacity-constraint explanations, which predict a monotonic relationship between state weakness and proxy reliance. Here, weak states delegate not just because they lack the capacity to act directly, but because revealing their type is politically costlier than creating a rival. Second, blowback should be systematically associated with state-sponsored delegation, specifically, because the act of delegation endogenously creates the political conditions, equilibrium beliefs, and demonstrated proxy competence, under which a challenge to the state becomes viable. Groups used by the state are particularly likely to challenge their patrons because the information structure of the delegation game creates conditions under which blowback is both possible and, in equilibrium, anticipated. Third, the terms of proxy contracts matter for political stability: proxies that receive generous shares of captured resources have weaker incentives to contest political power, suggesting a tradeoff between short-run resource extraction and long-run political stability.

The paper engages with several literatures. *Political agency and accountability*: A central theme in political economy concerns how voters discipline politicians through reelection incentives and how politicians respond strategically to electoral pressure (Besley, 2006, Ashworth, 2012). This literature typically studies information revelation about types through policy choices or observable outcomes. I extend these ideas to settings where incumbents choose the *mode of action*, and where the equilibrium interpretation of that choice endogenously determines the severity of political consequences. Delegation is not merely a principal-agent problem; it is a strategic choice whose political ramifications can depend on whether voters interpret it as a signal of weakness or as routine statecraft.

My mechanism is related to models of *strategic information transmission and media capture* (Besley and Prat, 2006). In those settings, politicians manipulate information flows to influence voter beliefs. In my setting, politicians manipulate information by choosing actions whose outcomes are more or less diagnostic of their type. The key innovation is that the informational properties of actions are themselves strategic objects.

Signaling and information in conflict: A growing literature studies how private information shapes conflict behavior (Fearon, 1995, Powell, 2006, Fey and Ramsay, 2011, Baliga and Sjöström, 2004, 2012). These papers typically consider bargaining between

adversaries with private information about resolve or capability. I study a slightly different problem: how an incumbent’s private information about its *own* strength shapes its choice of violence mode, with consequences for domestic rather than international audiences, which ultimately determine the incumbent regime’s survival.

Proxy warfare and non-state violence: An empirical literature documents patterns of state-proxy relationships, identifying correlates of militia formation, deployment, and behavior (Salehyan, 2010, Carey et al., 2015, Berman and Lake, 2019). This work provides rich descriptive evidence but lacks a unified theoretical framework for understanding when and why states delegate. This paper provides such framework with predictions about the conditions for violence delegation and the factors that determine blowback risk.

Endogenous political competition: Standard models of electoral competition and political accountability take the set of political contestants as given (Downs, 1957, Alesina, 1988, Persson et al., 1997, Lizzeri and Persico, 2001). A related strand on endogenous candidacy allows citizens to self-select into political competition (Osborne and Slivinski, 1996, Besley and Coate, 1997). Egorov and Sonin (2011) study a dictator’s choice of advisors who may become political threats, showing a competence-loyalty tradeoff: more capable agents are more useful but also more dangerous. My paper shares this “creating your own rival” logic but operates through a fundamentally different channel. In Egorov and Sonin (2011), the agent’s threat arises from exogenous competence; in my model, the proxy becomes a rival *endogenously* through the information revealed due to *proxy delegation*. Another closest antecedent to my mechanism is Powell (2013), who studies an infinite-horizon model in which a ruler decides whether and how quickly to consolidate power by disarming an armed opposition. In Powell’s framework, the ruler cannot credibly commit to future transfers: once the opposition’s military capacity is reduced, the ruler has an incentive to renege on promised concessions. This commitment problem over resource-sharing generates inefficient fighting, and the outcome depends on whether “contingent spoils” (economic gains) from monopolizing violence are large enough to justify a costly conflict rather than a peaceful buyout. My paper differs in two aspects. First, the driving friction is *informational*, not contractual: the state delegates to hide weakness from domestic audiences, not because it cannot commit to transfers. Second, blowback in my model operates through an electoral channel mediated by voter’s beliefs, whereas in Powell’s framework, rivalry arises from the coercive capacity that armed agents accumulate. More broadly, in my model, the proxy becomes a political rival *because of* delegation, the act of outsourcing violence creates the competitor.

The remainder of the paper proceeds as follows. Section 2 presents the model, defining players, information structure, timing, and payoffs. Section 3 characterizes equilibria, establishing existence conditions for separating, pooling, and semi-pooling configurations. Section 4 derives essential comparative statics, Section 5 discusses extensions, endogenous surplus sharing with optimal contracts, and an outside option. Finally, Section 6 concludes.

2 The Model

2.1 Players and Environment

There are two strategic players and one passive actor:

- A **State** (S), the incumbent political power with private information about its strength.
- A **Proxy** (P), a non-state violent organization that can be enlisted for violence.
- A **Target** (T), a passive actor controlling resources that can be captured through violence.

In addition, a representative **Voter** (V) observes outcomes and decides whether to retain the incumbent. There is a fixed political-economic surplus (the *pie*) that is normalized to size one. Initially, the target controls a share $\tau \in (0, 1)$ of the pie, while the state controls the remaining share $1 - \tau$. The proxy controls no share initially. Following violence outcomes, voters observe public signals, update their beliefs, and decide whether to retain the incumbent or replace it. The game unfolds over three main stages:

Stage 0. Nature: Nature draws the state's strength $\theta \in \{\theta_L, \theta_H\}$ and the proxy's political activation cost c_A , both of which are private information to the respective players.

Stage 1. Attack Mode: The state observes θ and chooses how to attack the target.

Stage 2. Violence: The attack outcome is realized and publicly observed.

Stage 3. Political Contest: If the proxy was used and succeeded, it may challenge the state; voters decide who governs.

I now describe each component of the model in detail.

2.2 Information Structure

2.2.1 State Strength

At the beginning of the game, nature draws the state's strength, which is a binary variable $\theta \in \{\theta_L, \theta_H\}$, where θ_L denotes a *weak* state, and θ_H denotes a *powerful* state. I normalize $\theta_L = 0$ and $\theta_H = 1$. The prior probability that the state is powerful is $\Pr(\theta = \theta_H) = p \in (0, 1)$. The realization of θ is the state's private information. The proxy and voters share the common prior p .

2.2.2 Proxy Strength

While the state's strength is private, the proxy's strength is publicly known, denoted by $\gamma = \frac{1}{1+\lambda}$. Here, $\lambda > 0$ measures how incompetent or ineffective the proxy is. Thus, a higher λ means an ineffective and weaker proxy. Note that the proxy's strength γ lies in $(0, 1)$. A fully competent proxy ($\lambda \rightarrow 0$) has $\gamma \rightarrow 1$; an incompetent proxy ($\lambda \rightarrow \infty$) has $\gamma \rightarrow 0$.

2.2.3 Proxy's Private Information

The proxy has its own source of private information. If the proxy is used and the attack succeeds, the proxy may choose to *activate* politically (challenge the state). Political activation entails a cost c_A (drawn by the nature) that is the proxy's private information:

$$c_A \sim H \text{ on } [\underline{c}, \bar{c}], \quad 0 < \underline{c} < \bar{c}, \quad (2.1)$$

where H is a continuously differentiable *CDF* with full support and density $h > 0$. I assume $\underline{c} < 1 - k\tau < \bar{c}$, ensuring $\pi \in (0, 1)$, where k is the post-violence share offered by the state to the proxy (discussed in Section 2.7.1).

This two-sided private information, the state over its strength, and the proxy over its activation cost, generates a genuine strategic uncertainty on both sides of the delegation relationship and the proxy blowback.

2.3 Stage 1: Attack Mode

The state observes its type θ and chooses an attack mode $m \in \{D, P\}$, where:

- $m = D$ (Direct attack): The state conducts a violent attack using its own forces, incurring a cost $c_D \geq 0$.
- $m = P$ (Proxy attack): The state delegates violence to the proxy, transferring resources $O = c_P$ to cover the proxy's operational cost $c_P > 0$.

An attack necessarily occurs; there is no option to refrain from violence.² This reflects settings where the state has already committed to addressing the target. If the proxy is used and the attack succeeds, the state retains a share $(1 - k)$, and the proxy receives a share $k \in (0, 1)$ of the captured surplus τ , where k is a fixed and known parameter of the proxy contract.³

2.4 Stage 2: Violence and Outcomes

The outcome of the attack depends on the mode chosen and, in the case of a direct attack, on the state's type.

2.4.1 Direct Attack

If the state chooses $m = D$, the attack succeeds with a probability that depends on the state strength:

$$p_D(\theta) = \begin{cases} 1 & \text{if } \theta = \theta_H, \\ \alpha & \text{if } \theta = \theta_L, \end{cases} \quad (2.2)$$

²In Section 5.2, I discuss a no attack option.

³In Section 5.1, I discuss endogenous k with optimal contracts.

where $\alpha \in (0, 1)$. Thus, a powerful state always succeeds with a direct attack. A weak state, on the other hand, succeeds with probability $\alpha < 1$. If a direct attack succeeds, the share τ (initially held by the target) is transferred from the target to the state. If the attack fails, no redistribution occurs. Crucially, the dependence of success probability on θ makes direct attack outcomes *informative* about state strength.

2.4.2 Proxy Attack

If the state chooses $m = P$, the proxy attacks the target. The attack succeeds with probability $p_P = \rho\gamma = \frac{\rho}{1+\lambda}$, where $\rho \in (0, 1)$ is a fixed operational effectiveness capturing the proxy's baseline capability in conducting violence (e.g., quality of local intelligence, organizational cohesion, terrain familiarity). The proxy success probability is *independent* of state strength. It depends only on the proxy's inherent characteristics (ρ, γ) , not on the state's type θ . This is crucial and implies that proxy attack outcomes are uninformative about θ , implying that the proxy delegation decouples attack outcomes from state strength. The resource transfer c_P covers the proxy's operational costs (logistics, ammunition, personnel) but does not affect its probability of success. If the proxy attack succeeds, the share τ is captured from the target. If the attack fails, no redistribution occurs.

Note that since $\rho < 1$ and $\lambda > 0$, we always have $p_P < 1 (= p_D(\theta_H))$. Thus, the proxy is strictly less effective than a powerful state, and a powerful state has no capacity constraint.

2.5 Public Signals

After the violence outcome is realized, voters observe a public signal consisting of both the mode of attack and whether it succeeded:

$$\sigma \in \Sigma = \{(D, S), (D, F), (P, S), (P, F)\}, \quad (2.3)$$

where the first component indicates the attack mode and the second indicates success (S) or failure (F). Voters use this signal to update their beliefs about the state's type before the political contest.

2.6 Belief Updating

Let $s_\theta \in \{D, P\}$ denote the state's strategy (attack mode) as a function of type. Voters update their beliefs about the state's type using Bayes' rule. Define:

$$\Theta_D = \{\theta : s_\theta = D\}, \quad (2.4)$$

$$\Theta_P = \{\theta : s_\theta = P\}. \quad (2.5)$$

Definition 1 (Posterior Belief). *The posterior belief after signal σ is:*

$$\mu(\sigma) = \Pr(\theta = \theta_H \mid \sigma). \quad (2.6)$$

Given the normalization $\theta_H = 1$, $\theta_L = 0$, the posterior expected strength equals the posterior belief, implying $\mathbb{E}[\theta \mid \sigma] = \mu(\sigma)$. The following lemmas establish the key informational asymmetry between attack modes.

Lemma 2.1 (Posteriors under $m = P$). *If the state plays $m = P$, the attack outcome is uninformative about θ :*

$$\mu(P, S) = \mu(P, F) = \Pr(\theta_H \mid \theta \in \Theta_P). \quad (2.7)$$

Lemma 2.2 (Posteriors under $m = D$). *If both types play $m = D$, the posteriors satisfy:*

$$\mu(D, S) = \frac{p}{p + (1 - p)\alpha}, \quad (2.8)$$

$$\mu(D, F) = 0. \quad (2.9)$$

If only the weak type plays $m = D$: $\mu(D, S) = \mu(D, F) = 0$. If only the powerful type plays $m = D$: $\mu(D, S) = 1$, and $\mu(D, F)$ is off-path.

The proofs of both lemmas are straightforward applications of Bayes' rule.

Corollary 1 (Informativeness). *Direct attacks are informative: $\mu(D, S) \neq \mu(D, F)$ generically. Proxy attacks are uninformative: $\mu(P, S) = \mu(P, F)$.*

This asymmetry is the model's informational foundation. As proxy attacks do not update beliefs about state strength, whereas direct attacks do, the choice of attack mode is itself a signal, giving rise to the strategic considerations analyzed in the equilibrium.

2.7 Stage 3: Political Contest

Stage 3 is reached after the signal σ is publicly observed. Political consequences arise through two distinct channels, depending on the mode of violence.

If the state attacked directly ($m = D$), voters observe the outcome and evaluate the incumbent. Since direct attack outcomes are informative about state strength (Corollary 1), voters update their beliefs and apply the retention rule: the state is retained if $\mu(\sigma) \geq \gamma$ and displaced otherwise. The proxy plays no active role in this process; it was not involved and has not demonstrated competence. The political risk after direct attack is voter-driven: it arises because the state's own performance is publicly observable, and poor performance generates evidence of weakness that voters act on through domestic political processes.

If the state delegated to the proxy ($m = P$) and the attack succeeded ($\sigma = (P, S)$), a qualitatively different political process unfolds. The proxy has demonstrated military effectiveness, captured resources, and gained public attention with visible successful outcome, providing the foundation for mounting a credible political challenge. Whether this challenge materializes depends on the proxy's private activation cost (Section 2.7.1). If the proxy activates, voters adjudicate the contest using the retention rule $\mu(\sigma) \geq \gamma$. If the proxy does not activate, the state retains power.

If the proxy attack failed ($\sigma = (P, F)$), the proxy lacks the demonstrated competence

and captured resources to mount a viable challenge, and the state retains power. Proxy outcomes are uninformative about state strength (Lemma 2.1), so the proxy's failure does not generate public evidence of state weakness in the way that direct failure does.

The two channels reflect a substantive asymmetry between *direct accountability* and *proxy-mediated challenge*. When the state acts directly, its competence is publicly tested, and poor performance triggers voter-driven political consequences through the retention rule. When the state delegates, it maintains distance from the operation's outcome; political risk arises from the proxy's endogenous decision to contest power following a successful operation instead of voter accountability. This asymmetry is central to the model's mechanism: direct attack exposes the state to passive political risk (voter evaluation), while delegation exposes it to active political risk (proxy challenge).

2.7.1 Proxy's Activation Decision

If the proxy is used and succeeds, it must decide whether to accept the fixed share offered by the state or to fight the political contest. Formally, conditional on $\sigma = (P, S)$, the proxy privately observes c_A and chooses:

- **Inactivity** (I): Accept share $k \in (0, 1)$ of the captured surplus, yielding payoff $k\tau$.
- **Activation** (A): Challenge the state for control, incurring cost c_A .

2.7.2 Electoral Rule

Voters evaluate the incumbent by comparing their posterior belief about state strength to the proxy's known strength. This rule governs both voter-driven evaluation after direct attack and the adjudication of a proxy-initiated political challenge after delegation:

$$\text{Voters retain state} \iff \mu(\sigma) \geq \gamma. \quad (2.10)$$

- If $\mu(\sigma) \geq \gamma$: State wins, retains full pie.
- If $\mu(\sigma) < \gamma$: Proxy wins, gains full pie.

The microfoundation of this electoral rule follows from the voter's payoff, which is presented in Section 2.8.3. The voter's retention decision is straightforward: given posterior $\mu(\sigma)$, retaining the state is optimal if and only if $\mu(\sigma) \geq \gamma$, and the optimal action is unique whenever $\mu(\sigma) \neq \gamma$; at $\mu(\sigma) = \gamma$ the voter is indifferent, I break the tie in favor of retention. The electoral rule is therefore deterministic:⁴ the voter's response to any signal, on or off the equilibrium path, is an element of $\{\text{retain}, \text{displace}\}$. This convention is maintained throughout the equilibrium analysis. The refinement test in Appendix A.9 follows the standard formulation of D1, quantifying over the voter's full best-response correspondence,

⁴I adopt the deterministic rule for tractability to isolate the informational mechanism as cleanly as possible. Under a probabilistic contest technology (e.g., the proxy wins with probability increasing in γ/μ or $\mu < \gamma$), the proxy's activation threshold would become belief-dependent, and the stark outcome where any activating proxy defeats a weak delegator (Remark 3) would be attenuated. The qualitative structure would still survive, but the all-or-nothing electoral outcome would be replaced by graded risk.

which admits randomization at the indifference belief $\mu = \gamma$. The test, therefore, admits a wider set of responses than the deterministic rule maintained in equilibrium play.

Definition 2 (Political Prize). *If the proxy wins the political contest, it replaces the state entirely as the governing authority and controls the full post-attack surplus:*

$$w = (1 - \tau) + \tau = 1. \quad (2.11)$$

This political prize definition essentially implies that the political victory confers complete governing authority. The proxy, having demonstrated both military effectiveness through the successful attack and political viability through the electoral contest, displaces the incumbent and inherits control over the entire political-economic surplus. Symmetrically, an incumbent displaced through voter evaluation, whether after a failed or a militarily successful direct attack, forfeits all governing authority and retains no share of the surplus, including any resources just captured, so its terminal share is zero with attack costs being sunk. This sharpens the state's tradeoff between delegation and political survival, especially when the political contest is decisive.⁵

2.7.3 Proxy's Activation Threshold

Since the proxy privately observes its political activation cost, even though its strength is publicly observable, the decision to fight a political contest remains unpredictable from the state's perspective. Even for the proxy, the political win may not be worthwhile if the political prize comes at a very high cost, c_A .

Lemma 2.3 (Proxy Activation Condition). *Let $\mu^* = \mu(P, S)$. The proxy activates iff:*

$$c_A < \hat{c}(\mu^*) \equiv \begin{cases} w - k\tau = 1 - k\tau & \text{if } \mu^* < \gamma, \\ -k\tau & \text{if } \mu^* \geq \gamma. \end{cases} \quad (2.12)$$

Since $c_A > \underline{c} > 0$ and $-k\tau < 0$, the proxy never activates when $\mu^* \geq \gamma$.

Proof. If $\mu^* \geq \gamma$: activation leads to a certain political loss. Payoff is $0 - c_A < 0$, while inactivity gives $k\tau > 0$. If $\mu^* < \gamma$: activation leads to a certain political victory. Payoff is $w - c_A = 1 - c_A$. Proxy activates iff $1 - c_A > k\tau$, i.e., $c_A < 1 - k\tau$. $\square$

Definition 3 (Activation Probability). *When $\mu^* < \gamma$, the probability that the proxy activates politically is:*

$$\pi = H(1 - k\tau). \quad (2.13)$$

When $\mu^* \geq \gamma$, $\pi = 0$.

The activation probability π , captures the likelihood of blowback conditional on proxy success and is determined by the distribution of the proxy's private activation cost.

⁵In practice, power transitions may be partial or shared.

2.8 Payoffs

I now derive the expected payoffs for each player. These payoffs depend on the attack mode, violence outcome, and the realization of the political contest.

2.8.1 State's Payoff

Under direct attack ($m = D$): The attack costs c_D and succeeds with probability $p_D(\theta)$. Voters observe the outcome and evaluate the incumbent using the retention rule $\mu(\sigma) \geq \gamma$. Success leads to resource capture τ ; failure leaves the distribution unchanged.

Assumption 1 (Prior Strength). *The prior probability that the state is powerful exceeds the proxy's political strength:*

$$p > \gamma. \quad (2.14)$$

This restriction focuses the analysis on environments where voters who observe no informative signal would retain the state over the proxy, so that political displacement requires evidence of weakness.⁶

Remark 1 (Retention after Direct Attack). *Under Assumption 1, the Bayesian posterior after direct success satisfies $\mu(D, S) > p > \gamma$ whenever θ_H plays D with probability one. To see this, note that $\mu(D, S) = p/(p + (1-p)q\alpha)$ for some $q \in [0, 1]$ representing the probability that the weak type plays D , and since $(1-p)q\alpha < 1-p$, we have $p + (1-p)q\alpha < 1$, so $\mu(D, S) > p > \gamma$. Consequently, a state that succeeds under direct attack is always retained. Conversely, $\mu(D, F) = 0 < \gamma$ whenever failure is on-path (since only the weak type can fail), so a state that fails under direct attack is always displaced. When direct attack is off-path, $\mu(D, F) = 0$ remains forced by consistency, while $\mu(D, S)$ is unrestricted by Bayes' rule; Appendix A.9 pins it down whenever the restriction is payoff-relevant.*

By Remark 1, the powerful state succeeds with certainty and is retained. The weak state succeeds with probability α and is retained, but fails with probability $1 - \alpha$ and is displaced, losing its initial share $(1 - \tau)$:

$$U_S^D(\theta_H) = 1 - c_D, \quad (2.15)$$

$$U_S^D(\theta_L) = \alpha \cdot (1 - c_D) + (1 - \alpha) \cdot (-c_D) = \alpha - c_D. \quad (2.16)$$

The powerful state's payoff is unchanged from a model without voter evaluation: it always succeeds and is always retained. The weak state's payoff, however, reflects the political cost of direct action. Compared to a setting without voter evaluation, the weak state's expected payoff is lower by $(1 - \alpha)(1 - \tau)$: with probability $1 - \alpha$ the attack fails, voters observe $\mu(D, F) = 0 < \gamma$, and the state loses its initial share $1 - \tau$ through political displacement. This is the cost of direct accountability: the informativeness of direct attack

⁶With unfavorable priors, $p \leq \gamma$, even uninformative pooling on proxy may generate on-path blowback, producing a qualitatively different equilibrium structure, and is not analyzed here.

(Corollary 1) has payoff consequences because voter beliefs are tied to political survival.

Under proxy attack ($m = P$): The state pays c_P and the outcome depends on both the proxy's military success and, if successful, the subsequent political contest. Let $\mu^* = \mu(P, S)$ denote the posterior belief upon observing a successful proxy attack. There are two cases:

Case 1: $\mu^* \geq \gamma$ (no blowback). The proxy never activates (Lemma 2.3). With probability p_P the attack succeeds and the state retains share $(1 - k)\tau$:

$$U_S^P = (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (2.17)$$

Case 2: $\mu^* < \gamma$ (blowback possible). With probability $(1 - p_P)$, the attack fails and the state retains $(1 - \tau) - c_P$. With probability p_P , the attack succeeds. Now with probability $(1 - \pi)$, the proxy stays inactive; with probability π , the proxy activates and wins (since $\mu^* < \gamma$), leaving the state with $-c_P$:

$$\begin{aligned} U_S^P &= (1 - p_P)[(1 - \tau) - c_P] \\ &\quad + p_P(1 - \pi)[(1 - \tau) - c_P + (1 - k)\tau] \\ &\quad + p_P \cdot \pi \cdot [-c_P] \end{aligned} \quad (2.18)$$

$$= (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (2.19)$$

Denoting the net expected gain with proxy success by $B \equiv (1 - \pi)(1 - k)\tau - \pi(1 - \tau)$, the payoff is $U_S^P = (1 - \tau) - c_P + p_P B$. This expression is intuitive and it allows us to capture probable scenarios when the delegation is likely to destroy the value. The state's net expected gain from proxy success can be expressed as the difference between two competing forces: the expected resource capture when the proxy remains inactive, $(1 - \pi)(1 - k)\tau$, and the expected loss of the state's initial endowment when blowback occurs, $\pi(1 - \tau)$. Delegation is beneficial only when the former exceeds the latter. Conversely, delegation destroys value for the state when:

$$(1 - k) < \frac{\pi}{1 - \pi} \cdot \frac{1 - \tau}{\tau}. \quad (2.20)$$

The left-hand side captures the state's share of the spoils; the right-hand side is the product of the odds of blowback and the ratio of the state's initial endowment to the contested resources. Delegation is more likely to destroy value when blowback is probable (π high), the state's share of captured surplus is small (k high), and the state has more to lose from political displacement (τ low).

2.8.2 Proxy's Payoff

If $m = D$: $U_P = 0$ (the proxy is not involved). If $m = P$ and the attack fails: $U_P = O - c_P = c_P - c_P = 0$ (the transfer $O = c_P$ covers exactly the operational cost). If $m = P$ and attack succeeds:

- Inactive: $U_P = k\tau$.

- Activates and wins: $U_P = w - c_A = 1 - c_A$.
- Activates and loses: $U_P = -c_A$.

2.8.3 Voter's Payoff

Voters prefer to be governed by a stronger entity. Formally, the representative voter's payoff equals the strength of whoever holds power:

$$U_V = \begin{cases} \theta & \text{if the state governs,} \\ \gamma & \text{if the proxy governs.} \end{cases} \quad (2.21)$$

Given posterior belief $\mu(\sigma) = \Pr(\theta = \theta_H \mid \sigma)$, the voter's expected payoff from retaining the state is $\mathbb{E}[\theta \mid \sigma] = \mu(\sigma)$, while the expected payoff from electing the proxy is γ . Hence, voters optimally retain the state if and only if $\mu(\sigma) \geq \gamma$. This microfounds the electoral rule specified in Section 2.7.2.

2.9 Strategies and Equilibrium

The representative voter follows a deterministic retention rule. If the state uses a direct attack, for (D, S) , the incumbent is retained, but for (D, F) , $\mu(D, F) = 0$ and the incumbent is replaced by a baseline alternative. If the state delegates to the proxy ($m = P$) and the attack succeeds, the proxy decides whether to challenge the state. If the proxy remains inactive (I), the incumbent state is retained. If the proxy actively challenges (A), the voter retains the incumbent if and only if the state's expected strength meets or exceeds the proxy's known capability, that is $\mu(\sigma) \geq \gamma$.

Having described the model environment and voter's rule, I now define the solution concept. The equilibrium concept is *Perfect Bayesian Equilibrium* (PBE). Formally, strategy and the equilibrium are defined as follows.

Definition 4 (Strategy). *A (behavioral) strategy for the state is a mapping $s : \{\theta_L, \theta_H\} \rightarrow [0, 1]$, where $s(\theta)$ is the probability of choosing $m = D$. A pure strategy sets $s(\theta) \in \{0, 1\}$ for each θ . A strategy for the proxy is a mapping $\beta : \Sigma \times [\underline{c}, \bar{c}] \rightarrow \{I, A\}$. Voters follow a deterministic rule to retain the incumbent state if and only if $\mu(\sigma) \geq \gamma$.*

I focus on pure strategies; Section 3.5 analyzes the semi-pooling case where θ_L mixes.

Definition 5 (Perfect Bayesian Equilibrium). *A Perfect Bayesian Equilibrium (PBE) consists of strategies (s^*, β^*) and beliefs μ^* such that:*

1. *For each θ , $s^*(\theta)$ maximizes the state's expected payoff.*
2. *For each σ and c_A , $\beta^*(\sigma, c_A)$ maximizes the proxy's expected payoff.*
3. *$\mu^*(\sigma)$ is derived via Bayes' rule for on-path signals. For off-path signals, beliefs are unrestricted (any $\mu \in [0, 1]$ is admissible)*

3 Equilibrium Analysis

I now characterize the Perfect Bayesian Equilibria of this game. Note that the voter follows the deterministic retention rule. With binary types, there are four candidate pure-strategy profiles for the state. I also examine whether semi-pooling equilibria, in which one type mixes, can arise.⁷ Before analyzing each equilibrium candidate, Lemma 3.1 establishes the proxy’s best response, which applies across all equilibrium configurations.

3.1 Proxy’s Best Response

The proxy’s equilibrium strategy follows directly from the activation condition derived in Lemma 2.3. Since the proxy moves after observing the public signal and its private cost, its optimal behavior can be stated in closed form.

Lemma 3.1 (Proxy’s Equilibrium Strategy). *In any equilibrium, the proxy’s strategy is:*

$$\beta^*(\sigma, c_A) = \begin{cases} A & \text{if } \sigma = (P, S), \mu(P, S) < \gamma, c_A < 1 - k\tau, \\ I & \text{otherwise.} \end{cases} \quad (3.1)$$

Proof. Immediate from Lemma 2.3. □

The lemma establishes that the proxy’s behavior is fully determined by the posterior belief $\mu(P, S)$ and the realized cost c_A . The state’s equilibrium problem therefore reduces to choosing an attack mode, taking as given that delegation creates blowback risk whenever $\mu(P, S) < \gamma$. I now analyze each candidate equilibrium in turn, beginning with the separating equilibrium.

3.2 Separating Equilibrium: $\theta_H \rightarrow D, \theta_L \rightarrow P$

The separating equilibrium is the most substantively important configuration. The powerful state signals strength through direct action, while the weak state delegates to the proxy, exploiting the type-independent success probability of proxy violence. This is the equilibrium that generates the “snakes in the backyard” dynamic: weak states rationally create political rivals by delegating violence. Note that in the separating equilibrium, delegation does not conceal the state’s type from voters; voters correctly infer that delegators are weak. The strategic logic is that weak states prefer the type-independent lottery of proxy violence with probabilistic blowback to the type-dependent lottery of direct violence. The signaling structure operates on both branches of the state’s decision. On the direct-attack branch, the informativeness of violence outcomes has payoff consequences through voter evaluation: failure reveals weakness and triggers political displacement (Remark 1). On the delegation branch, equilibrium beliefs endogenize blowback severity

⁷Off-path beliefs are disciplined by the D1 criterion, developed in Appendix A.9; the uniqueness statements in Section 3.6 refer to pure-strategy PBE satisfying D1.

and generate the multiplicity region (Section 3.6) where identical fundamentals support both a blowback and a blowback-free equilibrium. When α is low, the direct-attack lottery is materially inferior, both because the expected capture $\alpha\tau$ is small and because the political cost of failure $(1 - \alpha)(1 - \tau)$ is large, making delegation attractive despite the blowback risk it creates.

Proposition 3.1 (Separating Equilibrium Characterization). *Consider the strategy profile $(s(\theta_H) = D, s(\theta_L) = P)$ with off-path belief $\mu(D, F) = 0$, which is forced by consistency since only θ_L can fail under direct attack. This is a PBE if and only if the following conditions are simultaneously satisfied:*

$$c_D - c_P \leq \tau - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (3.2)$$

$$c_D - c_P \geq \alpha - (1 - \tau) - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (3.3)$$

where $\pi = H(1 - k\tau)$ is the blowback probability.

Proof. In Appendix A.1. □

Condition (3.2) ensures that the powerful state prefers direct attack over delegation. If θ_H deviates to P , it faces the equilibrium belief $\mu(P, S) = 0$ and the attendant blowback risk. Since θ_H succeeds with certainty under D and captures τ with guaranteed retention, this holds when $c_D - c_P$ is not too large. Condition (3.3) ensures that the weak state prefers delegation despite blowback. Under direct attack, the weak state's expected payoff is $\alpha - c_D$: it succeeds with probability α and is retained, but fails with probability $1 - \alpha$ and is displaced. When α is low, the direct-attack lottery is unattractive both because expected capture is small and because political displacement is likely. Together, the two conditions require $c_D - c_P$ to lie in an interval of width $1 - \alpha$, which is positive for all $\alpha < 1$. Lower α widens this interval: the separating equilibrium is sustained over a larger range of cost differentials when the weak state is more militarily disadvantaged.

3.3 Pooling on Proxy: $\theta_H \rightarrow P, \theta_L \rightarrow P$

When a direct attack is sufficiently costly, both types may prefer to delegate. In this case, the posterior equals the prior, and the proxy's activation decision depends on how the prior compares to proxy strength.

Proposition 3.2 (Pooling on Proxy Equilibrium). *Consider the strategy profile $(s(\theta_H) = P, s(\theta_L) = P)$. Under this profile:*

$$\mu(P, S) = \mu(P, F) = p. \quad (3.4)$$

This is a PBE if Assumption 1 holds and:

$$c_D - c_P \geq \tau - p_P(1 - k)\tau (= \bar{\Delta}_P) \text{ i.e., } c_D \text{ is sufficiently high favoring delegation} \quad (3.5)$$

Proof. In Appendix A.2 □

When both types delegate, the proxy attack conveys no information about θ , so the posterior remains at the prior p . If $p > \gamma$, voters perceive the state as stronger than the proxy, and the proxy never activates regardless of c_A . This is essentially a *blowback-free* delegation. The equilibrium is therefore blowback-free: delegation is “safe” for both types. However, this requires two conditions. First, the state must be perceived as strong enough *a priori* ($p > \gamma$). Second, a direct attack must be costly enough that even a powerful state, which would succeed with certainty, prefers the expected payoff from delegation. Notice that this condition coincides with $\Delta \geq \bar{\Delta}_P$ (defined in Section 3.6), while the separating equilibrium requires $\Delta \leq \bar{\Delta}_S$ (condition (3.2)). Since $\bar{\Delta}_P < \bar{\Delta}_S$, there is an overlap region $\bar{\Delta}_P < \Delta < \bar{\Delta}_S$ where both equilibria coexist, provided $\underline{\Delta} < \bar{\Delta}_P$ (which holds under Assumption 2; see Section 3.6).

3.4 Pooling on Direct: $\theta_H \rightarrow D, \theta_L \rightarrow D$

At the other extreme, when direct attack is cheap, both types may prefer to act directly, forgoing delegation entirely.

Proposition 3.3 (Pooling on Direct Equilibrium). *Consider $(s(\theta_H) = D, s(\theta_L) = D)$. Under this profile:*

$$\mu(D, S) = \frac{p}{p + (1 - p)\alpha}, \quad (3.6)$$

$$\mu(D, F) = 0. \quad (3.7)$$

This is a PBE if:

$$c_D - c_P \leq \alpha - (1 - \tau) - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] = \underline{\Delta} \quad (3.8)$$

under the off-path belief $\mu(P, S) = 0$, which is the unique D1-consistent assignment whenever D1 has bite and is not needed otherwise (Proposition A.1(a)).

Proof. In Appendix A.3 □

Under pooling on direct, both types attack using their own forces. The posterior $\mu(D, S)$ is interior; success raises the probability that the state is powerful, but does not fully reveal type since both types may succeed. The posterior $\mu(D, F) = 0$ because only the weak type can fail. The no-deviation condition requires that even the weak state, which benefits most from delegation’s informational obscurity, finds direct attack preferable. This holds when $c_D - c_P$ is sufficiently low relative to the weak state’s net direct return $\alpha - (1 - \tau)$. Note that pooling on direct requires $\Delta \leq \underline{\Delta}$, while pooling on proxy generically survives D1 only for $\Delta \geq \bar{\Delta}_P$ (Proposition A.1(c)). Since $\underline{\Delta} < \bar{\Delta}_P$ under Assumption 2, the two equilibria generically cannot coexist among pure-strategy PBE satisfying D1.

3.5 Semi-Pooling Equilibrium

I now analyze equilibria in which one type plays a pure strategy while the other mixes. The most natural candidate is a *semi-pooling* equilibrium where the powerful state always attacks directly, while the weak state randomizes between direct and proxy attack. In this configuration, types partially pool on direct attack.

Proposition 3.4 (Semi-Pooling Equilibrium). *Consider a strategy profile where θ_H plays D with probability 1, and θ_L plays D with probability $q \in (0, 1)$ and P with probability $1 - q$. Under this profile:*

$$\mu(D, S) = \frac{p}{p + (1 - p)q\alpha}, \quad (3.9)$$

$$\mu(D, F) = 0, \quad (3.10)$$

$$\mu(P, S) = \mu(P, F) = 0. \quad (3.11)$$

This constitutes a PBE if and only if:

$$c_D - c_P = \alpha - (1 - \tau) - p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)], \quad (3.12)$$

where $\pi = H(1 - k\tau)$ is the activation probability.

Proof. In Appendix A.4 □

Remark 2 (Non-Genericity). *The semi-pooling equilibrium exists if and only if condition (3.12) holds exactly. This is a knife-edge condition with measure zero in the parameter space $(c_D, c_P, \alpha, \tau, k, \lambda)$. When the condition fails:*

- *If $c_D - c_P < \alpha - (1 - \tau) - p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]$: the weak state strictly prefers D , leading to pooling on direct attack.*
- *If $c_D - c_P > \alpha - (1 - \tau) - p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]$: the weak state strictly prefers P , leading to the separating equilibrium (Proposition 3.1).*

Remark 3 (Indeterminacy of the Mixing Probability). *When condition (3.12) holds, the mixing probability q is indeterminate: any $q \in (0, 1)$ is consistent with equilibrium. This arises because, while voters evaluate the state following a direct attack, the posterior upon success $\mu(D, S) > p > \gamma$ guarantees political retention for any $q \in (0, 1)$. Consequently, the exact value of the mixing probability affects the posterior belief but not the electoral outcome, rendering it payoff-irrelevant and unable to be pinned down by incentive compatibility.*

For completeness, I also analyze the fourth candidate equilibrium (Appendix B): powerful states delegate while weak states attack directly. Proposition B.1 in Appendix B formally proves the non-existence of this type of equilibrium.

The failure of the reverse separating equilibrium reflects the severity of direct accountability in the model. Under this profile, only θ_H delegates, making delegation “safe” (no blowback) since $\mu(P, S) = 1 > \gamma$. Conversely, direct attack is played only by θ_L . Under strict Bayesian updating, any direct attack on-path perfectly reveals

weakness ($\mu(D, S) = 0 < \gamma$), guaranteeing political displacement even upon military success. Consequently, both types face the exact same payoff from a direct attack (certain displacement) and the exact same payoff from delegation. For the equilibrium to hold, both types must be exactly indifferent between these options, reducing the equilibrium to a non-generic, measure-zero parameter knife-edge. A weak state cannot be induced to choose certain political displacement when a safe, pooling option is available.

3.6 Equilibrium Mapping

Collecting the results from Propositions 3.1 - 3.4 and Appendix B, the equilibrium structure is governed by a single sufficient statistic: the cost differential $\Delta \equiv c_D - c_P$ between direct and delegated violence. I now formalize how the equilibrium regions partition the Δ -line. Define three threshold values of Δ :

$$\underline{\Delta} \equiv \alpha - (1 - \tau) - p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)], \quad (3.13)$$

$$\overline{\Delta}_P \equiv \tau - p_P(1 - k)\tau, \quad (3.14)$$

$$\overline{\Delta}_S \equiv \tau - p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]. \quad (3.15)$$

Assumption 2 (Sufficient Type Separation). *The military gap between state types exceeds the blowback wedge (Γ):*

$$1 - \alpha > \Gamma. \quad (3.16)$$

where $\Gamma \equiv p_P \pi(1 - k\tau) > 0$

The left side of the inequality measures the relative military advantage of the strong state (the exact difference in success probabilities). The right-side, blowback wedge (Γ), measures the difference in expected cost between facing blowback (under separating beliefs) and avoiding it (under pooling beliefs). It captures the payoff consequence of voter expectations: when voters interpret delegation as a signal of weakness, the state suffers an expected loss of Γ relative to the case where delegation is uninformative. This assumption requires that the weak state is sufficiently militarily disadvantaged relative to the political cost of blowback. When the assumption fails, the separating equilibrium still exists, but it coexists with a pooling equilibrium; the “snakes in the backyard” mechanism remains operative but no longer generates a unique prediction. Under the retention rule, the width of the separating region depends solely on the military capability gap $1 - \alpha$, not on the contested resource τ . The contested resource affects the threshold levels but not the range of cost differentials supporting separation.

Under Assumption 2, the three thresholds are strictly ordered: $\underline{\Delta} < \overline{\Delta}_P < \overline{\Delta}_S$. To see this, note that $\overline{\Delta}_S - \underline{\Delta} = 1 - \alpha > 0$ and $\overline{\Delta}_S = \overline{\Delta}_P + \Gamma$, so $\overline{\Delta}_P < \overline{\Delta}_S$. The remaining ordering $\underline{\Delta} < \overline{\Delta}_P$ follows from $\overline{\Delta}_P - \underline{\Delta} = (1 - \alpha) - \Gamma > 0$, which holds by Assumption 2.

The following proposition provides a complete equilibrium characterization.⁸

Proposition 3.5 (Equilibrium Characterization). *Suppose Assumptions 1 and 2 hold, and restrict attention to pure-strategy PBE satisfying D1 (Appendix A.9). The cost differential Δ partitions the parameter space into four generic regions:*

- (i) $\Delta < \underline{\Delta}$: **Pooling on Direct** is the unique such equilibrium.
- (ii) $\underline{\Delta} < \Delta < \bar{\Delta}_P$: **Separating** is the unique such equilibrium.
- (iii) $\bar{\Delta}_P < \Delta \leq \bar{\Delta}_S$: Both the **Separating** and **Pooling on Proxy** equilibria coexist.
- (iv) $\Delta > \bar{\Delta}_S$: **Pooling on Proxy** is the unique such equilibrium.

The three thresholds are measure-zero boundaries at which adjacent equilibria coincide: at $\Delta = \underline{\Delta}$ the Pooling on Direct, Separating, and Semi-Pooling equilibria coexist and at $\Delta = \bar{\Delta}_P$ and $\Delta = \bar{\Delta}_S$ the Separating and Pooling on Proxy equilibria coexist.

Proof. In Appendix A.5 □




Figure 1: Equilibrium regions as a function of the cost differential $\Delta = c_D - c_P$, under Assumption 1 and Assumption 2. Regions refer to pure-strategy PBE satisfying D1. The unique separating region (ii) has width $1 - \alpha - \Gamma$. The multiplicity region (iii) has width $\Gamma = p_P \pi (1 - k\tau)$.

Note that the unique separating region (ii) has width $1 - \alpha - \Gamma$. As blowback risk increases (π rises), the unique separating region shrinks, and the multiplicity region expands. The “snakes in the backyard” mechanism is most sharply predictive when blowback is *moderate*, large enough to matter but not so large as to make voter expectations decisive. The multiplicity region $(\bar{\Delta}_P, \bar{\Delta}_S)$ exhibits a self-fulfilling prophecy structure. In this range, two self-consistent equilibria coexist, distinguished not by fundamentals but by how voters interpret the state’s choice of violence mode. The expectations are self-validating. If voters *expect* delegation to signal weakness, only weak states

⁸Under the electoral rule of Section 2.7.2, the restriction to pure strategies for θ_H is generically without loss. Since the delegation payoff is type-independent, any profile in which θ_H mixes requires exact indifference between the two modes: $\Delta = \bar{\Delta}_P$ or $\Delta = \bar{\Delta}_S$ when direct success leads to retention (without and with blowback after delegation, respectively), and the knife-edge of Proposition B.1 when it leads to displacement, where consistency of beliefs rules out blowback. Each is a measure-zero restriction, as is the θ_L -mixing condition (3.12).

delegate, blowback follows, and the expectation is confirmed. If voters *expect* delegation to be uninformative, both types delegate, no blowback occurs, and the expectation is again confirmed. Whether the state creates a political rival depends not on the conflict environment, which is identical across equilibria, but on the prevailing interpretation of proxy delegation.

This multiplicity is robust to standard refinements. In the pooling-on-proxy equilibrium, the only off-path signal is a direct attack outcome. Since $\Delta > \bar{\Delta}_P$ implies that the equilibrium payoff strictly exceeds the maximum direct-attack deviation payoff for both types under any off-path belief, neither the Intuitive Criterion nor D1 restricts off-path beliefs after D (Proposition A.1(b)). The separating equilibrium has no ambiguous off-path beliefs. Both equilibria, therefore, survive D1. Elsewhere on the Δ -line, however, D1 is not idle. For $\Delta < \bar{\Delta}_P$, pooling on proxy can be sustained as a PBE by the off-path belief that any direct attacker is weak, but every such profile generically fails D1: a direct attack is unambiguously more tempting for the strong type (deviation payoffs $x - c_D$ versus $\alpha x - c_D$ at retention probability x), so D1 attributes an off-path direct attack to θ_H , and the implied retention makes the strong type's deviation strictly profitable (Proposition A.1(c)). This elimination delivers the uniqueness claims of Proposition 3.5 in regions (i) and (ii). Within region (iii), D1 has no selective power. The multiplicity is a consequence of the belief-dependent blowback mechanism. The signaling structure generates substantive content beyond the material lottery comparison that drives the separating equilibrium on its own. Equilibrium selection depends on how domestic audiences coordinate their interpretation of proxy violence, an inherently political determination, external to the model, that could vary across countries, historical periods, or institutional environments.

Thus, whether delegation produces blowback depends on the strategic environment as well as the prevailing political narrative about what delegation means. Media coverage, opposition rhetoric, or historical precedent that frames proxy use as a sign of weakness can lock in the separating equilibrium and its attendant political instability, while a narrative in which delegation is routine statecraft sustains the blowback-free pooling equilibrium over the same parameter range. The model thus provides a formal foundation for the scenarios in which proxy relationships produce wildly different political outcomes across otherwise similar contexts.

4 Comparative Statics

The equilibrium analysis identifies three configurations, pooling on direct, separating, and pooling on proxy, governed by the cost differential $\Delta = c_D - c_P$ relative to thresholds that depend on α , τ , k , λ , and the distribution H . I now examine how changes in key parameters affect both the blowback probability within the separating equilibrium and the equilibrium regions themselves. Throughout this section, I focus on the separating

equilibrium ($\theta_H \rightarrow D$, $\theta_L \rightarrow P$) unless stated otherwise.⁹

4.1 Blowback Probability

The central endogenous object in the separating equilibrium is the blowback probability $\pi = H(1 - k\tau)$. Recall that the proxy activates if and only if $c_A < 1 - k\tau$: the proxy challenges the state when its private cost of political entry is below the net gain from winning power ($w = 1$) minus what it would earn by staying inactive ($k\tau$).

Proposition 4.1 (Comparative Statics on Blowback Probability). *In the separating equilibrium, the blowback probability $\pi = H(1 - k\tau)$ satisfies:*

1. *The blowback probability (π) is decreasing in the proxy's offered share of captured surplus k .*
2. *The blowback probability (π) is decreasing in the target's share of the pie (τ).*

Proof. The activation threshold is $\hat{c} = 1 - k\tau$. Since H is continuously differentiable with density $h > 0$:

$$\begin{aligned}\frac{\partial \pi}{\partial k} &= h(1 - k\tau) \cdot (-\tau) < 0, \\ \frac{\partial \pi}{\partial \tau} &= h(1 - k\tau) \cdot (-k) < 0.\end{aligned}$$

□

The intuition for both results operates through the proxy's opportunity cost of activation. A higher k means the proxy receives a larger share of captured resources if it remains inactive; this raises the opportunity cost of challenging the state, reducing blowback. A higher τ means more is at stake in the violence stage; since the proxy's inactive payoff is $k\tau$, a larger τ increases what the proxy certainly foregoes by activating. This again makes the outside option of staying inactive more attractive relative to the uncertain payoff from political entry. These results have a natural policy interpretation. States can reduce blowback risk by offering proxies more generous terms, a larger share k of the spoils. However, this comes at the cost of the state's own surplus from successful proxy operations: the state retains only $(1 - k)\tau$. There is thus a straightforward tradeoff between short-run resource extraction and long-run political stability.

4.2 The Role of Weak State's Military Capability

The parameter α , the weak state's probability of success under direct attack, plays a central role in determining whether the separating equilibrium exists.

⁹See [Online Appendix](#) for more comparative statics

- Proposition 4.2** (Effect of Weak State's Capability). *1. The lower threshold $\underline{\Delta} = \alpha - (1 - \tau) - p_P B$ is strictly increasing in α with unit slope: $\frac{\partial \underline{\Delta}}{\partial \alpha} = 1$.
2. The upper threshold $\overline{\Delta}_S = \tau - p_P B$ is independent of α .
3. The width of the separating equilibrium region, $\overline{\Delta}_S - \underline{\Delta} = 1 - \alpha$, is strictly decreasing in α . As $\alpha \rightarrow 0$, the separating equilibrium region is maximally wide with width 1. The region is non-empty for all $\alpha < 1$.
4. The pooling-on-proxy threshold ($\overline{\Delta}_P$) is independent of α . The multiplicity width ($\overline{\Delta}_S - \overline{\Delta}_P = \Gamma$) is invariant to α , while the unique-separating ($\overline{\Delta}_P - \underline{\Delta} = (1 - \alpha) - \Gamma$) width is strictly decreasing in α .*

The proof is straightforward with simple differentiation. This result is central to the paper's narrative. The separating equilibrium arises precisely when weak states are militarily disadvantaged. A weak state with low α gains little from direct violence and faces a high probability of political displacement: direct attack yields expected payoff $\alpha - c_D$, where both the expected capture $\alpha\tau$ and the probability of retaining power α are low. Delegation offers a type-independent lottery at the cost of probabilistic blowback. As α increases, the direct-attack lottery improves on both margins, eventually making delegation unnecessary.

Item (4) refines the predictive content of the α comparative static. As α increases, the unique-separating width $(1 - \alpha) - \Gamma$ shrinks at the same rate as the total separating width $1 - \alpha$, while the multiplicity width Γ is invariant. Consequently, for states with higher α , a larger fraction of the parameter range supporting the separating equilibrium involves coexistence with pooling-on-proxy, where the prediction that weak states delegate is not uniquely determined. In the unique-separating region ($\underline{\Delta} < \Delta < \overline{\Delta}_P$), weak-state delegation is the sole outcome among pure-strategy PBE satisfying D1; in the multiplicity region ($\overline{\Delta}_P < \Delta < \overline{\Delta}_S$), the same fundamentals also support a blowback-free pooling equilibrium.

Note that the width of the separating region, $1 - \alpha$, depends solely on the military capability gap between state types and is invariant to the contested resource τ . The contested resource affects the level of both thresholds, shifting the separating region along the Δ -line, but not its width. In the absence of voter evaluation after direct attack, the width would be $\tau(1 - \alpha)$ and would shrink as τ falls. The invariance arises because the weak state's entire payoff is tied to the probability of military success: under direct attack, the weak state earns 1 with probability α and 0 with probability $1 - \alpha$ (net of costs), so the payoff gap between types is $1 - \alpha$ regardless of how the surplus is divided. Moreover, $\alpha \rightarrow 1$ shuts down the informational asymmetry that drives separation; the separating interval width equals $1 - \alpha = 0$: no generic separating equilibrium exists.

4.3 Proxy Inefficiency

The parameter λ simultaneously affects the proxy's military effectiveness (p_P), its political strength (γ), and the equilibrium structure.

Proposition 4.3 (Effect of Proxy Inefficiency). *As proxy inefficiency λ increases:*

1. *Proxy political strength $\gamma = \frac{1}{1+\lambda}$ and attack success probability $p_P = \frac{\rho}{1+\lambda}$ both decrease.*
2. *The width of the separating region $\bar{\Delta}_S - \underline{\Delta} = 1 - \alpha$ is invariant to λ , but both thresholds $\underline{\Delta}$ and $\bar{\Delta}_S$ shift with the sign of B :*

$$\frac{\partial \underline{\Delta}}{\partial \lambda} = \frac{\partial \bar{\Delta}_S}{\partial \lambda} = \frac{\rho B}{(1 + \lambda)^2}. \quad (4.1)$$

When $B > 0$ (delegation has positive net value per success), both thresholds increase: for a fixed Δ , pooling on direct becomes more likely. When $B < 0$, both thresholds decrease.
3. *As $\lambda \rightarrow \infty$: $p_P \rightarrow 0$, $\underline{\Delta} \rightarrow \alpha - (1 - \tau)$, $\bar{\Delta}_S \rightarrow \tau$, and the multiplicity region width $\Gamma = p_P \pi(1 - k\tau) \rightarrow 0$. The military value of delegation vanishes, and the separating region converges to $(\alpha - (1 - \tau), \tau)$.*
4. *The multiplicity width $\Gamma = p_P \pi(1 - k\tau)$ is strictly decreasing in λ . The unique-separating width $(1 - \alpha) - \Gamma$ is strictly increasing in λ : a less effective proxy narrows the multiplicity region, where equilibrium selection depends on voter expectations, and expands the region where the model delivers the unambiguous prediction that weak states delegate.*

Proof. (1) follows from the definitions. For (2), since $\underline{\Delta} = \alpha - (1 - \tau) - p_P B$ and $p_P = \rho/(1 + \lambda)$: $\frac{\partial \underline{\Delta}}{\partial \lambda} = -B \cdot \left(-\frac{\rho}{(1+\lambda)^2}\right) = \frac{\rho B}{(1+\lambda)^2}$. The same applies to $\bar{\Delta}_S = \tau - p_P B$. The width $\bar{\Delta}_S - \underline{\Delta} = 1 - \alpha$ does not contain λ . (3) follows by substituting $p_P \rightarrow 0$ into the threshold definitions. (4) $\Gamma = \rho \pi(1 - k\tau)/(1 + \lambda)$, so $\frac{\partial \Gamma}{\partial \lambda} = -\rho \pi(1 - k\tau)/(1 + \lambda)^2 < 0$. Moreover, $\frac{\partial}{\partial \lambda}[(1 - \alpha) - \Gamma] = -\frac{\partial \Gamma}{\partial \lambda} > 0$. $\square$

Remark 4 (Weak Proxy Paradox). *A very weak proxy (λ large) might seem to pose less blowback risk since $\gamma \rightarrow 0$ makes the proxy less electorally threatening. However, in the separating equilibrium, the proxy's electoral threat is invariant to its strength; the posterior $\mu(P, S) = 0$ always loses to any $\gamma > 0$, regardless of how small γ is. The relevant effect of high λ is therefore not on the electoral outcome but on the value of delegation: a weak proxy is militarily ineffective ($p_P \rightarrow 0$), reducing the state's incentive to delegate in the first place. The paradox, therefore, is that the relevant margin for blowback risk is not the proxy's political weakness but its military competence. States should worry less about how strong a proxy might become politically and more about whether the proxy is effective enough to justify delegation in the first place.*

5 Extensions

5.1 Endogenous Surplus Sharing

In the baseline model, the proxy's share k of captured surplus is exogenous. I now endogenize this sharing rule by allowing the state to make a take-it-or-leave-it offer to the proxy following a successful proxy attack. This extension formalizes the tradeoff identified in the comparative statics: a larger k reduces blowback but also reduces the state's surplus from delegation.

5.1.1 Modified Timing

The timing of the game is modified as follows. After a successful proxy attack ($\sigma = (P, S)$), a new sub-stage is inserted:

- **Stage 2.5 (Surplus Division):** The state makes a take-it-or-leave-it offer $k \in [0, 1]$ to the proxy, specifying the proxy's share of the captured surplus τ .

The proxy then observes k together with its private cost c_A and decides whether to accept (yielding payoff $k\tau$) or reject and activate politically (yielding payoff $1 - c_A$ if it wins the political contest). If the proxy rejects but activation is unprofitable ($c_A \geq 1$), the proxy receives 0; however, since $k\tau > 0$, such types always accept the offer. The proxy's participation in the operation is trivially satisfied: the operational transfer covers c_P exactly, so the proxy bears zero net cost and receives a weakly positive expected payoff from any subsequent offer $k > 0$. All other aspects of the game remain unchanged.

5.1.2 The Proxy's Acceptance Decision

The proxy's decision rule is identical in structure to the baseline. Given an offer k , the proxy accepts if and only if the offer exceeds the net gain from political activation:

$$k\tau \geq 1 - c_A \iff c_A \geq 1 - k\tau. \quad (5.1)$$

The proxy rejects the offer and activates with probability $\pi(k) = H(1 - k\tau)$. This is the same activation probability as in the baseline model. The difference is that k is now a choice variable.

5.1.3 The State's Optimal Offer

I focus on the separating equilibrium, where the proxy's activation decision is the only relevant consideration (since $\mu(P, S) = 0 < \gamma$ and the proxy wins with certainty if it activates). At Stage 2.5, the cost c_P is already sunk. The state chooses k to maximize its expected continuation payoff after proxy success:

$$V_S(k) = \underbrace{[1 - H(1 - k\tau)]}_{\text{prob. proxy accepts}} \cdot \underbrace{(1 - k\tau)}_{\text{state's surplus if accepted}} + \underbrace{H(1 - k\tau)}_{\text{prob. proxy activates}} \cdot 0. \quad (5.2)$$

The second term reflects that if the proxy activates and wins, the state loses the entire pie.¹⁰ Note that although c_P is relevant for payoffs, it is irrelevant for the optimization decision as it appears on both sides.

Proposition 5.1 (Optimal Surplus Sharing). *Suppose H has a strictly increasing hazard rate on $[\underline{c}, \bar{c}]$ and satisfies $\underline{c} \cdot h(\underline{c}) < 1$, and $\bar{c} \leq 1$. Then the state's optimal offer k^* is uniquely characterized as follows. Define $\hat{c}^*$ as the unique solution to:*

$$\frac{1 - H(\hat{c}^*)}{h(\hat{c}^*)} = \hat{c}^*, \quad (5.3)$$

where $\hat{c}^* \in (\underline{c}, \bar{c})$. The optimal sharing rule is:

$$k^* = \frac{1 - \hat{c}^*}{\tau}, \quad (5.4)$$

provided $k^* \leq 1$ (i.e., $\hat{c}^* \geq 1 - \tau$). If $\hat{c}^* < 1 - \tau$, the state sets $k^* = 1$, offering the entire captured surplus to the proxy.

Proof. In Appendix A.6 □

The characterization has a natural economic interpretation. Condition (5.3) equates the *inverse hazard rate* of the proxy's activation cost distribution. The inverse hazard rate $(1 - H(\hat{c}))/h(\hat{c})$ measures the “mass of remaining types” per unit density at the margin: it captures how many additional proxy types the state can buy off by marginally increasing its offer. Thus, the state is directly choosing the cutoff proxy type $\hat{c}$ it wants to appease. The state balances two forces: a higher offer (lower $\hat{c}$) buys off more proxy types, reducing blowback, but also reduces the state's surplus from each bought-off type. The optimum trades off the marginal reduction in blowback against the marginal decrease in surplus, precisely the condition captured by (5.3). Formally, this is equivalent to a monopolist's pricing problem: the state sets a price $\hat{c}$ for political peace, facing demand $1 - H(\hat{c})$ from proxy types willing to accept that price, and maximizes expected revenue $\hat{c} \cdot [1 - H(\hat{c})]$.

Corollary 2 (Uniform Distribution). *If $c_A \sim \text{Uniform}[0, 1]$ ¹¹, then $\hat{c}^* = \frac{1}{2}$, and:*

1. *The optimal sharing rule is $k^* = \frac{1}{2\tau}$, which is interior (i.e., $k^* \leq 1$) iff $\tau \geq \frac{1}{2}$.*
2. *When the interior solution obtains ($\tau \geq \frac{1}{2}$), the equilibrium blowback probability is $\pi^* = \frac{1}{2}$. At the corner ($\tau < \frac{1}{2}$, $k^* = 1$), the blowback probability is $\pi^* = 1 - \tau > \frac{1}{2}$.*
3. *The state's optimal continuation payoff after proxy success is $V_S^* = \frac{1}{4}$ for interior solution and $V_S^* = \tau(1 - \tau) < \frac{1}{4}$ at the corner.*

Proof. For $H(x) = x$ on $[0, 1]$, the FOC becomes $(1 - \hat{c})/1 = \hat{c}$, yielding $\hat{c}^* = 1/2$. Then $k^* = (1 - 1/2)/\tau = 1/(2\tau)$, $\pi^* = H(1/2) = 1/2$, and $V_S^* = (1/2)(1/2) = 1/4$. For $\tau < 1/2$ the corner binds at $\hat{c} = 1 - \tau$, giving $\pi^* = 1 - \tau$ and $V_S^* = \tau(1 - \tau)$. □

¹⁰The post-success timing is without loss of generality. In the state's ex ante expected payoff from delegation, $(1 - p_P)(1 - \tau) + p_P \cdot V_S(k) - c_P$, only $V_S(k)$ depends on k . Maximizing over k ex ante and post-success therefore yields the identical optimal offer k^* ; no commitment problem arises from the timing of the surplus division.

¹¹Uniform $[0, 1]$ is the boundary case $\underline{c} = 0$ of the maintained support. Lemma 2.3 requires only $c_A \geq 0$, no conflict arises with any other assumption.

The uniform case illustrates a sharp result: when the interior solution obtains ($\tau \geq 1/2$), the state accepts exactly a 50% blowback probability regardless of τ . When $\tau < 1/2$, the state hits the corner $k^* = 1$ and the blowback probability becomes $\pi^* = 1 - \tau$, which exceeds 50% and depends on τ . What varies with τ is the generosity of the offer: a larger contested resource allows the state to buy off the proxy with a smaller share.

5.1.4 Properties of the Optimal Sharing Rule

Proposition 5.2 (Comparative Statics on Optimal Sharing). *Under the conditions of Proposition 5.1:*

1. For $k^* < 1$ (interior solution), the blowback probability $\pi^* = H(\hat{c}^*)$ depends only on the distribution H and not on target value τ , proxy capability λ , or any other model parameter. At corner $k^* = 1$, $\pi^* = H(1 - \tau)$ depends on τ .
2. For $k^* < 1$, the optimal share $k^* = (1 - \hat{c}^*)/\tau$ is strictly decreasing in τ : when the contested resource is larger, the state can afford to offer a smaller share. At the corner $k^* = 1$, the share is constant.
3. The state's ex ante expected payoff from delegation (in the separating equilibrium) with optimal surplus sharing is:

$$U_S^P(\theta_L; k^*) = (1 - p_P)(1 - \tau) + p_P \cdot V_S^* - c_P, \quad (5.5)$$

where, for $k^* < 1$, $V_S^* = [1 - H(\hat{c}^*)] \cdot \hat{c}^*$ is a constant determined entirely by H .

Proof. Proof in Appendix A.7 □

For the interior solution, Property (1) states that the *probability of blowback* is invariant to the economic environment once the state optimizes its offer. The state always chooses the same threshold type $\hat{c}^*$ to separate accepting and rejecting proxies, and hence the same blowback probability. What adjusts in response to changes in τ is the actual share k^* offered, not the equilibrium level of political instability. At corner, $k^* = 1$ is fixed and the blowback probability $H(1 - \tau)$ absorbs the variation in τ .

5.1.5 Implications for Equilibrium Existence

The separating equilibrium conditions in the baseline model (Proposition 3.1) required $c_D - c_P$ to lie in an intermediate range determined by the fixed k . With endogenous surplus sharing, the weak state's payoff from delegation becomes $U_S^P(\theta_L; k^*) = (1 - p_P)(1 - \tau) + p_P V_S^* - c_P$. The equilibrium conditions, appropriately modified, become:

$$c_D - c_P \leq \tau - p_P V_S^* + p_P(1 - \tau), \quad (5.6)$$

$$c_D - c_P \geq \alpha - (1 - \tau) - p_P V_S^* + p_P(1 - \tau). \quad (5.7)$$

The interval width remains $1 - \alpha$, identical to the baseline formulation (Proposition 3.1). Endogenizing k shifts the equilibrium region along the Δ -line, since the optimized expected

state payoff from delegation differs from the fixed- k payoff, but it does not expand or shrink the separating region: the width depends solely on the military capability gap $1 - \alpha$, and the informational tradeoff between revealing type and creating a rival is unchanged. Only the proxy's compensation adjusts optimally.

5.2 Endogenous Violence - The No-Attack Option

In the baseline model, the state must attack; it chooses only the mode. This captures settings where the state has already committed to addressing the target. In practice, however, states can also choose inaction: tolerating the target's control over τ and avoiding the costs and risks of violence altogether. I now introduce this option and characterize when the “snakes in the backyard” logic operates, and when it breaks down.

5.2.1 Modified Strategy Space and the Status Quo

The state's strategy set is expanded to:

$$m \in \{D, P, \emptyset\}, \quad (5.8)$$

where $m = \emptyset$ denotes *inaction*. Under *inaction*, no violence occurs, no resources are redistributed, and no public signal of state capacity is generated. Crucially, because inaction avoids the public shock of a violent operation, it maintains the status quo: it *does not trigger* a Stage 3 political contest or voter evaluation. The state retains its initial endowment $(1 - \tau)$ and incurs no cost regardless of voter beliefs. Thus, the payoff from inaction is:

$$U_S^\emptyset(\theta) = 1 - \tau \quad \text{for both } \theta \in \{\theta_L, \theta_H\}. \quad (5.9)$$

5.2.2 The Participation Condition

The introduction of inaction imposes an additional constraint on each equilibrium: each type must prefer its equilibrium action to doing nothing. For the powerful state, direct attack is preferred to inaction whenever $1 - c_D \geq 1 - \tau$, i.e.:

$$c_D \leq \tau. \quad (5.10)$$

This is a straightforward condition saying that the cost of a direct attack must be lower than the share of the pie held by the target (T). If $c_D > \tau$, the powerful state also prefers inaction, so no violence occurs regardless of type.

For the weak state in the active separating equilibrium, the relevant comparison is between delegation and inaction. The *net expected value of delegation* (say $\mathcal{V}$) is given by:

$$\mathcal{V} \equiv p_P [(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] - c_P. \quad (5.11)$$

Since the weak state's payoff from delegation is $(1 - \tau) + \mathcal{V}$, it prefers delegation to inaction if and only if $\mathcal{V} \geq 0$.

Proposition 5.3 (Active vs. Passive Separating Equilibrium). *Suppose $c_D \leq \tau$.*

1. If $\mathcal{V} \geq 0$: the active separating equilibrium $(\theta_H \rightarrow D, \theta_L \rightarrow P)$ from Proposition 3.1 survives, subject to the baseline incentive compatibility conditions (3.2)-(3.3) and the additional requirement $c_D \leq \tau$.
2. If $\mathcal{V} < 0$ and $c_D \geq \alpha - (1 - \tau)$: a passive separating equilibrium $(\theta_H \rightarrow D, \theta_L \rightarrow \emptyset)$ exists, supported by the off-path belief $\mu(P, S) = 0$.

Proof. In Appendix A.8 □

Now, the participation condition $\mathcal{V} \geq 0$ can be rewritten as:

$$\underbrace{p_P(1 - \pi)(1 - k)\tau}_{\text{expected capture if proxy stays inactive}} \geq \underbrace{c_P + p_P\pi(1 - \tau)}_{\text{delegation cost + expected blowback loss}}. \quad (5.12)$$

The requirement $c_D \leq \tau$ for the powerful state resonates with the baseline separating bounds. From Proposition 3.1, the upper bound for the separating equilibrium is $c_D - c_P \leq \tau - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]$, which can be rewritten exactly as $c_D \leq \tau - \mathcal{V}$. When the weak state's participation constraint holds ($\mathcal{V} \geq 0$), it mathematically guarantees that $\tau - \mathcal{V} \leq \tau$. Therefore, whenever the baseline separating equilibrium holds and proxy delegation is profitable, the powerful state's participation constraint $c_D \leq \tau$ is automatically satisfied.

The “snakes in the backyard” mechanism is operative only when the expected spoils from delegation exceed its combined costs. The mechanism shuts down, and the weak state prefers inaction, when: (i) blowback is highly likely increasing the expected loss $p_P\pi(1 - \tau)$, (ii) the contested resource is small (τ low), reducing the gains from violence while increasing the state's exposure to political displacement, (iii) delegation is too expensive (c_P high), or (iv) when the state offers generous rewards (k high).

The no-attack option introduces a participation constraint that bounds the domain of the paper's central mechanism. The active separating equilibrium, in which delegation generates endogenous blowback, arises only when the expected value of proxy violence justifies the attendant political risk. When it does not, the model predicts a qualitatively different outcome: powerful states reveal themselves through action while weak states reveal themselves through inaction. Both equilibria are informationally separating, but they differ in their consequences. The active equilibrium produces endogenous political instability; the passive equilibrium produces political stagnation. This distinction reveals potential empirically testable consequences. The model predicts that proxy delegation should be observed in environments with valuable targets, reasonably effective proxies, and moderate blowback risk. When conditions are unfavorable, volatile political environments, low-value targets, or unreliable proxies, weak states should refrain from violence altogether rather than gamble on delegation. This is one of the potential reasons why some weak states with access to potential proxies nevertheless do not use them.

6 Conclusion

This paper develops a formal model explaining why states empower domestic violent actors despite the risk of future political competition. The key mechanism is that direct violence mode has type-dependent success probabilities, while delegated violence does not. A powerful state succeeds with certainty, direct attack failure is off-path for the strong type and perfectly reveals weakness; delegation decouples outcomes from state strength. Since voters condition political support on perceived strength, weak states face incentives to delegate even at the cost of creating political rivals. The separating equilibrium captures the self-undermining logic of delegation: rather than facing the high probability of military failure and political exposure that accompany direct action, weak states opt for the proxy's type-independent military effectiveness, accepting the endogenous creation of a political competitor. The model also characterizes pooling equilibria in which both types choose the same mode and shows that semi-pooling equilibria are non-generic, and a reverse separating configuration is ruled out.

The analysis yields several empirically relevant insights. First, proxy delegation should be more prevalent among states for which revealing weakness is politically costly, rather than merely among those lacking military capacity. Second, blowback is endogenous to the delegation decision; it can be mitigated through more generous resource-sharing arrangements, though at the expense of the state's surplus. The model also highlights a weak proxy paradox: inherently less capable proxies pose no less political risk, despite voters knowing that delegators are weak; weak proxies reduce the military value of delegation, eventually making it unattractive. These predictions distinguish the theory from existing explanations based on efficiency, deniability, or capacity constraints, providing testable implications for state-proxy relationships and subsequent political transitions. The equilibrium mapping identifies a multiplicity region where identical material fundamentals sustain both stable delegation and endogenous blowback, driven solely by voter beliefs. In this overlapping space, proxy deployment destabilizes the state only if domestic audiences interpret delegation as a signal of weakness rather than as usual statecraft. Expanding the strategy space to include inaction establishes the strict boundaries of this mechanism. When the targeted resources are insufficient to offset the expected cost of proxy activation, weak states separate via inaction rather than delegation, choosing the status quo over the risk of rival creation.

Several extensions merit future investigation. Replacing the deterministic electoral rule with a probabilistic contest technology would moderate the all-or-nothing structure of blowback while preserving the core signaling mechanism. Modeling the political contest as non-electoral would formally accommodate violent state takeovers by armed groups. Finally, a dynamic extension in which the proxy accumulates political capital over repeated interactions could capture the escalation of rivalry and reputational concerns in political

competition observed in empirical settings such as Colombia, Lebanon, Sudan, and Palestine.

A Proofs

A.1 Proof of Proposition 3.1

Proof. Beliefs: Note that under the separating strategy:

- Only θ_H plays D , and θ_H always succeeds. So $\mu(D, S) = 1$.
- The signal (D, F) is off-path (impossible under this strategy, since θ_H never fails and θ_L does not play D). Since only θ_L can fail under direct attack, $\mu(D, F) = 0$ is forced by consistency; no refinement is involved.
- Only θ_L plays P . So $\mu(P, S) = \mu(P, F) = 0$.

Blowback: Since $\mu(P, S) = 0 < \gamma$ (as $\gamma > 0$), the proxy will activate if $c_A < 1 - k\tau$. The activation probability is $\pi = H(1 - k\tau) \in (0, 1)$.

Payoffs and incentive compatibility:

Powerful state (θ_H): The powerful state succeeds with certainty under direct attack. Since $\mu(D, S) = 1 > \gamma$ (Remark 1), it is retained. Its on-path payoff is:

$$U_S^D(\theta_H) = 1 - c_D. \quad (\text{A.1})$$

If θ_H deviates to P , voters observe a proxy signal and hold equilibrium belief $\mu(P, S) = 0$. The deviation payoff, accounting for blowback risk, is:

$$U_S^P(\theta_H; \mu^* = 0) = (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]. \quad (\text{A.2})$$

For θ_H not to deviate, we require $U_S^D(\theta_H) \geq U_S^P(\theta_H; \mu^* = 0)$:

$$\begin{aligned} 1 - c_D &\geq (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \\ c_D - c_P &\leq \tau - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \end{aligned}$$

This is condition (3.2). This is the *exact* no-deviation condition: the deviation payoff is fully determined because the proxy signal is on-path under the separating strategy ($\mu(P, S) = 0$ follows from Bayes' rule), so no off-path beliefs are involved in this condition.

Weak state (θ_L): The weak state's on-path payoff under P with $\mu(P, S) = 0$:

$$U_S^P(\theta_L; \mu^* = 0) = (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]. \quad (\text{A.3})$$

If θ_L deviates to D , two outcomes are possible. With probability α , the attack succeeds: voters observe (D, S) and update $\mu(D, S) = 1 \geq \gamma$ (since only θ_H is supposed to play D), so the state is retained and captures τ , yielding payoff $1 - c_D$. With probability $1 - \alpha$, the attack fails: voters observe the off-path signal (D, F) and update $\mu(D, F) = 0 < \gamma$,

so the state is displaced, losing its initial share $(1 - \tau)$ and yielding payoff $-c_D$. The deviation payoff is therefore:

$$U_S^D(\theta_L) = \alpha(1 - c_D) + (1 - \alpha)(-c_D) = \alpha - c_D. \quad (\text{A.4})$$

For θ_L not to deviate, we require $U_S^P(\theta_L; \mu^* = 0) \geq U_S^D(\theta_L)$:

$$(1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \geq \alpha - c_D$$

$$c_D - c_P \geq \alpha - (1 - \tau) - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]$$

This is condition (3.3). Note that this condition relies on the off-path belief $\mu(D, F) = 0 < \gamma$, which is forced by consistency: since θ_H succeeds with certainty under direct attack, any observed failure must come from a deviating θ_L .

Interval width: The two conditions require $c_D - c_P$ to lie in the interval:

$$\left[\alpha - (1 - \tau) - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)], \quad \tau - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \right]. \quad (\text{A.5})$$

This interval has a width:

$$\tau - [\alpha - (1 - \tau)] = 1 - \alpha > 0, \quad (\text{A.6})$$

So it is non-empty for all $\alpha < 1$. $\square$

A.2 Proof of Proposition 3.2

Proof. Beliefs and blowback: If both types play P , then $\mu(P, S) = \mu(P, F) = p$ (prior). Since $\mu(P, S) = p > \gamma$ (Assumption 1), the proxy never activates. The payoff for both types is:

$$U_S^P = (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (\text{A.7})$$

Deviation incentives: Consider a deviation to D . The signals (D, S) and (D, F) are off-path. The belief after failure is forced by consistency: since only θ_L can fail, $\mu(D, F) = 0 < \gamma$ and failure ensures displacement. The belief after success is unrestricted by Bayes' rule; I evaluate deviations under the deviation-favorable assignment $\mu(D, S) \geq \gamma$ (success $\implies$ retention), which yields the tightest no-deviation condition, so that when it holds, the profile is a PBE under every off-path belief. Proposition A.1(b) confirms that D1 leaves this belief unrestricted for $\Delta \geq \bar{\Delta}_P$, and Proposition A.1(c) shows that generically no off-path belief sustains the profile under D1 for $\Delta < \bar{\Delta}_P$. Now, θ_H 's deviation payoff is $U_S^D(\theta_H) = 1 - c_D$. For θ_H not to deviate:

$$(1 - \tau) - c_P + p_P(1 - k)\tau \geq 1 - c_D \quad (\text{A.8})$$

$$c_D - c_P \geq \tau - p_P(1 - k)\tau \quad (\text{A.9})$$

Similarly, if θ_L deviates to D , it succeeds with probability α (retained) and fails with probability $1 - \alpha$ (displaced, yielding $-c_D$). θ_L doesn't deviate if:

$$(1 - \tau) - c_P + p_P(1 - k)\tau \geq \alpha - c_D \quad (\text{A.10})$$

$$c_D - c_P \geq \alpha - (1 - \tau) - p_P(1 - k)\tau \quad (\text{A.11})$$

Hence, from equations (A.9) and (A.11), we need:

$$c_D - c_P \geq \max \{ \tau - p_P(1 - k)\tau, \alpha - (1 - \tau) - p_P(1 - k)\tau \}.$$

Since $\alpha < 1$, we have $\alpha - 1 < 0$, which implies $\alpha - (1 - \tau) < \tau$. Therefore, the strong state's condition (A.9) is the binding constraint, yielding condition (3.5): $c_D - c_P \geq \tau - p_P(1 - k)\tau$. Note that if $p \leq \gamma$: the proxy activates with probability $\pi > 0$ on the equilibrium path. The state's equilibrium payoff becomes $(1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]$, which is strictly lower than the no-blowback payoff. While this does not immediately rule out the equilibrium, it alters the deviation conditions and, more importantly, the blowback-free pooling characterization no longer applies. $\square$

A.3 Proof of Proposition 3.3

Proof. Note that $m = P$ is off-path.

Payoffs on path:

θ_H : $U_S^D(\theta_H) = 1 - c_D$ (always succeeds, retained).

θ_L : $U_S^D(\theta_L) = \alpha(1 - c_D) + (1 - \alpha)(-c_D) = \alpha - c_D$ (retained on success, displaced on failure).

Deviation to P : If a type deviates to P , signals (P, S) and (P, F) are off-path. Let off-path belief be $\mu(P, S) = \tilde{\mu}$.

Case 1: $\tilde{\mu} > \gamma$ (no blowback on deviation). Deviation payoff is:

$$U_S^P(\text{no blowback}) = (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (\text{A.12})$$

For θ_L not to deviate:

$$\alpha - c_D \geq (1 - \tau) - c_P + p_P(1 - k)\tau, \quad (\text{A.13})$$

$$c_D - c_P \leq \alpha - (1 - \tau) - p_P(1 - k)\tau. \quad (\text{A.14})$$

Case 2: $\tilde{\mu} \leq \gamma$ (blowback on deviation). Deviation payoff reduced by blowback risk:

$$U_S^P(\text{blowback}) = (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]. \quad (\text{A.15})$$

For θ_L not to deviate:

$$\alpha - c_D \geq (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)], \quad (\text{A.16})$$

$$c_D - c_P \leq \alpha - (1 - \tau) - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)]. \quad (\text{A.17})$$

Case 2 is strictly weaker (larger RHS since blowback reduces the deviation payoff), so it supports the equilibrium for a wider range of Δ .

Off-path belief selection: Since the deviation payoff from P is type-independent but $U_S^D(\theta_H) > U_S^D(\theta_L)$, D1 uniquely selects $\mu(P, S) = 0$ whenever it has bite, and no restriction is needed otherwise (Proposition A.1(a)). In either case the equilibrium is supported by $\mu(P, S) = 0$; since $0 < \gamma$, Case 2 applies, yielding the binding constraint $\Delta \leq \underline{\Delta}$. The powerful type's no-deviation condition is strictly implied by the weak type's since $\alpha < 1$. $\square$

A.4 Proof of Proposition 3.4

Proof. Beliefs: Under the proposed strategy:

- Signal (D, S) : Both types may play D . The powerful type plays D and succeeds with probability 1; the weak type plays D with probability q and succeeds with probability α . By Bayes' rule:
$$\mu(D, S) = \frac{p \cdot 1}{p \cdot 1 + (1 - p) \cdot q \cdot \alpha} = \frac{p}{p + (1 - p)q\alpha}. \quad (\text{A.18})$$
- Signal (D, F) : Only θ_L can fail, so $\mu(D, F) = 0$.
- Signals (P, S) and (P, F) : Only θ_L plays P , so $\mu(P, S) = \mu(P, F) = 0$.

Blowback: Since $\mu(P, S) = 0 < \gamma$, the proxy will challenge the state if it activates. By Lemma 2.3, the proxy activates whenever $c_A < 1 - k\tau$, so the activation probability is $\pi = H(1 - k\tau)$.

Payoffs:

Strong state (θ_H): On-path payoff with direct attack is $U_S^D(\theta_H) = 1 - c_D$. If the strong state deviates to P , voters observe a proxy signal and believe $\mu(P, S) = 0$ (since only weak types are supposed to play P). The state, therefore, faces full blowback risk, and the deviation payoff is:

$$U_S^P(\theta_H) = (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (\text{A.19})$$

To prevent this deviation, we need

$$1 - c_D \geq (1 - \tau) - c_P + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (\text{A.20})$$

$$c_D - c_P \leq \tau - p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] \quad (\text{A.21})$$

Weak state (θ_L): On-path payoff under direct attack (retained on success, displaced on

failure):

$$U_S^D(\theta_L) = \alpha - c_D \quad (\text{A.22})$$

Weak state's payoff under proxy attack with $\mu(P, S) = 0$:

$$U_S^P(\theta_L) = (1 - \tau) - c_P + p_P(1 - \pi)(1 - k)\tau - p_P\pi(1 - \tau) \quad (\text{A.23})$$

Indifference condition for mixing: For θ_L to mix with interior probability $q \in (0, 1)$, we require indifference:

$$U_S^D(\theta_L) = U_S^P(\theta_L) \quad (\text{A.24})$$

$$\alpha - c_D = (1 - \tau) - c_P + p_P(1 - k)(1 - \pi)\tau - p_P\pi(1 - \tau) \quad (\text{A.25})$$

$$c_D - c_P = \alpha - (1 - \tau) - p_P[(1 - k)(1 - \pi)\tau - \pi(1 - \tau)]. \quad (\text{A.26})$$

For θ_H to strictly prefer D : Since the proxy attack payoff is independent of the state's type (p_P does not depend on θ), we have $U_S^P(\theta_H) = U_S^P(\theta_L)$. Using the indifference condition (A.26):

$$U_S^P(\theta_H) = U_S^P(\theta_L) = U_S^D(\theta_L) = \alpha - c_D. \quad (\text{A.27})$$

We need $U_S^D(\theta_H) > U_S^P(\theta_H)$:

$$1 - c_D > \alpha - c_D \quad (\text{A.28})$$

$$1 > \alpha. \quad (\text{A.29})$$

Since $\alpha \in (0, 1)$, the inequality holds strictly, confirming that θ_H strictly prefers D . Note that this also implies condition (A.21) is automatically satisfied whenever the indifference condition (A.26) holds. $\square$

A.5 Proof of Proposition 3.5

Proof. I record, for each of the four pure-strategy profiles, when it exists as a PBE and when it survives D1. The separating equilibrium exists if and only if $\Delta \in [\underline{\Delta}, \overline{\Delta}_S]$ (Proposition 3.1), and, whenever it exists, it survives D1: its only off-path signal is (D, F) , whose belief $\mu(D, F) = 0$ is forced by consistency, leaving no belief free for the refinement to restrict.

The pooling-on-proxy profile exists as a PBE for $\Delta \geq \overline{\Delta}_P$ (Proposition 3.2); for $\Delta < \overline{\Delta}_P$ it can be sustained as a PBE only by off-path beliefs with $\mu(D, S) < \gamma$, and every such profile fails D1 for generic parameters (Proposition A.1(c)), while for $\Delta \geq \overline{\Delta}_P$ D1 places no restriction on the off-path belief (Proposition A.1(b)). Hence, apart from the knife-edge $\Delta = \overline{\Delta}_P - 1$, pooling on proxy survives D1 if and only if $\Delta \geq \overline{\Delta}_P$.

The pooling-on-direct equilibrium exists if and only if $\Delta \leq \underline{\Delta}$, supported by the off-path belief $\mu(P, S) = 0$ (Proposition 3.3), which D1 uniquely selects whenever it has bite and leaves unrestricted otherwise (Proposition A.1(a)); it therefore survives D1 on its entire existence range. The reverse-separating profile is not a PBE for generic parameters (Proposition B.1).

Under Assumption 2, $\underline{\Delta} < \overline{\Delta}_P < \overline{\Delta}_S$. In region (i), $\Delta < \underline{\Delta}$ excludes separating and pooling on proxy does not survive D1 (Proposition A.1(c)); pooling on direct is the unique pure-strategy PBE satisfying D1. In region (ii), $\Delta > \underline{\Delta}$ excludes pooling on direct and pooling on proxy does not survive D1 (Proposition A.1(c)); separating is the unique such equilibrium. In region (iii), both $\Delta \in [\underline{\Delta}, \overline{\Delta}_S]$ and $\Delta \geq \overline{\Delta}_P$ hold, and both profiles survive D1 (Proposition A.1(b)), so both equilibria coexist. In region (iv), $\Delta > \overline{\Delta}_S$ excludes separating, and $\Delta > \underline{\Delta}$ excludes pooling on direct, leaving pooling on proxy as the unique such equilibrium. The boundary claims follow from the closed existence sets: at $\Delta = \underline{\Delta}$ the pooling-on-direct and separating conditions hold with equality, and semi-pooling obtains by Proposition 3.4; at $\Delta = \overline{\Delta}_P$ and $\Delta = \overline{\Delta}_S$ the separating and pooling-on-proxy conditions are simultaneously satisfied, since $\underline{\Delta} < \overline{\Delta}_P < \overline{\Delta}_S$.

□

A.6 Proof of Proposition 5.1

Proof. Substituting $\hat{c} = 1 - k\tau$ (so that $k = (1 - \hat{c})/\tau$ and the state's surplus conditional on acceptance is $\hat{c}$), the state's unconstrained problem becomes:

$$\max_{\hat{c}} V_S(\hat{c}) = [1 - H(\hat{c})] \cdot \hat{c}. \quad (\text{A.30})$$

Note that $V_S(\underline{c}) = [1 - H(\underline{c})] \cdot \underline{c} = \underline{c} > 0$ and $V_S(\bar{c}) = 0$, and V_S is continuous, so a maximum exists on $[\underline{c}, \bar{c}]$. Moreover, if $1 - \underline{c} \cdot h(\underline{c}) > 0$, then $V'_S(\underline{c}) > 0$, so the maximum cannot occur at the left boundary, and hence an interior maximizer exists. The FOC is:

$$\frac{dV_S}{d\hat{c}} = [1 - H(\hat{c})] - \hat{c} \cdot h(\hat{c}) = 0, \quad (\text{A.31})$$

which yields condition (5.3).

Uniqueness. Rewrite the FOC as:

$$\frac{1 - H(\hat{c})}{h(\hat{c})} = \hat{c}. \quad (\text{A.32})$$

The LHS is the inverse hazard rate, which is strictly decreasing under the increasing hazard rate (IHR) assumption. The RHS is strictly increasing. Hence, there is at most one intersection in $(\underline{c}, \bar{c})$. Existence follows from the intermediate value theorem: under the condition $1 - \underline{c} \cdot h(\underline{c}) > 0$, at $\hat{c} = \underline{c}$, the LHS equals $(1 - H(\underline{c}))/h(\underline{c}) > \underline{c}$ and at $\hat{c} = \bar{c}$, the LHS equals $0 < \bar{c}$. Therefore, there exists a unique $\hat{c}^* \in (\underline{c}, \bar{c})$ satisfying the FOC.

Second-order condition. Under IHR, the inverse hazard rate is decreasing. Hence, using $V_S'(\hat{c}) = h(\hat{c}) \left(\frac{1-H(\hat{c})}{h(\hat{c})} - \hat{c} \right)$, at the critical point $\hat{c}^*$, we have

$$\left. \frac{d^2 V_S}{d\hat{c}^2} \right|_{\hat{c}=\hat{c}^*} = h'(\hat{c}^*) \left(\frac{1-H(\hat{c}^*)}{h(\hat{c}^*)} - \hat{c}^* \right) + h(\hat{c}^*) \left(\frac{d}{d\hat{c}} \left[\frac{1-H(\hat{c})}{h(\hat{c})} \right] \Big|_{\hat{c}=\hat{c}^*} - 1 \right) \quad (\text{A.33})$$

Since the first term vanishes by the FOC, this simplifies to $h(\hat{c}^*) \left(\frac{d}{d\hat{c}} \left[\frac{1-H(\hat{c})}{h(\hat{c})} \right] \Big|_{\hat{c}=\hat{c}^*} - 1 \right) < 0$ as $h(\hat{c}^*) > 0$. Hence, $\hat{c}^*$ is a strict global maximum.

Optimal Sharing Rule and Corner Solution. The unconstrained global maximizer $\hat{c}^*$ implies an optimal offer $k^* = (1 - \hat{c}^*)/\tau$. Since the state's choice is restricted to $k \in [0, 1]$, this requires $\hat{c} \geq 1 - \tau$. Because $\hat{c}^*$ is the unique global maximum and $V_S(\hat{c})$ is strictly decreasing for $\hat{c} > \hat{c}^*$, the constrained optimum depends on τ . If $\hat{c}^* \geq 1 - \tau$, the interior solution $k^* = (1 - \hat{c}^*)/\tau \leq 1$ is feasible. If $\hat{c}^* < 1 - \tau$, the constraint binds, and the state's optimal feasible choice is the boundary $\hat{c} = 1 - \tau$, yielding the corner solution $k^* = 1$. $\square$

A.7 Proof of Proposition 5.2

Proof. (1) The characterizing equation (5.3) involves only H and h , no other model parameter appears; at corner the cutoff is $\hat{c} = 1 - \tau$, so $\pi^* = H(1 - \tau)$. (2) For $k^* < 1$, $k^* = (1 - \hat{c}^*)/\tau$, so $\partial k^*/\partial \tau = -(1 - \hat{c}^*)/\tau^2 < 0$ since $\hat{c}^* < 1$. (3) The state's ex ante payoff from delegation, anticipating the optimal post-success offer, is:

$$\begin{aligned} U_S^P(\theta_L; k^*) &= (1 - p_P) \underbrace{[(1 - \tau) - c_P]}_{\text{attack fails}} + p_P \underbrace{[V_S^* - c_P]}_{\text{attack succeeds}} \\ &= (1 - p_P)(1 - \tau) + p_P \cdot V_S^* - c_P. \end{aligned} \quad (\text{A.34})$$

For $k^* < 1$, $\hat{c}^*$ depends only on H , so V_S^* is a constant with respect to all parameters except H ; at corner $V_S^* = [1 - H(1 - \tau)](1 - \tau)$ varies with τ . $\square$

A.8 Proof of Proposition 5.3

Proof. Part 1 (Active Separating Equilibrium): The equilibrium strategies are $s(\theta_H) = D$ and $s(\theta_L) = P$. From Proposition 3.1, we know θ_H prefers D over P , and θ_L prefers P over D under conditions (3.2) and (3.3). We only need to check deviations to inaction ($\emptyset$).

For θ_H : $U_S^D(\theta_H) = 1 - c_D$. Deviation to $\emptyset$ yields $1 - \tau$. θ_H does not deviate iff $1 - c_D \geq 1 - \tau \implies c_D \leq \tau$.

For θ_L : $U_S^P(\theta_L) = (1 - \tau) + p_P[(1 - \pi)(1 - k)\tau - \pi(1 - \tau)] - c_P = (1 - \tau) + \mathcal{V}$.

Deviation to $\emptyset$ yields $1 - \tau$. θ_L does not deviate iff $(1 - \tau) + \mathcal{V} \geq 1 - \tau \implies \mathcal{V} \geq 0$.

Thus, the active separating equilibrium survives if $\mathcal{V} \geq 0$ and $c_D \leq \tau$.

Part 2 (Passive Separating Equilibrium): The equilibrium strategies are $s(\theta_H) = D$ and $s(\theta_L) = \emptyset$. *Beliefs:* On-path, $\mu(D, S) = 1$ (only θ_H plays D and succeeds). The signal $\emptyset$ is on-path for θ_L , so Bayes' rule yields $\mu(\emptyset) = 0$. However, by institutional rule, inaction does not trigger political displacement, so θ_L retains $1 - \tau$. The off-path belief is assigned as $\mu(P, S) = 0$. Off-path signal (D, F) yields $\mu(D, F) = 0$.

Deviations for θ_H : On-path payoff is $1 - c_D$. Deviation to $\emptyset$ yields $1 - \tau$. θ_H prefers D over $\emptyset$ iff $c_D \leq \tau$. Deviation to P under off-path belief $\mu(P, S) = 0$ yields full blowback risk: payoff $(1 - \tau) + \mathcal{V}$. Since $c_D \leq \tau$ implies $1 - c_D \geq 1 - \tau$, and $\mathcal{V} < 0$ implies $1 - \tau > (1 - \tau) + \mathcal{V}$, θ_H strictly prefers D over P .

Deviations for θ_L : On-path payoff is $1 - \tau$. Deviation to P yields $(1 - \tau) + \mathcal{V}$. Since $\mathcal{V} < 0$, θ_L strictly prefers $\emptyset$ over P . Deviation to D yields success with probability α (retained, capturing τ) and failure with probability $1 - \alpha$ (displaced, losing $1 - \tau$). The expected deviation payoff is $\alpha(1 - c_D) + (1 - \alpha)(-c_D) = \alpha - c_D$. θ_L prefers $\emptyset$ over D iff $1 - \tau \geq \alpha - c_D \implies c_D \geq \alpha - (1 - \tau)$.

Thus, the passive separating equilibrium exists when $\mathcal{V} < 0$ and $c_D \geq \alpha - (1 - \tau)$. $\square$

A.9 D1 Refinement

The refinement is applied to off-path *actions* $m' \in \{D, P\}$. For every belief $\tilde{\mu} \neq \gamma$, the voter's best response is unique and pure; at $\tilde{\mu} = \gamma$, the voter is indifferent between retaining and replacing the incumbent. Consequently, at that belief, the best-response correspondence is the entire interval $[0, 1]$: any response retaining the incumbent with probability r yields expected utility $r\tilde{\mu} + (1 - r)\gamma = \tilde{\mu} = \gamma$ (Section 2.7.2).

For an off-path action m' , let $D(\theta, m')$ denote the set of responses under which type θ strictly prefers deviating to m' over its equilibrium payoff, and $D^0(\theta, m')$ the set under which it is exactly indifferent. The D1 criterion (Banks and Sobel, 1987, Cho and Kreps, 1987) requires that beliefs following m' place zero probability on type θ whenever there exists another type θ' with

$$D(\theta, m') \cup D^0(\theta, m') \subsetneq D(\theta', m'),$$

where the inclusion is proper. In particular, if no type strictly gains from the deviation under any response, no proper inclusion holds, and D1 imposes no restriction.

The two off-path actions have the following response structure. After a deviation to D , the belief following failure is $\mu(D, F) = 0$ for every belief about the deviator's identity, since $p_D(\theta_H) = 1$ implies that only θ_L can fail; displacement after (D, F) is therefore forced, and a response is summarized by the retention probability $x \in [0, 1]$ after (D, S) . The deviation payoffs are $x - c_D$ for θ_H and $\alpha x - c_D$ for θ_L . After a deviation to P , a response is summarized by the retention probability $y \in [0, 1]$ in the political contest following (P, S) , together with the proxy's uniquely determined best-reply cutoff. The

induced deviation payoff $u(y)$ is the same for both types, since p_P , the proxy's activation rule, and the contest outcome do not depend on θ . It is continuous in y , is maximized at the no-blowback response $y = 1$, where $u(1) = (1 - \tau) - c_P + p_P(1 - k)\tau$, and satisfies $u(0) = (1 - \tau) - c_P + p_PB$, the blowback payoff.¹²

Proposition A.1 (D1 Refinement). *(a) In the pooling-on-direct equilibrium (Proposition 3.3), D1 uniquely selects $\mu(P, S) = 0 < \gamma$ whenever $\Delta > \underline{\Delta} - \Gamma$. For $\Delta \leq \underline{\Delta} - \Gamma$, deviation to P is unprofitable for both types under every response; D1 imposes no restriction, and none is needed, since the equilibrium is then supported by any off-path belief.*

(b) In the pooling-on-proxy equilibrium, for $\Delta \geq \bar{\Delta}_P$, D1 places no restriction on beliefs after the off-path direct attack. In particular, in the multiplicity region $\Delta \in (\bar{\Delta}_P, \bar{\Delta}_S)$, D1 eliminates neither the separating nor the pooling-on-proxy equilibrium.

(c) For $\Delta < \bar{\Delta}_P$ with $\Delta \neq \bar{\Delta}_P - 1$, the pooling-on-proxy profile does not survive D1: any off-path belief sustaining it as a PBE has $\mu(D, S) < \gamma$, D1 deletes θ_L after the deviation to D , and the resulting belief $\mu(D, S) = 1 \geq \gamma$ makes θ_H 's deviation strictly profitable. Combining (b) and (c), the pooling-on-proxy equilibrium generically survives D1 if and only if $\Delta \geq \bar{\Delta}_P$.

Proof. Part (a): In the pooling-on-direct equilibrium, $m = P$ is off-path, and the equilibrium payoffs are $U_S^D(\theta_H) = 1 - c_D > \alpha - c_D = U_S^D(\theta_L)$. Since the deviation payoff $u(y)$ is type-independent, $u(y) \geq 1 - c_D$ implies $u(y) > \alpha - c_D$, so $D(\theta_H, P) \cup D^0(\theta_H, P) \subseteq D(\theta_L, P)$ for every parameter configuration. The inclusion is proper if and only if some response yields a payoff strictly between $\alpha - c_D$ and $1 - c_D$. The existence condition of Proposition 3.3 gives $u(0) \leq \alpha - c_D$, and $u(y)$ is continuous with maximum $u(1)$, so such a response exists if and only if

$$u(1) = (1 - \tau) - c_P + p_P(1 - k)\tau > \alpha - c_D \iff \Delta > \alpha - (1 - \tau) - p_P(1 - k)\tau = \underline{\Delta} - \Gamma,$$

where the last equality uses $p_PB + \Gamma = p_P(1 - k)\tau$. Hence, for $\Delta \in (\underline{\Delta} - \Gamma, \underline{\Delta}]$, D1 deletes θ_H and uniquely selects $\mu(P, S) = 0 < \gamma$. For $\Delta \leq \underline{\Delta} - \Gamma$, every response yields at most $u(1) \leq \alpha - c_D < 1 - c_D$, so $D(\theta, P) = \emptyset$ for both types, no proper inclusion holds, and D1 is silent; no restriction is needed, since deviation to P is then unprofitable under every response.

Part (b): Suppose both types play P , with common equilibrium payoff $U_{eq} =$

¹²Facing retention probability y , the proxy activates iff $c_A < (1 - y) - k\tau$, so the activation probability is $\pi_y = H((1 - y) - k\tau)$ and $u(y) = (1 - \tau) - c_P + p_P[(1 - \pi_y)(1 - k)\tau + \pi_y(y - (1 - \tau))]$, using the fact that a state that wins the contest retains the full pie (Section 2.7.2). Direct computation gives $u(1) - u(y) = p_P\pi_y(1 - k\tau - y) \geq 0$, since $\pi_y > 0$ requires $y < 1 - k\tau$; continuity follows from that of H . The span across responses is $u(1) - u(0) = p_P\pi(1 - k\tau) = \Gamma$: the blowback wedge is exactly the payoff spread that off-path beliefs can generate after a proxy deviation.

$(1 - \tau) - c_P + p_P(1 - k)\tau$, and consider the off-path action D . Define

$$z \equiv U_{\text{eq}} + c_D = \Delta + (1 - \tau) + p_P(1 - k)\tau, \quad z < 1 \iff \Delta < \bar{\Delta}_P.$$

The deviation payoffs are $x - c_D$ and $\alpha x - c_D$, so $D(\theta_H, D) = \{x \in [0, 1] : x > z\}$ and $D(\theta_L, D) = \{x \in [0, 1] : \alpha x > z\}$. For $\Delta \geq \bar{\Delta}_P$ we have $z \geq 1$, so both sets are empty: no response makes any type strictly prefer the deviation (at $\Delta = \bar{\Delta}_P$, θ_H is exactly indifferent under $x = 1$, so $D^0(\theta_H, D) = \{1\}$, and still no proper inclusion holds). D1 therefore places no restriction on $\mu(D, S)$, and the equilibrium exists under any off-path belief in this range, since even the most deviation-favorable response, $x = 1$, yields $1 - c_D \leq U_{\text{eq}}$. The separating equilibrium has no ambiguous off-path beliefs: its unique off-path signal (D, F) carries the structurally forced belief $\mu(D, F) = 0$. Both equilibria therefore survive D1 in the multiplicity region.

Part (c): Fix $\Delta < \bar{\Delta}_P$, so $z < 1$. If an off-path belief with $\mu(D, S) \geq \gamma$ were assigned, θ_H 's deviation would yield $1 - c_D > U_{\text{eq}}$ and the profile would fail immediately; any sustaining belief therefore has $\mu(D, S) < \gamma$, under which a deviator is displaced after either outcome and receives $-c_D$, so the profile is a PBE if and only if $z \geq 0$. Assume $z > 0$. At $z = 0$, the non-generic knife-edge $\Delta = \bar{\Delta}_P - 1$ of Proposition B.1, both types have $D(\theta, D) = (0, 1]$ and $D^0(\theta, D) = \{0\}$, so no proper inclusion holds in either direction: D1 is silent and the profile survives. For $z < 0$, no such PBE exists. Then $D(\theta_L, D) \cup D^0(\theta_L, D) = \{x \in [0, 1] : \alpha x \geq z\}$, and since $z/\alpha > z$, every element of this set exceeds z , so it is included in $D(\theta_H, D) = (z, 1]$. The inclusion is proper: the interval $(z, \min\{z/\alpha, 1\})$ is non-empty because $0 < z < 1$ and $\alpha < 1$, and its elements tempt θ_H but not θ_L . The reverse deletion fails, since $x = z$ lies in $D^0(\theta_H, D)$ but not in $D(\theta_L, D)$. D1 therefore deletes θ_L , so beliefs after the deviation concentrate on θ_H : $\mu(D, S) = 1 \geq \gamma$, and the voter retains after success. Given this response, θ_H 's deviation yields $1 - c_D > U_{\text{eq}}$, a strict improvement, so no pooling-on-proxy profile with $\Delta < \bar{\Delta}_P$ survives D1. Combining with part (b), the pooling-on-proxy equilibrium survives D1 if and only if $\Delta \geq \bar{\Delta}_P$, apart from the knife-edge $z = 0$. $\square$

B The Reverse Separating Equilibrium (Non-existence):

$\theta_H \rightarrow P$, and $\theta_L \rightarrow D$

Proposition B.1 (Non-Existence of Reverse Separating Equilibrium). *The strategy profile $(s(\theta_H) = P, s(\theta_L) = D)$ is non-generic and cannot constitute a PBE except on a measure-zero set of parameters.*

Proof. Beliefs: Under the proposed strategy, all signals are on-path (or subsets of on-path strategies). Bayes' rule perfectly pins down beliefs:

- Only θ_H plays P . Hence $\mu(P, S) = \mu(P, F) = 1$.
- Only θ_L plays D . Hence $\mu(D, S) = \mu(D, F) = 0$.

Blowback: Since $\mu(P, S) = 1 > \gamma$ for any $\gamma < 1$, the proxy never activates. The powerful state faces zero blowback risk. **On-path payoffs:** For θ_H (playing P with no blowback):

$$U_S^P(\theta_H) = (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (\text{B.1})$$

For θ_L (playing D): Under strict application of Bayes' rule, observing (D, S) perfectly reveals the state is weak ($\mu(D, S) = 0 < \gamma$). Therefore, even if the weak state succeeds militarily, voters displace it. Its payoff from playing D is thus displacement regardless of the military outcome, $U_S^D(\theta_L) = -c_D$.

θ_L 's deviation incentive: If θ_L deviates to P , it produces an on-path signal. Voters believe $\mu = 1 > \gamma$, so no blowback occurs. The deviation payoff is:

$$U_S^P(\theta_L; \mu = 1) = (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (\text{B.2})$$

For θ_L not to deviate:

$$-c_D \geq (1 - \tau) - c_P + p_P(1 - k)\tau. \quad (\text{B.3})$$

θ_H 's deviation incentive: If θ_H deviates to D , it succeeds with probability 1, producing the on-path signal (D, S) . Since voters use Bayes' rule, they believe the state is weak ($\mu(D, S) = 0 < \gamma$) and displace it. The deviation payoff is $-c_D$. For θ_H not to deviate:

$$(1 - \tau) - c_P + p_P(1 - k)\tau \geq -c_D. \quad (\text{B.4})$$

Compatibility of conditions: For the equilibrium to exist, both (B.3) and (B.4) must hold simultaneously. This implies:

$$c_D - c_P = -[(1 - \tau) + p_P(1 - k)\tau]. \quad (\text{B.5})$$

This requires exact equality, meaning the equilibrium only exists on a measure-zero knife-edge in the parameter space and is therefore non-generic. Since $\tau \in (0, 1)$ and $p_P, k > 0$, the right-hand side is strictly negative, which requires $c_D < c_P - (1 - \tau)$. Thus, a robust Reverse Separating Equilibrium cannot exist. Moreover, since $\Delta < 0 \implies c_D < c_P$, when the no-attack option is available (Section 5.2), both types strictly prefer inaction to their equilibrium actions at these parameter values, so the knife-edge equilibrium is dominated. $\square$

References

- Acemoglu, D., Robinson, J. A., and Santos, R. J. (2013). The monopoly of violence: Evidence from colombia. *Journal of the European Economic Association*, 11(suppl_1):5–44.
- Ahram, A. I. (2011). *Proxy warriors: the rise and fall of state-sponsored militias*. Stanford University Press.
- Alesina, A. (1988). Credibility and policy convergence in a two-party system with rational voters. *The American Economic Review*, 78(4):796–805.
- Ashworth, S. (2012). Electoral accountability: Recent theoretical and empirical work. *Annual Review of Political Science*, 15(1):183–201.
- Baliga, S. and Sjöström, T. (2004). Arms races and negotiations. *The Review of Economic Studies*, 71(2):351–369.
- Baliga, S. and Sjöström, T. (2012). The strategy of manipulating conflict. *American Economic Review*, 102(6):2897–2922.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society*, pages 647–661.
- Berman, E. and Lake, D. A. (2019). *Proxy wars: Suppressing violence through local agents*. Cornell University Press.
- Besley, T. (2006). *Principled agents?: The political economy of good government*. Oxford University Press.
- Besley, T. and Coate, S. (1997). An economic model of representative democracy. *The quarterly journal of economics*, 112(1):85–114.
- Besley, T. and Persson, T. (2010). State capacity, conflict, and development. *Econometrica*, 78(1):1–34.
- Besley, T. and Persson, T. (2011). Pillars of prosperity: The political economics of development clusters. In *Pillars of Prosperity*. Princeton University Press.
- Besley, T. and Prat, A. (2006). Handcuffs for the grabbing hand? media capture and government accountability. *American economic review*, 96(3):720–736.
- Carey, S. C., Colaresi, M. P., and Mitchell, N. J. (2015). Governments, informal links to militias, and accountability. *Journal of Conflict Resolution*, 59(5):850–876.
- Carey, S. C., Mitchell, N. J., and Lowe, W. (2013). States, the security sector, and the

- monopoly of violence: A new database on pro-government militias. *Journal of Peace Research*, 50(2):249–258.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221.
- Downs, A. (1957). An economic theory of democracy. *Happer & Row*.
- Egorov, G. and Sonin, K. (2011). Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association*, 9(5):903–930.
- Fearon, J. D. (1995). Rationalist explanations for war. *International organization*, 49(3):379–414.
- Fey, M. and Ramsay, K. W. (2011). Uncertainty and incentives in crisis bargaining: Game-free analysis of international conflict. *American Journal of Political Science*, 55(1):149–169.
- Hughes, G. and Tripodi, C. (2009). Anatomy of a surrogate: historical precedents and implications for contemporary counter-insurgency and counter-terrorism. *Small Wars & Insurgencies*, 20(1):1–35.
- Kalyvas, S. N. (2006). *The logic of violence in civil war*. Cambridge University Press.
- Lizzeri, A. and Persico, N. (2001). The provision of public goods under alternative electoral incentives. *American Economic Review*, 91(1):225–239.
- Lyall, J. (2010). Are coethnics more effective counterinsurgents? evidence from the second chechen war. *American Political Science Review*, 104(1):1–20.
- Mazzei, J. (2009). *Death Squads or Self-Defense Forces? How Paramilitary Groups Emerge and Challenge Democracy in Latin America*. University of North Carolina Press, Chapel Hill, NC.
- Mitchell, N. J., Carey, S. C., and Butler, C. K. (2014). The impact of pro-government militias on human rights violations. *International Interactions*, 40(5):812–836.
- Myerson, R. B. (2008). The autocrat’s credibility problem and foundations of the constitutional state. *American Political Science Review*, 102(1):125–139.
- North, D. C., Wallis, J. J., and Weingast, B. R. (2009). *Violence and social orders: A conceptual framework for interpreting recorded human history*. Cambridge University Press.
- Osborne, M. J. and Slivinski, A. (1996). A model of political competition with citizen-candidates. *The Quarterly Journal of Economics*, 111(1):65–96.

- Persson, T., Roland, G., and Tabellini, G. (1997). Separation of powers and political accountability. *The Quarterly Journal of Economics*, 112(4):1163–1202.
- Powell, R. (2006). War as a commitment problem. *International organization*, 60(1):169–203.
- Powell, R. (2013). Monopolizing violence and consolidating power. *The Quarterly Journal of Economics*, 128(2):807–859.
- Salehyan, I. (2010). The delegation of war to rebel organizations. *Journal of conflict resolution*, 54(3):493–515.
- Stanton, J. A. (2015). Regulating militias: Governments, militias, and civilian targeting in civil war. *Journal of Conflict Resolution*, 59(5):899–923.
- Tilly, C. (1985). War making and state making as organized crime. In Evans, P. B., Rueschemeyer, D., and Skocpol, T., editors, *Bringing the State Back In*, pages 169–191. Cambridge University Press, Cambridge, UK.
- Weber, M. (1946). Politics as a vocation. In Gerth, H. H. and Mills, C. W., editors, *From Max Weber: Essays in Sociology*, pages 77–128. Oxford University Press, New York, NY.