Informative Pre-election Attacks and Post-Electoral Reprisals

Gavin Kiger* Miguel R. Rueda†

February 5, 2026

Abstract

Why do political parties use violence against voters before elections? Beyond suppressing opposition turnout, pre-electoral violence can coerce participation in favor of the perpetrating party through two reinforcing mechanisms. First, it signals a party's willingness and capacity to impose further punishment after unfavorable outcomes. Second, by reducing turnout in places where results are observed, it raises the pivotality of each remaining vote in determining whether reprisals occur. We formalize this logic in an incomplete-information model in which voters are uncertain about parties' violent capabilities. The model explains why violence is prevalent where voting follows ethnic lines, how the granularity of electoral reporting shapes both pre- and post-election violence, and why parties sometimes attack swing areas or strongholds of any party.

*Graduate student, Department of Political Science, Emory University. gavin.kiger@emory.edu

†Associate Professor, Department of Political Science, Emory University. miguel.rueda@emory.edu. Tarbutton Hall 315, 1555 Dickey Drive, Atlanta, GA 30322.

Ballots have long been cast in the shadow of violence. Historical accounts describe clashes between gangs and voter intimidation during elections in the late Roman Republic (Troxler, 2008), the American South (Kousser, 1974), and Western Europe in the 19th century (Hoppen, 1994). More recently, between 1990 and 2008, over half of elections in sub-Saharan Africa involved violent incidents (Taylor et al., 2017, p. 403), and from 1989 to 2017, more than 24,000 people were killed in election-related events worldwide (Fjelde and Höglund, 2022, p. 174). While some of this violence reflects public unrest over alleged fraud, it is frequently a calculated tool used to influence voter behavior. Pre-election attacks on the population, whether carried out by state forces, armed groups, or others affiliated to political parties, are indicative of this rationale. How do the perpetrators benefit from these attacks?

One possibility is that pre-election violence prevents turnout among opposition supporters (Chaturvedi, 2005; Collier and Vicente, 2012; Gonzalez-Ocantos et al., 2020). Mixed evidence on whether parties target low-support areas (e.g., Mares and Young, 2016; Wahman, 2024), however, suggests that the attacks can serve other purposes. A particularly understudied logic holds that fear of future reprisals induced by pre-election violence is a powerful force influencing not only turnout but also vote choice (Travaglianti, 2014). An incident during Colombia's 2002 elections in the town of San Onofre illustrates the mechanism. Right-wing paramilitary leader Cadena forced hundreds of residents to attend a campaign rally for Muriel Benito, a candidate for the House of Representatives. According to witnesses, at the end of the event, Cadena threatened to kill two councilmen in attendance and others at random if Muriel failed to win. These threats were highly credible. Cadena's reputation for carrying out assassinations against those who opposed him was well established. Unsurprisingly, Muriel received unusually high vote totals in the region.[1]

In other contexts, the credibility of similar threats may be weaker, leading perpetrators to stage additional displays of violence before the election. For example, during Kenya's 1992 elections, Kalenjin and Maasai "Warrior" groups attacked Kikuyu migrants in the Rift Valley. These assaults

[1] *Así se desarrolló la oscura conexión entre paramilitares y políticos en Sucre* [This is how the dark link between paramilitaries and politicians was developed], El Tiempo, November 11, 2006.

were intended to *show* that multiparty democracy would bring harm to those who supported any party other than the Kenyan African Union (Burchard, 2015, p. 60). While secret ballots were in place on these occasions, voters understood that local vote totals could trigger collective reprisals.

In this paper, we provide a theoretical framework to study terror induced by pre-electoral attacks targeting localities for which vote totals are available. We examine how attacks raise voters' expectations of future punishments that are conditional on aggregated electoral outcomes of small groups. The model helps reconcile seemingly inconsistent findings about the relationship between political preferences and the geography of electoral violence, clarifies how the levels at which electoral results are published affect voters' safety, and accounts for the observed higher incidence of violence in elections where voting occurs along ethnic lines.

We analyze the strategic interaction between a campaign operative and a small subset of the electorate. This group of voters may face post-electoral reprisals if the operative's candidate performs poorly in that locality. The baseline model captures a low-information environment in which voters' candidate preferences are private information and voters are uncertain about whether the operative has the capacity to carry out such reprisals. Both voters and the operative know the distribution of political preferences, and the operative observes only the group's aggregate vote totals, not individual choices. At the start of the game, the operative decides whether to engage in pre-electoral violence, which reduces turnout at a cost. Operatives capable of imposing post-election punishments face a lower cost of pre-election violence than those who are not. Before voting, voters use the occurrence of a pre-electoral attack to update their beliefs about whether a poor electoral outcome will trigger post-election punishment.

The model highlights two key benefits of pre-electoral violence for the attacker. First, in some equilibria, such violence increases voters' belief that the campaign operative is capable of carrying out post-electoral reprisals, prompting more opposition supporters to switch their votes out of fear. Second, because post-election punishment is based on group-level vote totals (due to ballot secrecy), reducing turnout through violence increases the pivotality of each remaining vote in determining whether punishment will be triggered. This heightened sense of individual responsibility

further incentivizes opponents to vote for the attacker's candidate. In essence, pre-election violence amplifies the perceived threat of post-election punishment by making each vote more consequential in smaller, targeted groups. Importantly, these two mechanisms are mutually reinforcing: the deterrent effect of increased vote pivotality is stronger when voters believe the operative will follow through on threats.

Examining how the fear of post-election punishments shapes voting allows us to interpret existing empirical patterns and contribute to ongoing debates in the literature. A central point of contention concerns where pre-electoral violence is most likely to occur: in opposition strongholds or in swing areas. Some argue that pre-electoral violence is primarily used to demobilize opposition supporters, and thus should be concentrated in areas where the opposition is strong (Rauschenbach and Paula, 2019). Others emphasize the strategic value of violence in swing regions, where the cost of losing due to nonviolent competition is high, and where violence may also serve as a mobilization tool (Daxecker et al., 2024; Robinson and Torvik, 2009). Furthermore, Wahman (2024) argues that violence can target the perpetrator's or competitor's strongholds to enforce or challenge territoriality—the perception of a party's local invincibility and dominance in the locality. Reviewing the literature, Mares and Young (2016) note that the empirical evidence remains inconclusive about whether violence systematically targets voters based on partisan preferences.

Our model explains pre-election violence in party strongholds, opposition areas, and swing districts. Targeting depends on two factors: whether multiple parties are capable of harming voters, and how much parties know about individual political preferences. When there is only one violent party that lacks information about individual voters' preferences, pre-election attacks are more likely where the opposition support is weak. When multiple parties can coerce but still only observe the distribution of preferences, the logic shifts: a party may target swing areas or even its own strongholds to counter the fear induced by their rivals' intimidation. By placing supporters "between a rock and a hard place," such attacks deter defection, since voting for the opposition no longer reduces the risk of harm.

When we allow the party operatives to observe individual prefered candidates and not just the distribution of preferences, the benefits of pre-electoral violence increase: the operative can now selectively target opposition supporters, reducing their turnout while avoiding targeting her supporters. Because an opponent only has a positive probability of voting for the operative's candidate out of fear, while supporters vote for that candidate with certainty, observability of individual preferences makes it easier to sustain all equilibria in which such attacks occur. Importantly, it also makes pre-election violence possible even where the opposition no longer has weak attachments to their candidate, in contrast to what we find when only the distribution of preferences is known. This finding can account for the higher incidence of pre-electoral violence in contexts where voting follows ethnic lines (Enamorado and Kosterina, 2022; Kuhn, 2015; Müller-Crepon, 2022; Travaglianti, 2014). When voter preferences are strongly correlated with observable individual characteristics, such as race, language, or religious customs, party operatives can identify and selectively target opposition supporters, making pre-election violence a more attractive campaign tactic.

The model also illustrates a complementary mechanism that makes pre-election violence more prevalent where ballots are cast primarily for co-ethnics. In the baseline model, voters make decisions exclusively focused on the effects of post-election reprisals on their own safety. Relaxing this assumption, by allowing voters to care about the well-being of others who share their political preferences, also increases the likelihood of pre-election violence. In this way, voters' stronger willingness to protect co-ethnics amplifies the coercive power of pre-election violence: by voting for the attacker, individuals seek to minimize expected future harm not only to themselves but also to members of their ethnic group. This logic, along with the observability of preferences mechanisms, contrasts with explanations that emphasize the scarcity of persuadable swing voters as the reason why violent tactics are common in ethnically polarized settings (Kuhn, 2015).

This paper contributes to the literature on political violence and elections that examines how elections shape overall levels of violence (e.g., Cox, 2009; Fearon, 2011; Little, 2012; Luo and Rozenas, 2018). Some models view elections as reducing violence by substituting for conflict or

lowering bargaining uncertainty, while others show that elections can raise violence by revealing incumbent vulnerability and triggering coups or uprisings. The informational content of election results, especially about incumbent strength, also creates incentives for election manipulation to shape public beliefs (e.g., Gehlbach and Simpser, 2015; Little, 2012; Simpser, 2013). Our model builds on this perspective but shifts the focus: rather than altering beliefs about incumbent strength, we study how pre-election violence changes voters' expectations about localized post-election consequences they may face, which may be independent of the election winner's identity.

Other formal work studies how electoral violence shapes outcomes. Ellman and Wantchekon (2000) and Wantchekon (1999) analyze how threats of post-election disruption affect platforms and voter behavior, but not the logic of pre-election violence. Work by Chaturvedi (2005), Robinson and Torvik (2009), and Collier and Vicente (2012) treats pre-election violence as a campaign tool without modeling its effects through voters' expectations. Our model provides microfoundations for Collier and Vicente (2012)'s targeting of weak opposition areas, while showing this strategy can be suboptimal when multiple parties can use violence or when voters individual preferences are observable. Brancati and Penn (2022) and Hassan and O'Mealia (2018) focus instead on agency problems between government or party officials and lower-level operatives, and on how these relationships shape the use of electoral violence, rather than on the agency relationship between party operatives and voters that we analyze here. Rundlett and Svolik (2016) likewise examine agency problems between incumbents and political operatives, but emphasize collective action and uncertainty about the incumbent's popularity as drivers of fraudulent behavior.

Prior empirical studies link electoral violence to expectations of unfavorable outcomes (Hafner-Burton et al., 2014), majoritarian rules and institutional constraints (Daxecker and Rauschenbach, 2023; Fjelde and Höglund, 2016; Müller-Crepon, 2022), and recent civil war (Ishiyama et al., 2022), as well as to voter characteristics such as socioeconomic status and prior exposure to violence (Bratton, 2008; Gutiérrez-Romero and LeBas, 2020). While this work underscores the strategic nature of violence, it largely ignores the informational role of pre-election attacks and how turnout reductions can shift remaining voters toward the attacker.

Finally, this paper also connects to work on how the allocation of rewards and punishments depends on the level at which votes are reported or constituencies are defined (Gingerich and Medina, 2013; Larreguy et al., 2016; Rueda, 2015; Rueda et al., 2025; Smith and de Mesquita, 2012). Prior studies show that more disaggregated results facilitate clientelism by improving monitoring of brokers (Larreguy et al., 2016) or increasing the pivotality of individual votes for continued benefits (Rueda, 2015; Smith and de Mesquita, 2012). We build on this latter mechanism, showing how pre-election violence, by depressing turnout, raises the pivotality of remaining voters and strengthens incentives to switch support to avoid punishment. We extend this literature by emphasizing the informational role of violence in making threats more credible, and how this effect depends on the size and composition of the monitored group.

A simple model

Consider a group of N voters, a subset of the electorate, in an election where two candidates, A and B, compete. Of these voters, N^A prefer candidate A, while the remaining ones prefer B. Each voter i receives an expressive private utility of γ_i when voting for her preferred candidate. Although γ_i values are unobserved by others, they are commonly known to be independent realizations from a continuous distribution F with support on $[0, 1]$. The election is conducted with a secret ballot: individual votes are not observable, but the group's aggregate vote totals are. We should think of the group of voters as those registered to vote at the same polling station whose vote totals will be published. Because voters are non-pivotal in large electorates, payoffs linked to election outcomes are independent of individual actions (Dekel et al., 2008); accordingly, we follow the literature and abstract from instrumental considerations (e.g., Casas, 2018).

A campaign operative working for A threatens the group of voters with post-electoral reprisals if A performs poorly in that particular locality. The operative may either have the capability to carry out this threat, denoted by $\overline{\omega}$, or lack such capability, denoted by $\underline{\omega}$. We refer to the former as *strong* and the latter as *weak*. Strong types possess the means, such as an organized group

of thugs, to inflict significant physical harm on voters, and are willing to do so. Weak types, by contrast, are either unable or unwilling to follow through on their threats. The operative's type is private information, but voters know that the fraction of strong types is $\mu \in (0, 1)$. This uncertainty reflects that even when a party has a reputation for violence, due to a history of insurgency or recent patterns of repression (Ishiyama et al., 2022), voters may not know whether that capability will be deployed in their locality. For now, we assume that candidate B's operatives do not engage in violence.

After the election, a strong operative carries out a reprisal with probability $\beta\left(1 - \frac{V^A}{\tilde{N}}\right)$. The parameter $\beta \in [0, 1]$ captures the likelihood that external constraints such as monitors, peacekeepers, or media scrutiny do not prevent the attack. The term $1 - \frac{V^A}{\tilde{N}}$ is the probability the attack is initiated, where $\tilde{N}$ is the post-threat group's size and V^A denotes A's votes; thus, reprisals are more likely when A performs poorly. In this baseline, local results (not national outcomes) determine the risk of punishment, though we later consider general outcome-contingent reprisals.

Operatives know that N^A voters prefer candidate A, but do not know the preferred candidate of each voter. We use this conservative assumption to isolate a rationale for pre-electoral violence that is not linked to changing the composition of the electorate (the main mechanism explored in the existing literature). Moreover, the assumption is consistent with voters' efforts not to reflect their voting intentions in violent campaigns. For example, supporters of the United Party for National Development (UPND) in Zambia 2016 often wore green campaign clothing from the Patriotic Front (PF) or abstained from using their party red; a so-called *watermelon strategy*, to avoid harassment from PF cadres (Wahman, 2024, p. 91).

Post-election reprisals do not affect all voters equally. Conditional on a reprisal, each voter is unaffected with probability $\delta \in [0, 1]$, which captures the extent of punishment, with higher values indicating less widespread harm. While β reflects whether outside factors prevent an attack by A, δ measures how broadly harm is applied if an attack occurs. Although conceptually distinct, β and δ play similar strategic roles in the baseline model. Keeping them separate, however, allows us to study two cases later: one where overall election results, rather than local outcomes, determine

Figure 1. Timeline

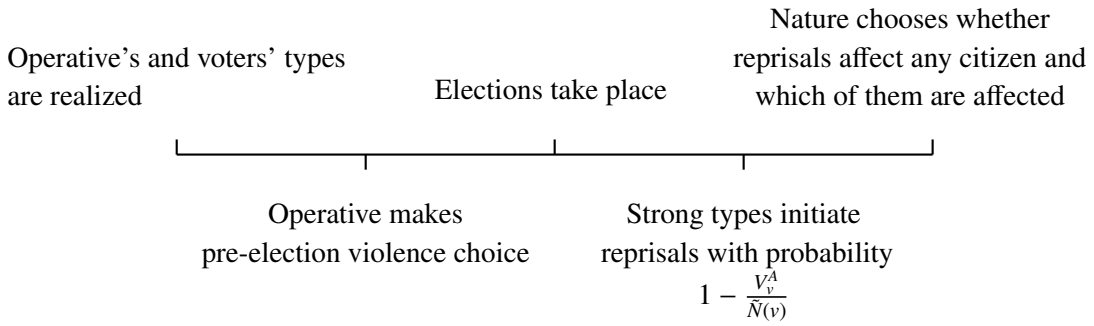

whether reprisals occur, and another where the severity of reprisals depends on the number of defectors.

In addition to caring about expressing their political preferences at the ballot, voters want to avoid post-electoral reprisals. At the beginning of the game, all voters have a utility unit. If a voter is affected during the reprisals, the voter loses it. Operatives, on the other hand, want to maximize the vote share of their candidate in this locality.[2]

Before the election, the operative can choose whether to commit a violent act against voters. This represents a show of force that could enhance the credibility of the post-electoral reprisal threat. The cost of engaging in pre-electoral violence for the strong type is $c_{\overline{\omega}}$, and that of the weak type is $c_{\underline{\omega}}$ with $0 \leq c_{\overline{\omega}} < c_{\underline{\omega}}$. That is, it is easier for strong type operatives to exert pre-election violence than for weak types. If the operative chooses to engage in pre-electoral violence, the number of voters is reduced by $K \in \{1, ..., N-2\}$,[3] with the voters that are eliminated chosen at random. The random selection of eliminated voters is consistent with the assumption that preferences are not observable. This turnout reduction is inline with a body of observational and quasi-experimental studies that document the demobilizing effects of pre-electoral violence (e.g., Alacevich and Zejcirovic, 2020; Collier and Vicente, 2014; Condra et al., 2018; Ley, 2018; van Baalen, 2024). We can interpret the reduction in turnout as the combined result of fear that election day might also

[2]In the Appendix, we describe changes in the results when the operative maximizes the number of votes.

[3]If $K = N - 1$ there will be certainty about the voting behavior of the remaining voter.

Table 1. B voter's expected payoffs after pre-electoral violence choice of v

Operative:	Strong $\overline{\omega}$		Weak $\underline{\omega}$
	Reprisal	No Reprisal	No Reprisal
Probability:	$1 - \frac{V_v^A}{\tilde{N}(v)}$	$\frac{V_v^A}{\tilde{N}(v)}$	1
A	$\beta\delta + (1-\beta)$	1	1
B	$\beta\delta + (1-\beta) + \gamma_i$	$1 + \gamma_i$	$1+\gamma_i$

be violent (von Borzyskowski et al., 2022), or backlash from voters against the use of violence (Rosenzweig, 2023).[4] After pre-electoral violence choices are made, all of those remaining in the group vote simultaneously, and nature decides whether there is a reprisal and who is affected by it. The timing is illustrated in Figure 1.

Table 1 summarizes the expected payoffs for a B supporter who votes for either A or B, conditional on the pre-electoral violence choice $v \in \{V, NV\}$, where V denotes violence and NV no violence. Consider first the payoff from voting for A (first row). If the operative is strong, a reprisal occurs with probability $1 - \frac{V_v^A}{\tilde{N}(v)}$. Conditional on a reprisal, it succeeds with probability β and fails with $1 - \beta$. If there is successful reprisal, the voter avoids harm with probability δ (retaining utility 1) or is harmed with probability $1 - \delta$; if no reprisal occurs—because A's vote share is high or the operative is weak, the voter retains the utility unit. The same arguments apply when the voter chooses B, except she also receives γ_i.

This is a dynamic game of incomplete information. An equilibrium in this game is a set of optimal strategies sustained by a belief system and such a belief system. A (pure) strategy for an operative is a function, $\chi : \{\underline{\omega}, \overline{\omega}\} \longrightarrow \{V, NV\}$, that maps the operative's type to a choice of pre-electoral violence. A strategy for a voter is function $\phi^j : [0,1] \times \{V, NV\} \longrightarrow \{A, B\}$, that maps the voter's type and observed operative's pre-electoral violent action to a voting choice for voters who support candidate $j \in \{A, B\}$. A belief system for voters is a probability distribution over the operatives' types, $\mu(.|v)$, derived via Bayes Rule given strategies when possible.

[4] We discuss the implications of the alternative turnout-enhancing attack in an extension.

Analysis

Note that A supporters have no reason to vote against their preferences. By voting for A, they reduce the chances of a post-electoral punishment and vote for their preferred candidate. B supporter i, on the other hand, will vote for A whenever

$$(1)\quad \mu_v\left(\beta\left(1-\frac{V^A_{-i,v}+1}{\tilde{N}(v)}\right)\delta+1-\beta\left(1-\frac{V^A_{-i,v}+1}{\tilde{N}(v)}\right)\right)\geq \mu_v\left(\beta\left(1-\frac{V^A_{-i,v}}{\tilde{N}(v)}\right)\delta+1-\beta\left(1-\frac{V^A_{-i,v}}{\tilde{N}(v)}\right)\right)+\gamma_i,$$

where $V^A_{-i,v}$ denotes (expected) others' votes for A and $\mu_v \equiv \mu(\overline{\omega}|v)$, the probability of facing a strong operative after observing pre-electoral violence choice v. We can deduce that the B supporters' net safety benefit for voting against their preferences is

$$(2)\qquad \Psi_0(v) \equiv \mu_v\frac{\beta(1-\delta)}{\tilde{N}(v)}.$$

The following result characterizes equilibrium voting behavior in the subgames that occur after A's operative chooses whether to engage in pre-electoral violence. All proofs are in the appendix.

Proposition 1. *In all equilibria, the strategies of voters after pre-electoral violence choice, v, are:*

1. *Candidate B supporter i votes for candidate A whenever* $\gamma_i \leq \Psi_0(v)$, *and votes for her preferred candidate otherwise.*

2. *Candidate A supporters vote for candidate A.*

The result implies that supporters of candidate B will vote for candidate A when their expressive attachment to B is weaker than the anticipated safety gain of supporting A. As expression 2 shows, this safety benefit increases with the likelihood of a successful reprisal (captured by the vulnerability parameter β) and with the expected severity of punishment, reflected in the risk of

personal harm $(1-\delta)$. This helps explain why violent political actors often target poor and rural areas, which typically face weaker government protection and receive less oversight from the media or observers (Gonzalez-Ocantos et al., 2020; Mares and Young, 2016).

The model indicates that pre-electoral violence is more effective in populations where a greater share of opposition voters have weak attachments to their preferred candidate. That is, where the distribution F places more mass on lower values of γ_i. This highlights a mechanism through which polarization, understood as a larger fraction of voters with strong attachments to their candidate, can reduce the effectiveness of coercive tactics. However, as we will show later, high polarization does not necessarily rule out the strategic use of pre-electoral violence under alternative assumptions about the information environment.

We also find that pre-electoral violence can increase the likelihood that B-supporters vote for A by reducing turnout in this precinct or polling station from N to $N-K$. As turnout falls, each vote becomes more pivotal in determining whether post-electoral reprisals occur, strengthening individual incentives to defect from one's preferred candidate to secure group safety. While the literature has emphasized turnout suppression among opposition supporters as the main objective of pre-electoral violence (Chaturvedi, 2005; Collier and Vicente, 2012; Gonzalez-Ocantos et al., 2020), our model highlights a distinct mechanism: shrinking the size of the monitored group can itself pressure remaining voters to support the attacker because their choices become more consequential in avoiding collective punishment. This effect does not rely on selectively targeting opponents; even indiscriminate violence that lowers turnout can yield local electoral gains. We refer to this logic as the *turnout mechanism*.

In addition to the turnout mechanism, Proposition 1 identifies a complementary channel through which pre-electoral violence benefits the perpetrator (captured by μ_v): it can shift beliefs by increasing the perceived likelihood that voters face a strong operative who will carry out post-electoral punishment. As this belief rises, so does the incentive to support the attacker to avoid future reprisals. We refer to this as the *terror mechanism*. Crucially, the turnout mechanism depends on threat credibility: if μ_v is close to zero, the greater pivotality created by lower turnout

has little effect on behavior. Together, these mechanisms are most effective when violence both reduces turnout and strengthens the credibility of post-election threats.

Terror: learning from pre-election attacks

The terror mechanism is more clearly illustrated by a separating equilibrium in which the only type of operative engaging in pre-electoral violence is the strong type. The next proposition characterizes this equilibrium.

Proposition 2. *There is a separating equilibrium in which the strong operatives engage in pre-electoral violence, and the weak do not, whenever*

$$c_{\underline{\omega}} \geq \left(1 - \frac{N^A}{N}\right) F\left(\frac{\beta(1-\delta)}{N-K}\right) \geq c_{\overline{\omega}}.$$

Voters follow the strategies described in Proposition 1 and believe the operative is strong after observing pre-electoral violence and weak when they are not attacked before the election.

The proposition shows that the net benefits of engaging in pre-electoral violence must be larger than the cost for the strong types but not higher than the costs for the weak types to sustain a separating equilibrium. The net benefits of pre-electoral violence, the expression between the costs parameters, represent the fraction of B supporters voting for A after such an attack. That is, the share of voters that would not have voted for A in the absence of threats, $1 - \frac{N^A}{N}$, multiplied by the ex-ante probability of them voting against their preferences, $F\left(\frac{\beta(1-\delta)}{N-K}\right)$. In sum, the benefits of pre-election attacks are the largest in areas where there are more *weak* supporters of the opposition. The finding can rationalize the observation that voters who live in opposition areas are often more fearful of violence than those living in areas where the incumbent has more support (Daxecker and Rauschenbach, 2023; Rauschenbach and Paula, 2019). An important caveat, however, is that the benefits of pre-election violence could be low even in an area where 100% of voters would have voted for the opposition. This could happen if all of those voters have a strong attachment to their

favorite candidate, making it harder for them to switch their votes for more levels of protection for voters, β, or intensities of pre- and post-attacks, K and $1-\delta$.

An important observation about Proposition 2 is that voters learn whether they are facing a strong type after observing the operative's choice of pre-electoral violence. If violence is observed, voters infer that they are dealing with a strong type. Conversely, if no violence occurs, they interpret any threats as coming from a weak operative and therefore view them as non-credible.

This is, however, not the only situation where the terror mechanism operates. When the costs of pre-electoral violence for the weak type are lower than $\left(1-\frac{N^A}{N}\right)F\left(\frac{\beta(1-\delta)}{N-K}\right)$, but not too low, there is a semi-separating equilibrium in which the weak operatives sometimes engages in pre-electoral violence while the strong continue choosing pre-election violence always. In this way, when voters observe pre-election violence, their assessment of the operative being strong will increase,[5] but without achieving certainty. Similarly, there could be an informative equilibrium where the strong type sometimes mimics a weak one by not attacking pre-election, and the weak one never attacks. Naturally, this situation arises when the costs of pre-electoral violence for the strong are high enough but not too high.

When the weak type's cost of engaging in pre-electoral violence are very low, weak operatives follow the strong types by engaging in pre-electoral violence always. When this happens, pre-electoral violence choices no longer reveal information about the type of operative voters face and violence is no longer useful to terrorize voters into believing worse outcomes will follow. This, however, does not mean pre-election violence will not be observed. Pre-election violence is still advantageous because the turnout mechanism is operating.

The equilibrium map in Figure 2 captures the previous discussion.[6] As the costs of pre-election violence increase for both types, we move from equilibria where strong and weak types attack

[5]The posterior probability of facing a strong operative after violence, μ_V, is $\frac{\mu}{\mu+(1-\mu)q_{\underline{\omega}}}$, where $q_{\underline{\omega}}$ is the probability of the weak operative attacking pre-election. This posterior is larger than μ when $q_{\underline{\omega}} < 1$.

[6]The figure represents a case where μ is small so $F\left(\frac{\beta(1-\delta)}{N-K}\right) - F\left(\mu\frac{\beta(1-\delta)}{N}\right) \geq F\left(\mu\frac{\beta(1-\delta)}{N-K}\right)$

Figure 2. Equilibrium Map

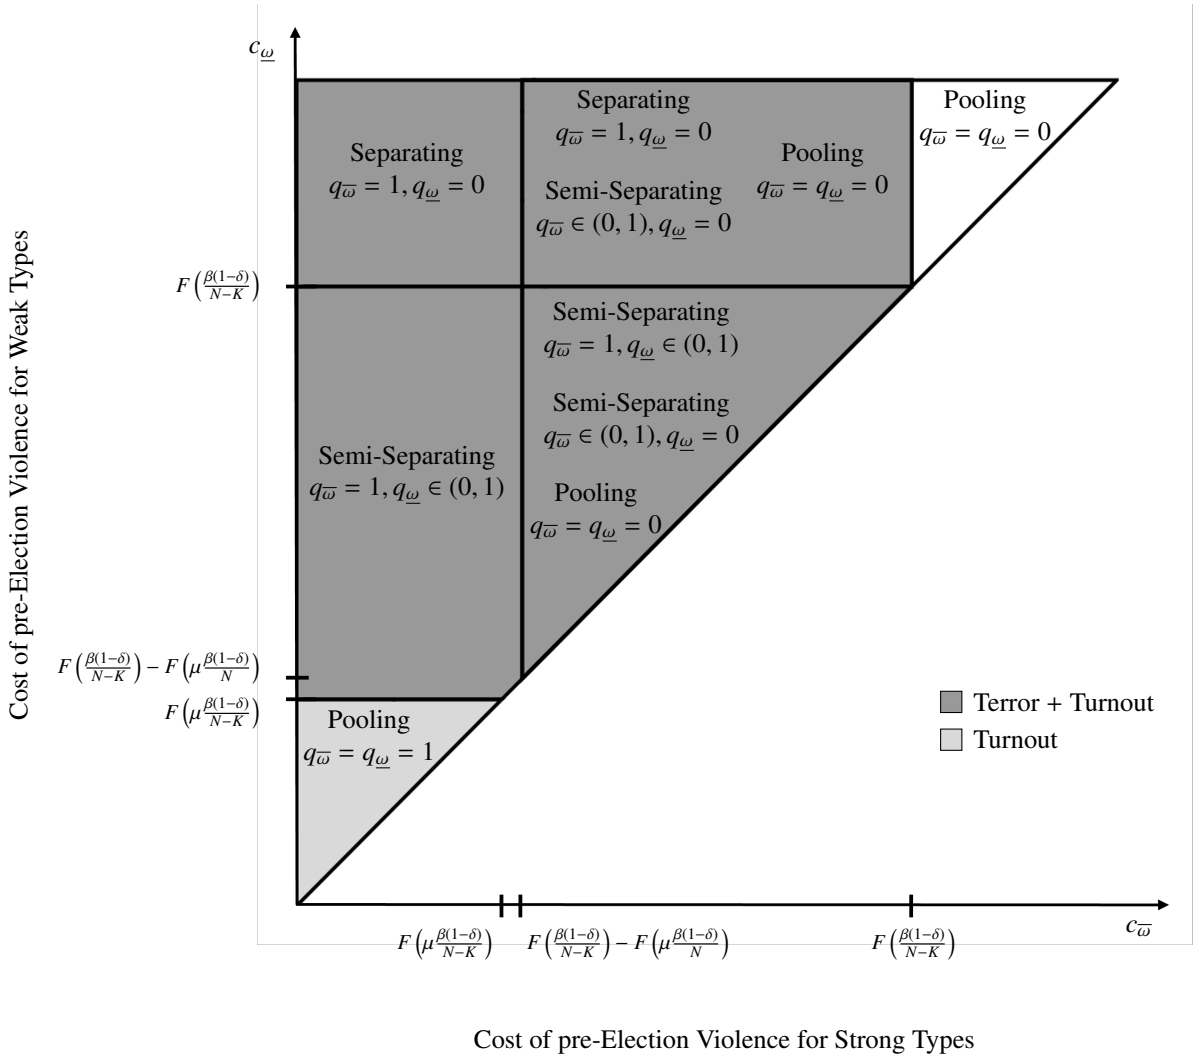

Note: the scale in both axis has been normalized by the share of B supporters, $1 - \frac{N^A}{N}$.

voters before the election and only the turnout mechanism operates (left south of the figure) to one where neither of them can engage in such attacks (right north). The appendix fully characterizes all the equilibria.

Comparative statics in the vulnerability parameter, β, further clarify the incentives at play. Consider the case where the strong type always uses pre-election violence ($c_{\overline{\omega}} = 0$), shown in

Figure 3. Comparative statics, voter vulnerability, β, when the strong type attacks pre-election ($c_{\overline{\omega}} = 0$).

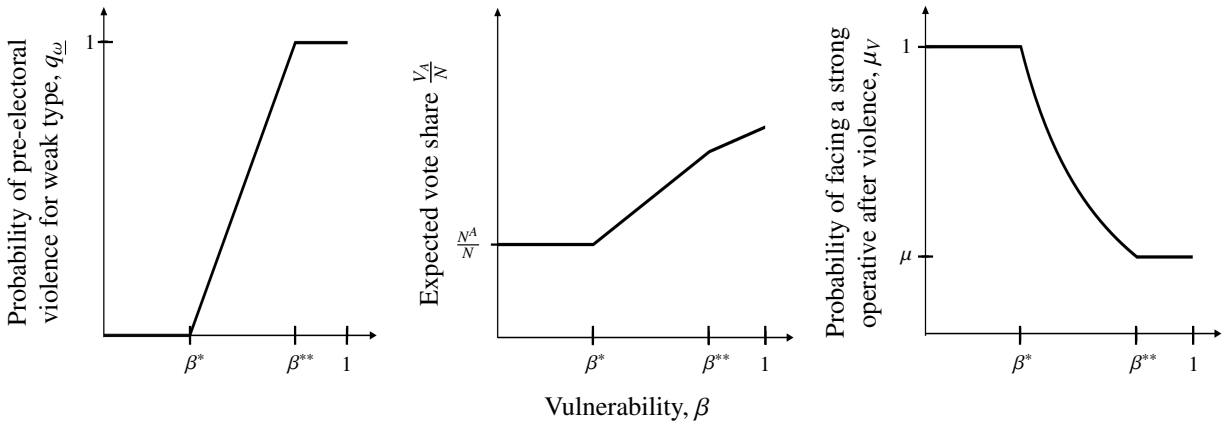

Figure 3.[7] When voter vulnerability is low ($\beta < \beta^*$), a weak operative does not attack, only A supporters vote for A, and voters know the operative is weak ($\mu_{NV} = 0$). For intermediate vulnerability ($\beta^* \leq \beta \leq \beta^{**}$), the weak type sometimes mimics the strong by attacking pre-election; fear and turnout suppression push A's vote share above $\frac{N^A}{N}$, and beliefs jump to $\mu_V > 0$, but as voters anticipate this mimicry, μ_V declines with β, producing the downward-sloping belief curve. For high vulnerability ($\beta > \beta^{**}$), the weak type always attacks, A's vote share continues to rise only through turnout effects (at a slower rate), and violence no longer signals the type ($\mu_V = \mu$).

A separate implication of the way pre-electoral violence affects the effectiveness of the threat of post-electoral violence is that the more intense the pre-electoral violence is (larger Ks), the smaller the chances of observing post-electoral violence. This is because if the reduction of turnout pre-election is large, the turnout mechanism will be stronger, making a larger share of B supporters switch their votes, which reduces the likelihood of a post-electoral reprisal.

[7] In Figure 3, the cutoffs β^* and β^{**} are $F^{-1}\left(\frac{c_\omega}{\left(1-\frac{N^A}{N}\right)}\right)\frac{N-K}{1-\delta}$ and $\frac{\beta^*}{\mu}$, respectively, where we use the fact that F is invertible in the support of the distribution.

Remark 1. *Conditional on observing pre-electoral violence in any equilibrium, the likelihood of observing post-electoral violence is weakly decreasing in the intensity of pre-electoral violence, K.*

Note that this result does not imply that areas with pre-electoral violence necessarily experience less post-electoral violence than those without it. In a separating equilibrium where only strong types attack pre-election, the absence of such violence guarantees no post-election reprisals, whereas its presence may be accompanied by future punishments.

The result also suggests an alternative empirical test of whether pre-election violence affects vote choice, complementing standard correlations between violence and party vote shares. If such violence increases a party's vote share, then in areas with pre-election attacks we should observe a negative relationship between the same party's pre- and post-election violence.

Visibility of preferences, group regarding preferences, and ethnic voting

The baseline model assumes the operative does not observe individual voters' preferred candidates, reflecting settings where parties are weak, attachments are fluid, and voters may conceal preferences under the threat of violence. This is less appropriate when preferences align with observable traits, as in many multi-ethnic societies where religion, language, or other visible attributes strongly correlate with vote choice. We now assume candidate preferences are observable, while maintaining that the intensity of preferences, γ_i, remains private information.

When voter preferences are observable, the benefits of pre-electoral violence increase for the operative, since the attacks can be precisely targeted at those who would otherwise vote against the operative's candidate. Under this assumption, the share of A supporters voting increases, as none are mistakenly targeted. Meanwhile, the number of B supporters who could switch to voting for A decreases. However, the gain in votes from A supporters outweighs the loss of switched votes from B supporters, as A supporters vote for A with certainty, whereas B supporters only do so with probability $F\left(\mu_V \frac{\beta(1-\delta)}{N-K}\right)$. The observability of preferences thus makes any equilibrium involving pre-electoral violence easier to sustain and increases the likelihood of such violence. The following proposition summarizes this dicussion.

Proposition 3. *In all equilibria, the benefits of engaging in pre-electoral violence when the operative observes each voter's preferred candidate are*

$$\frac{N^A}{N-K} + \frac{N-N^A-K}{N-K} F\left(\mu_V \frac{\beta(1-\delta)}{N-K}\right), \tag{3}$$

which are larger than those obtained when she does not.

Moreover, this observability premium increases with the share of A supporters, the intensity of the pre-electoral attack, and the share of B supporters with high expressive attachment to B.

The second statement in the proposition underscores how the observability of voters' preferred candidates makes pre-electoral violence particularly useful when the operative's candidate enjoys broad support. In such cases, if preferences were unobservable, the impact of violence and the fear it generates would be limited to a smaller subset of opponents. Similarly, when the operative believes that B supporters are likely to vote according to their preferences despite intimidation (for example, due to strong centralized protections against coercion or strong expressive attachments), then knowing individual preferences becomes even more valuable. It allows the operative to selectively demobilize those opponents for whom threats are least likely to change behavior.

In addition to the possibility that operatives can identify individual voters' preferred candidates, a second feature of elections in which ethnic identities are salient and tied to political preferences is that individuals might care relatively more about the welfare of others in their group. If we extend the baseline model to allow voters' utility functions to include the utility of others who share their preferred candidate weighted by $\zeta \in (0,1)$, the likelihood of B supporters voting for A increases. This is regardless of whether they care about all supporters of their candidate or just those who voted after a pre-election attack; although in the latter case, we require the weight placed on others in the utility function, ζ, to be small.[8] Voters in this setting comply not only to avoid personal

[8] In particular, $\zeta < \frac{1}{(N-N^A-1)\left(1-\frac{K}{N-1}\right)}$.

punishment, but also to protect fellow group members. Naturally, such collective concerns further strengthen the incentives to employ pre-electoral violence.

Proposition 4. *In all equilibria, the probability of voting against their preferred candidate when threatened with post-electoral punishments is higher when voters care about the welfare of others who prefer the same candidate than when they do not.*

Empirical research has shown that ethnic voting is less prevalent among ethnic minorities when there is a fear of intimidation (Enamorado and Kosterina, 2022), suggesting that violence can be an effective tool for altering vote choices in such contexts. Furthermore, using various indicators of pre-electoral violence, Kuhn (2015) finds a robust positive correlation between these measures and ethnic voting in a sample of sub-Saharan African countries. The rationale offered is that when ethnic identity or group affiliation outweighs campaign appeals, candidates are less inclined to rely on persuasion and instead turn to coercive strategies. The model gives an additional explanation: ethnic voting increases the strategic appeal of pre-electoral violence because it reduces the risk of inadvertently targeting the operative's own supporters. At the same time, it enhances the effectiveness of threats by making it more likely that voters will choose to vote against their group's candidate in order to avoid reprisals that may harm other members of their ethnic group.

Ethnic voting may be associated with strong expressive attachments to the preferred candidate (i.e., a low probability under F that voters derive little utility from voting sincerely), which tends to limit the effectiveness of pre-electoral violence because expressive costs can outweigh security benefits. However, Proposition 3 shows that even when few opponents are willing to switch (i.e., the second term in 3 is near zero), violence can remain profitable. By targeting only opponents, the first term in 3 stays positive and increases with the number of eliminated opponents K. Thus, while the baseline model predicts more violence where many weak opponents exist, ethnic voting can sustain violence even in opposition areas with strong attachments.

Terror competition

Up to this point, we have considered a setting in which only one party possesses the capacity to harm voters. The baseline model, in this sense, captures scenarios where one party exercises exclusive control over coercive resources in a given region, or where opposition parties deliberately refrain from employing intimidation tactics. However, it is unfortunately common to observe contexts in which multiple parties can threaten voters.

In this section, we allow both candidates (through their operatives) to threaten post-electoral violence if electoral outcomes are unsatisfactory. Each operative observes their own type—either *strong* or *weak*—but not the type of the opposing operative. As before, a weak operative is unable to carry out post-electoral punishment. Voters, in turn, do not observe the type of either party's operative. It is common knowledge, however, that the fraction of strong types is μ.

A strong operative working for candidate A will initiate post-electoral reprisals with probability $\beta\left(1 - \frac{V^A_{v_A,v_B}}{\tilde{N}(v_A,v_B)}\right)$, while a strong operative working for candidate B will do so with probability $\beta\frac{V^A_{v_A,v_B}}{\tilde{N}(v_A,v_B)}$. Thus, if voters face strong operatives from both parties, a post-election punishment is certain to occur, but whether it is carried out by A or B depends on the local electoral outcome. Conditional on A's operative conducting a successful reprisal, the probability that a voter escapes harm is δ_A; if the reprisal comes from B's operative, that probability is δ_B. We assume, without loss of generality, that $\delta_B > \delta_A$, meaning that party A's strong operative is *advantaged* in the sense of being capable of carrying out more extensive post-election punishments.

We further assume that operatives decide whether to engage in pre-electoral violence without knowledge of the other party's decision. The pre-electoral violence implemented by party j's operative reduces turnout by K_j, where $N > K_A + K_B$. The costs associated with pre-electoral violence remain as in the baseline model for both operative types, and voters observe which party is responsible for the pre-electoral attack.

The next proposition characterizes a fully separating equilibrium in which only strong operatives from both parties engage in pre-electoral violence.

Proposition 5. *The following strategies and beliefs constitute an equilibrium.*

1. *Advantaged party operative*

   Strong types engage in pre-electoral violence, and weak types do not whenever

$$c_{\underline{\omega}} \geq \mu\left(\frac{N^A}{N}F\left(\frac{\beta(1-\delta_B)}{N-K_B}\right)+\left(1-\frac{N^A}{N}\right)F\left(\frac{\beta(\delta_B-\delta_A)}{N-K_A-K_B}\right)\right)+(1-\mu)\left(1-\frac{N^A}{N}\right)F\left(\frac{\beta(1-\delta_A)}{N-K_A}\right) \geq c_{\overline{\omega}};$$

2. *Disadvantaged party operatives*

   Strong operatives engage in pre-electoral violence and weak types do not whenever,

$$c_{\underline{\omega}} \geq \mu\left(1-\frac{N^A}{N}\right)\left(F\left(\frac{\beta(1-\delta_A)}{N-K_A}\right)-F\left(\frac{\beta(\delta_B-\delta_A)}{N-K_A-K_B}\right)\right)+(1-\mu)\frac{N^A}{N}F\left(\frac{\beta(1-\delta_B)}{N-K_B}\right) \geq c_{\overline{\omega}};$$

3. *Voters*

   The safety benefit of a voter voting for their least preferred candidate j are given by

$$\Psi_0^j(v_j, v_{-j}) \equiv \begin{cases} \frac{\beta(1-\delta_j)}{N-K_j} & \text{if } v_j = V \text{ and } v_{-j} = NV, \\ \frac{\beta(\delta_{-j}-\delta_j)}{N-K_j-K_{-j}} & \text{if } v_j = v_{-j} = V, \\ 0 & \text{Otherwise.} \end{cases}$$

   Therefore,

   (a) *Voter i votes for her least preferred candidate j after j's operative is the only one engaging in pre-electoral violence whenever $\gamma_i \leq \Psi_0^j(V_j, NV_{-j})$ and votes for her preferred candidate otherwise. Supporters of party j, vote for party j.*

   (b) *When both parties engage in pre-electoral violence, a supporter of party B i will vote for A whenever $\gamma_i \leq \Psi_0^A(V, V)$ and vote for B otherwise. Supporters of A vote for A.*

(c) When there is no pre-electoral violence, all voters vote for their preferred candidate.

4. *Beliefs*

   Voters believe the operative is strong after observing pre-electoral violence. If they do not observe pre-electoral violence, they believe they face a weak operative.

The operatives' strategies reveal that pre-electoral violence accomplishes two tasks: 1) it induces voters who weakly support the opposing candidate to switch their vote in favor of the aggressor; and 2) it deters some voters from defecting to the opposing party out of fear of retaliation from their own party. By engaging in pre-electoral violence, B's operative prevents $\left(F\left(\frac{\beta(1-\delta_A)}{N-K_A}\right) - F\left(\frac{\beta(\delta_B-\delta_A)}{N-K_A-K_B}\right)\right)$ of B supporters to defect to A. Crucially, this implies that an operative may find it optimal to employ pre-electoral violence even in areas where her candidate would have performed well in the absence of any intimidation. Specifically, if the perceived probability of facing a strong operative from candidate A is high, the disadvantaged strong operative working for B may use pre-electoral violence to prevent A from poaching supporters in B's stronghold.

While much of the literature examines whether electoral violence is more likely in competitive areas or party strongholds (e.g., Hafner-Burton et al., 2014; Robinson and Torvik, 2009), it has not, to our knowledge, considered how the mutual capacity of multiple parties to use pre-electoral violence interacts with voters' preferences. Wahman (2024), using data from the 2016 Zambian and 2014 Malawian elections, shows that competitive areas were less likely to experience pre-election violence. In these contexts, the main competing parties—the PF and UNPD in Zambia, and the Democratic Progressive Party and the People's Party in Malawi—used violence and intimidation as campaign tactics to defend territorial strongholds and to contest them. Our analysis complements this account by highlighting how parties may also deploy violence in their own strongholds when they fear losing voters to their opponents' intimidation.

Interestingly, terror competition may compel a strong but disadvantaged (low delta) B operative to carry out a pre-electoral attack even when doing so risks pushing more of its supporters to vote for A. This outcome arises when voters expect that A will also engage in pre-election violence. To

see why, consider that a pre-electoral attack by B signals the presence of a strong operative. As a result, some of B's own supporters may become fearful of triggering a post-election punishment by their own party if B underperforms, and thus refrain from voting for A. This effect is reflected by the shift in the numerator in Ψ_0^A, from $\beta(1-\delta_A)$ to $\beta(\delta_B-\delta_A)$. At the same time, if both A and B launch pre-electoral attacks, turnout declines further than if only A's operative had attacked. This raises the salience of each individual vote in determining the collective punishment that will be more effectively applied by A (since $\delta_B > \delta_A$), making voters more likely to vote for A. If the latter effects dominate the former, a B operative might push more B supporters to vote for A when attacking pre-election.[9] When does this happen? If the fraction of A supporters is large and the likelihood of encountering a strong A operative is low, a strong B operative could try to take A supporters by signaling its type, despite the (low) risk of pushing its supporters away.

Endogenous post-election reprisals, voters' vulnerability, and turnout enhancing pre-election attacks

We now return to the case where only A threatens voters and study the implications of three assumptions of the baseline model: 1) the probability that an individual voter is harmed in a post-election reprisal, conditional on such reprisals occurring, is independent of how poorly candidate A performed; 2) post-election reprisals depend solely on local election outcomes, not on who wins the election, and 3) pre-election attacks reduce turnout.

We start by relaxing the first assumption allowing for the possibility that the severity of post-election punishment (and not just its occurrence) is a function of the group's electoral support for candidate A. This can be the result of voters anticipating that especially poor results for A could lead to harsher reprisals. A natural way to formalize this idea is by replacing δ with $\frac{V_v^A}{\tilde{N}(v)}$, so that the probability of remaining unharmed in a post-electoral reprisal increases with the share of votes that A receives from the group. If this is the case, a B supporter would vote for A whenever,

[9] In this case, $\left(F\left(\frac{\beta(1-\delta_A)}{N-K_A}\right) - F\left(\frac{\beta(\delta_B-\delta_A)}{N-K_A-K_B}\right)\right) < 0$.

$$\mu_v\left(\beta\left(1-\frac{V^A_{-i,v}+1}{\tilde{N}(v)}\right)\frac{V^A_{-i,v}+1}{\tilde{N}(v)}+1-\beta\left(1-\frac{V^A_{-i,v}+1}{\tilde{N}(v)}\right)\right)\geq\mu_v\left(\beta\left(1-\frac{V^A_{-i,v}}{\tilde{N}(v)}\right)\frac{V^A_{-i,v}}{\tilde{N}(v)}+1-\beta\left(1-\frac{V^A_{-i,v}}{\tilde{N}(v)}\right)\right)+\gamma_i.$$

If q is the ex-ante probability that a B supporter votes for A, we can show that the expected net benefit for a B supporter of voting against its preferences after pre-election violence choice v is

$$\Psi_1(q;v)\equiv\frac{\mu_v\beta}{\tilde{N}(v)}\left(2\left(1-\frac{N^A}{N}-\frac{N-N^A-1}{N}q\right)-\frac{1}{\tilde{N}(v)}\right). \tag{4}$$

Unlike in the baseline model, a voter's expectations about how other B supporters will vote now affect the perceived benefit of voting insincerely: Ψ_1 depends on q, whereas Ψ_0 does not. Since Ψ_1 is decreasing in q (assuming nonzero vulnerability or a positive probability of facing a strong type), higher expectations that others will vote for A lower the individual incentive to do so. This reflects free riding: if enough B supporters defect, reducing the likelihood or severity of reprisals, a given voter may feel protected and vote sincerely. The following proposition formalizes this result.

Proposition 6. *When the intensity of post-electoral reprisals decreases in the vote share of the attacker,*

1. *B supporters have less incentives to vote for A the more they expect other B supporters to vote for A.*
2. *There is a unique equilibrium in the voting subgame in which:*
   (a) *A B supporter i votes against her preferences whenever $\gamma_i \leq \Psi_1(q_v^*;v)$, where q_v^* is the solution to $F(\Psi_1(q;v))=q$ after pre-electoral violence choice v, and*
   (b) *A supporters vote for A.*
3. *The equilibrium probability of B supporters voting for A is decreasing in the number of A supporters.*

The final observation from the proposition also highlights the possibility of free-riding on A sympathizers' votes, which shield a B voter from more severe post-election reprisals. The model thus clarifies a separate incentive to concentrate pre-electoral violence in opposition strongholds: in areas where support for the threatening party is low, opposition voters anticipate harsher reprisals (and a higher likelihood of the reprisal taking place), pushing them to vote for their attacker.

For the remainder of this section, we return to the case where δ is exogenous, but consider an alternative scenario in which post-election reprisals occur only if candidate A wins the election. The alternative setting reflects the idea that a strong operative loyal to A may be unable to carry out reprisals if candidate B wins and mobilizes state forces to protect B voters.

In this version of the model, we retain all baseline elements but introduce an external (outside the group of N) group of M voters. We can think of M as the size of the rest of the electorate. A fraction $\alpha \in [0, 1]$ of these M voters are assumed to vote for A. A post-election reprisal targeting the group of $\tilde{N}(v)$ voters occurs only if 1) candidate A wins the election, and 2) a strong operative initiates a reprisal in response to poor local support among the $\tilde{N}(v)$ voters. Thus, the probability of a post-election reprisal following a pre-election violence choice v is given by $\frac{\alpha M + V_v^A}{M + \tilde{N}(v)} \left(1 - \frac{V_v^A}{\tilde{N}(v)}\right)$.

As in the case where the intensity of post-election reprisals was endogenous, a B supporter's expectations about how others will vote continue to influence her decision. However, unlike in the endogenous δ case, here a B supporter is *more* likely to vote for A if she believes others will do the same. The reason is that the greater the likelihood that B supporters vote for A, the more likely A is to win the election, which would enable the operative to carry out reprisals if the local results are not favorable enough in that particular locality.

Proposition 7. *When post-election reprisals materialize only when candidate A wins the election,*

1. *B supporters have more incentives to vote for A the more they expect other B supporters to vote for A*

2. *There is an equilibrium in the voting subgame in which:*

(a) A B supporter i votes against her preferences whenever $\gamma_i \leq \Psi_2(\tilde{q}_v; v)$, where $\tilde{q}_v$ is a solution to $F(\Psi_2(q; v)) = q$ after pre-electoral violence choice v, and

$$\Psi_2(q; v) \equiv \frac{\mu_v(1-\delta)}{\tilde{N}(v) + M}\left(2\left(\frac{N^A}{N} + \frac{N - N^A - 1}{N}q\right) + \frac{\alpha M + 1}{\tilde{N}(v)} - 1\right),$$

(b) A supporters always vote for A.

3. *For large electorates (as M approaches infinity), the equilibrium in the voting subgame is unique, and the equilibrium probability of a B supporter voting for A, $\tilde{q}_v$ is $F\left(\mu_v \frac{\alpha(1-\delta)}{\tilde{N}(v)}\right)$.*

Without additional restrictions on the distribution of attachment intensities, equilibrium uniqueness in the voting subgame is not guaranteed given the upward-sloping benefits-of-switching function, Ψ_2. As the electorate grows while the fraction of A supporters outside the locality remains fixed, however, the model converges to the baseline case as described in the third observation of the proposition. In this limit, the fraction of A supporters, α, replaces the vulnerability parameter, β. In short, the baseline model is a special case of a broader model in which the election outcome determines whether a local group of voters can be punished for their behavior.

Finally, recall that we have assumed turnout is reduced with pre-election violence. This was motivated by a large body of literature that documents reductions in turnout as a consequence of violent campaigns. Nevertheless, it is possible that pre-election violence increases voter turnout, when operatives force people to vote, and others have shown some support for such expectations (e.g., Daxecker et al., 2024). To study this situation we take the baseline model but now assume that after a pre-election attack, there are K additional voters, with their preferences drawn from the original distribution.

A pre-election attack by a A operative now has two opposing effects on B supporters. The terror mechanism can still operate, raising voters' beliefs that they face a strong type after observing an attack. At the same time, because turnout increases, a B supporter is less likely to vote for A since her vote becomes less pivotal in triggering post-election reprisals for the monitored group. This implies that when a party lacks information about individual preferences, mobilizing voters

through violence removes a key channel for inducing defection, making turnout deterrence preferable to coercive mobilization. This holds *regardless* of the preference distribution, because all gains from attacks come from opponents who switch. Mobilization through violence is beneficial only when preferences are observable, allowing the operative to ensure that the K additional voters are B supporters who vote for B with probability 1. Further discussion appears in the appendix.

Conclusions

We develop a theoretical framework that clarifies multiple ways in which pre-election violence can yield electoral gains. Attacks signal a party's willingness to punish poor performance, while suppressing turnout in monitored groups raises the influence of remaining votes on triggering future reprisals. Together, fear and heightened pivotality in small monitored groups push voters toward the perpetrating party—even under secret ballots and beyond simple opponent demobilization (or supporters mobilization).

The model explains why parties may target swing districts, opposition strongholds, or even their own bases. Targeting depends not only on expected vote choice but on the strength and observability of voter preferences, and on whether multiple parties can credibly threaten voters. Weak attachments make voters more responsive to fear, while strong but easily identifiable individual preferences allow violence to selctively suppress turnout. The possibility that multiple parties can coerce voters further shapes targeting: attacking swing areas or strongholds may be optimal if the intimidation offsets rival parties' coercive efforts.

Although our analysis is local, electoral violence can generate broader spillovers affecting general election dynamics. Attacks in one area may signal a candidate's willingness to punish dissent or disregard rules, reshaping expectations and voting behavior elsewhere. Understanding how parties account for these geographic externalities and choose optimal campaign strategies that account for local agency problems, like those studied here, is an important direction for future research.

Competing Interests: The authors declare none.

Acknowledgements: We thank Dan Gingerish, Maggie Penn, Jay Shon, Scott Tyson, and seminar participants at Vanderbilt, PEPSI conference University of Essex, and the ASPEN Political Economy conference for their feedback. All remaining errors are our own.

AI statement The authors made use of ChatGPT (GPT-5, OpenAI) to assist with English language editing. All analysis and interpretation were made by the authors.

References

Alacevich, C. and D. Zejcirovic (2020). Does violence against civilians depress voter turnout? Evidence from Bosnia and Herzegovina. *Journal of Comparative Economics 48*(4), 841–865.

Brancati, D. and E. M. Penn (2022). Stealing an election: Violence or fraud? *Journal of Conflict Resolution 67*(5), 858–892.

Bratton, M. (2008). Vote buying and violence in Nigerian election campaigns. *Electoral Studies 27*(4), 621–632.

Burchard, S. M. (2015). *Electoral Violence in Sub-Saharan Africa: Causes and Consequences*. Boulder, Colorado: FirstForumPress.

Casas, A. (2018). Distributive politics with vote and turnout buying. *American Political Science Review 112*(4), 1111–1119.

Chaturvedi, A. (2005). Rigging elections with violence. *Public Choice 125*(1/2), 189–202.

Collier, P. and P. C. Vicente (2012). Violence, bribery, and fraud: The political economy of elections in Sub-Saharan Africa. *Public Choice 153*(1-2), 117–147.

Collier, P. and P. C. Vicente (2014). Votes and violence: Evidence from a field experiment in Nigeria. *The Economic Journal 124*(574), F327–F355.

Condra, L. N., J. D. Long, A. C. Shaver, and A. L. Wright (2018). The logic of insurgent electoral violence. *American Economic Review 108*(11), 3199–3231.

Cox, G. W. (2009). Authoritarian elections and leadership succession, 1975-2004. Presented at the Annual Meeting of the American Political Science Association, Toronto.

Daxecker, U., A. Deglow, and H. Fjelde (2024). Voter intimidation as a tool of mobilization or demobilization? evidence from West Bengal, India. *Journal of Conflict Resolution 69*(7-8), 1232–1257.

Daxecker, U. and M. Rauschenbach (2023). Election type and the logic of pre-election violence: Evidence from Zimbabwe. *Electoral Studies 82*, 102583.

Dekel, E., M. Jackson, and A. Wolinsky (2008). Vote buying: General elections. *Journal of Political Economy 116*(2), 351–380.

Ellman, M. and L. Wantchekon (2000). Electoral competition under the threat of political unrest. *The Quarterly Journal of Economics 115*(2), 499–531.

Enamorado, T. and S. Kosterina (2022). Surrounded and threatened: How neighborhood composition reduces ethnic voting through intimidation. *Political Science Research and Methods 10*(1), 68–81.

Fearon, J. D. (2011). Self-enforcing democracy. *The Quarterly Journal of Economics 126*(4), 1661–1708.

Fjelde, H. and K. Höglund (2016). Electoral institutions and electoral violence in sub-Saharan Africa. *British Journal of Political Science 46*(2), 297–320.

Fjelde, H. and K. Höglund (2022). Introducing the deadly electoral conflict dataset (DECO). *Journal of Conflict Resolution 66*(1), 162–185.

Gehlbach, S. and A. Simpser (2015). Electoral manipulation as bureaucratic control. *American Journal of Political Science 59*, 212–224.

Gingerich, D. W. and L. F. Medina (2013). The endurance and eclipse of the controlled vote: A formal model of vote brokerage under the secret ballot. *Economics & Politics 25*(3), 453–480.

Gonzalez-Ocantos, E., C. K. de Jonge, C. Meléndez, D. Nickerson, and J. Osorio (2020). Carrots and sticks: Experimental evidence of vote-buying and voter intimidation in Guatemala. *Journal of Peace Research 57*(1), 46–61.

Gutiérrez-Romero, R. and A. LeBas (2020). Does electoral violence affect vote choice and willingness to vote? Conjoint analysis of a vignette experiment. *Journal of Peace Research 57*(1), 77–92.

Hafner-Burton, E. M., S. D. Hyde, and R. S. Jablonski (2014). When do governments resort to election violence? *British Journal of Political Science 44*(1), 149–179.

Hassan, M. and T. O'Mealia (2018). Uneven accountability in the wake of political violence: Evidence from Kenya's ashes and archives. *Journal of Peace Research 55*(2), 161–174.

Hoppen, T. K. (1994). Grammars of electoral violence in nineteenth-century England and Ireland. *The English Historical Review CIX*(432), 597–620.

Ishiyama, J., M. C. Marshall, and B. Stewart (2022). Are former rebel parties more likely to engage in electoral violence in Africa? *Journal of Elections, Public Opinion and Parties 33*(2), 278–299.

Kousser, J. M. (1974). *The Shaping of Southern Politics: Suffrage Restriction and the Establishment of the One-Party South, 1880–1910*. New Haven: Yale University Press.

Kuhn, P. (2015). Do contentious elections trigger violence? In P. Norris, R. W. Frank, and F. M. i Coma (Eds.), *Contentious Elections: From Ballots to Barricades*, pp. 89–110. New York and London: Routledge.

Larreguy, H., J. Marshall, and P. Querubin (2016). Parties, brokers, and voter mobilization: How turnout buying depends upon the party's capacity to monitor brokers. *American Political Science Review 110*(1), 160–179.

Ley, S. (2018). To vote or not to vote: How criminal violence shapes electoral participation. *Journal of Conflict Resolution 62*(9), 1963–1990.

Little, A. T. (2012). Elections, fraud, and election monitoring in the shadow of revolution. *Quarterly Journal of Political Science 7*(3), 249–283.

Luo, Z. and A. Rozenas (2018). Strategies of election rigging: Trade-offs, determinants, and consequences. *Quarterly Journal of Political Science 13*(1), 1–28.

Mares, I. and L. Young (2016). Buying, expropriating, and stealing votes. *Annual Review of Political Science 19*, 267–288.

Müller-Crepon, C. (2022). Local ethno-political polarization and election violence in majoritarian vs. proportional systems. *Journal of Peace Research 59*(2), 242–258.

Rauschenbach, M. and K. Paula (2019). Intimidating voters with violence and mobilizing them with clientelism. *Journal of Peace Research 56*(5), 682–696.

Robinson, J. A. and R. Torvik (2009). The real swing voter's curse. *American Economic Review 99*(2), 310–15.

Rosenzweig, S. C. (2023). *Voter Backlash and Elite Misperception: The Logic of Violence in Electoral Competition*. Cambridge University Press.

Rueda, M. R. (2015). Buying votes with imperfect local knowledge and a secret ballot. *Journal of Theoretical Politics 27*(3), 428–456.

Rueda, M. R., G. Grossman, and S. Ge (2025). Do more disaggregated electoral results deter aggregation fraud? *British Journal of Political Science 55*, e140.

Rundlett, A. and M. W. Svolik (2016). Deliver the vote! Micromotives and macrobehavior in electoral fraud. *American Political Science Review 110*(1), 180–197.

Simpser, A. (2013). *Why Governments and Parties Manipulate Elections: Theory, Practice, and Implications*. Political Economy of Institutions and Decisions. Cambridge University Press.

Smith, A. and B. B. de Mesquita (2012). Contingent prize allocation and pivotal voting. *British Journal of Political Science 42*(2), 371–392.

Taylor, C. F., J. C. Pevehouse, and S. Straus (2017). Perils of pluralism: Electoral violence and incumbency in sub-Saharan Africa. *Journal of Peace Research 54*(3), 397–411.

Travaglianti, M. (2014). Threatening your own. electoral violence within ethnic groups in Burundi and beyond. New York University, Doctoral Dissertation.

Troxler, H. (2008). Electoral abuse in the late Roman Republic. USF Tampa Graduate Theses and Dissertations.

van Baalen, S. (2024). Polls of fear? Electoral violence, incumbent strength, and voter turnout in Côte d'Ivoire. *Journal of Peace Research 61*(4), 595–611.

von Borzyskowski, I., U. Daxecker, and P. M. Kuhn (2022). Fear of campaign violence and support for democracy and autocracy. *Conflict Management and Peace Science 39*(5), 542–564.

Wahman, M. (2024). *Controlling Territory, Controlling Voters: The Electoral Geography of African Campaign Violence*. New York: Oxford University Press.

Wantchekon, L. (1999). On the nature of first democratic elections. *Journal of Conflict Resolution 43*(2), 245–258.