Tug of War: The Heterogeneous Effects of Outbidding Between Terrorist Groups

Casey Crisman-Cox[a]* and Michael Gibilisco[b]

[a]Department of Political Science, Texas A&M University, College Station, TX, USA
[b]Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, USA
*Corresponding author. Email: c.crisman-cox@tamu.edu

Abstract We introduce a dynamic game of outbidding where two groups use violence to compete in a tug-of-war fashion for evolving public support. We fit the model to the canonical outbidding rivalry between Hamas and Fatah using newly collected data on Palestinian public support for these groups. Competition has heterogeneous effects, and we demonstrate that intergroup competition can discourage violence. Competition from Hamas leads Fatah to use more terrorism than it would in a world where Hamas abstains from terrorism, but competition from Fatah can lead Hamas to attack less than it otherwise would. Likewise, making Hamas more capable or interested in competing increases overall violence, but making Fatah more capable or interested discourages violence on both sides. These discouragement effects of competition on violence emerge through an asymmetric contest, in which we find that Fatah uses terrorism more effectively to boost its support, although Hamas has lower attack costs. Expanding on these results, we demonstrate that outbidding theory is consistent with a positive, negative, or null relationship between measures of violence and incentives to compete.

Outbidding is an explanation for terrorism that posits that competing antigovernment groups use violence to garner popular support at the expense of their rivals. In this story, terrorism signals resolve or capacity to a population that is uncertain about which group best represents its interests. Popularity and attention are in turn critical for groups' recruitment numbers, financial resources, political influence, and day-to-day operations.[1] It is a unique theory of terrorism because "the enemy is only tangentially related to the strategic interaction," and therefore outbidding "provides a potential explanation for terrorist attacks that continue even when they seem unable to produce any real results."[2] Because scholars are still debating the degree to which terrorism helps groups achieve their long-term political objectives (such as military

1. Crenshaw 1981; Fortna 2015; Polo and González 2020.
2. Kydd and Walter 2006, 77.

International Organization 79, Winter 2025, pp. 81–116

doi:10.1017/S0020818324000390

victories or government concessions), outbidding provides an important explanation for observed variation in terrorism and intrastate violence.[3]

Following Bloom's foundational work, researchers generally hypothesize greater violence when groups have stronger incentives to compete.[4] Conrad and Greene concisely summarize a key mechanism underlying outbidding theory: "Since competition directly and indirectly threatens the resource base necessary to sustain the organization and ensure its effectiveness, it follows that terrorist organizations should make tactical choices in an effort to increase their share of resources within a competitive environment."[5] When looking for evidence of outbidding, scholars therefore regress measures of violence on proxies for incentives to compete, such as the number of groups in a conflict, using time-series cross-sectional data and test for a positive association.[6] Within this framework, Findley and Young find no relationship between competition and violence; Chenoweth, Cunningham, Bakke, and Seymour, and Wood and Kathman find a positive relationship; Polo and Welsh find a negative one; and others highlight more conditional findings.[7]

Previous research designs have two substantial weaknesses when seeking evidence for or against outbidding theories. First, they require proxies for competitive incentives, but directly evaluating the strength of these proxies is difficult, especially when commonly used measures (such as the number of terrorist groups) may be confounded by other aspects of the conflict (such as state strength). Second, they assume that the *encouragement effect* of competition is consistent with outbidding, while suggesting that the *discouragement effect* is not.[8] Yet both are implications of competition in contest models.[9] On the one hand, enhanced competition can *encourage* violence, because if one group becomes more competitive, others may fight harder to keep up. On the other hand, enhanced competition can *discourage* violence, because if one group becomes more competitive, others may recognize their disadvantage and reduce their violence. This creates a feedback loop where even the most competitive group uses little violence because it anticipates no pushback. By associating outbidding only with an encouragement effect, previous research designs have overlooked the discouragement effect and how the two countervailing effects can wash out in the aggregate, masking evidence of outbidding.[10]

3. For the debate, see Abrahms 2006; Fortna 2015; Getmansky and Zeitzoff 2014; Gould and Klor 2010; Thomas 2014.

4. Bloom 2004.

5. Conrad and Greene 2015, 547.

6. There is disagreement on how to measure competitive incentives and on whether to measure the extent or the intensity of terrorism. Nemeth 2014; Polo and Welsh 2024.

7. Chenoweth 2010; Conrad and Greene 2015; Conrad and Spaniel 2021; Cunningham, Bakke, and Seymour 2012; Findley and Young 2012; Nemeth 2014; Polo and Welsh 2024; Wood and Kathman 2015.

8. As Conrad, Greene, and Phillips 2024 put it, "[the outbidding] argument *is* that intergroup competition leads to more violence" (1895, emphasis added).

9. Chaudoin and Woon 2018; Dechenaux, Kovenock, and Sheremeta 2015.

10. To be clear, we follow Conrad and Spaniel 2021 and Kydd and Walter 2006 in using "outbidding" to refer to a theory where groups use costly terrorism to increase their popularity relative to another group. But unlike past works, we do not assume that outbidding is consistent only with the encouragement effect.

In this paper, we show how scholars can estimate the effects of competition on violence and better quantify the degree to which outbidding explains terrorism data. Our key departure from previous work is the structural approach. Broadly, the goal of the structural approach is to construct an outbidding model, estimate its parameters and equilibrium from observed data, and study the properties of the fitted model.[11] This approach has three main benefits in the context of outbidding. First, we flexibly estimate groups' incentives to compete, sidestepping the need for proxies. Second, we use the fitted model to quantify the substantive effects of competition on violence by asking counterfactual questions such as "What would happen if one group expected no violence from its rival?" and "How would violence change if a group's competitive incentives increased?" Third, we can check how well outbidding fits the data, because we can see whether nonsensical parameter estimates arise and explicitly analyze model fit and comparison.

To do these three things, we focus on the canonical outbidding example: the rivalry between Hamas and Fatah. Narrowing the scope has several benefits. Theoretically, in a two-group rivalry, we model outbidding as a dynamic contest wherein each side uses terrorism to pull public opinion toward itself and away from its opponent in a tug-of-war fashion. Empirically, given the rivalry's lengthy history, we compile monthly survey data on aspects of Palestinian public opinion from 1994 to 2018. The data provide fine-grained details on how Palestinians view the conflict and the two groups, which we use to measure their relative popularity. Substantively, because it is the canonical (and theory-generating) example of outbidding,[12] it is of first-order importance to understand whether discouragement effects emerge in this rivalry. If they do, then work extrapolating to other environments should not treat the discouragement effect as a mere theoretical curiosity when looking for evidence of outbidding.

Our main result is that we identify and quantify two discouragement effects. First, we compare the estimated equilibrium rates of terrorism to those from counterfactual scenarios in which each group never anticipates violence from its rival. Comparing how a group behaves with and without violence from its rival is one way to examine group behavior in competitive and noncompetitive environments, respectively. We find that competition from Hamas has an encouragement effect on Fatah's use of violence (as expected per the outbidding literature), where Fatah is 34 percent more violent in equilibrium than when it expects Hamas to never attack. In contrast, we find that competition from Fatah can deter violence by Hamas. During the Oslo era, between 1994 and 2001, Hamas is 4 percent less violent in equilibrium than when it expects Fatah to never use violence. That is, competition from Fatah depresses Hamas's use of violence even during the time when the two groups are publicly vying for support from the Palestinians. This is the unexpected discouragement effect. After the Oslo era, we again find an encouragement

11. Canen and Ramsay 2024.
12. Bloom 2004; Jaeger et al. 2015.

https://doi.org/10.1017/S0020818324000390

effect where Hamas uses more violence because of competition from Fatah. What distinguishes these two periods in our framework is the groups' relative popularity: Fatah is much more popular than Hamas before the Second Intifada. When one group, especially a strong one, has a commanding lead in a dynamic contest, discouragement effects can appear.

Second, we conduct comparative statics that demonstrate how equilibrium rates of violence change as a group becomes more or less competitive—that is, has stronger or weaker incentives to compete. Whereas the first set of counterfactuals fixes the behavior of one group, this second set illustrates how the behavior of both groups changes as incentives to compete change. In contest models, groups have stronger competitive incentives when they place greater value on their popularity, have lower costs of attacking, or become more effective at using terrorism to attract support. We find that making Hamas more competitive along any of these three dimensions increases the probability that either group uses terrorism. This is the encouragement effect expected by the outbidding literature, where increasing the competitiveness of an actor leads to an increase in violence not only by the group in question but by all groups involved. If Fatah becomes more competitive along any of these dimensions, however, both groups' propensities for terrorism decrease. This is the unexpected discouragement effect of outbidding.

Because we use the structural approach, our theory provides an explicit explanation for the results rooted in asymmetric contests. Although we find that Hamas has lower costs for terrorism and places higher value on public support than Fatah, Fatah is more effective than Hamas at garnering support through attacks. That is, Fatah's attacks result in larger pro-Fatah shifts in public opinion than Hamas's attacks provide for Hamas.[13] Because Fatah is substantially more capable of moving public opinion with violence, *if* its incentives to compete increase, then it is more willing to take on the immediate costs of violence to move popular opinion more quickly. Hamas then struggles to compete with Fatah's greater efficiency and reduces its use of terrorism, creating an equilibrium feedback loop where Fatah also uses less violence.

We acknowledge that this explanation for why the discouragement or the encouragement effect appears is limited by our data and theory. Specifically, we treat each group's incentives to compete as latent, to-be-estimated preferences. In our model, these preferences are described by exogenous parameters and determine equilibrium behavior. Thus we cannot systematically explain *why* these specific preference estimates arise and *why* they are asymmetric. Nonetheless, we demonstrate how the parameters can be identified given the observed data and describe how they are consistent with explanations in other work. Likewise, we cannot identify substantive features of conflicts that cause these asymmetries to be large enough to generate discouragement

13. This finding is robust to time-varying controls and different codings of attacks. It also holds when instrumenting group attacks with past weather conditions. See Appendix D.

effects. To answer these questions, we would need to compare group preferences across conflicts or explicitly model the determinants of the competitive incentives.

This is not the first paper to find a discouragement effect between group competition and terrorism. Polo and Welsh also document such an effect when considering how a rebel group's decision to attack soft rather than hard (that is, civilian rather than military) targets is affected by group competition, which they proxy using the number of attacks on other groups.[14] They find that, as groups attack other groups in the same civil conflict, the proportion of their attacks against soft targets decreases. Where we diverge is in considering whether such an effect supports or refutes outbidding theory, as they argue that such a finding "emphasizes the strategic limitations of outbidding."[15] This conclusion may not be warranted for three reasons. First, we find that group popularity increases after terrorist attacks.[16] This finding holds when we restrict ourselves to attacks on civilians.[17] Second, the discouragement effect is entirely consistent with outbidding theory: it appears when we fit an outbidding model to the theory's generating case study.

Third, we can use our analysis to provide *some* evidence on the strength of outbidding *in this specific case* that neither requires indirect proxies for competition nor assumes away heterogeneous effects. The first piece of evidence is that outbidding implies restrictions on our model's parameters—for example, groups should value popularity. We do not impose these restrictions, and our estimates satisfy these restrictions in our analysis and robustness checks. The second is that we compare our model to a version where competition does not arise because either the groups do not care about popularity or attacks do not affect popularity. We reject this no-competition model. The third is that we compare our outbidding model to an alternative tit-for-tat model, which we fit to the same attack data. Using a non-nested model fit test, we find that the outbidding model better fits the data. To be clear, we are not claiming that outbidding is the best explanation of, or explains all of, the observed terrorism.[18] Instead, the exercise demonstrates that, given our data, outbidding theory cannot be easily dismissed, and that the theory can be compared explicitly to others when scholars use the structural approach.

This discussion highlights an important implication for the conflict literature more broadly. Our results imply that reduced-form correlations between proxies for competition and violence, like those reported in time-series cross-section regressions, cannot falsify outbidding because it is consistent with a positive, negative, or null relationship between competition and violence. Moreover, these correlations risk hiding evidence of outbidding because the encouragement and discouragement effects run in opposite

14. Polo and Welsh 2024.

15. Ibid., 4.

16. Jaeger et al. 2015 find similar results using different polling data from the conflict. Outside of the Israeli–Palestinian conflict, Polo and González 2020 find indirect evidence by examining the relationship between terrorism and whether or not violence occurs along in-group/out-group cleavages.

17. See Table D.4 in the online appendix.

18. Indeed, we find some Hamas attacks in the mid-1990s that occur even when our model predicts low Hamas attack probabilities. These attacks were attributed to spoiling motives by Kydd and Walter 2002.

directions. Although the contest literature has theoretically characterized the conditions under which discouragement effects appear,[19] it is unclear whether encouragement or discouragement effects would dominate in any given case or how conflict scholars would know. In investigating these questions, we use the structural approach. While this approach is not necessary, a close connection between theory and data will be needed in future work to determine the extent that outbidding explains the relationship between competition and violence outside of our case.

Finally, we provide a general methodological approach to studying the effects of competition in dynamic contests in and outside of international relations. In intrastate conflict, outbidding also appears among separatist groups in Northern Ireland and militant leftists in Colombia, which are cases with straightforward applications of our methods. In the interstate setting, arms races can be cast as a country using military investments to favorably adjust its security environment vis-à-vis a rival.[20] With time-series data on countries' decisions to acquire arms and on the evolution of military power, scholars can estimate an identical dynamic contest and use similar counterfactuals to quantify the substantive effects of competition on the balance of power. Trade wars and major-power competition for influence and protégés can also be conceptualized as a tug of war. A growing political-economy literature estimates contest-like models, but these are either one-shot games[21] or include only one long-term player.[22] Thus, our paper helps scholars study empirical contests in a wider array of scenarios.

Model

Hamas (H) and Fatah (F) compete over a infinite number of periods indexed by $t \in \mathbb{N}$. In our data, a period corresponds to a calendar month. Period t's interaction explicitly depends on a publicly observed state variable $s^t \in \mathcal{S}$ that represents the popularity of Fatah relative to Hamas among Palestinians.[23] The set of states $\mathcal{S} = \{s_1, s_2, \ldots, s_K\} \subseteq \mathbb{R}$ comprises $K \geq 3$ equally spaced popularity levels, where $k > k'$ if and only if $s_k > s_{k'}$. We say Fatah is relatively more popular in state s than in state s' if $s > s'$, and vice versa for Hamas. In other words, smaller (larger) states represent periods when Hamas (Fatah) is more relatively popular.

In each period t, Hamas and Fatah choose whether to commit a terrorist attack ($a_i^t = 1$) or not ($a_i^t = 0$), where $i = \mathrm{H}, \mathrm{F}$ indexes the group.[24] Given an action

19. Kirkegaard 2012; Konrad and Kovenock 2005; Stein 2002.
20. Fearon 2011; Powell 1993.
21. Kang 2016; Kenkel and Ramsay 2024; Köning et al. 2017.
22. Iaryczower, López-Moctezuma, and Meirowitz 2024.
23. We focus on relative popularity because outbidding theories implicitly assume that benefits are "primarily relative or positional—i.e., the value of the resources gained depends on how much of that resource the group's competitors possess." Gibilisco, Kenkel, and Rueda 2022, 9.
24. We model actions as binary for two reasons. Theoretically, discrete-choice models have well-understood properties (Pesendorfer and Schmidt-Dengler 2008; Su and Judd 2012). Empirically, these groups rarely attack more than once month: Fatah attacks more than once (twice) a month in 2.7 percent

profile $a^t = (a^t_{\mathrm{H}}, a^t_{\mathrm{F}})$, per-period payoffs are $u_i(a^t_i, s^t; \theta) + \varepsilon^t_i(a^t_i)$. The term $\varepsilon^t_i \in \mathbb{R}^2$ is a vector of action-specific payoff shocks that are private information to group i, where $\varepsilon^t_i(a^t_i)$ refers to the $(a^t_i + 1)$th element of the vector ε^t_i. The shock $\varepsilon^t_i(a_i)$ is an independent and identically distributed (IID) draw from a standard type-one extreme value (T1EV) distribution.[25] The shocks account for unobserved factors temporarily affecting the costs and benefits of terrorism and ensure that choices within each period are stochastic.

The term $u_i(a^t_i, s^t; \theta)$ is the systematic component of group i's per-period payoff and consists of popularity benefits and attack costs:

$$u_i(a^t_i, s^t; \theta) = \underbrace{\beta_i s^t}_{\text{popularity benefit}} + \underbrace{\kappa_i a^t_i}_{\text{attack cost}}, \tag{1}$$

where $\theta = (\beta_{\mathrm{H}}, \beta_{\mathrm{F}}, \kappa_{\mathrm{H}}, \kappa_{\mathrm{F}})$. Because $\beta_i s^t$ captures i's benefit from relative popularity level s^t, we expect $\beta_{\mathrm{H}} < 0$ and $\beta_{\mathrm{F}} > 0$—that is, groups want more favorable support. This is one incentive for groups to compete. Likewise, κ_i denotes i's *cost* of attacking, which is another competitive incentive, and we expect $\kappa_i < 0$. Note that these inequalities are theoretical expectations from the outbidding literature. We do not impose them as *a priori* restrictions, but we explicitly test them after estimating the unobserved competitive incentives.

In contrast, outbidding theories do not offer explicit expectations about the relative magnitudes of β_i and κ_i across actors. It could be that Hamas cares more about relative popularity than Fatah ($|\beta_{\mathrm{H}}| > |\beta_{\mathrm{F}}|$) because Fatah has outside support from Israel and the United States, so it might care less about local Palestinian support. A similar argument suggests the opposite, however, because Hamas has outside support from Iran, Syria, and Qatar during our time frame. While intuition suggests that Hamas has lower attack costs given the differences in the groups' use of violence, outbidding theories do not have explicit predictions about relative attack costs. The model accommodates all possibilities and allows us to quantify the differences post-estimation.

The sequence of the game in period t is as follows.

1. Group i observes s^t and ε^t_i.
2. Groups simultaneously choose whether to attack: $a^t_i \in \{0, 1\}$.[26]
3. Payoffs are accrued.

(0.7%) of observations, and Hamas more than once (twice) a month in 26 percent (16%) of observations—see Figure A.1 in Appendix A.

25. This assumption induces easy-to-use logit choice probabilities over actions and is a common simplifying assumption in structural models. Crisman-Cox and Gibilisco 2018; Frey, López-Moctezuma, and Montero 2023; Rust 1994.

26. Simultaneous choice is a standard assumption in the contest literature and a useful simplification. To estimate a sequential model we would need to specify a particular group to move first. We cannot infer such an ordering from the observed data, however, because the group that attacks first may be different from the group that had the first opportunity to attack.

https://doi.org/10.1017/S0020818324000390

4. Transition to period $t+1$.

As the game transitions from period t to $t+1$, popularity evolves according to an AR-1 process with a mean that depends on the chosen actions and state. Given today's support and attack decisions (a^t, s^t), we define the mean of tomorrow's support s^{t+1} as

$$\mu[a^t, s^t; \gamma] = \gamma_0 + \gamma_1 s^t + \sum_i (\gamma_{i,1} + \gamma_{i,2} s^t) a_i^t. \tag{2}$$

The term $(\gamma_{i,1} + \gamma_{i,2} s^t)$ represents group i's ability to use terrorist attacks to increase its support—what we call i's *effectiveness* of attacks, which is the third competitive incentive in the model. Outbidding theories expect $\gamma_{H,1} < 0$ and $\gamma_{F,1} > 0$; that is, attacks by group i pull popular support in i's preferred direction. These inequalities are theoretical expectations but are not imposed in estimation. As with the payoff parameters, outbidding does not have explicit expectations about the relative magnitudes of $\gamma_{F,1}$ and $\gamma_{H,1}$ (that is, about which group is more effective at using terrorism), but the model accommodates either possibility.

Note that Equation 2 allows the effects of i's attacks $(\gamma_{i,1} + \gamma_{i,2} s^t)$ to depend on the current popularity level, s^t. *A priori*, it is not clear whether group i's attacks should be more or less effective as its popularity increases. On the one hand, higher popularity could make its attacks more effective due to support from the local population, implying that $\gamma_{i,2} > 0$; but on the other hand there is less of the population to be won over, implying that $\gamma_{i,2} > 0$.

In period $t+1$, the probability that $s^{t+1} = s'$ given action profile a^t and state s^t is $f(s'; a^t, s^t, \gamma)$. We specify f using a discretized normal distribution:

$$f(s'; a^t, s^t, \gamma) = \begin{cases} \Phi\left(\dfrac{s' + d - \mu[a^t, s^t; \gamma]}{\sigma}\right) - \Phi\left(\dfrac{s' - d - \mu[a^t, s^t; \gamma]}{\sigma}\right) & s' \in \{s_2, \ldots, s_{K-1}\} \\ \Phi\left(\dfrac{s_1 + d - \mu[a^t, s^t; \gamma]}{\sigma}\right) & s' = s_1 \\ 1 - \Phi\left(\dfrac{s_K - d - \mu[a^t, s^t; \gamma]}{\sigma}\right) & s' = s_K \end{cases} \tag{3}$$

where Φ is the standard normal cumulative distribution function, σ is the standard deviation parameter, and $2d = s_2 - s_1$ is the distance between the relative popularity levels. The parameters $\gamma = (\gamma_0, \gamma_1, \gamma_{H,1}, \gamma_{H,2}, \gamma_{F,1}, \gamma_{F,2}, \sigma)$ describe the transitions of the game, and we estimate them later. We choose this specification because γ can be estimated using standard techniques for continuous AR-1 models even though popularity levels are discrete.[27]

27. Tauchen 1986.

Equilibria

Given a sequence of states, actions, and payoff shocks $\{s^t, a_i^t, \varepsilon_i^t\}_{t=1}^{\infty}$, group i's total payoffs are $\sum_{t=1}^{\infty} \delta^{t-1}[u_i(a_i^t, s^t) + \varepsilon_i^t(a_i^t)]$, where $\delta \in [0, 1)$ is a fixed, common discount factor. Discount factors are difficult to identify in dynamic discrete-choice models.[28] Following Rust and others,[29] we estimate the model at several discount factors and fix the discount factor at $\delta = 0.999$, which gives the highest log-likelihood.[30] This matches anecdotal descriptions of the groups that highlight their long time horizons.[31]

Markov equilibria in discrete dynamic games with per-period private-information payoff shocks have a straightforward characterization.[32] Dropping references to time, let $v_i(a_i, s)$ denote i's net-of-shock expected utility from choosing action a_i in state s and continuing to play the game. The vector $v_i = (v_i(a_i, s))_{(a_i,s)\in\{0,1\}\times\mathcal{S}}$ collects these values for each (a_i, s) pair. Given a vector of expected utility values v_i and a vector of random shocks $\varepsilon_i = (\varepsilon_i(0), \varepsilon_i(1))$, group i chooses action a_i in state s if and only if

$$a_i = \underset{a_i\in\{0,1\}}{\operatorname{argmax}} \; \{v_i(a_i, s) + \varepsilon_i(a_i)\}.$$

Thus, v_i implicitly specifies a cutoff strategy for i, where i chooses to attack in state s if and only if $v_i(1, s) - v_i(0, s) > \varepsilon_i(0) - \varepsilon_i(1)$, where we sidestep the zero-probability event that i is indifferent. Because $\varepsilon_i(0)$ and $\varepsilon_i(1)$ are IID draws from a standard T1EV distribution, i chooses a_i in state s with probability

$$P(a_i, s; v_i) = \frac{\exp\{v_i(a_i, s)\}}{\exp\{v_i(0, s)\} + \exp\{v_i(1, s)\}}. \tag{4}$$

Let g denote the joint density of the action-specific payoff shocks, ε_i. Group i's continuation value for state s is

$$\begin{aligned} V_i(s, v_i) &= \int \max_{a_i\in\{0,1\}} \{v_i(a_i, s) + \varepsilon_i(a_i)\} g(\varepsilon_i) d\varepsilon_i \\ &= \log(\exp\{v_i(0, s)\} + \exp\{v_i(1, s)\}) + C, \end{aligned}$$

where C is Euler's constant. The second equality in Equation 5 follows from McFadden because g is the joint density of two IID standard T1EV random

28. Abbring and Daljord 2020.
29. Rust 1994; Frey, López-Moctezuma, and Montero 2023.
30. Our results are robust for $\delta \geq 0.975$. See Appendix I.
31. A reporter describes it as follows: "It's sometimes shocking to sort of hear what their timeline is. And they'll say … that justice is on our side and that we're doing the right thing. And if we're not able to do it, maybe our children will do it or maybe our grandchildren will do it. But they have this very long-term view of where this is going." ("Why Hamas Keeps Fighting and Losing," *New York Times*, 26 May 2021, available at <https://www.nytimes.com/2021/05/26/podcasts/the-daily/gaza-hamas-israel-war.html>.)
32. Pesendorfer and Schmidt-Dengler 2008, in Theorem 1, prove existence of Markov equilibria in a class of games that subsumes our game.

https://doi.org/10.1017/S0020818324000390

variables.[33] Consider a profile $v = (v_i, v_j)$ of action-state expected utility values. Group i's iterative net-of-shock expected utility of action a_i in state s, denoted $\mathcal{V}_i(a_i, s, v; \theta, \gamma)$, is

$$\mathcal{V}_i(a_i, s, v; \theta, \gamma) = \underbrace{u_i(a_i, s; \theta)}_{\substack{i\text{'s payoff} \\ \text{today}}} + \delta \underbrace{\left[\sum_{a_j} P(a_j, s; v_j) \underbrace{\sum_{s' \in \mathcal{S}} f(s'; a_i, a_j, s, \gamma) V_i(s', v_i)}_{\substack{i\text{'s } \textit{expected continuation} \\ \text{value given } a_j}} \right]}_{\text{iterated expectation over } j\text{'s } \textit{action}}. \tag{6}$$

An equilibrium is a profile v that satisfies the fixed-point condition

$$v = \mathcal{V}(v; \theta, \gamma) \equiv \times_i \times_{(a_i, s)} \mathcal{V}_i(a_i, s, v; \theta, \gamma). \tag{7}$$

Equations 4 to 7 characterize equilibria as a system of $4K$ equations, where K is the number of relative popularity levels. In words, starting with i's net-of-shock, action-specific expected utilities, Equations 4 and 5 return i's choice probabilities and continuation values, respectively. Then Equation 6 updates i's net-of-shock action-specific expected utilities, holding fixed i's continuation values and j's choice probabilities. An equilibrium is a fixed point in Equation 7. In Appendix B, we consider a symmetric example, use Equation 7 to compute equilibria, and then study their substantive properties and comparative statics.

Remarks

Three remarks on the model. First, because this is a model of outbidding, it explains variation in violence via intergroup competition. Such a spartan approach is critical for our argument: outbidding produces heterogeneous relationships between competition and violence, and one such relationship is the discouragement effect. Furthermore, this discouragement effect appears in the canonical case of outbidding, and other forces are not necessary to generate discouragement effects. Adding more moving parts to the analysis—while potentially interesting in future work—only obfuscates this central result. Thus, the model does not include other motives for terrorism considered by Kydd and Walter.[34]

We therefore prioritize matching the model to outbidding theories and abstract away from many other features that appear in the broader literature. Overwhelmingly, outbidding theories focus on terrorist groups' use of violence.[35] Conrad and Spaniel also use a contest to build a theory of outbidding.[36] Groups

33. McFadden 1978, 82.
34. Kydd and Walter 2006.
35. Conrad, Greene, and Phillips 2024; Conrad and Spaniel 2021; Polo and Welsh 2024.
36. Besides the structural approach, our departure from Conrad and Spaniel 2021 is that we consider a dynamic and asymmetric contest.

may use other tools, like public goods provision, to boost their support, but studies considering how groups provide public goods focus on their competition with the government, not rival groups.[37] Likewise, a key assumption in most outbidding theories is that the government is tangential.[38] Nonetheless, in Appendix D we discuss whether and how unobserved government actions can affect our parameter estimates and explicitly control for government actions when considering the robustness of our approach. Likewise, because this is a model of outbidding, it abstracts away from some of the specifics of the Fatah-Hamas rivalry, such as principal–agent problems within the groups.

The model is not inconsistent with other features affecting terrorism and public support. It includes exogenous shocks to the costs of attacking, and relative popularity evolves stochastically, capturing outside forces. We put assumptions on these features, but they mimic standard assumptions in reduced-form work (for example, ε_i^t is drawn IID). As discussed by Canen and Ramsay, all quantitative empirical work requires models and assumptions to make causal claims, and this paper is no exception.[39] In subsequent sections, we examine model fit and the robustness of our inferences to different modeling assumptions. For the former, we test outbidding's theoretical expectations concerning groups' competitive incentives and explicitly compare the outbidding model to competing models. For the latter, numerous robustness checks show that our estimates and predictions are insensitive to excluding or including specific time frames where there are known shifts in the conflict. Our estimates of the groups' payoffs are unaffected by uncertainty in the estimated transition function f and by small changes in the discount factor. As mentioned, a key feature of our estimates is that Fatah is more effective than Hamas at using violence to increase public support. Appendix D is dedicated to demonstrating the robustness of this result to different terrorism measures, control variables, and (unobserved) omitted variables.

Second, we do not explicitly model the decisions of individuals who choose a group to support. This is a simplifying assumption that also appears in Conrad and Spaniel and structural models of dynamic elections.[40] Instead, individuals and their choices are captured by Equations 2 and 3, which describe how relative support evolves given the attack decisions of the two groups and their current popularity. Rather than microfounding this behavior, we calibrate it to data by estimating the relevant parameters of interest, γ. This allows us to sidestep additional assumptions detailing the preferences and decisions of local individuals, which could be quite complicated.[41] We do not assume that they want violence or that terrorism must increase support; indeed, the model accommodates the possibility that terrorism decreases support. When we estimate the $\gamma_{i,1}$ and $\gamma_{i,2}$ parameters, however, we find

37. Berman and Laitin 2008; Stewart 2018; Heger and Jung 2017; Wagstaff and Jung 2020.
38. Kydd and Walter 2006.
39. Canen and Ramsay 2024.
40. Conrad and Spaniel 2021; Iaryczower, López-Moctezuma, and Meirowitz 2024.
41. Ze'evi 2008.

that violence does increase relative support for the two groups (that is, $\gamma_{i,1} + \gamma_{i,2}s^t$ is estimated to be positive for Fatah and negative for Hamas at the observed relative popularity levels s^t). This result is consistent with previous empirical work,[42] Bloom's case study of the conflict,[43] and the theoretical mechanism underlying outbidding.[44]

Third, the model treats competitive incentives as exogenous actor-specific parameters. This matches other contest models, and we explicitly borrow their terminology of "value," "cost," and "effectiveness." On the one hand, this approach allows us to flexibly accommodate time-invariant, group-level heterogeneity. That is, these parameters can be composed of time-invariant, group-specific features, and we avoid additional functional-form assumptions by treating these as fixed effects. This is a strength because we anticipate systematic differences between Hamas and Fatah, such as different corruption levels, relationships with Israel, and capacities for violence.[45]

On the other hand, several interesting questions arise about the origins of these incentives—for example, why does one group have lower costs than another? We choose to prioritize a flexible model with actor-specific parameters rather than endogenizing the parameters because the contest literature anticipates that the discouragement effect appears when groups have asymmetric incentives.[46] Also, our data cover only two groups, so we cannot leverage cross-sectional variation to determine the covariates of the competitive incentives. In Appendices D and H, we explore whether these incentives change across time periods, but we find little variation. If future work fits the model to several different conflict environments, then we can compare how these incentives vary in a post-estimation exercise—in a manner similar to how Crisman-Cox and Gibilisco study the correlates of their estimated audience cost parameters.[47]

Data Sources and Measurement

Terrorism data are from the Global Terrorism Database (GTD), where we record terrorist attacks by Fatah/PLO and Hamas from January 1994 to December 2018.[48] The GTD records both suicide bombings, which are the focus of Bloom and Findley and Young, and other types of terrorism, such as rocket attacks, which are a greater part

42. Jaeger et al. 2015; Polo and González 2020.
43. Bloom 2004, chap. 2.
44. Kydd and Walter 2006.
45. Stewart 2018 and Tokdemir and Akcinaroglu 2016 record group provision of public goods and potentially find differences by group (depending on the specific measure) but not over time. Our model would also accommodate group-level heterogeneity in the provision of public goods.
46. Kirkegaard 2012; Siegel 2014.
47. Crisman-Cox and Gibilisco 2018.
48. We use data from Acosta and Ramos 2017 for December 1993, which is missing from the GTD. In Appendix H, we re-estimate the model using different time frames; our results are stable across subsamples.

of violence against Israelis in recent years.[49] Hamas engages in an average of roughly 1.5 attacks per month, while Fatah engages in an average of less than one per month (see Figure A.1).[50] To measure group *i*'s attack decision in month *t*, we use a dummy variable for whether the group committed any terrorist attacks in that month.[51]

The model's state variable represents the relative popularity of the two groups among Palestinians. To measure it, we treat relative popularity as a dynamic latent variable and use observed public-opinion variables as its indicators. To assemble the set of indicators, we use surveys from the Jerusalem Media and Communication Centre (JMCC) and the Palestinian Center for Policy and Survey Research (PCPSR).[52] The JMCC publishes two to six surveys per year using random samples of Palestinian adults. It conducts face-to-face interviews in randomly selected households from randomly selected neighborhoods throughout the West Bank and Gaza Strip; within each home, the subjects inside are selected using Kish tables. Each survey typically occurs over a few days but less than a week. The average sample size is 1,205, with a range between 815 and 1,920.[53] On average, 63 percent of respondents are from the West Bank. Given their rich data on Palestinian attitudes, these surveys also appear in other studies.[54]

The PCPSR runs two to nine surveys per year. It generally uses a multi-step selection process, randomly sampling from locations in proportion to the population from a list of all cities, towns, villages, and refugee camps in the West Bank and Gaza Strip. Once locations are selected, it samples individual blocks and then individual households. Each survey typically occurs over several days but less than a week. Sample sizes vary between 1,076 and 2,006, with a mean of 1,312.[55] West Bank respondents tend to make up 60 to 67 percent (mean, 63 percent) of the overall sample, with Gaza respondents making up the rest.

We search through every survey published by these centers between 1994 and 2018 to track Palestinian public opinion on both actors using three questions. The first (from the JMCC) asks which political or religious group the respondent trusts most. The second (from the PCPSR) asks which political party they support. The third (from the JMCC) asks which party they intend to vote for in the legislative elections. For each of these three questions we track the proportion of respondents who

49. Bloom 2004; Findley and Young 2012; Getmansky and Zeitzoff 2014.

50. Fatah's last attack was in 2009, but our estimates and model's predictions are insensitive to excluding data from later in the time frame (Appendix H).

51. First-stage results are not dependent on using either fatalities or fatalities/attack as the main measure of interest (Appendix D.2). Attacks do not appear to get more or less deadly over time (Appendix D.5).

52. Jerusalem Media and Communication Centre, N.d; Palestinian Center for Policy and Survey Research, N.d (known as the Center for Palestine Research and Studies until July 2000).

53. In most (but not all) of its polls, the JMCC breaks down answers geographically. In the West Bank, the average sample size is 764, with a range of 518 to 1,246. In the Gaza Strip, the average is 441, with a range of 297 to 674.

54. Clauset et al. 2010; Jaeger et al. 2012.

55. In the West Bank and Gaza Strip the ranges are 664–1,311 (mean of 857) and 390–695 (mean of 497), respectively. The surveys continue to report the results by region but stop reporting the regional sample sizes in 2009.

answer Hamas or Fatah.[56] All three questions are open-ended. Appendix C gives details on question wording and variation by geography.

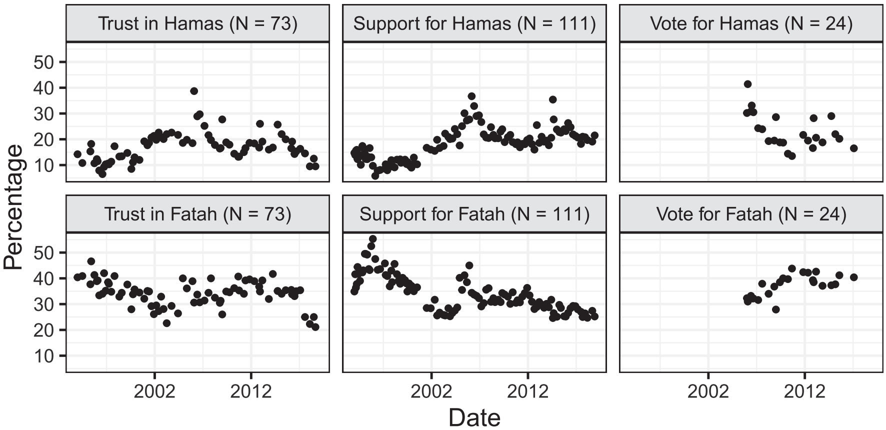

Notes: Left column: JMCC question, "Which political or religious faction do you trust the most?" Middle column: PCPSR question, "Which of the following political parties do you support?" Right column: JMCC question, "If Legislative Council elections were held today, which party would you vote for?" N is the number of months a question was asked.

FIGURE 1. *Survey responses over time*

Figure 1 graphs responses to these six survey questions over time. Generally, public attitudes toward the two groups are inversely related. During the 1990s and early 2000s, support for Fatah falls and support for Hamas rises. These trends level out a bit in later years, with Fatah regaining some support at the expense of Hamas. The surveys mostly correlate with each other in the expected directions (Table C.2), which suggests that they can be collapsed onto one dimension. To do this, we use a dynamic factor model that transforms these polling questions into a continuous representation $\tilde{s}^t$ of the theoretical state variable s^t. See Appendix C for details.

Having constructed the continuous state variable $\tilde{s}^t$, we assess its validity. All indicators load onto the factor in the expected directions (Table C.3). Figure 2 shows how the state variable evolves from 1994 to 2018. Fatah is favored in earlier periods; its relative popularity peaks during the 1996 Oslo II process (about 12.5 in January 1996). Hamas is at its most popular relative to Fatah in

56. We also examine the total percentage of people saying they trust/support either group and how this varies over time. Regressing these totals on time, observed terrorism, and their interactions, we find that fitted values range from 51.2 to 51.4 percent for trust and 50 to 55 percent for support over our sample period. The expected level of trust/support for these two actors is fairly stable over time, with the estimated conditional mean shifting by only a few percentage points. The notable, but one-off, exception is at the end of the Second Intifada, when there is a surge in Hamas support.

2006, in the aftermath of the general election in which Hamas took control of Gaza (about –10.9 in August 2006). The mean of this variable is −0.87; the median is about −3, and the standard deviation is 6.59 (with an interquartile range of −6.02 to 4.13). The continuous state variable is easily mapped back onto the original surveys, such that, on average, a one-unit increase in $\tilde{s}^t$ roughly corresponds to increases of 0.9, 1.5, and 2 percentage points in net trust, support, and intention to vote for Fatah over Hamas, respectively.

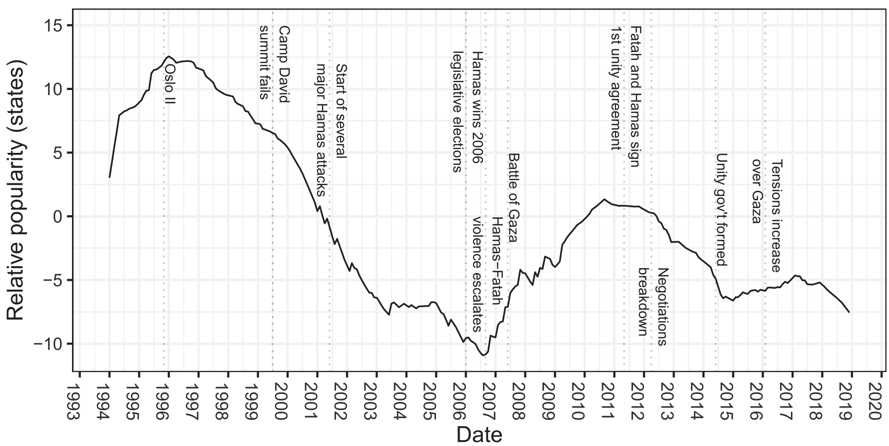

FIGURE 2. *Popularity of Fatah relative to Hamas over time*

Several important events are flagged in Figure 2, providing context and face validity to the idea that this variable captures the relative ups and downs of the two groups. The late 1990s are typically regarded as an important inflection point for these two groups, and that is clearly reflected here. Fatah sees its popular support erode as the peace process unravels. Furthermore, our measure has rich variation, with substantial ups and downs that went undetected in earlier measures of these groups' popularity.[57] In Appendix C we demonstrate that our latent measure of relative popularity is robust to different model specifications. The estimated state variables are closely correlated (0.87–0.99) across specifications.

Estimation and Identification

We use a two-step procedure, where we first estimate how relative support evolves (γ) and then estimate the groups' payoff parameters (β, κ).[58] To do this, first rewrite the

57. For example, Tokdemir and Akcinaroglu 2016 do not find differences in the popularity of Fatah versus Hamas after 1997.
58. As in Rust 1994.

https://doi.org/10.1017/S0020818324000390

AR-1 model in Equation 2 in terms of the continuous state variable $\tilde{s}^t$:

$$\tilde{s}^t = \gamma_0 + \gamma_1 \tilde{s}^{t-1} + \gamma_{H,1} a_H^{t-1} + \gamma_{H,2}(\tilde{s}^{t-1} a_H^{t-1}) + \gamma_{F,1} a_F^{t-1} + \gamma_{F,2}(\tilde{s}^{t-1} a_F^{t-1}) + \nu^t, \quad (8)$$

where a_F^{t-1} and a_H^{t-1} are binary indicators for whether Fatah and Hamas attack, respectively, and $\nu^t \sim N(0, \sigma^2)$.[59]

The first-step estimates are then used to construct the Markov transition probabilities, f. To discretize the continuous state $\tilde{s}^t$, we define the lowest and highest (most Hamas-friendly and most Fatah-friendly) states as the bottom and top 2.5th percentiles of $\tilde{s}^t$. Discrete states between these extremes are defined at equally spaced intervals with distance $2d = 0.05$. In the baseline model, $K = 440$. We then map the continuous measure $\tilde{s}^t$ into the discrete measure s^t by finding the closest discrete state.[60] Let $\mu[a, s; \hat{\gamma}]$ be the fitted values from the first model (reported in Table 1) for all possible combinations of action profiles with the discrete states. Plugging these fitted values and the estimated standard deviation $\hat{\sigma}$ into Equation 3 produces the transition probabilities.

We use a constrained maximum likelihood estimator (CMLE) to estimate the payoff parameters $\theta = (\beta, \kappa)$.[61] Specifically, let $Y = (s^t, a_H^t, a_F^t)_{t=1}^T$ denote the time series of observed data (relative popularity levels and attacks). We fix the transition probabilities using the first-step estimates, $\hat{\gamma}$, and the definition of f in Equation 3. The CMLE estimates $(\hat{\theta}, \hat{v})$ maximize the log-likelihood

$$L(v|Y) = \sum_{t=1}^{T} [\log P(a_H^t; s^t, v_H) + \log P(a_F^t; s^t, v_F)]$$

subject to the equilibrium constraint equations $v = \mathcal{V}(v; \theta, \hat{\gamma})$. For standard errors, we follow Silvey in using the bordered Hessian to compute the variance-covariance matrix and use a two-step correction described in Appendix F.[62]

The game can have multiple equilibria. The CMLE allows for this multiplicity, with its main identification assumption being that the data Y are generated from only one of these equilibria.[63] By treating the endogenous equilibrium quantities, v, as auxiliary parameters, the CMLE selects the values of v that best describe the data while still being an equilibrium of the model. In other words, the CMLE imposes an empirical selection rule: choose the equilibrium associated with the highest log-likelihood. This process is a computationally feasible alternative to an approach that computes *all* equilibria at *every* optimization step and then always chooses the equilibrium that maximizes the log-likelihood at that optimization

59. Unit root tests suggest that the state variable $\tilde{s}_t$ is not stationary. However, because $\tilde{s}^t$ and $\tilde{s}^{t-1}$ are cointegrated, OLS will produce superconsistent estimates. We also fit the model using the Engle-Granger error correction model (ECM) of for hypothesis testing.

60. Our estimates are robust to changes in the discretization process (Appendix J).

61. Crisman-Cox and Gibilisco 2018; Su and Judd 2012.

62. Silvey 1959.

63. Crisman-Cox and Gibilisco 2018; Su and Judd 2012.

https://doi.org/10.1017/S0020818324000390

step.[64] The CMLE imposes this same empirical selection rule, but without the infeasible requirement of repeatedly enumerating all equilibria.

TABLE 1. *Regressing relative popularity (state variable) on terrorist attacks*

	Dependent variable	
	State AR(1)	Δ state ECM
Hamas attack, $\gamma_{H,1}$	−0.21	−0.21
		(0.04)
Fatah attacks, $\gamma_{F,1}$	1.12	1.04
		(0.05)
lag state, γ_1	1.00	
Δ lag state		0.33
		(0.04)
Hamas attacks × lag state, $\gamma_{H,2}$	0.01	0.002
		(0.01)
Fatah attacks × lag state, $\gamma_{F,2}$	0.03	0.01
		(0.01)
constant, γ_0	−0.02	−0.01
		(0.02)
T	299	298
Adj. R^2	0.999	0.721
σ	0.216	0.183

Notes: Newey-West standard errors in parentheses. No standard errors are reported for the AR(1) model due to unit root. ECM refers to the Engle-Granger error correction model.

Along with the assumption that one equilibrium is generating the data, three empirical moments pin down our parameters of interest. We estimate γ through observed variation in the state variable over time. We know that each action profile has a positive probability of being played at each relative popularity level given the distributional assumptions on ε_i^t, and that the probability of transitioning from level s to level s' is positive for all s and s'. Thus f can be estimated nonparametrically from frequency estimators with a sufficiently long time frame, because the equilibrium path will eventually visit all states, and all action profiles will be played in every state. When the transition probabilities are known, the payoff parameters are identified by their relationship to the equilibrium constraint $\mathcal{V}$ in Equation 6. A group's attack cost is identified through its baseline propensity to attack regardless of the state, and its value of public support is identified by the variation in its propensity to attack across states.

Formal identification of the payoff parameters θ follows from Propositions 2 and 3 in Pesendorfer and Schmidt-Dengler.[65] The former is a necessary condition stating that up to K payoff parameters per actor can be identified. We seek to estimate

64. Su and Judd 2012, Proposition 1.
65. Pesendorfer and Schmidt-Dengler 2008.

two parameters per group using $K = 440$ states, which satisfies the necessary condition. The connection between K and identification raises a question about how sensitive the estimates are to discretizing relative popularity; in Appendix J we show that our estimates are robust to both small and major changes in this process. The latter is a more involved sufficient condition for identifying θ that depends on the equilibrium choice probabilities, which we can verify given our estimated equilibrium (Appendix E).

Although the groups' competitive incentives can be identified given data generated from an equilibrium of the game, another reasonable concern is how sensitive the estimated incentives are to forces outside the model, in particular to interventions from Israel. Here, we anticipate that the importance of Israeli actions depends on whether the incentive is group effectiveness or directly enters into the groups' payoff functions. For the former, we can compare our baseline estimates of $\gamma_{i,1}$ to those in robustness exercises where we control for Israeli actions or their proxies (such as the number of Palestinian fatalities or the time since the last Israeli election) or instrument group attacks with rainfall. Our estimates of groups' attack effectiveness are stable across specifications (Appendix D).

For the latter, the analysis is murkier because we are unable to conduct such robustness exercises. If we wanted to include Israeli interventions when estimating the groups' value of support and cost of attacking, then we would need monthly data on Israeli actions against the individual groups during our time frame. Furthermore, we would need to either estimate how these actions evolve according to relative popularity levels and attack decisions, or explicitly model the Israeli government as a third strategic actor. Given the scarcity of high-frequency data on how Israel responds to individual groups, and given that outbidding theories generally treat governments as tangential, we think that the appropriate first step is to structurally estimate an outbidding model without government interventions. Nonetheless, we anticipate that the groups' costs of attacking include both upfront costs (such as explosives) and strategic backlash from the Israeli government (such as border walls and air strikes). Also, if Israeli interventions are aimed at reducing the likelihood of attacks, then these interventions should target groups precisely when the tug-of-war predicts high attack probabilities. Thus, the observed probability of attacks would appear flatter as a function of relative popularity than in a world without interventions. This would attenuate our estimates of the groups' values of support, because these are identified by variation in changes in attack probabilities as a function of relative popularity.

Parameter Estimates

Table 1 shows the first-stage estimates. Attacks by Fatah and Hamas move the state space in the expected direction. Recall that estimates of $\gamma_{i,1}$ reflect each group's effectiveness at using terrorism to shift public support toward itself and away from its rival. In months when Hamas attacks, its relative popularity improves by an average of about 0.11 to 0.28 in the following month. When Fatah attacks, it can expect its relative popularity to improve by about 0.87 to 1.4 on average. As mentioned, the scale of $\tilde{s}^t$ can be roughly compared with the net level of trust in Fatah over Hamas, so on average, these magnitudes roughly reflect shifts in net levels of trust for Fatah over Hamas.[66] Both of these effects are statistically significant in the ECM model. These results provide evidence that groups are capable of outbidding and that acts of terrorism carry popularity benefits for the group committing them.[67]

We also find that Fatah's use of terrorism increases support for Fatah more effectively than Hamas's use of terrorism increases support for Hamas. Specifically, we reject the hypothesis that the groups are equally effective at moving public opinion ($H_0{:}\gamma_{\mathrm{H},1} + \gamma_{\mathrm{H},2} \cdot s + \gamma_{\mathrm{F},1} + \gamma_{\mathrm{H},2} \cdot s = 0$) at every level of relative popularity s using the estimates and Newey-West variance matrix from the ECM model.

One possible explanation is that attacks by Fatah provide more information to the public. This could be for a variety of reasons. For example, it is often seen as the more pro-peace actor[68] or, alternatively, as the more corrupt or possibly inept actor.[69] Attacks by Hamas are expected, so they do little to adjust public opinion. Attacks by Fatah, being more surprising, are more likely to shift the public's beliefs about how committed Fatah is to the Palestinian cause. Thus, even though attacks demonstrate the resolve of both groups, Fatah's attacks give it a larger boost in public opinion. This explanation is consistent with our parameter estimates, but it is of course a conjecture because it involves assumptions about the Palestinian population that we deliberately did not microfound. Future studies should consider the population side of the outbidding process to better explain these asymmetries in the response to terrorism.

In Appendix D, we show that these relationships are not driven by omitted economic or political factors, such as unemployment, Palestinian attitudes toward violence, the Second Intifada, Israeli election timing, or Palestinian fatalities from Israeli forces (which is one proxy for government actions). We also find no evidence that the groups are becoming more or less effective during our time frame (Table D.3). Overall, the relationships between attacks and shifts in public support are largely unchanged in either direction or magnitude across model specifications.

We also consider alternative measures of attacks. Whether we use attack counts, fatalities, or fatalities per attack, we find Fatah is more effective than Hamas (Table D.4). We also study plausibly exogenous variation in attacks driven by

66. These numbers can be multiplied by 1.5 or 2 to translate them into the average effect of terrorism on net support and net voting intention, respectively.

67. This is consistent with the results of Jaeger et al. 2015 and Polo and González 2020.

68. Kydd and Walter 2006.

69. Milton-Edwards and Farrell 2010 note that Fatah had a "reputation for sleaze and inefficiency" (238), and they quote Fatah campaign chief Nabil Saath as saying that Fatah "look[ed] like they [were] quarrelling and fighting over trivia" (252). While corruption perceptions may be part of this asymmetry, they do not fully explain it; it persists even when controlling for annual perceptions of corruption (see Appendix D.1).

extreme rainfall shocks in the Gaza Strip and the West Bank.[70] Our baseline estimates of $\gamma_{F,1}$ and $\gamma_{H,1}$ (in Table 1) are similar in direction and magnitude to those from an instrumental-variables analysis, although we do not want to over-interpret these results (see Appendix D.3 for details).

We also conduct formal sensitivity analyses for our estimates $\hat{\gamma}_{F,1}$ and $\hat{\gamma}_{H,1}$ (Appendix D.4). Here, we ask how strong omitted variables would have to be to make the asymmetry between these estimates disappear or to make the estimates null. We find that these unobserved effects would have to be implausibly large to explain away our findings.

Table 2 presents estimates for the two groups' value of popularity and cost of attack. The sign on each estimate is in the direction expected by outbidding theory and is statistically significant at conventional levels (one-sided tests). Both actors like being more popular than their opponent. It may be concerning that the β_i estimates are quite close to 0, but we reject the null hypothesis that both are 0. Furthermore, the estimates of β_i have strong impacts on the equilibrium attack probabilities despite their seemingly small magnitudes (see Appendix K). Interestingly, Hamas values its support more than Fatah, with $|\hat{\beta}_H|$ being an order of magnitude larger than $|\hat{\beta}_F|$. One possible explanation for this is that Fatah has more support from outside actors to consider than Hamas. While that would be consistent with our parameter estimates, our analysis cannot rule out other explanations.

Intuitively, we find that terrorism is less costly for Hamas than for Fatah, a finding which itself has several potential explanations. First, it could reflect different preferences between the two groups regarding violence. Second, Hamas has made a concerted effort to build its capacity for violence by developing infrastructure to acquire weapons and better train its members. Hence, the group would find it less costly to engage in violence than Fatah, which has devoted more resources to governance and engagement with the Israeli and US governments. Both explanations fit with the historical record, which typically depicts Hamas as a more extreme actor while Fatah is a more practical political entity.[71]

Beyond the face validity of the point estimates, we consider the robustness of the estimates in Table 2. In Appendix F, we conduct a sensitivity analysis to demonstrate that they are stable across a range of plausible first-stage estimates. In Appendix H, we consider several shorter time frames that represent potential starts, stops, or change points in the Hamas-Fatah rivalry—for example, ending the analysis with the 2011 coalition agreement, or starting in 1997.[72] In Appendix J, we demonstrate

70. Köning et al. 2017 pursue a similar approach in studying groups' use of violence in the Second Congo War.

71. Fatah officially renounced terrorism as part of its push to be recognized as a legitimate political actor, so attacks likely carry additional reputation costs for violating this pledge. Schanzer 2003 notes that this was a fundamental constraint on Fatah's ability to respond violently when Hamas's popularity was increasing during the Roadmap to Peace era.

72. This is the first year included in Bloom 2004.

robustness to how we discretize our measure of relative popularity; our results are stable even with a small number of states ($15 \leq K \leq 22$).

TABLE 2. *Payoff estimates*

	Estimates	Standard errors BH	Standard errors Two-step
Hamas value of popularity, β_H	−0.0071	0.0042	0.0056
Fatah value of popularity, β_F	0.0005	0.0003	0.0004
Hamas attack cost, κ_H	−0.95	0.23	0.28
Fatah attack cost, κ_F	−2.46	0.28	0.40
Log-likelihood	−278.20		
T	300		

Note: Bordered-Hessian (BH) standard errors and two-step corrected standard errors.

Figure 3 graphs the groups' estimated attack probabilities over time—that is, $P(a_i = 1; s^t, \hat{v}_i)$. We also graph the relative popularity level s^t over time on the second horizontal axis for reference. Notice that Hamas has a higher probability of attacking than Fatah regardless of its relative popularity. Averaging over the observed states, Hamas attacks with probability 0.42, and Fatah with probability 0.11. This maps onto our estimates. Hamas cares more about its popularity than Fatah, and it has a comparatively lower attack cost, although Fatah uses terrorism more effectively to increase its support. Also, terrorism is particularly prevalent when Hamas is relatively popular, specifically during the Second Intifada and after the group wins legislative elections in 2006.

Model Fit and Comparison

Before considering the substantive implications of the estimated outbidding model, we consider how well it describes the data, both on its own terms and in comparison to alternative theories. In this section, our goal is not to test a particular hypothesis but to demonstrate the validity and usefulness of the model in explaining variation in the observed terrorism data.

For the first exercise, recall that the estimated competitive incentives match the direction posited by outbidding theories: attacking is costly, groups value support, and attacks increase relative support. These restrictions were not imposed during estimation, and we would be skeptical of outbidding's ability to explain the data if they did not hold. For example, how could outbidding be a consistent theoretical explanation if groups wanted to become less popular? We can also examine the states in which the model predicts attacks well; Figure 4 does so visually. Ideally, we should see more attacks when the equilibrium choice probabilities are higher, all else equal. For the most part, this is true: observed attacks fall mostly when relative popularity s is between −7 and −3, where the equilibrium choice probabilities peak.

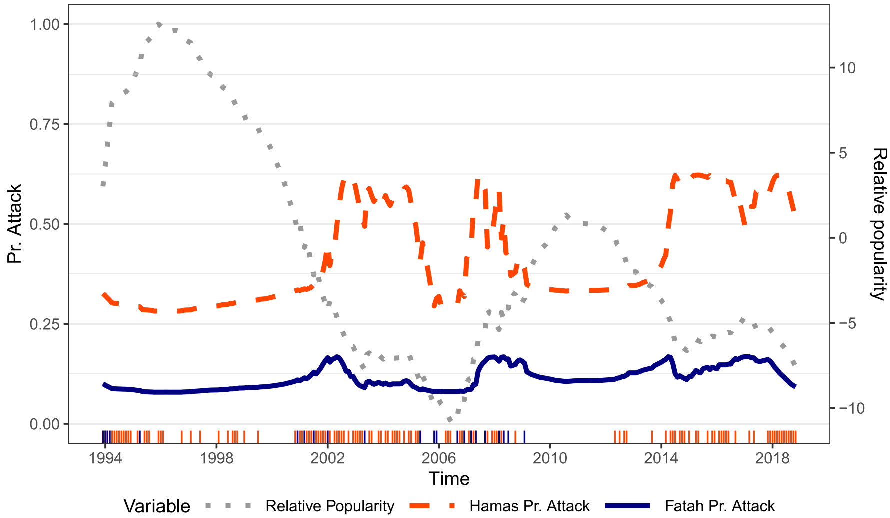

Notes: Horizontal axis denotes sample months/periods. Left vertical axis is the estimated probability that i attacks in month t—that is, $P(a_i = 1; s^t, \hat{v}_i)$, where s^t is the relative popularity level observed in period t and $\hat{v}_i$ is estimated from CMLE. For reference, s^t is also plotted on the right vertical axis. The rug plot indicates observed attacks.

FIGURE 3. *Estimated equilibrium probability of attacking over time*

Nonetheless, the data show a cluster of Hamas attacks around $s \in [8, 9]$, where attack probabilities are smaller. These attacks are difficult to attribute to outbidding as Hamas was close to its nadir of popularity, and the estimates suggest that, given the relatively few number of periods, this many attacks are unlikely. Thus, it seems reasonable to suspect that another theory of terrorism may better explain these attacks. Kydd and Walter argue that some of these attacks were part of an attempt by Hamas to undermine or "spoil" the Oslo process and drive a wedge between Fatah and Israeli negotiators.[73] If these attacks are indeed more associated with spoiling and less attributable to outbidding, then it is not surprising that they stand out in Figure 4. This analysis highlights an advantage of this structural approach: it allows us to identify observations that do not fit the theory's predictions.

It also raises a question: can other theories of terrorism better explain the data? It is beyond the scope of this paper, or any one paper, to consider and adjudicate among all theories of terrorism. Indeed, we think the field's understanding of the strategic forces behind terrorism will advance if scholars construct competing models of terrorism from different theories and estimate these models on the same data. This would let us compare the models and their associated theories regarding how well they explain observed terrorism. Given the lack of previous structural models of terrorism,

73. Kydd and Walter 2002.

we have no obvious prior model to compare ours to. Thus, we create some alternative structural models and acknowledge that, until more models are available, our comparison models are inherently ad hoc.

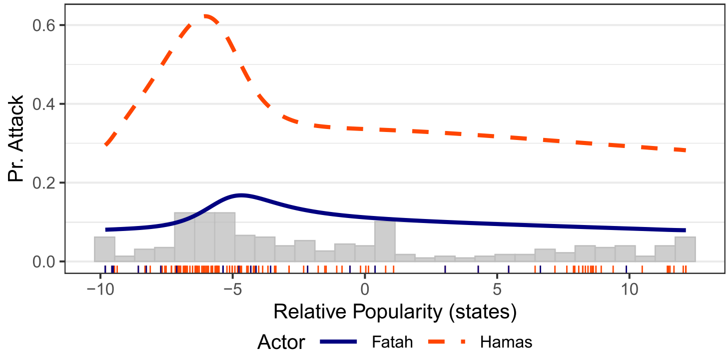

Notes: Estimated probability that each group attacks as a function of relative popularity. The horizontal axis includes a rug plot of observed attacks and a histogram of the observed states s^t; gray bars represent the density of the observed states.

FIGURE 4. *Estimated equilibrium attack probabilities as a function of state*

The first alternative model is a null model in which there is no competition between groups because they cannot (or do not care to) compete for popularity. We can nest such a model within our outbidding model by assuming $\gamma_{F,1} = \gamma_{F,2} = \gamma_{H,1} = \gamma_{H,2} = 0$. With this assumption, we cannot identify β. The only parameters left to fit the no-competition model are κ_F and κ_H—that is, groups are attacking without reference to relative popularity and only due to static incentives. Because this alternative model is nested, it can be compared to the main model using a standard likelihood ratio test. We reject the null hypothesis that the no-competition model fits as well as the main model (Table 3).

The second alternative is a non-nested model based on tit-for-tat retaliation.[74] For this model, when group i chooses to attack ($a_i^t = 1$) or not ($a_i^t = 0$) in each period, we make the following assumptions:

1. The new publicly observed state variable $r^t = (r_F^t, r_H^t) \in \{0, 1\} \times \{0, 1\}$ is a two-dimensional variable that represents whether each actor attacked in the previous period, with $r_i^t = 1$ meaning group i attacked in period $t - 1$.

74. We choose this model because it is (i) a dynamic model, so we can use similar tools to characterize equilibria and estimate its parameters; (ii) supported by news articles and scholarly work (Johannsen 2011; Brown 2012); and (iii) an alternative explanation for competition suggested by Michael Joseph, whom we thank for this suggestion. It has the same number of payoff parameters as the outbidding model.

2. The systematic utility function for group i is now

$$u_i(a_i^t, r^t; \tau, \kappa) = a_i^t(\underbrace{\kappa_i}_{\substack{\text{baseline}\\ \text{cost}}} + \underbrace{\tau_i r_{-1}^t}_{\substack{\text{retaliation}\\ \text{benefit}}}),$$

where κ_i is the baseline cost of attacking and τ_i is the additional benefit or cost a group receives when attacking in response to a previous attack by its rival. We collect these parameters into the vectors $\kappa = (\kappa_H, \kappa_F)$ and $\tau = (\tau_H, \tau_F)$.

TABLE 3. *Comparative model tests*

Alternative model	Test	Null distribution	Statistic	p value
No competition	Likelihood ratio	$\chi 2$ (6)	279.9	< 0.01
Tit for tat	Clarke's test	Binomial (300, 0.5)	182	< 0.01

As in the baseline model, we assume that groups' per-period payoffs depend on privately observed action-specific payoff shocks that are IID draws from the standard T1EV. Additionally, this tit-for-tat model is a discrete dynamic game, so we can use techniques almost identical to those used with the main model to characterize Markov equilibria, except with appropriate changes to the utility functions and the state transitions, which are now deterministic, as $r_i^t = a_i^{t-1}$. Moreover, we can use a CMLE to fit the model to the same GTD data to estimate κ and τ.[75] The goal is to compare this model with our outbidding model regarding how well they explain the attack data.[76]

The point estimates from the tit-for-tat model are all in the expected directions for a tit-for-tat theory: attacking is costly, but groups have an additional benefit if they attack in response to their rival (Appendix G.1). Comparing the tit-for-tat model against the outbidding model requires a non-nested model test. We use Clarke's test, which is a comparison of "pointwise" log-likelihood values.[77] The null hypothesis is that the two models are equally good; we reject this, finding that the outbidding model better fits the data (Table 3). Overall, we conclude that the outbidding model explains the data better than the tit-for-tat model.[78]

75. Because transitions are deterministic, we do not need to estimate how the state variable evolves. Results from Pesendorfer and Schmidt-Dengler 2008, 913, Eqs. 16–17, imply identification of the payoff parameters.

76. We also set the discount factor δ to 0.999 to match the main model, but the model fit and point estimates are unchanged for nearly any $\delta > 0$.

77. Clarke 2007.

78. Appendix G gives more information on model fit. It also includes a comparison of the outbidding model to reduced-form results from a vector autoregression model.

Substantive Effects of Competition on Violence

What is the substantive effect of competition on violence? Does heightened competition encourage or discourage violence? Answering these questions without a structural analysis is difficult because raw attack rates—even changes in attack rates—cannot be used as evidence for either discouragement or encouragement effects. If we see a group using violence more (or less) frequently in a given time frame, then the pattern could be explained by small (or large) attack costs, the equilibrium path visiting states in which a group is likely (or unlikely) to use violence, or merely small-sample bias arising via stochastic decisions. Instead of interpreting the data atheoretically, we use the fitted structural model to quantify how a group's use of violence *changes* as competition *changes* while holding everything else equal. To do this, we warp different aspects of competition in the fitted model while keeping other parts fixed and record how its predictions concerning the groups' use of violence would change in response. We consider this in two ways, by adjusting each actor's competitive behavior and then their competitive incentives.

Before proceeding, note that these exercises differ conceptually from those that would be computed from standard regression-based studies. Instead of focusing on observed correlations between competition and violence to quantify the effects of interest, we construct counterfactuals based on the theoretical model. Specifically, we take the competitive environment described by the estimated model and equilibrium as fixed and change specific features (re-solving for equilibrium) to find the effects of competition on violence. For example, the discouragement effects we will describe do not reflect an observed correlation between terrorist attacks and relative popularity. Rather, we are comparing the groups' estimated attack probabilities to their attack probabilities in the counterfactual world, holding fixed relative popularity.

This approach has two main benefits. First, it does not require untested proxy variables for competition. Second, we do not have to worry that unobserved confounders related to both competition and violence create a spurious result. The reason for this is that we control exactly what is changing in the model and change *only* parts of the model related to competition (individual group behavior or one of the competitive incentives). In other words, after fitting the model, we can tweak it in specific ways that are related to competition only and assess the changes. Evaluating the consequence of policies or behaviors that have never been observed is one strength of structural approaches because, without a model, it is unclear how to estimate the effects of such changes that are outside the data's support.[79]

The main concern with this approach is if the model poorly describes the data or if we have done a poor job estimating its primitives; we considered these concerns in the model fit section, and say more in Appendices D and F to J.[80] To be clear,

79. Canen and Ramsay 2024.

80. We can accommodate a richer set of robustness exercises for effectiveness of attacks ($\gamma_{i,1}$) than for the other competitive incentives (κ_i and β_i). The reason is that the former can be estimated using standard

https://doi.org/10.1017/S0020818324000390

these are not trivial concerns; however, we believe that the structural approach is complementary to more traditional empirical analyses because each comes with its own set of limitations and trade-offs.

Effect of Competitive Behavior on Violence

First, we compare how a group behaves with and without violence from its rival. That is, would Fatah use more or less violence if Hamas did not engage in terrorism, and vice versa? Specifically, we compare group i's estimated equilibrium attack probability (in Figure 3) to i's attack probability in its single-agent problem—that is, i's predicted use of violence if it expects its rival to never attack. Subtracting the latter from the former is one way to quantify the effect of competitive behavior on violence where the equilibrium attack probabilities represent violence in a competitive environment and the single-agent attack probabilities are from a noncompetitive environment. Figure 5 graphs these differences over time given the observed relative popularity s^t. Positive values indicate a positive effect of competition on violence, where a group's equilibrium probability of attacking is higher than its probability of attacking in its single-agent problem. Negative values indicate a negative effect.[81] Thus, one interpretation of the figure is that the value in month t with popularity level s^t indicates the effect on group i's immediate attack probability if group $-i$ were to stop using violence in all future periods.[82]

Before interpreting this figure, we provide some additional historical context for the period after the 2006 election. This era is characterized by various reconciliation attempts and agreements, with different levels of success, as well as several moments of tension. The first post-election spike in Figure 5, for example, appears during the 2007 Battle of Gaza and the consolidation of Hamas control in the Gaza Strip.[83] The relatively flat spot of this figure runs from 2009 to 2014, a period largely characterized by reconciliation talks, while the peak in 2014 is around the time of a failed coup attempt, when Hamas tried to unseat Fatah's leadership in the West Bank.[84]

Turning to the counterfactual, for Fatah, the values are entirely positive, indicating that Hamas encourages Fatah to use more violence than it would in the absence of competition. On average, competition from Hamas increases Fatah's use of violence

time-series regression techniques, but the latter requires a bespoke model, estimator, and identification conditions. On the one hand, this is a strength of our analysis, because Fatah's advantage in effectiveness is driving the estimated discouragement effects, and this result can be subjected to the most numerous robustness tests. On the other hand, this illustrates the main drawback of the structural approach: we cannot immediately add another control variable to the baseline model—we have mentioned the complications of adding the government as a third actor, for example.

81. Figure A.2 graphs the difference in attack probabilities as a function of relative popularity levels.

82. Rather than showing evidence either for or against outbidding, Figure 5 shows evidence of encouragement effects (positive numbers) or discouragement effects (negative numbers) for the two groups in different time periods.

83. Milton-Edwards and Farrell 2010, 272.

84. Ginsburg 2014.

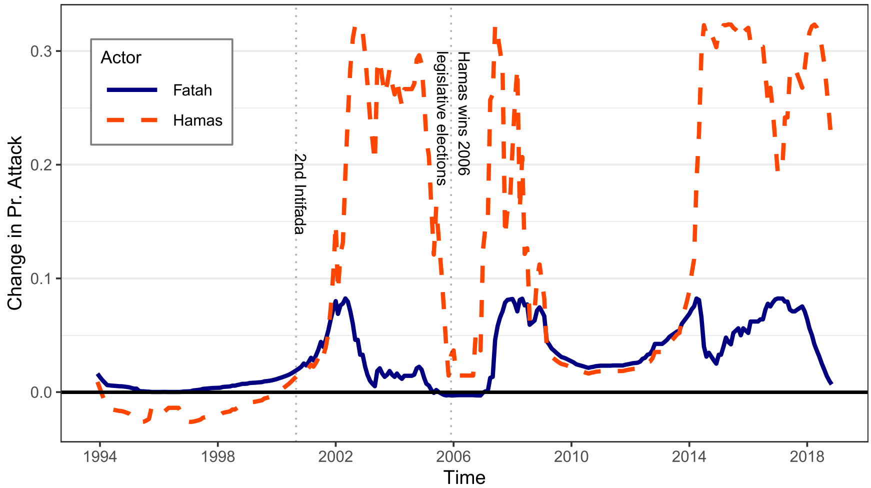

Notes: For each month t (horizontal axis), we compare group i's equilibrium probability of terrorism to the probability that would obtain if i expected its rival to never use violence, by subtracting the latter from the former, given the observed state s^t. Positive (negative) values indicate that competition increases (decreases) violence by group i in period t with state s^t.

FIGURE 5. *Effects of competitive behavior on violence*

by 34 percent relative to the counterfactual noncompetitive environment. This is the encouragement effect of competition on violence expected by the outbidding literature. Table 4 decomposes the effect into three time periods. Fatah's propensity for terrorism increases by about three percentage points due to competition from Hamas, especially after the start of the Second Intifada.

For Hamas the story is different, as heterogeneous effects exist. Competition from Fatah depresses Hamas's use of violence during the Oslo era, although we find a positive effect during and after the Second Intifada. During the Oslo era, Hamas's propensity for terrorism would increase by about one percentage point, on average, in the absence of competition from Fatah. This point estimate represents an average over this period, but if we consider the largest monthly effect, then we would predict a 9 percent increase in Hamas attacks were Fatah to commit no violence. Put differently, this corresponds to a 4 to 5 percent reduction in violence by Hamas during the Oslo era compared to its counterfactual single-agent problem where Fatah never attacks. This is the discouragement effect of competition on violence, where a group uses less violence in the competitive environment than in a noncompetitive one. Substantively, this change implies two or three more months with Hamas terrorism in the counterfactual world than in the observed data. While this is a relatively small effect, the potential devastation and loss of life associated with any given attack (particularly by Hamas) means that it is likely to be substantively meaningful.

TABLE 4. *Average effect of competitive behavior on violence in three eras*

	Jan. 1994 to Sept. 2000 (Oslo era)	Oct. 2000 to Jan. 2006 (2nd Intifada)	Feb. 2006 to Dec. 2018 (after 2006 election)
Hamas	−0.01	0.18	0.15
	(0.001)	(0.01)	(0.01)
Fatah	0.005	0.03	0.04
	(0.0004)	(0.002)	(0.001)

Note: Average difference between equilibrium and single-agent attack probabilities from different eras, with standard errors in parentheses.

These estimates suggest a competition-based explanation for the Oslo lull. Specifically, the popularity of the peace process of the 1990s boosted Fatah's standing among the Palestinian population. Figure 1 shows that Fatah frequently leads Hamas by thirty percentage points in terms of trust and support during this time. Accordingly, relative popularity is overwhelmingly in Fatah's favor relative to the rest of the sample (Figure 4). Hence, although Hamas has incentives to use violence, it also knows that the competition is very lopsided in Fatah's favor. Furthermore, Fatah is also more effective at using violence to increase its popularity, which depresses Hamas's use of violence.

This theoretical account has anecdotal support in some contemporary understandings of the conflict. As Kristianasen writes, "While the Oslo agreement consecrated Hamas's role as a new national resistance to Israel, it ushered in a reality that progressively would tie the movement's hands."[85] They further argue that Hamas had trouble remaining relevant during parts of this period due to Fatah's popularity, and that delays and discontentment with the peace process (that is, negative shocks to Fatah's public approval) helped Hamas remain relevant. Others affirm this trouble, pointing out that "Hamas was swimming against a tide of popular support" during this era,[86] and that in the mid-to-late 1990s, Hamas reduced their operations in the face of popular resistance.[87] While Hamas still pulled off several high-profile attacks during this time, its overall public support was low enough that it was unclear to contemporary observers whether the group would continue to be a relevant actor.[88]

This explanation requires two caveats. First, it is not the only possible explanation for this period of Fatah-Hamas interactions. As mentioned, the model does not include other key aspects of this relationship, such as the efforts of Hamas to spoil the peace, outside of a desire for local support. Second, because our results are derived from counterfactual comparisons to a world where Fatah never uses violence,

85. Kristianasen 1999, 20.
86. Milton-Edwards and Farrell 2010, 230.
87. Natil 2015, 38.
88. Kristianasen 1999, 33–34.

qualitative evidence cannot directly support the effects described here. Such evidence is inherently indirect, as the internal workings and strategic calculations of groups are often not well known in the real world and are completely unknown in the counterfactual world. However, the historical record does lend credence to the idea that Hamas was deterred by Fatah's popularity during this period; contemporary writers and conflict historians acknowledge that Fatah's popularity at that time had a notable effect on Hamas' strategic calculus.

We see in Figure 5 and Table 4 that the presence of a rival terrorist group can depress violence. With a rival that is an effective outbidder (Fatah), a group (Hamas) may use less violence than it normally would when it falls behind in the race for public opinion and sees the competition as increasingly difficult. As Figure 5 illustrates, this discouragement effect emerges in the Oslo era, when Fatah was more popular than Hamas and at its peak of popularity. Although some argue that increasing the number of terrorist groups—a common proxy for competitiveness—can decrease violence, the mechanisms underlying that argument do not appear in this setting. For example, Nemeth argues that increasing the number of ideologically similar groups should decrease violence through free-riding dynamics.[89] Hamas and Fatah are generally seen as ideologically opposed, however, and there are no free-riding incentives in the model. Another example comes from Conrad and Spaniel, who argue that the government may change its demands in response to a larger number of terrorist groups, leading to a negative correlation between this number and violence.[90] We find that endogenous government demands are not necessary for competition to reduce violence. Instead, the contest itself explains these results; the discouragement effect emerges when the popularity contest becomes lopsided in favor of the more effective actor.

Effect of Competitive Incentives on Violence

Now we examine how groups' incentives affect their attack probabilities. For example, how would overall violence levels change if group i became a more effective outbidder? The first counterfactual quantified the effects of competitive behavior on violence; this one illustrates the effects of competitive incentives on violence. To do this, we fix the transition parameters estimated from Table 1, the payoff parameters in Table 2, and the estimated equilibrium quantities. Then, for each group i, we change how effectively i can boost its popularity through terrorism by increasing and decreasing $\gamma_{i,1}$ by 1 percent. As the effectiveness of attacks changes, the equilibrium probabilities of attacks will change as well. Recall that $\gamma_{i,1}$ reflects the effectiveness of i in using terrorism to shift relative public opinion. An increase or decrease in $\gamma_{i,1}$ may reflect a change in tactics that the public may find more or less distasteful.

89. Nemeth 2014.
90. Conrad and Spaniel 2021.

https://doi.org/10.1017/S0020818324000390

Because multiple equilibria can exist, we cannot just vary $\gamma_{i,1}$, compute a new equilibrium, and compare choice probabilities under the old and new parameter values. Doing so would not guarantee that the new equilibrium bears any resemblance to the estimated one. Indeed, it may be possible to change behavior even though $\gamma_{i,1}$ does not change, by changing the selected equilibrium. To ensure that the counterfactuals fix the equilibrium that is selected by the data and the CMLE, we use a homotopy method to map equilibria as locally continuous functions of the parameters.[91] Appendix L gives details.

Figure 6 graphs these differences given the change in $\gamma_{i,1}$ and observed state s^t. Positive (negative) values indicate that violence from group i in observed state s^t increases (decreases) in the counterfactual scenario. As in the earlier exercise, one interpretation of the figure is that values in month t with popularity level s^t indicate the effect on the groups' immediate attack probabilities if i were to exogenously become more or less competitive.

Focusing on the effects of Hamas's competitive incentives, we find evidence of outbidding's expected encouragement effect: when Hamas has greater incentives to compete, violence by both groups increases. We estimate that a 1 percent increase in Hamas's effectiveness results in a one percentage-point increase in the frequency of terrorism by Hamas and a 0.1-percentage-point increase in the frequency of terrorism by Fatah. On average, this implies Hamas would increase its use of violence by 2 percent, and Fatah by 1 percent. These encouragement effects are even stronger when focusing on more recent observations, after the Oslo era.

Focusing on the effects of Fatah's competitive incentives, we find evidence of outbidding's unexpected discouragement effect: when Fatah has greater incentives to compete, violence by both groups decreases. We estimate that a 1 percent increase in Fatah's effectiveness results in a one percentage-point decrease in the frequency of terrorism by Hamas and a 0.2-percentage-point decrease in the frequency of terrorism by Fatah. On average, this implies both groups would decrease their violence by 2 percent if Fatah were to have greater incentives to compete via becoming 1 percent more effective at outbidding. Again, these discouragement effects are even stronger after the Oslo era.[92]

Figures A.3 and A.4 illustrate the same exercise for the value of support, β_i, and the costs of attacking, κ_i, respectively. The main takeaways are similar: when Hamas becomes more competitive, both sides attack more frequently (as expected by the outbidding literature); but when Fatah becomes more competitive, both sides tend to attack less frequently (in contrast to expectations in the outbidding literature).

These discouragement effects arise from asymmetric competition. Fatah has an advantage in its effectiveness in using terrorism to increase public support; that is, $|\gamma_{F,1}|$ is substantially larger than $|\gamma_{H,1}|$. When Fatah's incentive to compete increases, it more readily bears the upfront costs of attacks to increase future support. This affects Hamas's

91. Aguirregabiria 2012; Crisman-Cox and Gibilisco 2018.

92. Even with the estimated discouragement effect, we still note that Hamas is more likely than Fatah to attack in both the observed and the counterfactual world.

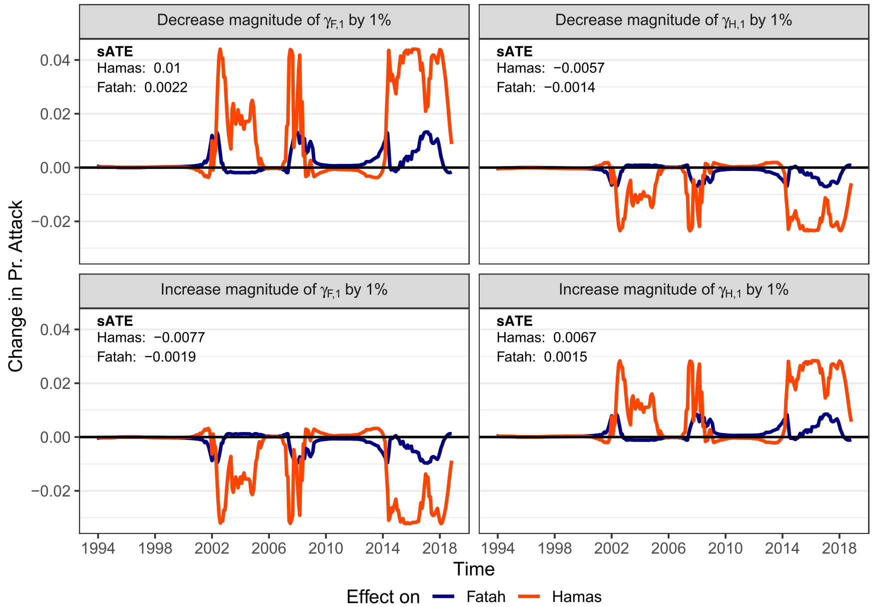

Notes: In each panel, we increase and decrease $\gamma_{i,1}$, for $i = H, F$, from its estimated value by 1%; all other parameters are held constant at their estimated values. We use a homotopy procedure to account for the potential presence of multiple equilibria; see Appendix L for details. Incentives to compete are greater when $\gamma_{i,1}$ is larger. The horizontal axis denotes the period/month t. The vertical axis is the difference between equilibrium attack probabilities (Figure 3) and counterfactual attack probabilities given the change in $\gamma_{i,1}$ and observed state s^t. Positive (negative) values indicate that violence by group i increases (decreases) in the counterfactual. Sample average treatment effects (sATE) are listed in the top-left corner of each pane.

FIGURE 6. *Relationship between terrorism and effectiveness of attacks*

equilibrium strategy. When Fatah becomes more aggressive, Hamas generally attacks less, because it cannot efficiently compete against the more aggressive and more capable Fatah. In equilibrium, this creates a feedback loop where Fatah uses less violence as Hamas uses less violence. These discouragement effects will be the strongest in periods (or states) when the model predicts that both groups will use substantial violence—that is, after Fatah loses substantial popularity (Figures 3 and 4)—because behavior in these periods (or states) will be the most sensitive to strategic incentives.

Discussion

The relationship between intergroup competition and terrorism is not as clear as previous scholarship suggests. Work looking for evidence of outbidding has focused on uncovering an encouragement effect in which greater competition leads to more violence. But we find that discouragement effects can also exist in a theory of outbidding where competition depresses violence. The key difference is the structural approach: we write down a model of outbidding, fit the model to data observed in the Fatah-Hamas rivalry, and then quantify the effects of competition on violence in the fitted model.

These heterogeneous effects matter for both researchers and policymakers. To see this, consider the effect of changes in the costs of terrorism, κ_i. For example, Israeli officials may want to pursue policies (such as barriers, trade restrictions, or violent reprisals) that make it harder for these groups to acquire arms or raise funds. Likewise, scholars would like to know how well a reduced-form study captures the relationship between competitive incentives (κ_i in this example) and the probability of violence. Increasing the costs of terrorism will decrease both groups' incentives to compete, and if we focus on just the encouragement effect, then we may expect that these changes will lead to less violence overall. However, with heterogeneous effects, the implications are less clear.

We illustrate the implications of changes in attack costs in Table 5. These counterfactuals follow the same procedure used to create Figure 6, only here we adjust κ_i by ± 0.13 for each actor individually (reflecting policy responses targeting a single group) and then for both actors (reflecting policy responses that affect both groups). This number translates into a roughly 5 percent and 15 percent change in the costs of terrorism for Fatah and Hamas, respectively. The values in this table are the probability of observing an attack by Hamas, Fatah, or either group, averaged over all values of state variable.

The first thing to note is that increasing only Hamas's attack costs has the desired effect: Hamas commits fewer attacks on average, and the overall rate of violence drops. The opposite effect appears when increasing only Fatah's attack costs: violence increases. But what happens when both groups are targeted? In this counterfactual, the encouragement and discouragement effects cancel out, and the overall attack probability is unchanged.

The implications for policy and research are clear. Simple tactics like trying to reduce terrorism by raising its cost may not have the desired effect in a competitive environment. Indeed, indiscriminate tactics that target all groups can leave the average probability of terrorism unchanged, as the competitive incentives cancel each other out. Boosting Fatah and targeting Hamas appears to be the most effective path away from terrorism in this conflict. For researchers, this heterogeneity should be concerning. Traditionally, outbidding scholars test their theories by regressing terrorist attacks on proxies for incentives to compete. What Table 5 makes clear, however, is that *even when these incentives are changing*, the overall effect might not be detected. Such a scenario can arise even when the competitive incentives are changing by the same amount in the same direction for all actors. Thus, standard approaches based on correlations between violence and proxies for competition cannot falsify the outbidding hypothesis. In this case, researchers regressing violence on the costs of outbidding may mistakenly conclude that outbidding is not a factor between these groups because when both actors become more or less competitive (via changes in κ_i) the overall probability of attacks is unchanged.

TABLE 5. *Average attack probabilities as costs* κH *and* κF *change*

		Pr(Hamas attacks)	Pr(Fatah attacks)	Pr(either attacks)
Baseline		0.37	0.11	0.43
Higher costs for	Hamas	0.33	0.10	0.40
	Fatah	0.46	0.10	0.51
	Both	0.36	0.10	0.43
Lower costs for	Hamas	0.44	0.12	0.50
	Fatah	0.35	0.11	0.42
	Both	0.38	0.11	0.44

In contrast, the structural approach provides a method for directly modeling competitive incentives, estimating the effects of changing these incentives, and quantifying how well outbidding explains the data. Doing this with outbidding theory, we uncover heterogeneous effects without relying on the commonly used, but untestable, proxies for competition. We are also able to assess the model for face validity and then explicitly consider the fit of the model both in terms of how well it explains violence on its own and in comparison to two alternative models that do not contain outbidding.

Naturally, our analysis raises new questions that cannot be answered in a single paper. For example, what explains the variation of competitive incentives across groups, and what substantive features of the conflict environment determine whether asymmetric incentives are strong enough to generate discouragement effects? Our model and data cannot answer these questions because we treat the groups' incentives as exogenous parameters to be estimated and our data consists of only two groups. Nonetheless, these questions *can* be answered if the model is applied to other cases of intergroup competition or generalized to include more than two groups. For example, contests among republican groups in Northern Ireland, leftist groups in Colombia, or Tamil groups in Sri Lanka are natural places to study outbidding. The main hurdle to studying alternative conflicts or more groups is the need for long-term public support data, but as intrastate conflict data become more fine-grained, we anticipate more applications outside the specific Hamas-Fatah rivalry.

Data Availability Statement

Replication files for this article may be found at <https://doi.org/10.7910/DVN/NDGZ8G>.

Supplementary Material

Supplementary material for this article is available at <https://doi.org/10.1017/S0020818324000390>.

References

Abbring, Jaap H., and Øystein Daljord. 2020. Identifying the Discount Factor in Dynamic Discrete Choice Models. *Quantitative Economics* 11 (2):471–501.

Abrahms, Max. 2006. Why Terrorism Does Not Work. *International Security* 31 (2):42–78.

Acosta, Benjamin, and Kristen Ramos. 2017. Introducing the 1993 Terrorism and Political Violence Dataset. *Studies in Conflict and Terrorism* 40 (3):232–47.

Aguirregabiria, Victor. 2012. A Method for Implementing Counterfactual Experiments in Models with Multiple Equilibria. *Economics Letters* 114 (2):190–94.

Berman, Eli, and David D. Laitin. 2008. Religion, Terrorism and Public Goods: Testing the Club Model. *Journal of Public Economics* 92 (10–11):1942–67.

Bloom, Mia M. 2004. Palestinian Suicide Bombing: Public Support, Market Share, and Outbidding. *Political Science Quarterly* 119 (1):61–88.

Brown, Nathan J. 2012. Gaza Five Years On: Hamas Settles In. Carnegie Endowment for International Peace. Available at <https://carnegieendowment.org/research/2012/06/gaza-five-years-on-hamas-settles-in>.

Canen, Nathan, and Kristopher Ramsay. 2024. Quantifying Theory in Politics: Identification, Interpretation and the Role of Structural Methods. *Journal of Theoretical Politics* 36 (4):301–327.

Chaudoin, Stephen, and Jonathan Woon. 2018. How Hard to Fight? Cross-Player Effects and Strategic Sophistication in an Asymmetric Contest Experiment. *Journal of Politics* 80 (2):585–600.

Chenoweth, Erica. 2010. Democratic Competition and Terrorist Activity. *Journal of Politics* 72 (1):16–30.

Clarke, Kevin A. 2007. A Simple Distribution-Free Test for Nonnested Model Selection. *Political Analysis* 15 (3):347–63.

Clauset, Aaron, Lindsay Heger, Maxwell Young, and Kristian Skrede Gleditsch. 2010. The Strategic Calculus of Terrorism: Substitution and Competition. *Cooperation and Conflict* 45 (1):6–33.

Conrad, Justin, and Kevin Greene. 2015. Competition, Differentiation, and the Severity of Terrorist Attacks. *Journal of Politics* 77 (2):546–61.

Conrad, Justin, Kevin T. Greene, and Brian J. Phillips. 2024. Introducing ViNSAR: Dyadic Data on Violent Non-state Actor Rivalry. *Journal of Conflict Resolution* 68 (9):1884–1908.

Conrad, Justin, and William Spaniel. 2021. *Militant Competition: How Terrorists and Insurgents Advertise with Violence and How They Can Be Stopped*. Cambridge University Press.

Crenshaw, Martha. 1981. The Causes of Terrorism. *Comparative Politics* 13 (4):379–99.

Crisman-Cox, Casey, and Michael Gibilisco. 2018. Audience Costs and the Dynamics of War and Peace. *American Journal of Political Science* 62 (3):566–80.

Cunningham, Kathleen Gallagher, Kristin M. Bakke, and Lee J.M. Seymour. 2012. Shirts Today, Skins Tomorrow: Dual Contests and the Effects of Fragmentation in Self-Determination Disputes. *Journal of Conflict Resolution* 56 (1):67–93.

Dechenaux, Emmanuel, Dan Kovenock, and Roman M. Sheremeta. 2015. A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments. *Experimental Economics* 18 (4):609–669.

Fearon, James D. 2011. Arming and Arms Races. Working paper, Stanford University. Available at <https://web.stanford.edu/group/fearon-research/cgi-bin/wordpress/wp-content/uploads/2013/10/arming3.pdf>.

Findley, Michael G., and Joseph K. Young. 2012. More Combatant Groups, More Terror? Empirical Tests of an Outbidding Logic. *Terrorism and Political Violence* 24 (5):706–721.

Fortna, Virginia Page. 2015. Do Terrorists Win? Rebels' Use of Terrorism and Civil War Outcomes. *International Organization* 69 (3):519–56.

Frey, Anderson, Gabriel López-Moctezuma, and Sergio Montero. 2023. Sleeping with the Enemy: Effective Representation Under Dynamic Electoral Competition. *American Journal of Political Science* 67 (4):915–31.

Getmansky, Anna, and Thomas Zeitzoff. 2014. Terrorism and Voting: The Effect of Rockets in Israeli Elections. *American Political Science Review* 108 (3):588–604.

Gibilisco, Michael, Brenton Kenkel, and Miguel R. Rueda. 2022. Competition and Civilian Victimization. *Journal of Conflict Resolution* 66 (4–5):809–835.

Ginsburg, Mitch. 2014. Israel Says It Foiled Hamas Plan for Massive Attacks on Israel, Coup Against PA. *Times of Israel*, 18 August. Available at <https://www.timesofisrael.com/israel-says-it-foiled-hamas-plan-for-coup-against-pa-in-west-bank/>.

Gould, Eric D., and Esteban F. Klor. 2010. Does Terrorism Work? *Quarterly Journal of Economics* 125 (4):1459–1510.

Heger, Lindsay L., and Danielle F. Jung. 2017. Negotiating with Rebels: The Effect of Rebel Service Provision on Conflict Negotiations. *Journal of Conflict Resolution* 61 (6):1203–1229.

Iaryczower, Matias, Gabriel López-Moctezuma, and Adam Meirowitz. 2024. Career Concerns and the Dynamics of Electoral Accountability. *American Journal of Political Science* 68 (2):696–713.

Jaeger, David A., Esteban F. Klor, Sami H. Miaari, and M. Daniele Paserman. 2012. The Struggle for Palestinian Hearts and Minds: Violence and Public Opinion in the Second Intifada. *Journal of Public Economics* 96 (3–4):354–68.

Jaeger, David A., Esteban F. Klor, Sami H. Miaari, and M. Daniele Paserman. 2015. Can Militants Use Violence to Win Public Support? Evidence from the Second Intifada. *Journal of Conflict Resolution* 59 (3):528–49.

Jerusalem Media and Communication Centre. N.d. Polls. Accessed November 2019. Available at <http://www.jmcc.org/en/Category/4/polls>.

Johannsen, Margret. 2011. A Balance of Fear: Asymmetric Threats and Tit-for-Tat Strategies in Gaza. *Journal of Palestine Studies* 41 (1):45–56.

Kang, Karam. 2016. Policy Influence and Private Returns from Lobbying in the Energy Sector. *Review of Economic Studies* 83 (1):269–305.

Kenkel, Brenton, and Kristopher W. Ramsay. 2024. The Effective Power of Military Coalitions: A Unified Theoretical and Empirical Model. *Journal of Politics*. Available at <https://www.journals.uchicago.edu/doi/abs/10.1086/732953?journalCode=jop>.

Kirkegaard, René. 2012. Favoritism in Asymmetric Contests: Head Starts and Handicaps. *Games and Economic Behavior* 76 (1):226–48.

Köning, Michael D., Dominic Rohner, Mathias Thoenig, and Fabrizio Zilibotti. 2017. Networks in Conflict: Theory and Evidence from the Great War of Africa. *Econometrica* 85 (4):1093–1132.

Konrad, Kai A., and Dan Kovenock. 2005. Equilibrium and Efficiency in the Tug-of-War. CESifo Working Paper No. 1564. Available at <http://dx.doi.org/10.2139/ssrn.844644>.

Kristianasen, Wendy. 1999. Challenge and Counterchallenge: Hamas's Response to Oslo. *Journal of Palestine Studies* 28 (3):19–36.

Kydd, Andrew H., and Barbara F. Walter. 2002. Sabotaging the Peace: The Politics of Extremist Violence. *International Organization* 56 (2):263–96.

Kydd, Andrew H., and Barbara F. Walter. 2006. The Strategies of Terrorism. *International Security* 31 (1):49–80.

McFadden, Daniel. 1978. Modelling the Choice of Residential Location. In *Spatial Interaction Theory and Planning Models*, edited by Anders Karlqvist, Lars Lundqvist, Folke Snickers, and Jörgen W. Weibull, 75–96. North-Holland.

Milton-Edwards, Beverly, and Stephen Farrell. 2010. *Hamas: The Islamic Resistance Movement*. Polity Press.

Natil, Ibrahim. 2015. *Hamas Transformation: Opportunities and Challenges*. Cambridge Scholars.

Nemeth, Stephen. 2014. The Effect of Competition on Terrorist Group Operations. *Journal of Conflict Resolution* 58 (2):336–62.

Palestinian Center for Policy and Survey Research. N.d. Index Survey Research. Accessed November 2019. Available at <https://www.pcpsr.org/en/node/105>.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. Asymptotic Least Squares Estimators for Dynamic Games. *Review of Economic Studies* 75 (3):901–928.

Polo, Sara M.T., and Belén González. 2020. The Power to Resist: Mobilization and the Logic of Terrorist Attacks in Civil War. *Comparative Political Studies* 53 (13):2029–2060.

Polo, Sara M.T., and Blair Welsh. 2024. Violent Competition and Terrorist Restraint. *International Organization* 78 (3):397–426.

Powell, Robert. 1993. Guns, Butter, and Anarchy. *American Political Science Review* 87 (1):116–32.

Rust, John. 1994. Structural Estimation of Markov Decision Processes. In *Handbook of Econometrics*, edited by Robert F. Engle and Daniel L. McFadden, 3081–3143. Elsevier.
Schanzer, Jonathan. 2003. The Challenge of Hamas to Fatah. *Middle East Quarterly* 10 (2):29–38.
Siegel, Ron. 2014. Asymmetric Contests with Head Starts and Nonmonotonic Costs. *American Economic Journal: Microeconomics* 6 (3):59–105.
Silvey, Samuel D. 1959. The Lagrangian Multiplier Test. *Annals of Mathematical Statistics* 30 (2):389–407.
Stein, William E. 2002. Asymmetric Rent-Seeking with More than Two Contestants. *Public Choice* 113 (3):325–36.
Stewart, Megan A. 2018. Civil War as State-Making: Strategic Governance in Civil War. *International Organization* 72 (1):205–226.
Su, Che-Lin, and Kenneth L. Judd. 2012. Constrained Optimization Approaches to Estimation of Structural Models. *Econometrica* 80 (5):2213–30.
Tauchen, George. 1986. Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters* 20 (2):177–81.
Thomas, Jakana. 2014. Rewarding Bad Behavior: How Governments Respond to Terrorism in Civil War. *American Journal of Political Science* 58 (4):804–818.
Tokdemir, Efe, and Seden Akcinaroglu. 2016. Reputation of Terror Groups Dataset: Measuring Popularity of Terror Groups. *Journal of Peace Research* 53 (2):268–77.
Wagstaff, William A., and Danielle F. Jung. 2020. Competing for Constituents: Trends in Terrorist Service Provision. *Terrorism and Political Violence* 32 (2):293–324.
Wood, Reed M., and Jacob D. Kathman. 2015. Competing for the Crown: Inter-rebel Competition and Civilian Targeting in Civil War. *Political Research Quarterly* 68 (1):167–79.
Ze'evi, Dror. 2008. Clans and Militias in Palestinian Politics. Crown Center for Middle East Studies, Brandeis University. Available at <https://www.brandeis.edu/crown/publications/middle-east-briefs/pdfs/1-100/meb26.pdf>.

Authors

Casey Crisman-Cox is Assistant Professor of Political Science at Texas A&M University. He can be reached at c.crisman-cox@tamu.edu.
Michael Gibilisco is Professor of Political Science at the California Institute of Technology. He can be reached at michael.gibilisco@caltech.edu.

Acknowledgments

Thanks to Federico Echenique, Alex Hirsch, Karam Kang, Brenton Kenkel, Gabriel López-Moctezuma, Jack Paine, Kris Ramsay, Jean-Laurent Rosenthal, Kirssa Cline Ryckman, Matt Shum, William Spaniel, Yi Xin, and audiences at Caltech's Applied Political Theory Group, Vanderbilt University's Conflict Workshop, Penn State's Political Science Colloquium, the CEPR-Warwick-Princeton-Yale Symposium on Political Economy, and APSA 2022 for comments and discussions. We also thank the editors and four anonymous reviewers for their feedback. Previous versions of this paper were circulated under the title "Tug of War: Dynamic Outbidding Between Terrorist Groups."

Key Words

Terrorism; outbidding; contests; dynamic games; structural estimation; Israel-Palestine

Date received: April 11, 2023; Date accepted: December 4, 2024