

Appendix for “Tug of War: The Heterogeneous Effects of Outbidding between Terrorist Groups”

Casey Crisman-Cox* Michael Gibilisco†

September 2024

Contents

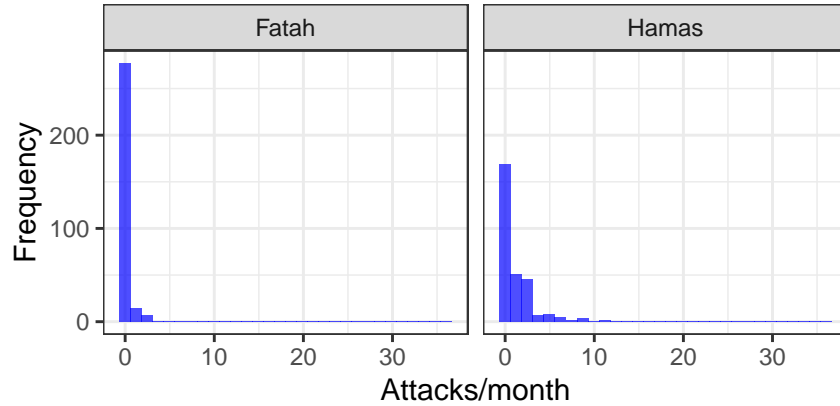
A	Additional figures	1
B	Numerical example	4
C	More details on the survey data and dynamic factor model	8
D	First-stage robustness	13
	D.1 Control variables and specification choices	13
	D.2 Measurement of the actions	19
	D.3 Instrumental variable analysis	19
	D.4 Sensitivity analysis	24
	D.5 Is an attack an attack an attack?	27
E	Identification	29
	E.1 Identification intuition	31
F	Standard errors and sensitivity of the second stage estimates	32
G	Model fit and comparisons	34
	G.1 Comparison to alternative theoretical models	34
	G.2 Comparison to reduced-form results	36
H	Robustness to different time spans	39
I	Choice of discount factor	40
J	Robustness to discretization	40
K	Interpreting estimates of β_i	46

*Texas A&M University. Email: c.crisman-cox@tamu.edu

†California Institute of Technology. Email: michael.gibilisco@caltech.edu

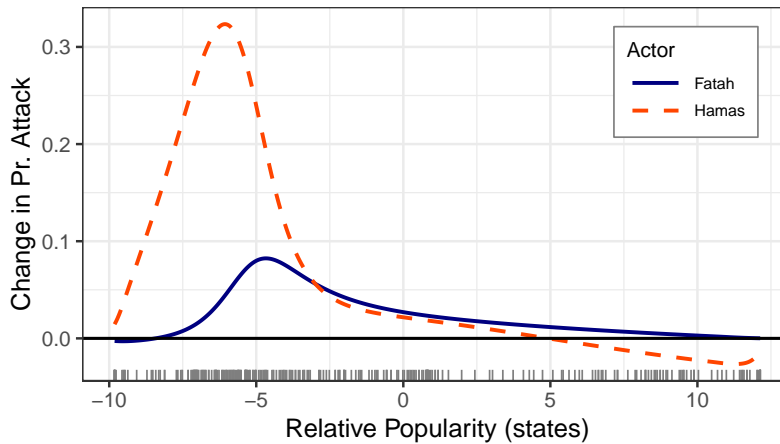
A Additional figures

Figure A.1: Empirical distribution of the number of GTD attacks in each month.



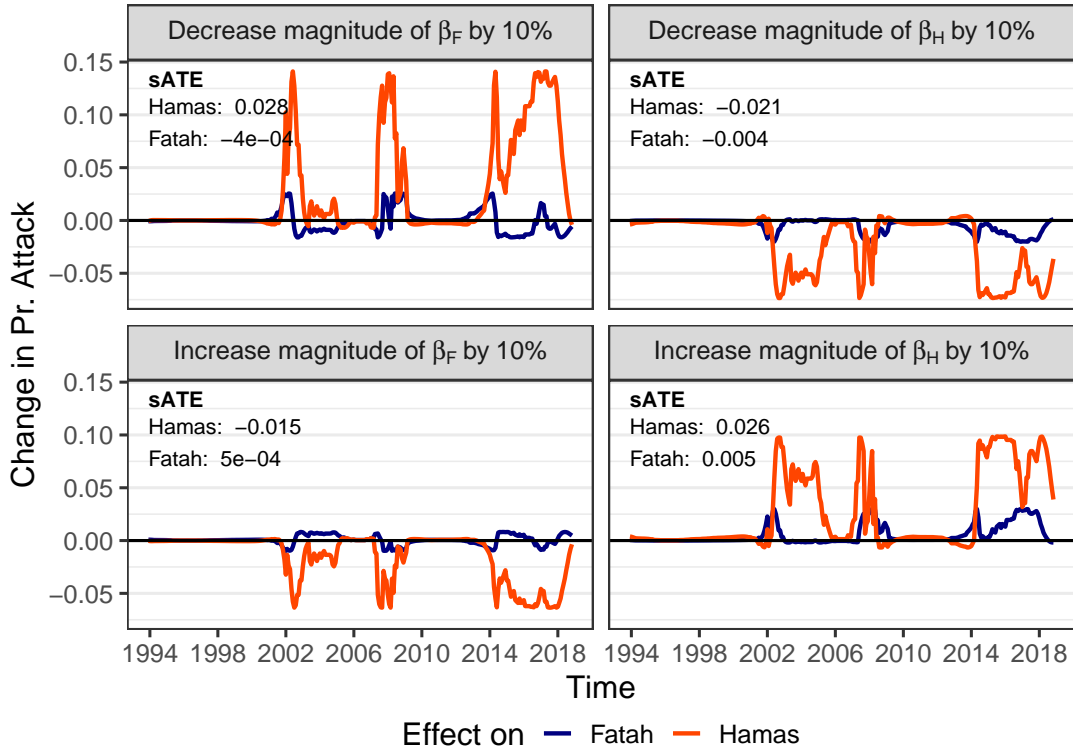
Note: For each month, we count the number of attacks attributed to Fatah (left) and Hamas (right) in the GTD and plot the distribution. The GTD is a standard source for recording terrorist acts, which are defined as either a threat or an attack that meets two of the following conditions: occurs outside the confines of legitimate warfare; is designed to signal to a larger audience than the immediate victims; and helps to attain a political, religious, or social goal (START 2019a, 6). For Hamas, the mean is 1.5, median is 0, and range is 0–36. For Fatah, the mean is 0.2, median is 0, and range is 0–15.

Figure A.2: Effects of competition on violence for all states s




Note: We compare group i 's equilibrium probability of terrorism in state s to the probability that would arise if i expects its rival to never use violence, by subtracting the latter from the former. Whereas Figure 5 graphs the difference over time conditional on the observed relative popularity s^t , Figure A.2 shows the difference as a function of all relative popularity levels s on the horizontal axis. Positive values indicate that competition increases violence by group i in state s ; negative values indicate that competition decreases violence by group i in state s . Rug plot denotes observed states s^t .

Figure A.3: Relationship between terrorism and value of support in observed states.



Note: In each panel, we increase and decrease the magnitude of β_i for $i = H, F$ from its estimated value by 10%; all other parameters are held constant at their estimated values. We use a procedure from Aguirregabiria (2012) to account for the potential presence of multiple equilibria—see Appendix L for details. Incentives to compete are greater when the value of support, β_i , is larger in magnitude. 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 β_i and observed state s^t . Positive (negative) values indicate that violence by group i increases (decreases) in the counterfactual.

Figure A.4: Relationship between terrorism and cost of attacking in observed states.



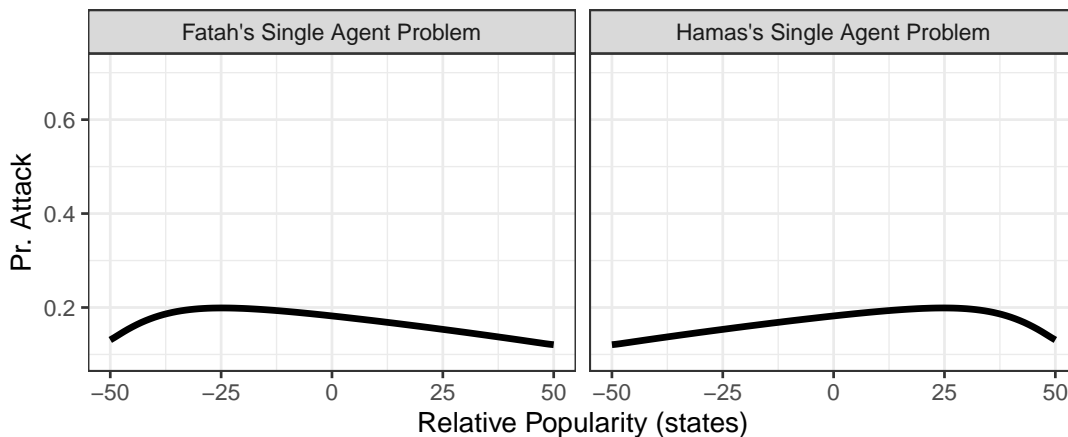
Note: In each panel, we increase and decrease the magnitude of κ_i for $i = H, F$ from its estimated value by 1%; all other parameters are held constant at their estimated values. We use a procedure from Aguirregabiria (2012) to account for the potential presence of multiple equilibria—see Appendix L for details. Incentives to compete are greater when attack costs, κ_i , are closer to zero. 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 κ_i and observed state s^t . Positive (negative) values indicate that violence by group i increases (decreases) in the counterfactual.

B Numerical example

To illustrate the strategic tensions in the model, we pick hypothetical values for the parameters and study the equilibria that arise. The specification is symmetric to aid interpretation, but our model is more general. The popularity levels are $\mathcal{S} = \{-S, -S + 1, \dots, S - 1, S\}$ where $S = 50$. For the payoff parameters, we set $\beta_F = \frac{1}{500} = -\beta_H$ and $\kappa_i = -2$. For the transitions, we assume $\mu[a^t, s^t; \gamma] = s^t - a_H^t + a_F^t$ and $\sigma = 2$. In other words, group i 's attacks shift the mean of tomorrow's expected relative popularity by one in its preferred direction. The current popularity level does not change the effectiveness of attacks ($\gamma_{i,2} = 0$).

To build intuition, Figure B.1 presents group i 's optimal attack probabilities when its rival never attacks, i.e., when i is the only relevant group. The probabilities range from 0.1 to 0.2. Notice i is most likely to attack when its relative popularity is weak (small state s for Fatah and large state s for Hamas), although end-point effects emerge because the state space \mathcal{S} is bounded. When the current state is at a boundary, one group's popularity cannot get worse, while the other's cannot get better tomorrow. This decreases the groups' incentives to attack. When a group is relatively unpopular, it has stronger dynamic benefits from using costly attacks to increase its popularity: attacking increases i 's future payoffs for some time and decreases its need to use costly attacks in the future. Thus, comparing across the two single-agent problems, the groups generally attack in different states—the correlation coefficient of their attack probabilities is $\rho = -0.13$.

Figure B.1: Attack probabilities without competition in the numerical example.



Note: Left panel graphs the probability that Fatah attacks (y -axis) as a function of the states (x -axis) in its single-agent dynamic programming problem, i.e., when Hamas never attacks. The right panel graphs the attack probabilities for Hamas's single-agent dynamic programming problem, i.e., when Fatah never attacks.

Turning to the strategic setting, we investigate equilibrium attack probabilities. To find equilibria, we repeatedly compute solutions to Equation 7 (given the parameters in this example) using the Newton-Raphson method and ten thousand different starting values. After knowing the solutions to Equation 7, we use Equation 4 to compute each group's

probabilities of attacking. Overall, we found three different solutions to Equation 7 given the fixed parameters in the symmetric example. Of course, additional equilibria might exist within this numerical example. Moreover, equilibrium multiplicity is a feature of this numerical example (i.e., the fixed parameters θ and γ above), but it is not a general feature of the model. That is, under other parameters a unique equilibrium exists.¹

Figure B.2: Computed equilibria in the numerical example.

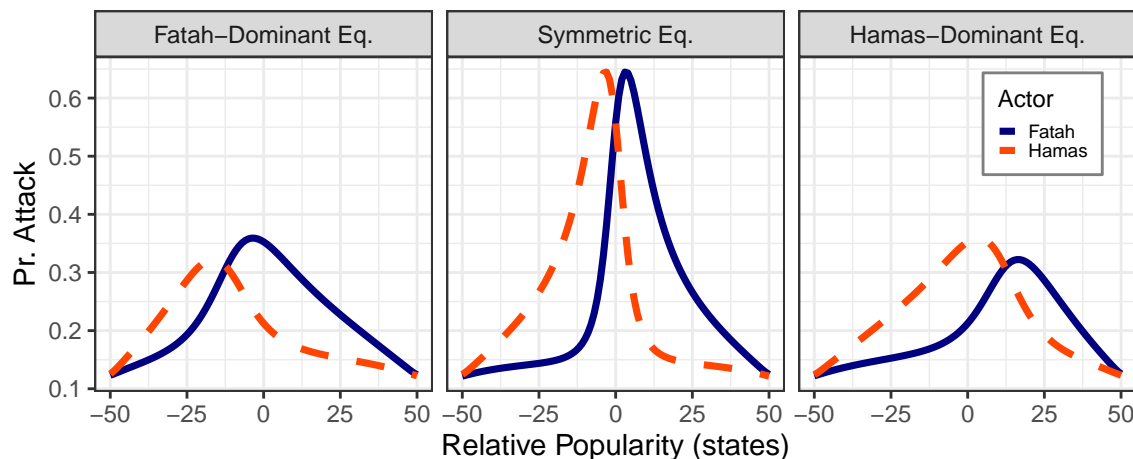
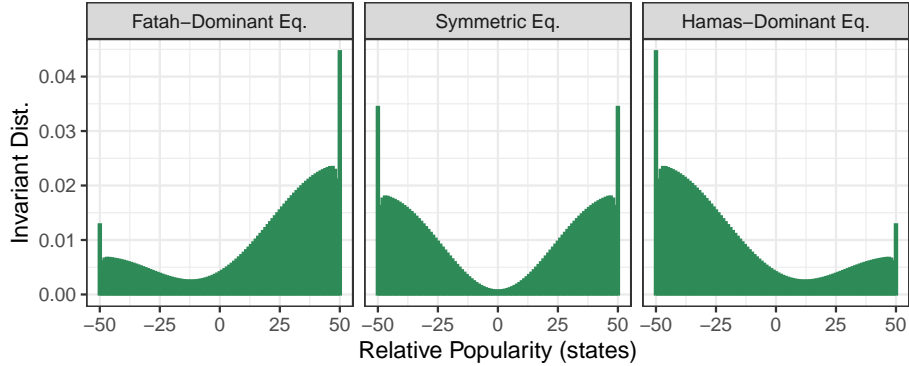


Figure B.2 graphs the attack probabilities for each of the three computed equilibria. In the symmetric equilibrium, terrorism is most fierce when the groups are equally popular, and group i attacks with the highest probability once it begins to be slightly more popular than its rival. The other two equilibria are asymmetric but are essentially the same with the actors and states flipped. In these equilibria, one actor (labeled dominant) is using violence with higher probability than the remaining actor for the majority of the state space. Figure B.3 graphs the invariant distribution for each of the three equilibria.

The example illustrates three features. First, violence between groups exhibits some strategic complementarities: attack probabilities are positively correlated across states. In the asymmetric equilibria, the correlation coefficient is $\rho = 0.41$, and in the symmetric equilibrium it is $\rho = 0.28$. These complementarities do not arise through the group’s per-period payoffs in Equation 1, because i ’s per-period payoff does not depend on its opponent’s action. In addition, the groups do not become more effective in certain states as $\gamma_{i,2} = 0$. Instead, the complementarities arise endogenously through tug-of-war dynamics in which competition can increase violence. Indeed, group i ’s attack probabilities in any of the three equilibria are larger than those in its single-agent problem. Thus, our dynamic model can

¹If $\beta_i = 0$, then a unique equilibrium exists. To see this, when $\beta_i = 0$, i ’s only benefit from attacking comes from the fixed costs κ_i and the iid shocks ε_i . Thus, i ’s optimal attack probabilities will be constant as a function of the state s and will not depend on the behavior of j . From the perspective of group j , they are essentially facing a single-agent dynamic program, which has a unique solution.

Figure B.3: Equilibrium invariant distributions in the numerical example.



Note: The invariant distribution indicates how likely it is for the equilibrium path to visit each of the relative popularity levels in the long run. The equilibrium path in an asymmetric equilibrium is more likely to visit popularity levels that are favorable to the dominant (more violent) actor. In the symmetric equilibrium, the invariant distribution is symmetric around zero. Spikes occur at the extreme values of the state space because the interaction can bunch at high and low values due to relative popularity levels being bounded.


endogenize the strategic complementarities for violence found in previous analyses using static games (Gibilisco, Kenkel and Rueda 2022).

Second, these complementarities are moderate, i.e., the attack probabilities are not perfectly correlated. In all equilibria, the state in which Hamas is most likely to attack is strictly less than the state in which Fatah is most likely to attack. This arises because, all else equal, Hamas wants to exert costly effort to attack at popularity levels where becoming more popular (decreasing the state variable) reduces Fatah’s likelihood to attack and vice versa. In contrast, Fatah wants to attack at levels where becoming more popular (increasing the state variable) reduces Hamas’s likelihood to attack. These incentives temper the potential strategic complementarities in the model.

Third, equilibrium rates of attacks are not perfect measures of competitive incentives. Consider the Fatah-dominant equilibrium. At the majority of popularity levels, Fatah is attacking with greater probability than Hamas, so one might conclude that Fatah has smaller attack costs or a greater value of popularity. The example is symmetric, however, and both groups have identical competitive incentives. Thus, incentives to compete do not directly map onto observed rates of violence as the relationship is mediated by a strategic interaction. As a result, reduced-form regressions using observed terrorism as the dependent variable may obscure some aspects of the outbidding process. Directly estimating the model’s parameters and equilibrium allows for a deeper exploration of how competition affects violence.

Figure B.4 graphs comparative statics to illustrate how changes in competitive incentives can affect violence. In this symmetric example, we find that increasing Hamas’s value for public support can increase, decrease, or have mixed effects on each actor’s willingness

Figure B.4: Comparative statics in the numerical example.



Note: We increase and decrease Hamas’s value of popularity, β_H , from the baseline numerical example, where $\beta_H = -0.002$, using a procedure from Aguirregabiria (2012) to account for the presence of multiple equilibria—see Appendix L for details. All other parameters are held fixed. The horizontal axis is the relative popularity levels (smaller values are more favorable to Hamas) and the vertical axis is the probability of an attack. Columns correspond to the equilibrium and rows denote which group’s probability of attacking is graphed. Darker red denotes stronger incentives to compete, i.e., β_H is more negative.

depending on the equilibrium.² These mixed comparative statics motivate fitting this model to data as it is not clear what parameter values or equilibrium are empirically relevant. Overall, the example illustrates that enhanced incentives to compete can either increase or decrease overall violence levels, raising two questions: Which effect dominates in the data? Do these effects vary across groups?

To conclude this section, a remark is in order: for several reasons, it is not possible to map the three equilibria in Figure B.1 to potential equilibria in the estimated model. As stated above, the model can have a unique equilibrium under some specifications. Furthermore, this example is symmetric but the fitted model is asymmetric. It is therefore unlikely a symmetric equilibrium exists in the fitted model. Moreover, the substantive labels given to the equilibria only make sense when comparing across equilibria. For example, the equilibrium in the left panel of Figure B.1 is Fatah-dominant because Fatah is using more violence than Hamas *and* other equilibria exists in which either the groups use equal

²We use a procedure from Aguirregabiria (2012) and Crisman-Cox and Gibilisco (2018) to conduct counterfactuals in the presence of multiple equilibria—see Appendix L for details.

Table C.1: Survey questions and frequency.

Source	Question	Frequency
JMCC	“Which political or religious faction do you trust the most?”	2-6 times/year
PCPSR	“Which of the following political parties do you support?”	2-9 times/year
JMCC	“If Legislative Council elections were held today, which party would you vote for?”	0-5 times/year starting in 2006

levels of violence or Hamas uses more violence than Fatah. Systematically comparing behavior across equilibria is generally impossible given our model, however. Doing so requires computing the set of equilibria, which is equivalent to finding all solutions to Equation 7. Equation 7 is a system of $4K$ nonlinear equations. Computing all solutions is only feasible when K is small enough so that grid searching is feasible or when the system satisfies others properties, e.g., it is polynomial.³

C More details on the survey data and dynamic factor model

Table C.1 lists the wordings of each survey question and the frequency at which it was asked. To illustrate, Figure C.1 presents screenshots of two polls that we record. For the JMCC, the August 2015 poll breaks down trust by geographic region, although not all JMCC polls do this. For the PCPSR, the December 2006 poll also breaks down support by geographic region, where the columns are, from left to right, are total, West Bank, and Gaza. In this PCPSR poll, there are 1270 adults: 830 in the West Bank and 440 in Gaza.

Figure C.2 shows how the aggregate survey responses vary between Gaza and the West Bank. The correlation coefficients summarize the graphs: attitudes toward each group are highly correlated across geographic areas. One noticeable pattern is that *both* groups are generally more popular in Gaza than in the West Bank although Hamas seems to get a larger boost. Also, in some months, the percentage for the entire survey is outside the interval for the percentages in the West Bank and Gaza, e.g., trust in Fatah in September 1995, a theoretical impossibility. In these months, the individual survey does not report break down averages by geographic area, and the impossibility is a result of the linear interpolation in Figure C.2.

Table C.2 shows the raw correlations between pairs of survey questions using the entire sample of respondents (i.e., both West Bank and Gaza). Fixing a group, trust and support are highly correlated (0.69 and 0.77 for Hamas and Fatah, respectively). Likewise, support

³In finite, normal-form games, the set of Nash equilibria can be represented as the set of solutions to a system of polynomial equations and inequalities, so tools exist to compute all equilibria, although this can still be computationally intensive. See Herings and Peeters (2010) for a recent discussion. Gibilisco and Montero (2022) exploit these properties and tools in their structural analysis of major-power interventions into civil war to quantify the effects of equilibrium selection. These tools are not applicable in our environment given that we are using private-information payoff shocks to rationalize the data, however.

Figure C.1: Examples of public opinion polls.

Q31- Which Political or religious faction do you trust the most?

	Total	West Bank	Gaza
	n=1199	n= 749	n=450
Fatah	35.4	34.8	36.2
Hamas	20.0	18.6	22.4
PFLP	3.2	3.1	3.3
Other Islamic factions	3.1	1.9	5.1
Others	3.5	4.0	2.7
I don't trust anyone	29.2	31.1	26.0
No answer	5.6	6.5	4.3

*This was an open-ended question no options were read to the interviewee

(a) JMCC: August 2015

Which of the following political parties do you support?

1) PPP	1.4	1.6	0.9
2) PFLP	3.1	3.2	3.0
3) Fateh	32.3	31.7	33.3
4) Hamas	29.3	25.4	36.0
5) DFLP	1.0	1.4	0.2
6) Islamic Jihad	1.6	1.8	1.1
7) Fida	0.2		0.5
8) National Initiative (Mubadara)	1.0	1.2	0.7
9) Independent Islamists	3.3	3.9	2.3
10) Independent Nationalists	4.3	4.2	4.6
11) None of the above	22.4	25.3	17.4
12) Other, specify	0.2	0.4	0.0

(b) PCPSR: December 2006

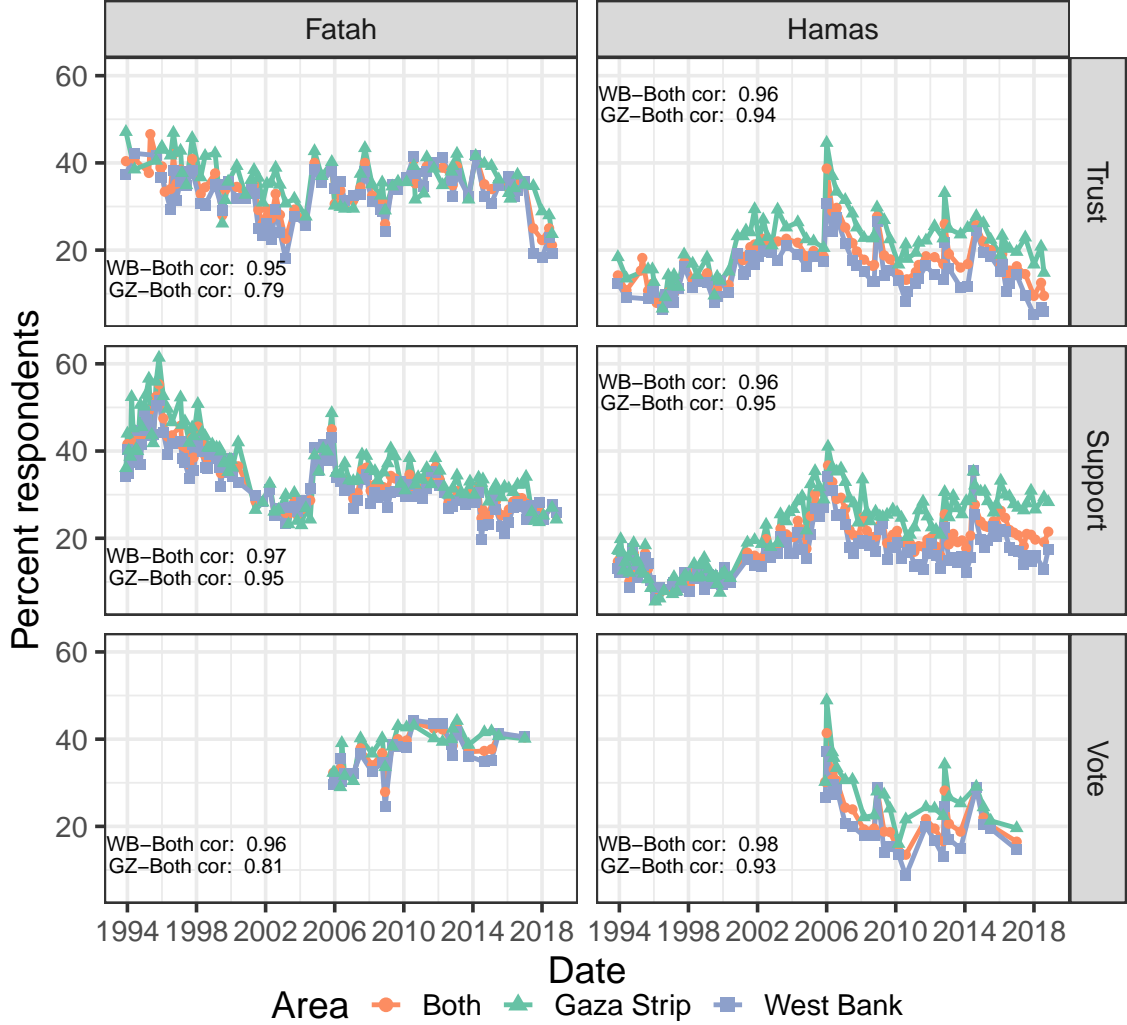
for one group is negatively correlated with support for the other. The remaining correlations are almost all in the expected directions, suggesting that the population does in fact trade off among supporting these two leading actors. The only exceptions are the negative correlation between voting for and supporting Fatah and the positive correlation between supporting Fatah and voting for Hamas. However, trust in Fatah correlates highly with both supporting and voting for Fatah, while voting for Hamas correlates highly with support and trust in Hamas.

To produce the latent state variable, we first let y^t be the column vector denoting the 6 survey values at time t such that

$$y^t = \begin{bmatrix} \% \text{ of Population that trusts Fatah} \\ \% \text{ of Population that trusts Hamas} \\ \% \text{ of Population that supports Fatah} \\ \% \text{ of Population that supports Hamas} \\ \% \text{ of Population that plans to vote for Fatah} \\ \% \text{ of Population that plans to vote for Hamas} \end{bmatrix}^t. \quad (\text{C.1})$$

Let z^t denote the vector of z -transformed surveys, where for survey $j = 1, \dots, 6$, $z_j^t =$

Figure C.2: Survey responses and geographic areas.



Note: We graph survey responses over time broken down by geography. The both category is the aggregate survey results in Figure 1. We also report the correlation coefficient between the local survey responses and the aggregate ones in each panel.

$\frac{y_j^t - \bar{y}_j}{\sqrt{\text{Var}(y_j)}}$. We construct a continuous state variable \tilde{s}^t as a function of past terrorist attacks and past population support using the dynamic factor model given by

$$z^t = \tilde{s}^t \omega + \xi^t, \quad (\text{C.2})$$

and

$$\tilde{s}^t = \rho \tilde{s}^{t-1} + \alpha_0 + a_H^{t-1} \cdot \alpha_H + a_F^{t-1} \cdot \alpha_F + \eta^t. \quad (\text{C.3})$$

Here, a_F^{t-1} and a_H^{t-1} record attacks by Fatah and Hamas, respectively while the α_F and α_H weights the impact of those attacks. Including the attacks in the measurement of s^t reflects the strategic interdependence between the future states and today's actions as represented in Equation 3. Note that simply including attacks in the measurement model does not

Table C.2: Correlations among survey responses.

	Trust in Hamas	Trust in Fatah	Support for Hamas	Support for Fatah	Vote for Hamas	Vote for Fatah
Trust in Hamas	1.00	-0.19	0.77	-0.43	0.98	-0.72
Trust in Fatah	-0.19	1.00	-0.27	0.67	-0.54	0.92
Support for Hamas	0.77	-0.27	1.00	-0.57	0.89	-0.83
Support for Fatah	-0.43	0.67	-0.57	1.00	0.56	-0.32
Vote for Hamas	0.98	-0.54	0.89	0.56	1.00	-0.73
Vote for Fatah	-0.72	0.92	-0.83	-0.32	-0.73	1.00

presuppose their relationship in the first-stage regressions below. The α_i parameters can take on any value, including zero. Likewise, α_0 is a constant term, ω is a length-6 column vector of factor weights, and ρ is an AR(1) term on the state variable. Finally, $\eta^t \sim N(0, 1)$ and $\xi^t \sim N(0, \mathbf{1})$ are random perturbations, where $\mathbf{1}$ is the identity matrix.

The parameter vector $\Theta = (\omega, \rho, \alpha)$ can be estimated using maximum likelihood using the MARSS package for R (Holmes, Ward and Scheuerell 2018). Starting with an initial guess of the parameters $\hat{\Theta}_1$, the estimator relies on the following EM for iteration k :

1. **Expectation step:** Generate expected values of \tilde{s}^t using a Kalman filter and current givens $\hat{\Theta}_k$, z_t , and a^{t-1} . During this step missing values in z^t are also imputed by a Kalman filter.
2. **Maximization step:** Using the generated values of \tilde{s}^t and imputed z^t , maximize the multivariate normal log-likelihood. This step outputs $\hat{\Theta}_{k+1}$
3. Repeat the EM steps until no improvement in the log-likelihood is gained.

Table C.3: ML estimates for the factor model.

Equation	Variable	Estimate
Factor Weights (ω)	Trust in Hamas	-0.09
	Trust in Fatah	0.06
	Support for Hamas	-0.11
	Support for Fatah	0.10
	Vote for Hamas	-0.07
	Vote for Fatah	0.07
AR(1) term (ρ)	Lagged DV	0.99
Additional inputs (α)	Constant	-0.01
	Hamas attack	-0.28
	Fatah attack	1.05

The estimates of (ω, α) are reported in Table C.3, while the estimates of \tilde{s}^t are presented in Figure 2 in the main text. Notice that the observed indicators all load onto the dynamic

factor in the expected direction: pro-Hamas responses have negative weights and pro-Fatah responses have positive weights.

We also consider the robustness of this measurement model by comparing the estimated states \tilde{s}^t from the following models.

1. Main specification (described above)
2. Fix $\rho = 1$. Modify Eq. C.3 s.t.

$$\tilde{s}^t = \tilde{s}^{t-1} + \alpha_0 + a_H^{t-1} \cdot \alpha_H + a_F^{t-1} \cdot \alpha_F + \eta^t.$$

3. Estimated and homoskedastic variance in ξ^t . Modify Eq. C.2 s.t. $\xi^t \sim N(0, \sigma_\xi^2 \mathbf{1})$.
4. Estimated and heteroskedastic variance. Modify Eq. C.2 s.t. $\xi^{t,j} \sim N(0, \sigma_{\xi,j}^2 \mathbf{1})$, where $j = 1, \dots, 6$ indexes surveys
5. A model with $\alpha = 0$. Modify Eq. C.3 s.t.

$$\tilde{s}^t = \tilde{s}^{t-1} + \eta^t.$$

6. Remove the “plans to vote for” surveys (which start later than the other four). Modify y^t and z^t to only contain the first four survey responses.
7. Using only survey data from the Gaza Strip
8. Using only survey data from the West Bank

Note that each of the robustness checks considers one change to the main specification (i.e., these are not cumulative changes to the factor analysis). Fitting these models gives us eight specifications, each of which produces its own estimate of our continuous measure of relative popularity \tilde{s}^t . In Table C.4 we present the correlation matrix of these different approaches. Overall, we see that these methods all produce remarkably similar estimates. The biggest difference from the main model comes from heteroskedastic version, where a separate variance term is estimated for each of the six surveys. However, the correlation here is still roughly 0.9. As such, we conclude that these deviations from the main specification result in little change to \tilde{s}^t .

Of special note are Models 7 and 8. Here we consider the two Palestinian regions separately rather than just looking at the overall support across the two territories. In the former, we only track the polling results from Gaza, while in the latter we only use results from the West Bank. Overall, the underlying survey answers across these two areas correlate highly and tend to track each other, as in Figure C.2. These similarities appear in the measurement step as well. State variables measured using only Gaza or West Bank opinion both correlate with the main specification at about 0.99. This similarity gives us additional confidence that using the overall-combined opinion data is reflective of overall Palestinian support and not covering up or removing interesting regional variation.

Table C.4: Correlations across measurement model specifications.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Model 1	1.00	1.00	0.99	0.87	1.00	0.99	0.99	0.99
Model 2	1.00	1.00	0.99	0.86	0.99	0.99	0.99	0.98
Model 3	0.99	0.99	1.00	0.88	1.00	0.98	0.97	0.99
Model 4	0.87	0.86	0.88	1.00	0.88	0.91	0.87	0.84
Model 5	1.00	0.99	1.00	0.88	1.00	0.99	0.98	0.99
Model 6	0.99	0.99	0.98	0.91	0.99	1.00	0.99	0.97
Model 7	0.99	0.99	0.97	0.87	0.98	0.99	1.00	0.95
Model 8	0.99	0.98	0.99	0.84	0.99	0.97	0.95	1.00

D First-stage robustness

D.1 Control variables and specification choices

We want to ensure that the relationship between attacks and relative popularity is not driven by alternative factors. Our first concern is to ensure that the γ estimates are robust to some specification choices such as the decision to include or exclude the state-action interactions or to the possibility of time-varying effectiveness of each group’s ability to use terrorism. Additionally, we want to be sure that variables omitted from the first-step AR-1 model, such as attitudes about violence against Israelis, Israeli aggression, or other economic and political factors, are not biasing our estimates of the transition parameters, γ . To measure attitudes and economic factors, we return to the surveys and use additional dynamic factor models to build latent variables that capture Palestinian attitudes about violence towards Israelis and their employment status. The other controls that we consider are directly measurable.

For the two latent variables (describing attitudes toward violence and unemployment), we again aggregate survey data using dynamic factor models. Both models use the same basic specification described in Equations C.2 and C.3, but with different surveys forming z^t and as such a different latent variable output (i.e., \tilde{v}^t and \tilde{u}^t for attitudes for violence and unemployment, respectively rather than \tilde{z}^t). For attitudes towards violence, we use survey questions that record 26 different responses to various aspects of the conflict/peace process. From the surveys run by JMCC we include:

- 4 responses about attitudes to a two-state solution
- 2 responses about attitudes towards peace negotiations
- 2 responses about attitudes towards military operations against Israeli targets
- 2 responses about attitudes towards suicide bombings against Israeli civilians
- 4 responses about optimism/pessimism regarding a peaceful settlement with Israel
- 3 responses about attitudes towards the Oslo peace process
- 2 responses about attitudes towards the 2nd Intifada

Table D.1: ML estimates for latent support for violence.

Variable	Est.
% Supporting two-state solution	-0.15
% Supporting a one shared state solution	0.02
% Supporting a one Islamic state solution	0.15
% Saying there is no solution	0.14
% Supporting a peace process	-0.17
% Opposing a peace process	0.17
% Supporting military action against Israel	0.17
% Opposing military action against Israel	-0.18
% Supporting suicide bombings	0.17
% Opposing suicide bombings	-0.17
% Very optimistic about peace	-0.16
% Optimistic about peace	-0.17
% Pessimistic about peace	0.13
% Very pessimistic about peace	0.17
% Strongly support Oslo	-0.11
% Support Oslo	-0.10
% Oppose Oslo	0.14
% Support the Intifada	0.09
% Oppose the Intifada	-0.09
% Who think peace is dead	0.10
% Who think the peace process is stalled	-0.03
% Who think peace is alive	-0.13
% Support armed attacks generally	0.17
% Support armed attacks against civilians	0.16
% Support armed attacks against soldiers	0.13
% Support armed attacks against settlers	0.12

- 3 responses about whether the current peace process is alive, dead, or unclear.

Additionally, we add 4 responses from surveys by PCPSR that record support for armed attacks against

- Israel generally
- Israeli civilians
- Israeli soldiers
- Israeli settlers in the West Bank.

Among these last 4 we note that support for armed attacks against Israel generally never drops below 20%. Likewise, the median support for attacks against soldiers/settlers/civilians is 91/88/49% and the IQR is 87–92/83–91/43–54%.

Many of these variable are correlated, but we avoid perfect correlations between combinations of factors by the virtue of “don’t know,” “no answer,” and similar non-answers. The high correlations across answers and across questions provide strong evidence that these

responses can be reduced into a latent measure. The factor weights are reported in Table D.1. All the surveys load in the expected way where surveys that should correlate with approval towards violent tactics load positively and surveys that correlate with approval of peaceful tactics and negotiated settlement load negatively. Overall, this give us strong assurance that the latent variable captures the Palestinian public’s underlying attitudes toward violence against Israel at any given month.

For the unemployment latent variable we combine four survey responses:

1. % of respondents telling pollsters they are unemployed to JMCC pollsters
2. % of respondents telling pollsters they are unemployed to PCPSR pollsters
3. Estimated true unemployment by PCPSR
4. Unemployment rate reported in Labor Force Surveys published by the PCBS

The results from the dynamic factor analysis measuring unemployment are reported in Table D.2. Here we see that all the unemployment rates load onto the latent dimension in the same direction, but with different weightings.

Table D.2: ML estimates for latent unemployment conditions.

Variable	Est.
Self reported unemployment rate (JMCC)	0.18
Self reported unemployment rate (PCPSR)	0.17
Estimated unemployment rate (PCPSR)	0.08
Estimated unemployment rate (PCBS)	0.21

The robustness checks that include more covariates or change model specification are reported in Table D.3 (which should be compared to Table 1 in the main text). The main thing to note is that the effects of attacks on relative popularity are similar across all specifications. In Model 1, we show that our conclusions persist even when we do not control for state-attack interactions.

In Model 2, we want to justify the time-invariant estimation of each group’s ability to use terrorism to shift public opinion. It is reasonable to suspect that each side’s ability to outbid may shift over time, and we want to know if our time-invariant specification fits the data. Note, that the interactions of lagged attacks and lagged states in the main models allow for some heterogeneity in each group’s ability estimates. Over the range of the state variables we find little evidence for strong heterogeneity. To directly check for time-varying effects in each group’s effectiveness, we interact a linear time trend with the main attack variables. Here we see that the main point estimates do not change much. Additionally, while we see slight evidence to suggest that both groups are getting more effective over time, neither of these interaction terms is statistically significant. A Wald test fails to reject the joint null

Table D.3: Robustness checks for the first-stage model: Specification changes

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Hamas attacks	-0.21 (0.03)	-0.17 (0.08)	-0.20 (0.04)	-0.23 (0.03)	-0.21 (0.02)	-0.19 (0.02)	-0.34 (0.05)	-0.28 (0.05)
Hamas attacks \times time		-0.0001 (0.0004)						
Fatah attacks	1.00 (0.08)	0.81 (0.18)	1.06 (0.05)	1.04 (0.04)	1.05 (0.04)	1.05 (0.03)	1.14 (0.10)	1.27 (0.15)
Fatah attacks \times time		0.003 (0.002)						
Δ Lag state	0.34 (0.04)	0.33 (0.05)	0.31 (0.04)	0.25 (0.04)	0.19 (0.04)	0.17 (0.04)	0.11 (0.04)	0.14 (0.04)
Δ unemployment			-0.23 (0.09)	-0.12 (0.08)	-0.08 (0.05)	-0.06 (0.06)	-0.07 (0.07)	-0.15 (0.07)
Δ support for violence				-0.25 (0.06)	-0.32 (0.05)	-0.35 (0.05)	-0.29 (0.06)	-0.19 (0.05)
Time		0.0004 (0.0002)						
Second <i>Intifada</i>					-0.17 (0.04)	-0.17 (0.04)	-0.18 (0.03)	0.43 (0.11)
Time since last election						-0.003 (0.001)	-0.005 (0.001)	-0.005 (0.001)
Palestinian fatalities by Israel							0.01 (0.01)	0.02 (0.01)
Corruption								-0.48 (0.10)
Infant mortality								-0.04 (0.01)
Corruption \times post-2006 election								0.30 (0.16)
Infant mortality \times post-2006 election								0.02 (0.01)
Constant	-0.01 (0.02)	-0.07 (0.05)	-0.01 (0.02)	0.003 (0.02)	0.02 (0.02)	0.08 (0.03)	0.11 (0.04)	0.53 (0.16)
State-attack interactions	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
T	298	298	298	298	298	298	212	200
adj. R^2	0.721	0.729	0.743	0.782	0.812	0.821	0.845	0.857
$\hat{\sigma}$	0.184	0.181	0.176	0.162	0.151	0.147	0.137	0.125

Note: Newey-West standard errors in parentheses.

hypothesis that both interactions are zero. Overall, this provides reasonable evidence that the time-invariant effectiveness imposed by the main model is not doing too much violence against the data.

In Models 3 and 4, we add in the latent variables for economic and attitudes about violence, respectively. Unsurprisingly, as the Palestinian public becomes more accepting of violence, there is a shift in popularity toward Hamas. Likewise, poor economic conditions favor Hamas’s popularity. In Models 5 and 6, we consider some political context with an indicator for whether the Second Intifada is ongoing and the time since the last Israeli election. In Model 7, we control for the logged number of Palestinian fatalities due to Israeli forces, as recorded by B’Tselem, an Israeli human rights organization, to proxy for aggression by the Israeli government against Palestinians.⁴

Finally, in Model 8, we consider how the groups may compete in terms of public goods or quality of governance. Since groups may compete in ways that either complement or substitute for violence, we would like to know that the first-stage estimates are not the result of these alternative venues of competition that exist outside of traditional outbidding. Ideally, we would have group-specific, monthly-level data on important public goods provisions along with perceptions of corruption. We know of no such group-specific measure of corruption that varies over time. In contrast, Tokdemir and Akcinaroglu (2016) and Stewart (2018, Appendix 5) investigate whether anti-government or rebel groups provide public goods. Using a variety of media, academic, and archival sources, both code whether Fatah and Hamas provide public goods at the yearly level—Stewart (2018), but not Tokdemir and Akcinaroglu (2016), includes the Palestinian Authority as a separate group. Specifically, Tokdemir and Akcinaroglu (2016) code an indicator if a group in a given year provides “free education, food, security, or health services to their constituency” (270). Stewart (2018) has two indicator variables: the first denotes whether a group provides educational public goods *and* health public goods, and the second denotes whether a group provides educational public goods *or* health public goods. The more conservative measure is their preferred measure. However, neither records any variation in public goods provisions over time for either Hamas or Fatah. That is, given a specific measure, if a group provides a public good in one year, then they provide that public good in all years recorded in that paper’s data. While some measures pick up on differences across groups (e.g., Stewart’s less conservative measure finds that Hamas provides public goods but Fatah does not), there are no differences over time.⁵

⁴Data are only available after the start of the Second Intifada. Available at <https://www.btselem.org/statistics>

⁵There are some differences across datasets. E.g., Tokdemir and Akcinaroglu (2016) code Fatah and Hamas as always providing public goods, but Stewart (2018) codes both groups as never providing public goods in their most conservative measure. In their less conservative measure, Hamas provides public goods, but Fatah does not. In both of Stewart’s measures, the Palestinian Authority is always providing public goods.

As another validation check, we also explored the data in Wagstaff and Jung (2020) who scrape LexisNexis for English-language print news coverage and subsequently count the words that correspond with service provision that are used in articles also mentioning a terrorist group in a specific year. For example, educational public goods are associated with the words “school” and “teacher,” and health public goods are associated with the words “hospital” and “clinic.” The code book mentions several important caveats when using their data. Specifically, “our measure is an imprecise indicator of the precise level of service provision and, as such, should not be used to draw conclusions about exacting levels, nor should it be used to compare across small numbers of groups” (Heger, Jung and Wagstaff 2018, 7). In addition, “certain groups in the dataset concentrate their violence against targets that may correlate with some service sectors. . . this is most likely the case for the religion, security, and infrastructure [public goods]” (Heger, Jung and Wagstaff 2018, 8). Given the first concern, we decided to focus on their dichotomous measure of public goods that codes whether or not a word is mentioned in the same article as a group in a specific year. Given the second concern, we decided to focus on educational and health public goods, which are the key dependent variables in Wagstaff and Jung (2020), which also matches Stewart (2018). As in the measures above, we find little variation over time. Specifically, for every year and each group (Fatah and Hamas), there is at least one article using a word for an educational public good (school or teacher) and one article using a word for a health public good (hospital or clinic). When we look at the component words, school, teacher, and clinic always appear in an article about each group in every year. The word hospital is rarer. It only appears in three years (1999, 2000, and 2004) and is associated with Hamas only.⁶ Overall, this data also supports our conclusion that data on whether these groups provide public goods does not vary over time in our sample period.

As a second best approach we use two measures from the World Bank: control of corruption and infant mortality. Both of these are pooled measures for both the West Bank and Gaza Strip.⁷ The former is constructed based on public perceptions of how well the Palestinian Authority (PA) fights, deters, or tolerates corruption. We multiply the recorded measure by -1 so that larger, more positive numbers reflect less control (i.e., more perceived corruption). The latter is the ordinary infant mortality rate and is used to proxy the quantity and quality of public services provided by the PA. In both cases, we interact the measure with an indicator for all dates after the 2006 election that saw Hamas enter the government. This strategy allows us to see how these measures affect relative popularity

⁶A LexisNexis (now called NexisUni) search for the terms “hospital” and “Hamas” during the year 1999 returned 8 articles, but these were about Hamas leaders falling ill and needing hospital treatments (February and December) or included a quote from a Hamas lawyer requesting a Hamas leader be freed from custody to receive medical treatment at a hospital (September). Similar issues appear in 2000, and 2004 has more than 125 articles that meet the search criteria.

⁷Additionally, both are measured at the yearly level. Control of corruption is estimated by the World Bank for the Palestinian Territories every other year from 1996–2002 and then annually from 2003 onward. Infant mortality is measured in deaths per 1,000 births and measured annually from 1975 onward.

both before Hamas enters government (i.e., the PA is entirely dominated by Fatah) and after.

What we see in this model is that higher perceptions of corruption are associated with shifts in popularity toward Hamas. This result makes sense as Fatah is typically the actor seen as more corrupt and is generally associated with corruption in the PA (Milton-Edwards and Farrell 2010). Note that after the 2006 election the effect of corruption attenuates; however, the overall effect is still negative suggesting that Fatah is still punished more when control of corruption is perceived as weak. These same trends appear in infant mortality, with higher levels being associated with shifts in relative popularity toward Hamas. Likewise, while the effect attenuates after Hamas enters government, Fatah still bears the brunt of increases in infant mortality. Most importantly, the main estimates of $\gamma_{F,1}$ and $\gamma_{H,1}$ are relatively unchanged, so we see that Fatah is still more effective at using violence to increase popularity than Hamas.

D.2 Measurement of the actions

Next, we want to ensure that our results are not dependent on our binary coding of attacks or the type of attacks considered. Regarding the former, we consider four alternative measurements for attacks: counts, binary with fatalities as a control, only fatalities, and fatalities per attacks. These are the first four models in Table D.4. In these models, the main results survive: (i) violence by actor i makes i more relatively popular in the following month and (ii) the effects of Fatah’s violence are larger in magnitude than the effects of Hamas’s violence.⁸

Regarding the latter, in the last two models in Table D.4, we use the GTD’s attack target information to subset the attacks into two different types: attacks against civilians or not, where the latter category includes attacks against the military, government, or a non-state actor. Here again we find that Fatah attacks are more effective at boosting public opinion than Hamas attacks regardless of the type of attack. The overall consistency across the six model specifications is reassuring as we do not want the first stage to be dependent on the specific measure of violence. Overall, these various robustness checks of the AR(1) model from Table 1 provide confidence in our using it in the first stage of the analysis as the estimates of γ .

D.3 Instrumental variable analysis

The next set of robustness checks considers the potential endogeneity of attack decisions. While we attempt to address some of these concerns with the control variables above, we understand that there may be additional unobserved confounders or other forms of endogeneity bias that may affect the estimates of γ . There are four potentially endogenous

⁸The only circumstances where we fail to reject the null of equally effective groups are when $\bar{s}^t < -8.7$ with the measures in Models 1, 3, and 4 from this table. So if Hamas is *extremely* popular (about 5% of the data), the two groups might be equally effective at moving their relative popularity level according to some of these alternative specifications.

Table D.4: Robustness checks for the first-stage model: Measurement changes.

Dependent variable: Attack measure:	Counts			Binary & fatalities			Fatalities			Δ State		
	Counts	Binary & fatalities	Fatalities	Fatalities/attack	Civilian targets	Non-civilian targets	Fatalities/attack	Civilian targets	Non-civilian targets	Fatalities/attack	Civilian targets	Non-civilian targets
Hamas attacks	-0.03 (0.01)	-0.20 (0.03)		-0.02 (0.01)	-0.15 (0.03)	-0.05 (0.04)						
Fatah attacks	0.42 (0.05)	1.03 (0.04)		0.67 (0.07)	1.03 (0.05)	0.98 (0.04)						
Hamas fatalities		0.001 (0.001)	-0.01 (0.002)									
Fatah fatalities		0.01 (0.01)	0.31 (0.05)									
Second <i>Intifada</i>	-0.19 (0.04)	-0.17 (0.04)	-0.19 (0.04)	-0.22 (0.05)	-0.16 (0.04)	-0.20 (0.04)						
Δ unemployment	-0.08 (0.06)	-0.06 (0.06)	-0.06 (0.06)	-0.07 (0.06)	-0.09 (0.06)	-0.04 (0.05)						
Δ support for violence	-0.32 (0.06)	-0.35 (0.05)	-0.32 (0.06)	-0.32 (0.07)	-0.35 (0.06)	-0.32 (0.06)						
Time since late election	-0.003 (0.001)	-0.003 (0.001)	-0.003 (0.001)	-0.003 (0.001)	-0.004 (0.001)	-0.004 (0.001)						
Δ Lag state	0.09 (0.05)	0.17 (0.04)	0.14 (0.06)	0.15 (0.07)	0.13 (0.05)	0.20 (0.06)						
Constant	0.08 (0.03)	0.08 (0.03)	0.07 (0.03)	0.07 (0.04)	0.09 (0.04)	0.07 (0.03)						
State-attack interactions	Yes	Yes	Yes	Yes	Yes	Yes						
T	298	298	298	298	298	298						
adj. R^2	0.52	0.82	0.4	0.44	0.56	0.55						
$\hat{\sigma}$	0.24	0.14	0.26	0.26	0.23	0.23						

Note: Newey-West standard errors in parentheses

variables we need to consider: Hamas attacks, Fatah attacks, and the interactions of attack choices with the lagged state variables. We follow the research design proposed in Köning et al. (2017) and consider deviations from average rainfall in the West Bank and Gaza Strip as plausible instrumental variables (IVs) for Hamas and Fatah attack decisions. We also interact these rainfall variables with an additional lag of the state variable (continuous relative popularity) to instrument for the interaction terms. Of course, there has been a lot of push back against the idea of using rainfall as an instrument (e.g., Mellon 2022; Sarsons 2015), but it remains a standard instrument in the political economy of conflict literature (e.g., Brückner and Ciccone 2011; Miguel, Satyanath and Sergenti 2004). The goal of this section is to illustrate that our baseline estimates of $\gamma_{F,1}$ and $\gamma_{H,1}$ in Table 1 are similar in size and magnitude to those from when we use extreme rainfall as an instrument on the groups’ attack decision. This gives us greater confidence in using those estimates moving forward to the second stage in which estimate the groups’ payoff parameters.

Extra heavy rainfall in the West Bank (Gaza), we suspect, reduces Fatah’s (Hamas’s) propensity to attack in the following month, while perhaps raising Hamas’s (Fatah’s) willingness to attack as they may see an opportunity to attack without a response. This reduction in attacks can come from either inclement conditions discouraging attacks directly or changes in the local agricultural market that make terrorism a less attractive option for potential recruits (Köning et al. 2017). The validity of the instrument depends on rainfall only affecting the relative popularity of the two groups through their attack decisions, and it is difficult to imagine other ways in which rain would directly shift the relative popularity of these groups. The main threat to validity would be any effect that flows from rainfall to the economic conditions and from there to relative popularity of the actors. As such, we also consider models using the control variables from Table D.4, which includes Palestinian unemployment.

To measure rainfall, we use data from Global Precipitation Climatology Centre (GPCC) as provided by the NOAA Physical Science Lab (Schneider et al. 2020). These data provide monthly rainfall totals on a 0.25×0.25 spatial grid. We aggregate these for the Gaza Strip and West Bank separately and record how many standard deviations each observed month is from the historical average rainfall for that month. These data are lagged one month behind the attack variables and interacted with the second lag of the state variable to instrument for the state-action interactions.

We also acknowledge that rainfall is likely to be a weaker instrument here than in other contexts for two reasons. First, the small size of Israel means that our two instruments are correlated. Deviations from monthly average rainfall in the Gaza Strip and the West Bank have a correlation coefficient of 0.89. In comparison, Köning et al. (2017) study the Democratic Republic of Congo and use the yearly average of rainfall in each group’s homeland as an instrument for the group’s propensity to attack. The Democratic Republic of Congo is more than 100 times the size of Israel; it is 400 times the size of Palestine.

As such, there is not as much variation in weather conditions across the two regions in our context. Second, a key mechanism connecting rainfall to violence is opportunity costs: extreme weather depresses agricultural productivity and therefore decreases the reservation wages that would be derived via agricultural activities. This makes it more difficult for groups to recruit. It is unclear how prominent this mechanism is in the Palestinian case, however. On the one hand, agriculture makes up a small percentage of Palestinian GDP, ranging from 10% in 1999 to 5% in 2010 according to the Center for Economic and Policy Research (CEPR) (2012). With such a small proportion of the formal economy devoted to agriculture, it could be the case that rainfall only has weak effects on the opportunity costs of violence. On the other hand, agriculture makes up an important part of the informal economy. According to the CEPR (2012), 13.4% of the labor force and over 90% of all informal employment is absorbed by the agricultural sector. Nonetheless, it is difficult to get reliable estimates of the size of the informal economy (Shabaneh 2008), although Awad and Alazzeah (2020) estimate it to be 28% of Palestinian GDP in 2010.

As an additional instrumental variable we also consider an indicator for whether Palestinian Islamic Jihad (PIJ) attacks, lagged a month behind Fatah and Hamas attack decisions and measured from the GTD. This variable leads to stronger first-stage relationships, although we are likely still dealing with a weak instrument. For validity, it must be the case that PIJ attacks only affect the relative popularity between Hamas and Fatah through their own attack decisions. In this context it means that a Palestinian citizen, upon observing an attack from the PIJ, does not update their relative preferences between Hamas and Fatah unless one of those actors also decided to commit violence.⁹ Including this additional instrument allows for overidentification tests which provide some evidence of instrument validity.

The structural equation is the ECM form of Eq. 8, given as

$$\tilde{s}^t - \tilde{s}^{t-1} = \gamma_0 + \tilde{\gamma}_1(\tilde{s}^{t-1} - \tilde{s}^{t-2}) + \gamma_{H1}a_H^{t-1} + \gamma_{H2}(\tilde{s}^{t-1} \times a_H^{t-1}) + \gamma_{F1}a_F^{t-1} + \gamma_{F2}(\tilde{s}^{t-1} \times a_F^{t-1}) + \xi^t.$$

Let $\mathbf{y}^t = (a_H^{t-1}, a_F^{t-1}, \tilde{s}^{t-1} \times a_H^{t-1}, \tilde{s}^{t-1} \times a_F^{t-1})'$ be the column vector of endogenous variables, then the first-stage equations are given as

$$\mathbf{y}^t = \pi_0 + \pi_1(\tilde{s}^{t-1} - \tilde{s}^{t-2}) + \sum_{\ell=2}^L \left(\pi_{\ell} r^{t-\ell} + \pi_{\ell+L-1} (r^{t-\ell} \times \tilde{s}^{t-2}) \right) + \sum_{m=2}^M \lambda_{m-1} a_P^{t-m} + \zeta^t,$$

where r denotes the extreme rainfall measure and a_P denotes the PIJ attack indicator. Additionally, each parameter and the error term ζ^t are length-four column vectors. For the five models, presented in Table D.5, we have $L = 2$ in Models 1-4 and $L = 5$ in Model 5.

⁹We can think of this as a kind of independence of irrelevant alternatives assumption. In this sense it matches the identification assumptions used in multinomial logit models of group support (e.g., Jaeger et al. 2015).

Likewise, $\lambda = 0$ in Models 1 and 3, $M = 2$ in Models 2 and 4, and $M = 5$ in Model 5. The additional controls enter both the structural and first-stage equations and instrument for themselves.

Table D.5: Robustness checks for the first-stage model: Instrumental variables analysis.

Dependent variable:	Δ State				
	(1)	(2)	(3)	(4)	(5)
Hamas attacks	-0.16 (0.19)	-0.17 (0.20)	-0.23 (0.17)	-0.26 (0.21)	-0.25 (0.18)
Fatah attacks	1.39 (0.64)	1.62 (0.66)	1.36 (0.53)	1.82 (0.62)	1.70 (0.80)
Constant	-0.02 (0.07)	-0.04 (0.08)	0.11 (0.03)	0.07 (0.05)	0.07 (0.04)
Interactions	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes
T	298	298	298	298	295
$\hat{\sigma}$	0.229	0.240	0.189	0.250	0.264
Cragg and Donald statistic	0.001	0.316	0.001	0.412	0.974
Number of instruments	4	5	4	5	20
Sargan-Hansen p value		0.758		0.632	0.141

Note: Estimates from Fuller’s k estimator with parameter 1. Bekker’s robust standard errors in parentheses

To consider the strength of the instruments we look at both the first-stage F statistics and the Cragg-Donald statistics, which is a multivariate measure of instrument strength when there is more than one endogenous variable. For the models using only rainfall, the first-stage F statistics range from about 0.2 to 3.3, which is quite weak. Likewise, rainfall seems to do better at predicting Hamas attacks than Fatah attacks, with extreme rain in Gaza reducing the probability that Hamas attacks the following month. Adding in PIJ attacks improves these F statistics so that they range from 0.5 to 6.7. Additionally, while the instruments are still stronger for Hamas attacks we can note that a PIJ attack is a positive and statistically significant predictor of both Hamas and Fatah attacks in the following month. Additionally, the Cragg-Donald statistics suggest that the instruments are much stronger when PIJ attacks are included.

While these values suggest that there are first-stage relationships here, the values are overall low. When weak instruments are present, two-stage least squares (2SLS) estimation of the IV model can be biased toward OLS even in large samples. In contrast, k -class estimators like the limited information maximum likelihood (LIML) estimator are less biased in overidentified models with weak instruments but this comes at the cost of fatter tails and an increased IQR in the sampling distribution. To make the best of our weak instruments we use a Fuller estimator that modifies the regular LIML estimator in a way that usually

reduces both bias and mean-squared error (Stock and Yogo 2005). Likewise, Stock, Wright and Yogo (2002) show that the rule-of-thumb first-stage F statistics can be much lower for the Fuller estimator (relative to 2SLS) to reduce the estimator’s bias relative to the bias from OLS, although they don’t provide any critical or suggested values for the case of four endogenous regressors. For this set of models we use Bekker standard errors which are consistent in the face of heteroskedasticity and many weak instruments (Bekker 1994).

To take advantage of the weak instrument asymptotics of the Fuller estimator, we fit an additional model that uses the rainfall deviations for the four months before the attack decision along with their interactions with the state variable. Likewise, this model also uses PIJ attack indicators for four months prior to the attack decision for a total of 20 instruments.

The IV estimates are presented in Table D.5; these estimates should be compared to those the first two rows of Table 1 in the main text. Models 1 and 2 have no additional control variables, while Models 3–5 use the same controls as the models in Table D.4. We generally see that the main estimates largely agree with the main models in terms of direction and magnitude. Likewise, Fatah is still more effective than Hamas at using attacks to increase its relative popularity. The fact that the estimates remain largely unchanged as we increase the number of instruments is good news given that the Fuller estimator is consistent in the number of weak instruments. Additionally, with the overidentified models we can use Sargan’s test of the null hypothesis that the instruments are valid. We fail to reject this null.

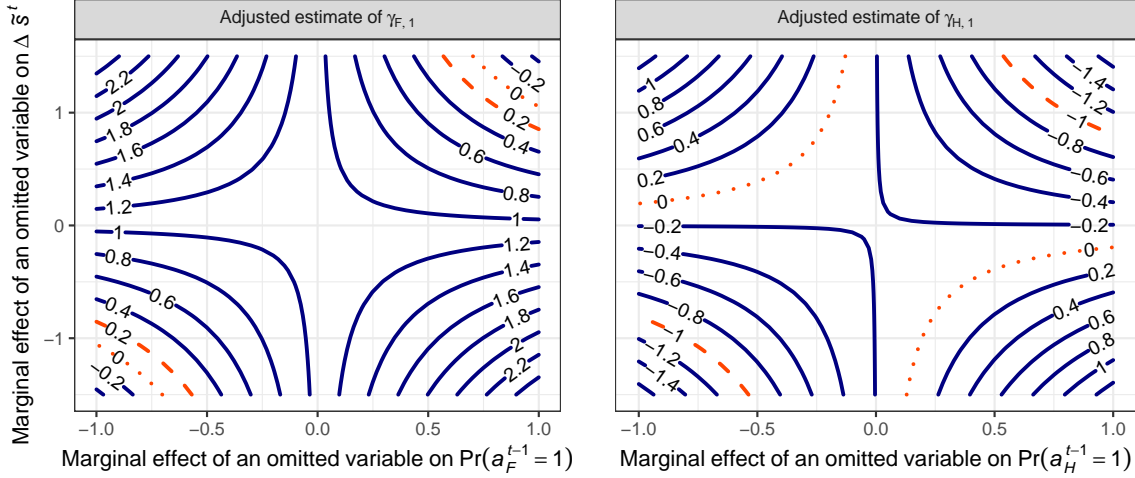
D.4 Sensitivity analysis

In this appendix, we conduct a formal sensitivity analysis with the following two questions in mind: how strong would an unobserved confounder need to be in order to (a) make the first-stage estimates symmetric (i.e., $\gamma_{H,1} = -\gamma_{F,1}$) or (b) suggest that actors are not effective at using violence $\gamma_{i,1} = 0$? The first question is particularly important because this asymmetry is driving the discouragement effect in Figures 5 and 6.

To answer these questions we use techniques from Cinelli and Hazlett (2020) who provide a suite of tools for assessing the sensitivity of regression coefficients to unobserved confounders. The results of this analysis are presented in Figure D.1, where we consider the sensitivity of the estimates in model 6 from Table D.3. We choose this model to consider as it already contains the most control variables without losing additional observations over the other ECM specifications.

Figure D.3 demonstrates that our first-stage results are robust. In the left-hand plot, we see that an unobserved confounder of Fatah attacks (a_F^{t-1}) would need to have an effect on the probability of a Fatah attack of at least positive (negative) 0.5 while simultaneously having a marginal effect on $\Delta \tilde{s}^t$ of at least positive (negative) 0.9 for the asymmetry result to disappear (denoted as the orange dashed-line). Regarding the former, note that 0.5 in

Figure D.1: Sensitivity of $\hat{\gamma}_{F,1}$ and $\hat{\gamma}_{H,1}$ to unobserved confounders.



Note: Contour lines reflect the adjusted estimate for $\gamma_{i,1}$ if an omitted variable with the specified marginal effects on $a_{i,t-1}^{t-1}$ and $\Delta \tilde{s}^t$ is added to the model. The orange, dashed line denotes the point where the model becomes symmetric ($\hat{\gamma}_{i,1} = -\hat{\gamma}_{j,1}$). The orange, dotted line denotes the point where i is no longer competitive ($\hat{\gamma}_{i,1} = 0$).

this case means a 50 percentage point increase in the probability that Fatah attacks. This is a sizable effect and would be an incredible predictor of Fatah violence, but how likely is it that such a confounder exists? Regressing a_F^{t-1} on the other covariates tells us that the strongest observed marginal effect is about 0.1 (Δ unemployment), followed by the effect of a Hamas attack, which raises the probability of a Fatah attack by about 0.05. Regarding the latter, note that in Table D.3 the only marginal effect size greater than 0.9 is the effect of Fatah attacks. All the other effects have a magnitude less than 0.5. In other words, an unobserved confounder needs to have about *10 times* the observed effect of Hamas violence on the probability of a Fatah attack while simultaneously affecting $\Delta \tilde{s}^t$ at the same magnitude of a Fatah attack. While the latter is quite plausible, the former is a very tall order. As seen in the figure, it would take an even stronger unobserved confounder for the estimated $\hat{\gamma}_{F,1}$ to be adjusted to 0.

Turning to the right-hand plot, we see that the estimate of $\gamma_{H,1}$ is possibly more sensitive to unobserved confounders, but overall we are still encouraged. First, we consider whether the asymmetry in effectiveness would be the result of possible attenuation bias in $\hat{\gamma}_{H,1}$. That is, there might be unobserved confounders that drive our estimate of $\hat{\gamma}_{H,1}$ closer to zero, which could explain the asymmetry in Table 1. As in the Fatah case, an omitted variable would need to change the probability of a Hamas attack of at least positive (negative) 0.5 while simultaneously having a marginal effect on $\Delta \tilde{s}^t$ of at least positive (negative) 0.9 to make the estimates symmetric. As previously mentioned, the latter magnitude is rare in our data, but perhaps plausible. The former, however, is still quite large. The best predictor of Hamas attacks among our observed covariates is the Second Intifada and even this only

raises the probability of a Hamas attack by about 0.28. As such, an omitted variable would need to simultaneously increase the probability of a Hamas attack by *twice* that of the Second Intifada and $\Delta\tilde{s}^t$ by the same magnitude of a Fatah attack in order to increase the size of $\hat{\gamma}_{H,1}$ enough to make these estimates symmetric. Overall, this may be more plausible than the conditions on $\hat{\gamma}_{F,1}$, but they are still quite strong.

Perhaps more concerning are the requirements on an observed confounder to shift $\hat{\gamma}_{H,1}$ to 0. Here, we see that the omitted variable needs to increase (decrease) the probability of a Hamas attack by only 0.1 while decreasing (increasing) $\Delta\tilde{s}^t$ by at least 0.1. In this case, a variable that has the same effect of at least 1/2 a Hamas attack on $\Delta\tilde{s}^t$ (≈ -0.1) and 2/3 of the effect of a Fatah attack on the probability that Hamas attacks (≈ 0.1) would be enough to move $\hat{\gamma}_{H,1}$ to 0.

To add additional context and further explore the sensitivity of $\hat{\gamma}_{H,1}$ in particular, we consider an additional question: How strong of a relationship with the unobservables is needed to shift $\hat{\gamma}_{i,1}$? To answer this question we consider robustness values (RVs) from Cinelli and Hazlett (2020), which are expressed in terms of partial R^2 . Recall that the partial R^2 of a variable a_i^{t-1} from a regression of $\Delta\tilde{s}^t$ on a_i^{t-1} , a_j^{t-1} , and other covariates x^t is given as

$$\text{partial } R^2(a_i^{t-1} | \Delta\tilde{s}^t, a_j^{t-1}, x^t) = \frac{\hat{\nu}'_0 \hat{\nu}_0 - \hat{\nu}' \hat{\nu}}{\hat{\nu}'_0 \hat{\nu}_0},$$

where the $\hat{\nu}$ and $\hat{\nu}_0$ are the residual vectors from regressions with and without a_i^{t-1} , respectively. As such, the partial $R^2(a_i^{t-1} | \Delta\tilde{s}^t, a_j^{t-1}, x^t)$ tells us the proportion of variation in $\Delta\tilde{s}^t$ that is explained by only a_i^{t-1} and not the other covariates. We present these results in Table D.6.

Table D.6: Robustness values (RVs) and sensitivity of $\gamma_{i,1}$.

	Fatah	Hamas
RV-symmetry	0.70	0.94
RV-null	0.76	0.44
R^2 other covariates	0.03	0.16
Max. partial R^2 of other covariates	0.01	0.05

The first two rows of Table D.6 show the RVs for $\gamma_{F,1}$ and $\gamma_{H,1}$. In the first row, we see that omitted confounders would have to explain at least 70% (94%) of the remaining variance in both a_F^{t-1} (a_H^{t-1}) and $\Delta\tilde{s}^t$ in order for the first-stage results to be symmetric. In the second row, we see that these values are 76% (44%) to shift $\hat{\gamma}_{F,1}$ ($\hat{\gamma}_{H,1}$) to 0. As above, we want to place these results in context of how likely confounders of these strengths may be. Obviously, we can never rule out the existence of such confounders, but a few points are encouraging. First, we note that all of the other observed covariates in this model only explain 3% and 16% of the variance in a_F^{t-1} and a_H^{t-1} , respectively. Second, the

covariates with the largest partial R^2 for a_F^{t-1} and a_H^{t-1} are only 1 and 5%, meaning that the unique contribution of each variable to variance explained is quite small. What this means as we would need to consider confounders that are many times stronger than any and all of observed variables to change the main results. As an example, even the most sensitive result ($\gamma_{H,1} < 0$) would require confounders that are twice as good the Second Intifada at predicting Hamas attacks (partial $R^2 = 0.05$) while being nearly $2\times$ stronger than Hamas attacks (partial $R^2 = 0.26$) at explaining changes in relative popularity. Such confounders may exist, but they would need to dwarf the explanatory power of the included variables. For the other three cases considered here, the orders of magnitude needed to create a symmetric or null result are large enough to strain the imagination.

Of course, the weakness of the above analysis is that, like nearly all formal sensitivity analysis, we can only focus on one variable at time (i.e., either a_F^{t-1} or a_H^{t-1} , but not both simultaneously). As such, the unobserved confounders considered here are only those that correlate with the variable under consideration but are orthogonal to the other (along with all other included variables). In the context of attack decisions, this may be a serious limitation. The instrumental variable analysis in Section D.3, however, does allow for these more general forms of unobserved confounding. We find that the main results persist there as well.

D.5 Is an attack an attack an attack?

While the results in Table D.4 indicate that the main first-stage results are robust to the alternative measures of attacks, we may be concerned about differences in terrorist attacks over time. For example, is it the case that attacks in 1994 are roughly similar to attacks in 2006? If attacks are getting deadlier over time, perhaps as the result of improved technology or other changes, then we may worry about how these differences affect the overall interpretation of the model. Relatedly, if Hamas attacks are particularly lethal during the Oslo era, then maybe the so-called Oslo lull is a mirage. We consider this question in Table D.7 where we regress fatalities against attacks interacted with measurements of time for each actor.¹⁰ Specifically, we consider both a linear time trend and dummies for the three main eras considered in the main text.

There are three takeaways we want to point out in this table. First, neither actor's attacks appear to have increasing lethality over time, as the coefficients on attacks interacted with a time trend are very small. For Fatah, the interaction effect is statistically indistinguishable from 0, while for Hamas the effect size suggests that Hamas attacks have slightly fewer fatalities on average over time with the marginal effect slightly decreasing each month. At first glance here we do not see anything to suggest that Hamas attacks were particularly lethal during the Oslo era.

¹⁰An alternative ANOVA analysis based on mean fatalities/attacks produces similar results.

Table D.7: Attack fatality over time

Dependent variable: Actor	Fatalities		
	Fatah	Fatah	Hamas
Attacks (count)	0.85 (0.22)	0.64 (0.10)	3.70 (0.70)
Attacks \times time	0.0016 (0.0013)		-0.01 (0.003)
Attacks \times Second Intifada – Election		1.28 (0.89)	0.91 (1.20)
Attacks \times Post-2006 election		0.45 (0.11)	-0.95 (0.68)
Time	-0.00004 (0.00006)		-0.001 (0.003)
Second Intifada – Election		0.02 (0.04)	1.49 (1.27)
Post-2006 election		-0.04 (0.02)	-0.83 (0.31)
Constant	-0.004 (0.03)	-0.01 (0.005)	0.52 (0.55)
Excluding June 2007 (Battle of Gaza)	No	No	No
T	300	300	300
adj. R^2	0.79	0.82	0.52

Note: Newey-West standard errors in parentheses

Second, when we consider how the effects change by conflict eras we see that each additional attack during the Second Intifada results in about 1 more fatality, on average, than attacks during the Oslo era, but these results are also indistinguishable from zero for both actors, suggesting that we do not have enough information to say that attacks were more fatal in the Second Intifada relative to the Oslo era. Likewise, for Hamas we fail to reject the null hypothesis that attacks during the Oslo era had the same number of fatalities on average as attacks after the 2006 election. We also fail to reject the joint hypothesis test that attacks during all three eras produce the same number of fatalities on average.

Third, for Fatah we do reject the hypothesis that attacks after the 2006 election produce the same number of fatalities, on average, as attacks during the Oslo era. However, the direction of this change depends on a very specific outlier: June 2007 during the so-called Battle for Gaza. The leverage of this one observation is about 0.94, which is roughly double that of the next most influential point (Cook’s distance of about 109). As such, it seems clear that this is a true outlier, so we report the Fatah results both with and without it. With(out) it, each additional attack during this era is associated with 0.45 more (0.37 fewer) fatalities, on average, versus attacks during the Oslo era. In both cases, we reject the null hypotheses that (a) there are no differences in average fatalities per attack between this era and the Oslo era, and (b) that each additional attack is associated with the same average increase in fatalities across all three eras. We conclude that Fatah attacks, with the exception of one particular month, have on average decreased slightly in their lethality from an average of 0.63 fatalities for each additional attack during the Oslo era to an average increase of 0.23 fatalities for each additional attack. Substantively, these differences are both still less than 1 death per attacks and do not suggest major changes in lethality over time.

Overall, there is nothing in these results to suggest that Hamas attacks in any era were particularly deadly. The average change in fatalities for each additional attack are roughly equivalent across eras. If anything Hamas attacks are at their least deadly during the final part of the data from 2006 onward. With the exception of June 2007, the same is true of Fatah attacks. Even in cases where differences over time are statistically significant, the substantive effects are quite small.

E Identification

With K states (ordered s_1, \dots, s_K) and two actions, each actor i has $4K$ possible payoffs to consider. Adopting notation from Pesendorfer and Schmidt-Dengler (2008), let $\pi_i(s_k, a_i, a_j)$ denote the systematic component of i ’s per-period payoff for choosing a_i in state s_k when its opponent takes action a_j . We can collect these payoffs into $4K \times 1$ matrix

of individual state-action payoffs

$$\mathbf{\Pi}_i = \begin{bmatrix} \pi_i(s_k, 0, 0) \\ \pi_i(s_k, 0, 1) \\ \pi_i(s_k, 1, 0) \\ \pi_i(s_k, 1, 1) \end{bmatrix}_{k=1}^K .$$

As noted in the main text, Proposition 2 from Pesendorfer and Schmidt-Dengler (2008) tells us that at most we can identify K payoffs. To identify up to K parameters, we need to impose at least $3K$ restrictions on i 's payoffs. Consider the following $3K$ restrictions:

$$\begin{aligned} \pi_i(s_k, 0, 0) - \pi_i(s_k, 0, 1) &= 0 && \text{for } k = 1, \dots, K \\ \pi_i(s_k, 1, 0) - \pi_i(s_k, 1, 1) &= 0 && \text{for } k = 1, \dots, K \\ \pi_i(s_{k+1}, 0, 0) - 2\pi_i(s_k, 0, 0) + \pi_i(s_k, 1, 0) &= 2d\beta_i + \kappa_i && \text{for } k = 1, \dots, K - 1 \\ \pi_i(s_1, 1, 0) - \pi_i(s_1, 0, 0) &= \kappa_i. \end{aligned}$$

The first two sets of restrictions reflect the fact that neither side considers its opponent's action within its utility function.¹¹ The third set of restrictions maps the unconstrained payoffs into part of the linear-in-the parameters utility function to capture the difference in utility for a change in both state and action. The last restriction considers actions within a state to separate κ_i from β_i . These last two lines formalize some of the intuition from the main text: identification of β_i depends on variation in both attacks and states, while variation in actions within a state separates κ_i from β_i .

The restrictions are all linear and can be written in matrix form as

$$\mathbf{R}\mathbf{\Pi}_i = r_i.$$

Here \mathbf{R} is the matrix of numeric constants that impose the above restrictions on the payoffs.

In addition to the necessary conditions given by Pesendorfer and Schmidt-Dengler (2008) in Proposition 2, we can also verify a sufficient condition for identification of θ , which they provide in their Proposition 3. The condition depends on the restrictions matrix \mathbf{R} and some indifference conditions that are functions of the equilibrium choice probabilities. The result says that given an equilibrium, a distribution of action-specific payoff shocks, and a discount factor, if a rank condition holds, then the payoff parameters are identified. To verify the condition for player i , we need the following matrices where $j \neq i$:

$$\underbrace{\mathbf{P}_i}_{K \times 4K} = \begin{bmatrix} \mathbf{p}_{ij}(s_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{p}_{ij}(s_K) \end{bmatrix}$$

¹¹Note that the above reflects a minimal set of linear restrictions to identify K parameters per player. We seek to identify only 2 and impose a total of $4K - 1$ payoff restrictions in estimation.

where

$$\mathbf{p}_{ij}(s_k) = \begin{bmatrix} P(0, s_k; v_i)P(0, s_k; v_j) \\ P(0, s_k; v_i)P(1, s_k; v_j) \\ P(1, s_k; v_i)P(0, s_k; v_j) \\ P(1, s_k; v_i)P(1, s_k; v_j) \end{bmatrix}',$$

$$\underbrace{\mathbf{P}_j}_{K \times 2K} = \begin{bmatrix} P(0, s_1; v_j) & P(1, s_1; v_j) & 0 & 0 \\ 0 & \ddots & \ddots & 0 \\ 0 & 0 & P(0, s_K; v_j) & P(1, s_K; v_j) \end{bmatrix},$$

and

$$\underbrace{\mathbf{P}^*}_{K \times 4K} = [-\mathbf{P}_j \quad \mathbf{P}_j].$$

We also need the matrix

$$\underbrace{\mathbf{X}_i}_{K \times 4K} = \mathbf{P}^* + \delta \mathbf{P}^* \mathbf{\Gamma} (\mathbf{I}_K - \delta \mathbf{P}^* \mathbf{\Gamma})^{-1} \mathbf{P}_i,$$

where $\mathbf{\Gamma}$ is the $4K \times K$ matrix of Markov transition probabilities, and \mathbf{I}_K is an identity matrix of size K . Note that \mathbf{X}_i depends on the equilibrium values v , the first-stage transition parameters γ , the discount factor δ , and the distributional assumption on the action-specific shocks ε , but not the payoff parameters β and κ . Pesendorfer and Schmidt-Dengler (2008) show that \mathbf{X}_i can be used to characterize indifferent types of player i , and they further show that a sufficient condition for the identification of player i 's payoff parameters is the full rank of the square matrix $\begin{bmatrix} \mathbf{X}_i \\ \mathbf{R} \end{bmatrix}$. We verify that this rank condition holds at the estimated equilibrium \hat{v} , fixed discount factor δ , and first-stage estimates $\hat{\gamma}$ for both Fatah and Hamas, thus satisfying the sufficient identification conditions.

E.1 Identification intuition

To follow up on the above analysis, we also want to briefly sketch how β_i and κ_i are separately identified from each other within a given equilibrium. Fix an equilibrium v and first-stage parameters γ . Now consider two separate solutions to the CMLE's constrained optimization problem $\theta = (\beta, \kappa) \neq \theta' = (\beta', \kappa')$ with associated Lagrange multipliers λ and λ' . Since these are both solutions to the CMLE problem, then both (θ, v, λ) and (θ', v, λ')

solve the first order conditions of the CMLE's Lagrangian $\mathcal{L}(\theta, v, \lambda; \gamma)$ such that

$$\begin{aligned}\frac{\partial}{\partial(v, \theta)} \mathcal{L}(\theta, v, \lambda; \gamma) &= \frac{\partial}{\partial(v, \theta')} \mathcal{L}(\theta', v, \lambda'; \gamma) = 0 \\ \frac{\partial}{\partial \lambda} \mathcal{L}(\theta, v, \lambda; \gamma) &= \frac{\partial}{\partial \lambda'} \mathcal{L}(\theta', v, \lambda'; \gamma) = 0 \\ \mathcal{V}(v; \theta, \gamma) - v &= \mathcal{V}(v; \theta', \gamma) - v = 0.\end{aligned}$$

Where this last line is the first-order condition with respect to the Lagrange multipliers. In words, this last line tells us that both θ and θ' have to satisfy the equilibrium constraint given the fixed v since they are both solutions. Note that θ only enters \mathcal{V} through the utility function, so this further simplifies to

$$\begin{aligned}[u_i(a_i, s; \theta)]_{i \in \{F, H\}, (a_i, s) \in \{0, 1\} \times \mathcal{S}} &= [u_i(a_i, s; \theta')]_{i \in \{F, H\}, (a_i, s) \in \{0, 1\} \times \mathcal{S}} \\ \begin{bmatrix} s\beta_F + \kappa_F a_F \\ s\beta_H + \kappa_H a_H \end{bmatrix}_{(a_i, s) \in \{0, 1\} \times \mathcal{S}} &= \begin{bmatrix} s\beta'_F + \kappa'_F a_F \\ s\beta'_H + \kappa'_H a_H \end{bmatrix}_{(a_i, s) \in \{0, 1\} \times \mathcal{S}} \\ \begin{bmatrix} (\beta_F - \beta'_F)s + (\kappa_F - \kappa'_F)a_F \\ (\beta_H - \beta'_H)s + (\kappa_H - \kappa'_H)a_H \end{bmatrix}_{(a_i, s) \in \{0, 1\} \times \mathcal{S}} &= 0.\end{aligned}$$

At this point, standard results from linear models apply. Since s and a_F are linearly independent, the only solution to the Fatah equations is $(\beta_F, \kappa_F) = (\beta'_F, \kappa'_F)$ and likewise for H (contradicting the assumption that $\theta' \neq \theta$).

Another way to see this is to rewrite the above such that

$$\begin{bmatrix} s \\ s \end{bmatrix} = \begin{bmatrix} a_F \frac{\kappa'_F - \kappa_F}{\beta_F - \beta'_F} \\ a_H \frac{\kappa'_H - \kappa_H}{\beta_H - \beta'_H} \end{bmatrix}_{(a_i, s) \in \{0, 1\} \times \mathcal{S}}.$$

Note that these equalities imply that the state variable is a fixed proportion of the two action variables and thus linearly dependent on both. Since s and a_i are linearly independent, we have reached a contradiction.

F Standard errors and sensitivity of the second stage estimates

In this appendix, we describe the standard errors reported for the two-step CMLE estimates and consider how sensitive the estimates of β_i and κ_i are to the first-stage estimates γ . The standard result on two-step estimation involving a maximum likelihood estimator comes from Murphy and Topel (1985). To use these results the way we do, we need the first-step estimates of $\hat{\gamma}$ to be normally distributed. This is not guaranteed because of the unit root estimated in the AR-1 model. To check for normality in the errors, we use the Kolmogorov-Smirnov nonparametric test and the Jarque-Bera moment-based test on the

residuals $\hat{\nu}^t$. Specifically, we use a parametric bootstrap to simulate the sampling distribution of $\hat{\gamma}$; we then use the same tests to test for normality for each estimated γ . In every case we fail to reject the null hypothesis of normality, which provides us with confidence in constructing two-step standard errors in this way.

Aguirregabiria and Mira (2007, Proposition 1) use Murphy and Topel's (1985) result to describe the asymptotic distribution of the two-step pseudo-likelihood estimator from Hotz and Miller (1993) and we follow the same approach here. Specifically, let $\theta_2 = (\beta, \kappa, v)$ be the set of parameters estimated in the second stage, then the two-step correction gives the variance of $\hat{\theta}_2$ as

$$\widehat{\text{var}}(\hat{\theta}_2) = \hat{\Sigma}_{\theta_2} + \hat{\Sigma}_{\theta_2} \left(\hat{\Omega} \hat{\Sigma}_{\gamma} \hat{\Omega}^T \right) \hat{\Sigma}_{\theta_2}.$$

Here $\hat{\Sigma}_{\theta_2}$ is the CMLE covariance matrix based on the bordered Hessian as described by Silvey (1959) which is

$$\hat{\Sigma}_{\theta_2} = \begin{bmatrix} H_{\theta_2} L(\hat{\nu}|Y) + J_{\theta_2} \mathcal{V}(\hat{\theta}_2; \hat{\gamma})^T & J_{\theta_2} \mathcal{V}(\hat{\theta}_2; \hat{\gamma}) & -J_{\theta_2} \mathcal{V}(\hat{\theta}_2; \hat{\gamma})^T \\ -J_{\theta_2} \mathcal{V}(\hat{\theta}_2; \hat{\gamma}) & \mathbf{0} & \mathbf{0} \end{bmatrix}^{-1},$$

where H_x and J_x respectively denote the Hessian and the Jacobian of a function with respect to x . The bordered Hessian standard errors are found using the square root of the diagonal of $\hat{\Sigma}_{\theta_2}$.

The remaining two matrices are related to the first-stage estimates $\hat{\gamma}$. The matrix $\hat{\Omega}$ describes how the CMLE's Lagrangian changes with respect to γ and θ_2 and is given by

$$\hat{\Omega} = \begin{bmatrix} J_{\theta_2} L^*(\hat{\nu}|Y, \hat{\gamma})^T J_{\gamma} L^*(\hat{\nu}|Y, \hat{\gamma}) + J_{\theta_2} \mathcal{V}(\hat{\theta}_2; \hat{\gamma})^T J_{\gamma} \mathcal{V}(\hat{\theta}_2; \hat{\gamma}) \\ \mathbf{0} \end{bmatrix}.$$

Here, L^* is the vector-valued log-likelihood of the entire data:

$$L^*(v|Y, \gamma) = (\log P(a_H^t; s^t, v_H) + \log P(a_H^t; s^t, v_H) + \log f(s^t; a^{t-1}, s^{t-1}, \gamma))_{t=1}^T.$$

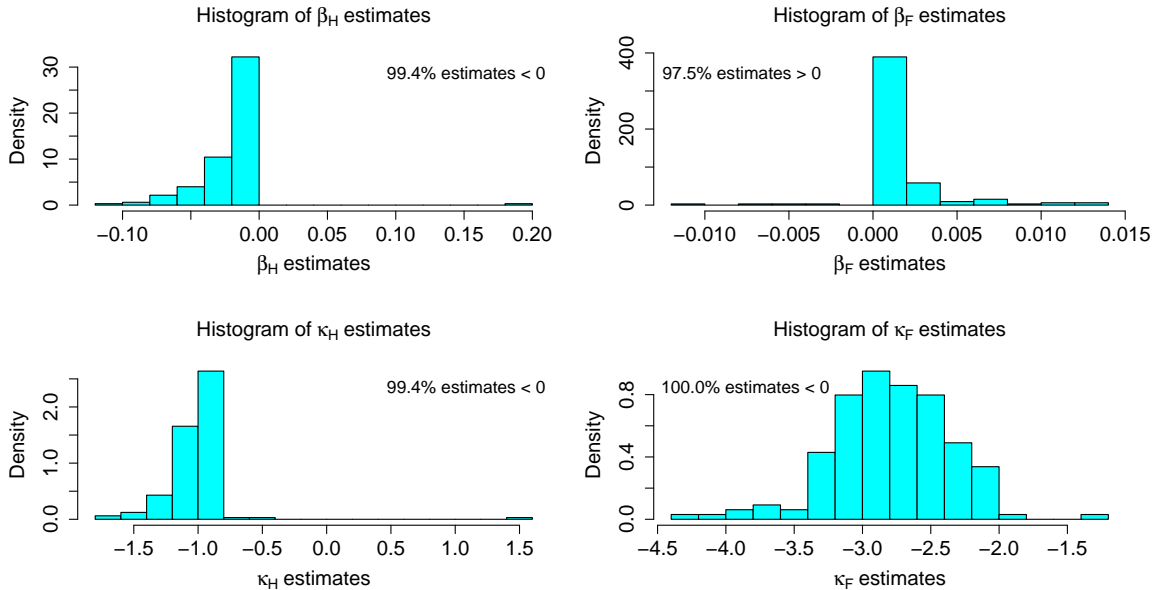
Note that for a given estimate of γ , the transition probabilities are fixed and so using either $L(v|Y)$ or $\sum_{t=1}^T L^*(v|Y, \hat{\gamma})$ as the CMLE's objective function will return the same constrained maximum likelihood estimates of θ_2 .¹² The final piece is the first-stage covariance matrix $\hat{\Sigma}_{\gamma}$, which we construct using the parametric bootstrap mentioned above.

Furthermore, we also want to know how sensitive the second-stage estimates are to changes in the first-stage estimates. To consider this we conduct a sensitivity analysis where for each iteration $b = 1, 2, \dots, B$ we conduct the following exercise:

1. Draw new values for continuous state variable \tilde{s}_b^t using the parameters from the first-stage model

¹²For completeness in L^* we impose $f(s^1; a^0, s^0, \gamma) = 1$ or $\log(f(s^1; a^0, s^0, \gamma)) = 0$.

Figure F.1: Sensitivity of payoff estimates to first-step estimates.



2. Re-fit the first-stage model to produce new estimates $\hat{\gamma}_b$.
3. Re-fit the second-stage model using $\hat{\gamma}_b$ and observed data $Y = (s^t, a^t)_{t=1}^T$. Save $\hat{\beta}_b, \hat{\kappa}_b$.
4. Repeat steps 1-3, B times.

This analysis allows us to consider how much variation there is in the second-step estimates under a range of plausible values of $\hat{\gamma}$. If the analysis is highly sensitive to the first-stage values, then we should see a large range of second-stage values. We are particularly interested in seeing how frequently the signs on the second-stage estimates change. We run the analysis for 500 iterations and save the results that report successful convergence. The results are reported in Figure F.1.

There are few points of interest in Figure F.1. Most importantly the signs on the estimates almost never change over the course of this experiment and are in the expected direction in more than 98% of simulations. The few cases that do not match the main results are clear outliers from the other cases. Overall, these histograms all peak around the point estimates reported in Table 2, which is a good sign that the exact estimates of γ used in the main model are not driving the main results.

G Model fit and comparisons

G.1 Comparison to alternative theoretical models

In this appendix, we continue our comparison of the main structural model to several alternative models. As mentioned in the main text, we consider a type of “null model” that we call the no-competition model. In the no-competition model, neither side can use

violence to move public opinion (i.e., outbidding cannot happen). Technically, this results in a model where $\gamma_{F,1} = \gamma_{F,2} = \gamma_{H,1} = \gamma_{H,2} = 0$. Under this assumption, we cannot identify β , so the only parameters used to fit the no-competition model are κ_F and κ_H . Intuitively, the no-competition model represents a world where costs/desire for violence (broadly defined) determine actions, but outbidding cannot be part of that because violence does not affect popularity. In the second comparison model, we consider the tit-for-tat model mentioned in the main text. The results from this comparison models are presented alongside the main model in Table G.1.

Table G.1: Comparing outbidding to other theories.

	Outbidding (main model)	No-competition	Tit-for-tat
β_H	-0.01 (0.004)		
β_F	0.0005 (0.0003)		
τ_H			0.75 (0.44)
τ_F			1.33 (0.49)
κ_H	-0.95 (0.23)	-0.25 (0.12)	-0.35 (0.13)
κ_F	-2.46 (0.28)	-2.54 (0.22)	-3.32 (0.42)
LL	-278.20	-284.18	-278.59
T	300	300	300

Standard errors from the bordered Hessian in parentheses

There are a few points of interest here. First, we see that there is some evidence in favor of tit-for-tat violence. Both τ estimates are positive, suggesting that groups receive some benefit for retaliatory violence. Interestingly, Fatah appears to place higher value on responding to violence from Hamas than vice-versa, while still facing larger costs to committing violence. A fuller exploration of these results is left to future work; of interest to us is the comparison of these two models, discussed in the main text.

In comparing the fitted model to the comparison models, we consider two exercises. First, in the main text, we directly compare the main model to the no-contest and tit-for-tat models using nested and non-nested tests, respectively. In both cases, we find support for outbidding model over the comparison model.

Second, we use each model's conditional choice probabilities to assess how many attacks by each actor we would expect to correctly predict (ePCP). For actor i , this value is given by

$$\text{ePCP}_i = \frac{1}{T} \sum_{t=1}^T \Pr(a_i^t; s^t),$$

where (a_i^t, s^t) are the observed action-state pairs from the data Y and \hat{v}_i are the equilibrium (net-of-shock) expected utilities estimated from the CMLE. We push on this comparison a little more, by also considering an overall measure that aggregates across actors

$$\text{ePCP} = \frac{1}{T} \sum_{t=1}^T \prod_{i=H,F} P(a_i^t; s^t, \hat{v}_i).$$

Table G.2: In-sample model fit.

	Main model	No competition	Tit-for-tat
ePCP– Hamas	0.54	0.51	0.51
ePCP– Fatah	0.83	0.86	0.87
ePCP– Overall	0.46	0.44	0.45

Table G.2 reports the results from both exercises. The three rows show us the expected percentage of actions correctly predicted by each model. The first two rows break this comparison down by actor. We see that the outbidding model is expected to correctly predict 54% of Hamas’ actions and 83% of Fatah’s. For the two comparison models, these values are 51% and 86-7%, respectively. Note that all three models do much better at predicting Fatah, because Fatah attacks are much rarer. This rarity means that lots of “no attack” predictions will be correct. Hamas attacks more often making it slightly trickier for the model to predict. Interestingly, we see that our model does a better job at explaining Hamas actions, while for Fatah the tit-for-tat model is slightly preferred.

When we consider the overall expected percent correctly predicted, we find that the main model is preferred. The explanatory gains in predicting Hamas actions are more than off-set by the decrease in understanding Fatah actions. This exercise opens the door for other, competing theories and models of these data that may improve on our outbidding-based approach.

As an additional, exploratory analysis, we also consider how the model fit tests in the main data break down across the eras reported in Table 4. These subsample results are reported in Table G.3. Here we see that in only one of the subsample time periods do we fail to reject the null hypothesis that the main model is preferred. Specifically, during the Second Intifada, we cannot distinguish between the models using this test. Nonetheless, note that this subsample is the smallest of the eras, and thus this test is the least powered of the ones considered.

G.2 Comparison to reduced-form results

In this appendix, we illustrate why reduced-form approaches cannot uncover evidence either for or against outbidding using the same data we analyze. Instead of examining model fit, we focus on why countervailing encouragement and discouragement effects mask outbidding’s presence.

Table G.3: Sub-sample model fit

Era	Obs	Clarke stat	Clarke p value	LR stat	LR p value
Oslo	81	50	< 0.05	80	< 0.05
2nd Intifada	64	34	0.27	30	< 0.05
Post-election	155	98	< 0.05	170	< 0.05

Note: Clarke (LR) test compares the main model to the tit-for-tat (no-competition) model within each time period

As previously noted, most reduced-form outbidding studies do not try to directly measure popular support and, as such, do not consider the effectiveness of attacks within an outbidding framework. A rare exception is Jaeger et al. (2015) who examine data from the Second Intifada and find that attacks from Hamas and Fatah make the groups more popular among Palestinians. Indeed, we find a similar result in Table 1 in the main text. Nonetheless, the authors do not attempt to quantify the effects of competition on violence.

The most direct translation from standard outbidding work would be to use a Poisson regression to regress the total number of terrorist attacks in a given month against some transformation of the state variable such that it reflects the tightness of the competitive environment, perhaps, negative distance from the mean popularity level. Applying such an approach to our data does return a positive coefficient, such that months when relative popularity is closer to the in-sample average have more violence, on average, than months when relative popularity is extreme. The estimate is statistically insignificant ($p > 0.80$) with small substantive implications, however. The results do not meaningfully change if one uses median level of relative popularity to compute the distance.

Thus, when examining the reduced-form correlation between relative popularity and attacks, it is not clear whether outbidding appears in the data. As such, we consider a more sophisticated approach based on vector auto-regression (VAR) of the state and action variables to give reduced form methods a better chance at estimating the relationships among these factors.

The VAR approach considers the states and actions as three outcomes in a system of equations. Given the unit root in \tilde{s}^t , we work in differences $\Delta\tilde{s}^t$, the VAR model is now specified as

$$\begin{bmatrix} \Delta\tilde{s}^t \\ a_H^t \\ a_F^t \end{bmatrix} = \Pi_0 + \sum_{\ell=1}^L \Pi_\ell \begin{bmatrix} \Delta\tilde{s}^{t-\ell} \\ a_H^{t-\ell} \\ a_F^{t-\ell} \end{bmatrix} + \begin{bmatrix} \nu_s^t \\ \nu_{a_H}^t \\ \nu_{a_F}^t \end{bmatrix}.$$

Here, Π_0 is a length-3 vector of constant terms. The remaining Π matrices are 3×3 coefficient matrices, again with each row representing the relationships between each outcome and lags of all other outcomes. Finally, the ν terms represent exogenous error terms for each equation. We use BIC comparisons to select $L = 2$ lags of the variables.

Table G.4: VAR regression results

	Δ States	Hamas Attack	Fatah Attack
$\Delta \tilde{s}_{t-1}$	0.91 (0.02)	0.04 (0.15)	-0.11 (0.08)
Hamas attacks $_{t-1}$	-0.27 (0.01)	0.29 (0.06)	0.07 (0.03)
Fatah attacks $_{t-1}$	1.02 (0.02)	-0.01 (0.11)	0.002 (0.06)
$\Delta \tilde{s}_{t-2}$	0.01 (0.02)	0.03 (0.09)	0.06 (0.05)
Hamas attacks $_{t-2}$	0.28 (0.01)	0.23 (0.07)	-0.03 (0.04)
Fatah attacks $_{t-2}$	-0.94 (0.03)	0.01 (0.18)	0.23 (0.10)
Const.	-0.01 (0.01)	0.21 (0.04)	0.03 (0.02)

Note: Least squares estimates with standard errors in parentheses.

The VAR results are presented in Table G.4, where each column represents one of the three equations, above. The estimates in the first column are similar to the ECM model in Table 1, which is unsurprising given that the setups are nearly identical. Again, we see that terrorism is effective at moving relative popularity in each group’s preferred direction. Turning to the other two columns, however, we see that there is little-to-no observed effect of popularity on either actor’s decision to use terrorism. Specifically, we see weak relationships where both groups are slightly more likely to attack if there is a positive change (pro-Fatah) in public opinion.

To make these marginal effects more clear, consider the impulse response function for these three variables, as presenting in Figure G.1. In this figure, the columns represent the effects of a 1 unit shock to that variable on the three outcomes (the rows). So the first, column shows the effects when $\Delta \tilde{s}_t$ increases by 1. The persistence in the state space is noted here, it takes more than six months for $\Delta \tilde{s}_t$ to return to its baseline, but there are no discernible effects on either actor. Overall, there is little relationship between changes in popularity and changes in attack probabilities when studying the marginal effects estimated by VARs.

However, the deceptively tricky part of reduced-form analysis, is that we have no idea if these weak relationships are evidence of outbidding or not! We still see evidence of each side’s ability to outbid (i.e., attacks from a group increase its popularity), but do these various null relationships provide evidence that groups are valuing popularity? Moreover, we see little evidence that changes in popularity affect violence here. As discussed above, both deterrence and emboldening effects are consistent with outbidding. But, we do not know, absent the structural model, if this null result suggests that: a) outbidding is a poor explanation or b) if heterogeneous effects are canceling each other out. Scholars finding

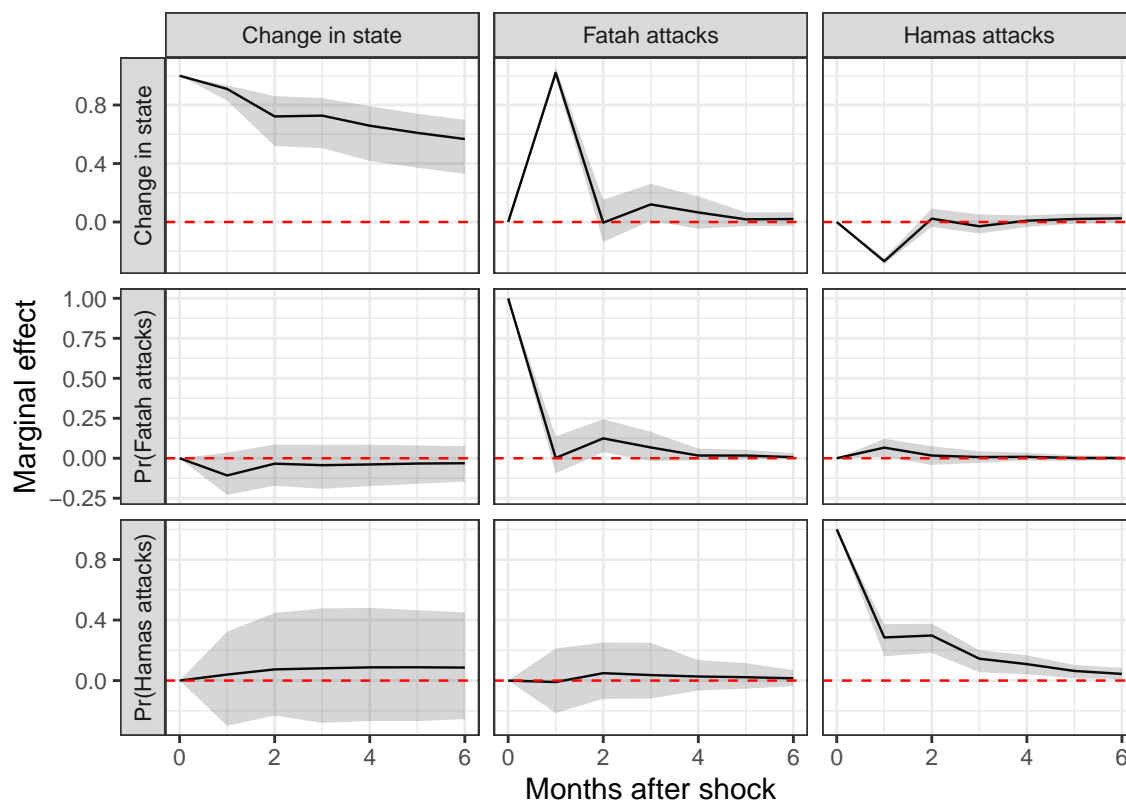


Figure G.1: Impulse response function from the VAR model, six months out. Shaded area is a bootstrapped 95% confidence interval

results like this may be tempted to report that they have found a null effect of popularity on terrorism, but our application of theory to these data reveal a much richer story of heterogeneous effects.

H Robustness to different time spans

In this appendix, we consider alternative time frames or subsamples of the data. The different time spans represent plausible break points in the Fatah-Hamas relationship, such that the underlying competition between the groups may have changed. As such we want to be sure that our estimates of the groups' preferences are robust to the exclusion of some of these later observations. Specifically, we consider the following time periods:

1. Until the formation of the Fatah-Hamas unity government in April 2014 (Jan. 1994–Mar. 2014)
2. Until the signing of the first Fatah-Hamas unity agreement in 2011 (Jan. 1994–Apr. 2011)
3. Until the last recorded Fatah attack in the GTD (Jan. 1994–Mar. 2009)
4. Until Hamas wins the 2006 legislative elections (Jan. 1994–Dec. 2005)
5. Until the start of the Second Intifada in 2001 (Jan. 1994–Aug. 2000)

6. From the same year that Bloom (2004) starts (Jan. 1997–Dec. 2018)
7. From 2001 (Jan. 2001–Dec. 2018)

Notice the first 5 subsamples end our overall sample at earlier dates, and the last two subsamples start the overall sample at a later date. The results of these short- T robustness checks are presented in Table H.1 along with a reproduction of the main model (1994-2018) for comparison. The CMLE failed to converge in the smallest sample so we employ an alternative estimation method called the nested-pseudo-likelihood estimator (Aguirregabiria and Mira 2007; Crisman-Cox and Gibilisco 2021), which did converge. In general, the estimates are very stable across samples.

I Choice of discount factor

In this appendix, we consider how our choice of discount factor affects our results. Specifically, we fix δ to 0 and then a few different values in the interval $[0.9, 1)$ and then reestimate the second-stage model at each value. Table I.1 shows the estimates and log-likelihoods of the second-stage model under different fixed values of the discount factor δ . The model with the best fit among these options is $\delta = 0.999$. As such we use this value in both the main model specification and the numerical examples.

Figure I.1 shows how the choice of discount factor affects the estimated equilibrium attack probabilities. Notice that when $\delta = 0.999$, the probabilities are identical to those in Figure 4 from the baseline analysis. The graphs in Figure I.1 demonstrate that for $\delta \in \{0.99, 0.999\}$ the attack probabilities are almost identical. Comparing across these two estimated models, the average difference in attack probabilities is 1.2 percentage points for both actors. The maximum difference is 2.4 percentage points for Hamas and 2.4 for Fatah. Technically, at a tolerance of 0.05 (i.e., 5 percentage points), the equilibrium choice probabilities are identical for $\delta \in \{0.99, 0.999\}$. Substantively, the predictions in the estimated model are roughly invariant to choosing a discount factor between $[0.975, 0.999]$. These similarities likely arise because discount factors are difficult to identify in dynamic discrete choice models generally (Abbring and Daljord 2020; Magnac and Thesmar 2002). When δ is difficult to identify, we would expect several values of δ to return essentially identical equilibrium attack probabilities, which is what we see in Figure I.1.

J Robustness to discretization

In this appendix, we show how our results are robust to the choices surrounding the discretization of the continuous state variable \tilde{s} into the discrete state variable s . There are two factors that can be adjusted here: the space between states ($2d$) and the cutpoint that signals the highest/lowest state. In the main model we use the as the 2.5th (97.5th) percentile of \tilde{s}^t as the cutpoint to denote the most Hamas (Fatah) friendly state. Likewise, we use $2d = 0.05$ as the distance between states, giving us $K = 440$.

Table H.1: Robustness to different time periods.

	Full Sample	2014 Agreement	2011 Agreement	2009 Last Fatah attack	2006 Elections	Second Intifada	Bloom's (2004) start year	From 2001
β_H	-0.01 (0.004)	-0.01 (0.005)	-0.01 (0.004)	-0.01 (0.004)	-0.02 (0.01)	-0.01 (0.002)	-0.02 (0.01)	-0.03 (0.01)
β_F	0.0005 (0.0003)	0.001 (0.0003)	0.0004 (0.0003)	0.0004 (0.0002)	0.0005 (0.0003)	0.0005 (0.43)	0.0003 (0.0001)	0.0002 (0.0001)
κ_H	-0.95 (0.23)	-0.90 (0.23)	-0.69 (0.22)	-0.40 (0.22)	-0.73 (0.25)	-0.89 (0.26)	-1.65 (0.33)	-1.94 (0.49)
κ_F	-2.46 (0.28)	-2.47 (0.31)	-2.31 (0.29)	-2.09 (0.29)	-2.33 (0.26)	-2.89 (3.64)	-2.29 (0.23)	-2.07 (0.25)
T	300	243	208	183	144	80	261	216
Start date	Jan 1994	Jan 1994	Jan 1994	Jan 1994	Jan 1994	Jan 1994	Apr 1997	Jan 2001
End date	Dec 2018	Mar 2014	Apr 2011	Mar 2009	Dec 2005	Aug 2000	Dec 2018	Dec 2018
LL	-278.20	-231.07	-206.89	-189.47	-125.40	-71.81 [†]	-228.05	-200.19

Note: [†]CMLE did not converge, estimates from nested-pseudo-likelihood (NPL) estimator. BH standard errors in parentheses.

Table I.1: Estimates at different discount factors

δ	0	0.9	0.925	0.95	0.975	0.99	0.999	0.9999
β_H		-0.51 (0.21)	-0.33 (0.15)	-0.18 (0.10)	-0.09 (0.03)	-0.03 (0.01)	-0.01 (0.00)	-0.01 (0.01)
β_F		-0.03 (0.08)	-0.02 (0.05)	-0.01 (0.03)	0.02 (0.01)	0.01 (0.00)	0.00 (0.00)	0.00 (0.00)
κ_H	-0.25 (0.12)	-1.17 (0.40)	-1.07 (0.39)	-0.91 (0.38)	-1.15 (0.27)	-1.05 (0.25)	-0.95 (0.23)	-2.09 (0.33)
κ_F	-2.54 (0.22)	-2.28 (0.70)	-2.28 (0.67)	-2.26 (0.63)	-2.76 (0.42)	-2.55 (0.34)	-2.45 (0.28)	-3.26 (0.36)
LL	-284.18	-281.03	-281.58	-282.47	-283.77	-280.83	-278.2	-307.4*

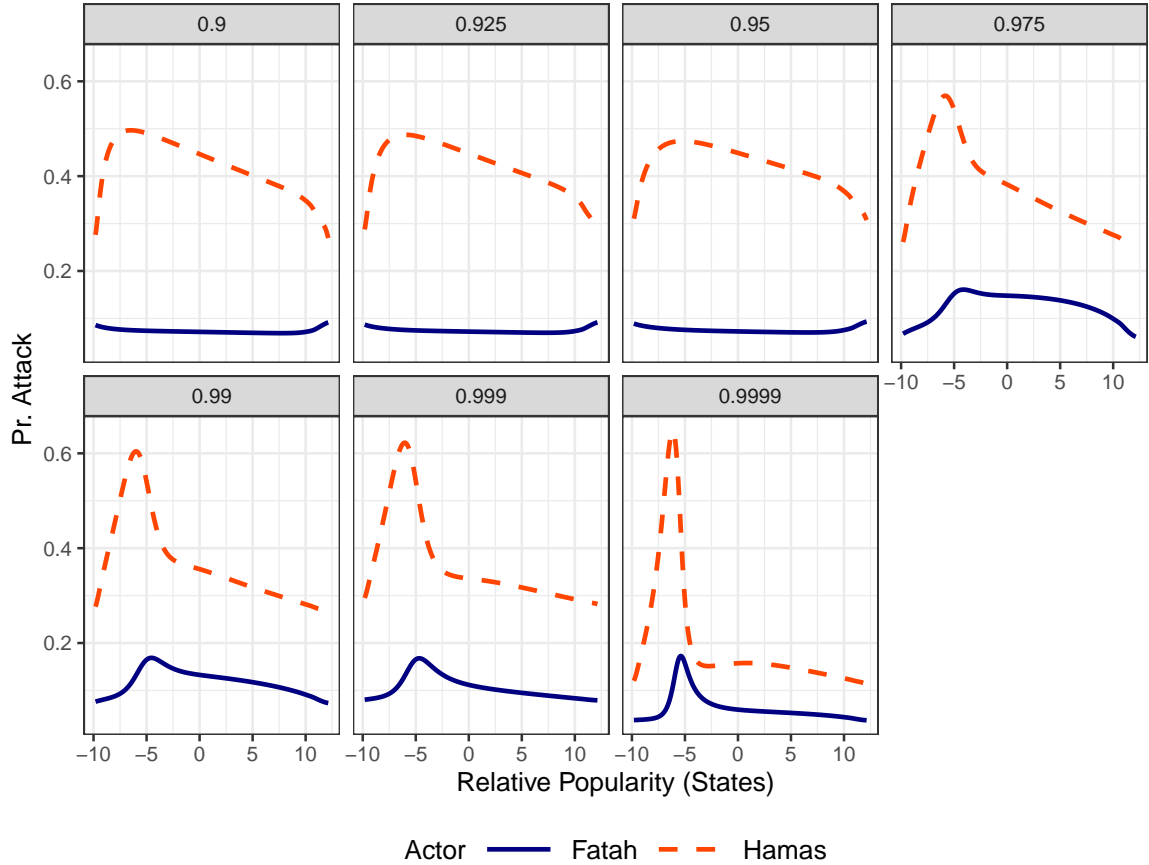
Note: *Model failed to converge. BH standard error in parentheses

Table J.1: Estimates at different discretizations of \tilde{s}^t

	0.025	0.025	0.025	0.05	0.05	0.05	0.05	0.05	0.075	0.075	0.075	0.075	0.075
End percentile	0.0125	0.025	0.0375	0.0125	0.025	0.0375	0.0125	0.025	0.0375	0.0125	0.025	0.0375	0.0375
K	909	879	863	455	440	432	455	440	432	303	293	293	288
β_H	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)	-0.01 (0.004)
β_F	0.0004 (0.0002)	0.0005 (0.0003)	0.001 (0.0003)	0.0004 (0.0002)	0.0005 (0.0003)	0.001 (0.0003)	0.0004 (0.0002)	0.0005 (0.0003)	0.001 (0.0003)	0.0004 (0.0002)	0.0005 (0.0003)	0.001 (0.0003)	0.001 (0.0003)
κ_H	-0.97 (0.23)	-0.95 (0.23)	-0.93 (0.23)	-0.97 (0.23)	-0.95 (0.23)	-0.93 (0.23)	-0.97 (0.23)	-0.95 (0.23)	-0.93 (0.23)	-0.97 (0.23)	-0.94 (0.23)	-0.93 (0.23)	-0.93 (0.23)
κ_F	-2.42 (0.27)	-2.46 (0.28)	-2.48 (0.29)	-2.42 (0.27)	-2.46 (0.28)	-2.48 (0.29)	-2.42 (0.27)	-2.46 (0.28)	-2.48 (0.29)	-2.42 (0.27)	-2.46 (0.28)	-2.48 (0.29)	-2.48 (0.29)
LL	-277.10	-278.15	-278.86	-277.08	-278.20	-278.87	-277.09	-278.20	-278.87	-277.09	-278.13	-278.13	-278.81

Note: BH standard errors in parentheses.

Figure I.1: Discount factors and equilibrium attack probabilities



Note: Graphs of the estimated equilibrium attack probabilities at 7 different discount factors. Setting $\delta = 0.999$ (second row, second column) is identical to Figure 4. The CMLE model failed to converge when $\delta = 0.9999$ as in Table I.1.

To assess the robustness of these two choices we consider all combinations of $2d \in \{0.025, 0.05, 0.075\}$ crossed with endpoint percentiles in $\{0.0125, 0.025, 0.0375\}$. The middle case $2d = 0.05$ and the endpoint percentile of 0.025, match the main results, the other models show how sensitive the results are to these choices. These new models are presented in Table J.1. Overall, we see that the main results are robust to changes in these discretization parameters.

The above analysis provides us with assurance that the main results are robust to changes in discretization choices around the original choice of $2d = 0.05$. However, it is worth noting the models in Table J.1 have many states relative to the number of observations. There may be a concern with estimating $v_i(a_i, s)$ for states are never visited or only visited once. As such, one may wonder whether our results are robust to a coarser discretization, although as mentioned in the main text, one of the justifications for using OLS to estimate γ is that we have many grid points. In Table J.2, we see that the point estimates are largely unchanged as we increase $2d$ even when K is quite small. When $K \leq 30$, all states are

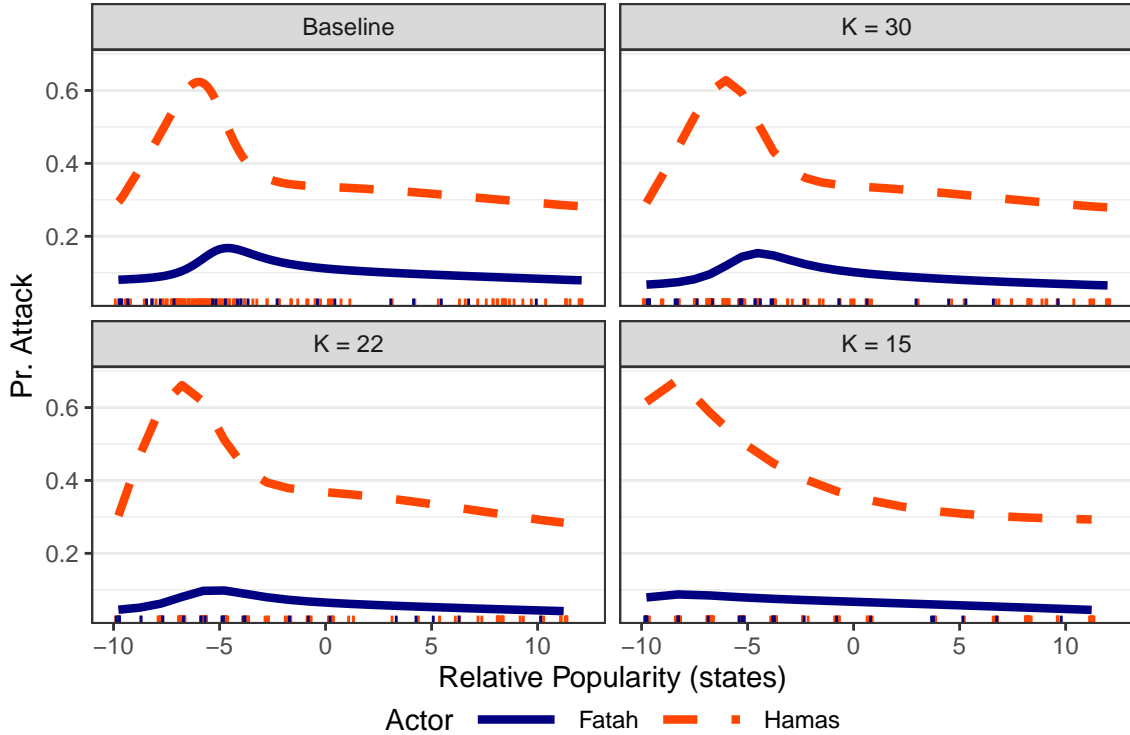
Table J.2: Estimates at coarse discretizations of \tilde{s} .

$2d$	0.05	0.75	1	1.5
End percentile	0.025	0.025	0.025	0.025
K	440	30	22	15
β_H	-0.01 (0.004)	-0.01 (0.005)	-0.02 (0.01)	-0.08 (0.02)
β_F	0.0005 (0.0003)	0.001 (0.0003)	0.001 (0.0003)	0.002 (0.002)
κ_H	-0.95 (0.23)	-0.94 (0.22)	-0.95 (0.23)	-0.89 (0.20)
κ_F	-2.46 (0.28)	-2.68 (0.31)	-3.09 (0.36)	-3.07 (0.62)
LL	-278.20	-277.12	-275.38	-274.42

Note: BH standard errors in parentheses.

visited at least twice. When $K = 15$, the value of popularity for each group increases in magnitude, so there is some change here.

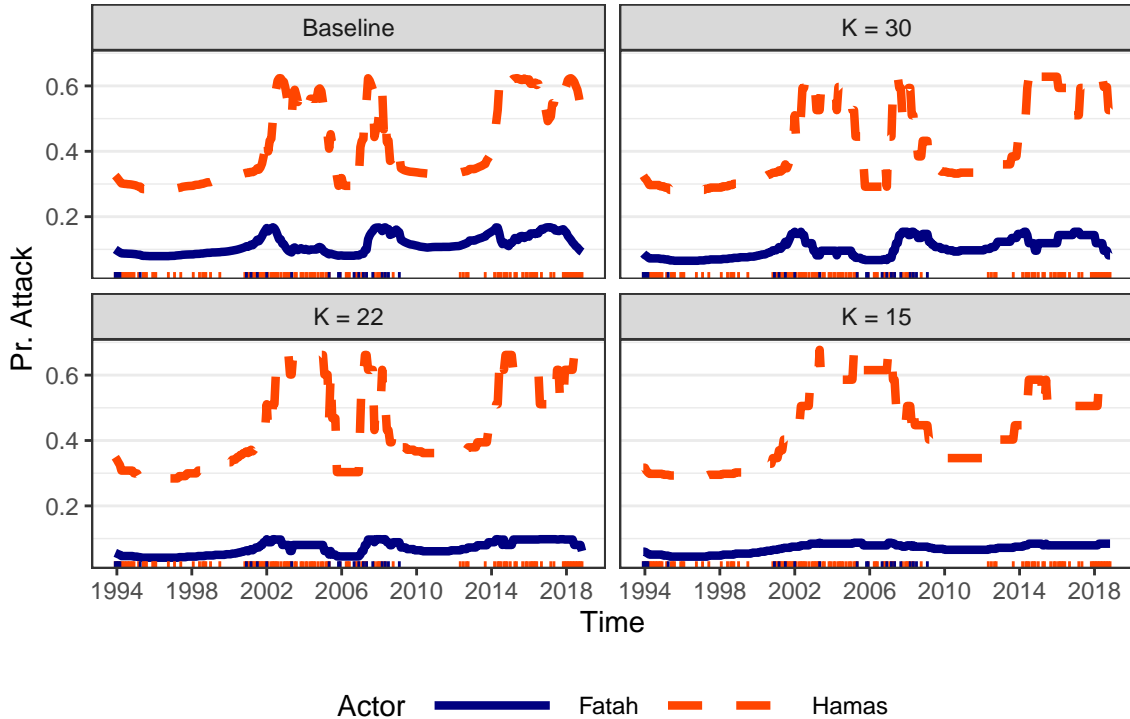
Figure J.1: Estimated attack probabilities over states with a coarse state space.



Note: Estimated probability that each group attacks as a function of the popularity level s . Hamas (Fatah) prefers smaller (larger) states. The horizontal axis includes a rug plot of observed attacks. Baseline panel replicates Figure 4 with $K = 440$.

Finally, we also investigate our estimated attack probabilities with a coarse state space. Figure J.1 graphs the attack probability as a function of the state, one for each model in

Figure J.2: Estimated attack probabilities over time with a coarse state space.



Note: Horizontal axis denotes sample months/periods. Left vertical axis is the estimated probability that i attacks in month t , i.e., $P(a_i = 1; s^t, \hat{v}_i)$ where s^t is the observed relative popularity level in period t and \hat{v}_i is estimated from the CMLE. The horizontal axis includes a rug plot of observed attacks. Baseline panel replicates Figure 3 with $K = 440$.

Table J.2; Figure J.2 graphs the attack probabilities as a function of time given the observed states s^t in the data. Again, the takeaway is the same. Our results are robust for even a small K , although attack probabilities change when $K = 15$, where the non-monotonic relationship between the probability of attacks and the state space only appears at the far left-end of the state space.

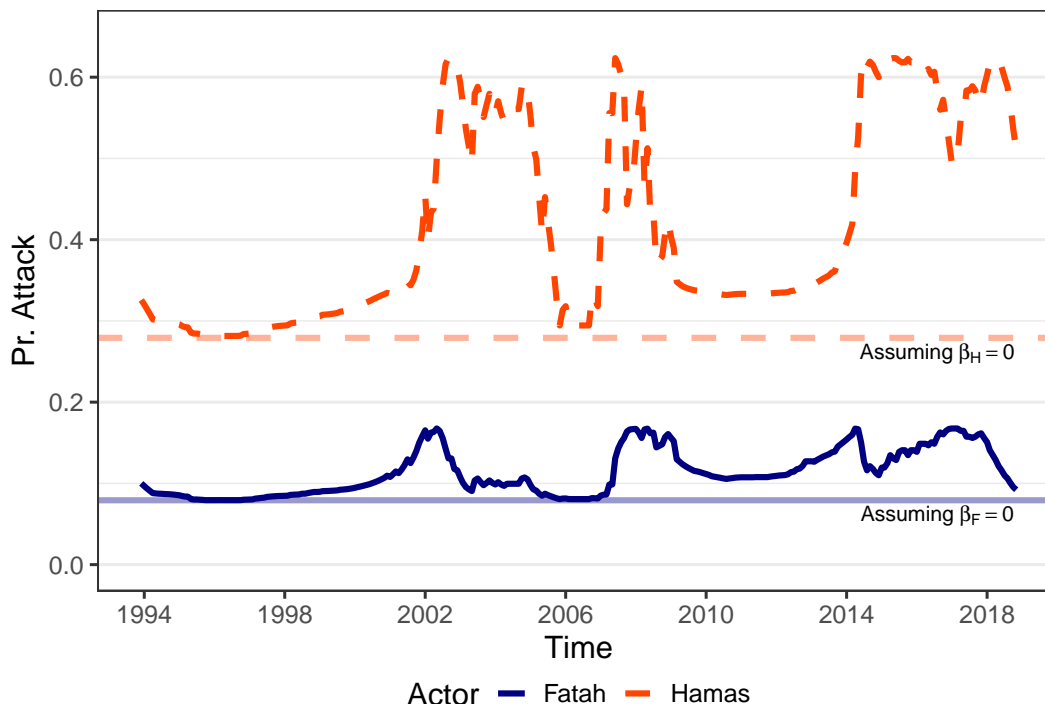
This reduced monotonicity is explained by two reasons. First, when K is small, we no longer see fewer attacks when Hamas is very popular, e.g., the year around when they win legislative elections. When K is small, observations from this lull are grouped with observations where Hamas is popular but still attacks, e.g., during the middle of the Second Intifada. Second, when K is small enough, the expected state is the current state regardless of whether a group attacks, even Fatah, the group that was estimated to be the most effective. As a result benefits of popularity β_i start to blow up to better explain the data.

K Interpreting estimates of β_i

It is difficult to directly interpret that magnitude of β_i for several reasons. First, the magnitude depends on the scale of relative popularity, so dividing or multiplying the ob-

served relative popularity levels by a constant will have a mechanical effect on the magnitude of β_i . Second, this is a dynamic model, so the effects of changing per-period incentives on the group’s propensity to attack is mediated by the group’s dynamic optimization problem. As such, the magnitude of β_i depends on the discount factor δ . When δ becomes larger, small changes in per-period incentives have larger effects on each group’s behavior.

Figure K.1: Illustrating the importance of the estimated values of popularity



Note: In the darker lines, we graph the estimated probability that i attacks in month t , i.e., $P(a_i = 1; s^t, \hat{v}_i)$ where s^t is the observed relative popularity level in period t and \hat{v}_i is estimated from the CMLE (as in Figure 3). In the transparent lines, we graph i ’s probability of attacking when $\beta_i = 0$ while holding all other parameters as fixed.

To better illustrate the importance of each group’s value of popularity, we compare group i ’s estimated equilibrium attack probabilities, i.e., those in Figure 3, to i ’s probability of attacking when $\beta_i = 0$, holding all other parameters fixed. Figure K.1 presents the comparison. As in Figure 3, the x -axis is time, and the y -axis is the probability of an attack. The darker lines show equilibrium attack probabilities for each group, i.e., $P(a_i = 1; s^t, \hat{v}_i)$ where s^t is the observed relative popularity level in period t and \hat{v}_i is estimated from the CMLE. The transparent lines denote the attack probabilities for each group when $\beta_i = 0$, holding fixed all other parameters, in particular κ_i , at their previously estimated values. Notice that when $\beta_i = 0$, i ’s attack probability is a constant because there are no longer strategic nor dynamic incentives to attack when i does not care about its relative popularity. As such, the probability that group i attacks is $\Pr(\varepsilon_i^t(0) - \varepsilon_i^t(1) < \hat{\kappa}_i)$, where $\hat{\kappa}_i$ is given in Table 2. This probability is roughly 0.28 for Hamas and 0.08 for Fatah.

Substantively, this leads to large effects. To see this, Hamas’s average propensity to attack in the estimated equilibrium, is 0.41, averaging over all observed states s^t in the data. Thus, if Hamas did not value its popularity ($\beta_H = 0$), then its propensity to attack would decrease by 14 percentage points, a roughly 33% decrease from the equilibrium baseline. Likewise, Fatah’s average propensity to attack in the estimated equilibrium is 0.11. Thus, if Fatah did not value its popularity ($\beta_F = 0$), then its propensity to attack would decrease by 3 percentage points, a roughly 30% decrease from the equilibrium baseline.

L Comparative statics with multiple equilibria

In this appendix, we detail how to compute comparative statics on the exogenous parameters using predictor-corrector homotopy method proposed in Aguirregabiria (2012). We use this method when examining how changes in exogenous parameters (β, κ, γ) affect equilibrium behavior.¹³ Because there are multiple equilibria, we cannot just vary the exogenous parameters, compute a new equilibrium, and compare choice probabilities under the old and new parameters values. Doing so would not guarantee that the new equilibrium bears any resemblance to the estimated one. For example, it is possible to uncover differences in choice probabilities when selecting among different equilibria even when exogenous features of the game do not change—see the numerical example in Section B as an example.

To describe the method, collect the exogenous parameters of interest in vector $\tilde{\theta} = (\beta, \kappa, \gamma)$, and suppose v is an equilibrium given exogenous parameters $\tilde{\theta}$, i.e., $\mathcal{V}(v|\tilde{\theta}) - v = 0$. Comparative statics refer to how the equilibrium v changes as we vary the parameters from $\tilde{\theta}$ to $\tilde{\theta}' \neq \tilde{\theta}$. When the Jacobian of \mathcal{V} (with respect to v) is not vanishing (i.e., the Jacobian matrix has full rank), then small changes in the exogenous parameters produce small changes in the endogenous equilibrium by the implicit function theorem. The rank condition on the Jacobian matrix can be verified given parameters $\tilde{\theta}$ and an equilibrium v . Furthermore, this insight can be used to conduct comparative statics in the presence of multiple equilibria. More specifically, for small changes in the parameter vector $\tilde{\theta}$, we first approximate changes in v using the implicit function theorem and linear interpolation. Next, we use these approximations as starting values in a Newton-Raphson algorithm that computes a new equilibrium after small changes in the parameters. Finally, we repeat this procedure until reaching the final parameter vector of interest $\tilde{\theta}'$. A minor difference between this routine and the specifics discussed in Aguirregabiria (2012) is that ours requires the computation of equilibria at each step along the line from $\tilde{\theta}$ to $\tilde{\theta}'$. While this does increase the computational burden of the procedure, it is feasible and will return an equilibrium upon convergence. This procedure is also used by Crisman-Cox and Gibilisco (2018).

Algorithm 1 presents the specifics of the procedure for reference and is an implementation of the predictor-corrector method applied to our model. Here, $n \in \mathbb{N}$ is a tuning

¹³Specifically, we use this procedure to create Figure 6 and Table 5 in the main text, Figures A.3, and A.4 in Appendix A, and Figure B.4 in Appendix B.

Algorithm 1: COMPARATIVE STATICS (CS) using a homotopy

Input: An initial coefficient vector $\tilde{\theta} = (\beta, \kappa, \gamma)$ and equilibrium v such that $\mathcal{V}(v|\tilde{\theta}) - v = 0$; new values $\tilde{\theta}'$; and a tuning parameter $n \in \mathbb{N}$. For the NEWTON solver, a convergence tolerance $\epsilon > 0$ and maximum number iterations $m \in \mathbb{N}$.

Output: An equilibrium v' under new parameters $\tilde{\theta}'$, i.e., $\mathcal{V}(v'|\tilde{\theta}') - v' = 0$

```

1  $\tilde{\theta}_{\text{old}} \leftarrow \tilde{\theta}$ 
2  $v_{\text{old}} \leftarrow v$ 
3 for  $i \leftarrow 1$  to  $n$  do
4    $\lambda \leftarrow \frac{i}{n}$ 
5    $\tilde{\theta}_{\text{new}} \leftarrow (1 - \lambda)\tilde{\theta} + \lambda\tilde{\theta}'$ 
6   slope  $\leftarrow - (J_v \mathcal{V}(v_{\text{old}} | \theta_{\text{old}}) - \mathbf{1})^{-1} J_{\tilde{\theta}} \mathcal{V}(v_{\text{old}} | \theta_{\text{old}})$ 
7   start  $\leftarrow v_{\text{old}} + [\theta_{\text{new}} - \theta_{\text{old}}] \mathbf{slope}^T$ 
8    $(v_{\text{new}}, \mathbf{success}) \leftarrow \text{NEWTON}(\mathbf{start}, \mathcal{V}(x | \theta_{\text{new}}) - x, \epsilon, m)$ 
9   if success then
10     $\theta_{\text{old}} \leftarrow \theta_{\text{new}}$ 
11     $v_{\text{old}} \leftarrow v_{\text{new}}$ 
12  else
13     $v_{\text{new}} \leftarrow \text{"Warning: Convergence Problems."}$ 
14    break
15  $v' \leftarrow v_{\text{new}}$ 
16 return  $v'$ 

```

parameter describing how many steps the algorithm takes as the parameters move away from $\tilde{\theta}$ to $\tilde{\theta}'$. When n is very large relative to $\|\tilde{\theta} - \tilde{\theta}'\|_2$, the algorithm is more careful to avoid switching equilibria when looking for counterfactual changes. In line 6, we invoke the implicit function theorem, where $\mathbf{1}$ is the identity matrix of size $4K$ and the variable **slope** stores $D_{\tilde{\theta}} v$. In line 7, we use linear interpolation to predict how the equilibrium changes as the exogenous parameters change from $\tilde{\theta}_{\text{old}}$ to $\tilde{\theta}_{\text{new}}$. In line 8, we use the Newton-Raphson method. As input, we give it the starting values in **start**, the function (that takes input $x \in \mathbb{R}^{4K}$) $\mathcal{V}(x | \theta_{\text{new}}) - x$, a convergence tolerance $\epsilon > 0$ and a maximum number of iterations m . As an output, it returns a pair including a potential solution given exogenous parameters $\tilde{\theta}_{\text{new}}$ and an indicator of a successful convergence. The Newton-Raphson algorithm has converged when $\|\mathcal{V}(v_{\text{new}} | \theta_{\text{new}}) - v_{\text{new}}\|_{\infty} < \epsilon$. In our experiments, we save the output of the Newton call (v_{new}) in each iteration to produce a continuous representation of the equilibrium, which we use to graphically verify that the output has indeed continuously traced the equilibrium as parameters vary from $\tilde{\theta}$ to $\tilde{\theta}'$.

References

- Abbring, Jaap H. and Øystein Daljord. 2020. "Identifying the Discount Factor in Dynamic Discrete Choice Models." *Quantitative Economics* 11(2):471–501.
- Aguirregabiria, Victor. 2012. "A Method for Implementing Counterfactual Experiments in Models with Multiple Equilibria." *Economics Letters* 114(2):190–194.
- Aguirregabiria, Victor and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75(1):1–53.
- Awad, Ibrahim and Wael Alazzez. 2020. "Using Currency Demand to Estimate the Palestine Underground Economy: An Econometric Analysis." *Palgrave Communications* 6(1):1–11.
- Bekker, Paul A. 1994. "Alternative Approximations to the Distributions of Instrumental Variable Estimators." *Econometrica* pp. 657–681.
- Bloom, Mia M. 2004. "Palestinian Suicide Bombing: Public Support, Market Share, and Outbidding." *Political Science Quarterly* 119(1):61–88.
- Brückner, Markus and Antonio Ciccone. 2011. "Rain and the Democratic Window of Opportunity." *Econometrica* 79(3):923–947.
- Center for Economic and Policy Research (CEPR). 2012. "Agriculture in Palestine: A Post-Oslo Analysis." Accessed February 2022 from <https://www.semanticscholar.org/paper/Agriculture-in-Palestine-%3A-a-post-Oslo-analysis/6170c6659e55c5ac2cbed8a263d20db7ca5a0879?p2df>.
- Cinelli, Carlos and Chad Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82(1):39–67.
- Crisman-Cox, Casey and Michael Gibilisco. 2018. "Audience Costs and the Dynamics of War and Peace." *American Journal of Political Science* 62(3):566–580.
- Crisman-Cox, Casey and Michael Gibilisco. 2021. "Estimating Crisis Signaling Games in International Relations: Problems and Solutions." *Political Science Research and Methods* 9(3):565–582.
- Gibilisco, Michael, Brenton Kenkel and Miguel R. Rueda. 2022. "Competition and Civilian Victimization." *Journal of Conflict Resolution* 66(4-5):809–835.
- Gibilisco, Michael and Sergio Montero. 2022. "Do Major-Power Interventions Encourage the Onset of Civil Conflict? A Structural Analysis." *The Journal of Politics* 84(2):944–959.
- Heger, Lindsay L., Danielle F. Jung and William A. Wagstaff. 2018. "Longitudinal Indicators of Service Provision by Terrorist and Insurgent Organizations." *On Earth Future Research*.
- Herings, P Jean-Jacques and Ronald Peeters. 2010. "Homotopy Methods to Compute Equilibria in Game Theory." *Economic Theory* 42(1):119–156.
- Holmes, E.E. E.J. Ward and M.D. Scheuerell. 2018. "Analysis of Multivariate Time-Series Using the MARSS Package." User guide. Version 3.10.10.

- Hotz, V. Joseph and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60(3):497–529.
- 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):528–549.
- Jerusalem Media and Communication Centre (JMCC). N.d. "Polls." Accessed November 2019 from <http://www.jmcc.org/polls.aspx>.
- 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.
- Magnac, Thierry and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70(2):801–816.
- Mellon, Jonathan. 2022. "Rain, Rain, Go Away: 192 Potential Exclusion-Restriction Violations for Studies Using Weather as an Instrumental Variable." Available at SSRN 3715610: <https://dx.doi.org/10.2139/ssrn.3715610>.
- Miguel, Edward, Shanker Satyanath and Ernest Sergenti. 2004. "Economic Shocks and Civil Conflict: An Instrumental Variables Approach." *Journal of Political Economy* 112(4):725–753.
- Milton-Edwards, Beverly and Stephen Farrell. 2010. *Hamas: The Islamic Resistance Movement*. Malden, MA: Polity Press. Kindle.
- Murphy, Kevin M. and Robert H. Topel. 1985. "Estimation and Inference in Two-Step Econometric Models." *Journal of Business & Economic Statistics* 3(4):370–379.
- National Consortium for the Study of Terrorism and Responses to Terrorism (START). 2019. "Codebook: Inclusion Criteria and Variables." Global Terrorism Database. Accessed December 2019.
- Palestinian Center for Policy and Survey Research (PCPSR). N.d. "Index Survey Research." Accessed November 2019 from <https://www.pcpsr.org/en/node/105>.
- Palestinian Central Bureau of Statistics (PCBS). N.d. "Press Releases." Accessed November 2019 from http://www.pcbs.gov.ps/pcbs_2012/PressEn.aspx.
- Pesendorfer, Martin and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *The Review of Economic Studies* 75(3):901–928.
- Sarsons, Heather. 2015. "Rainfall and Conflict: A Cautionary Tale." *Journal of Development Economics* 115:62–72.
- Schneider, Udo, Andreas Becker, Peter Finger, Elke Rustemeier and Markus Ziese. 2020. "GPCC Full Data Monthly Product Version 2020 at 0.25°: Monthly Land-Surface Precipitation From Rain-Gauges Built on GTS-Based and Historical Data." 10.5676/DWD_GPCC/FD_M_V2020_025.
- Shabaneh, Luay. 2008. "Measurement of Informal Sector in Adverse Conditions: Relevance of Methods." Accessed February 2019 from https://www.pcbs.gov.ps/Portals/_PCBS/Researchs/inform_sector.pdf.

- Silvey, Samuel D. 1959. "The Lagrangian Multiplier Test." *The Annals of Mathematical Statistics* 30(2):389–407.
- Stewart, Megan A. 2018. "Civil War as State-Making: Strategic Governance in Civil War." *International Organization* 72(1):205–226.
- Stock, James, Jonathan Wright and Motohiro Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business & Economic Statistics* 20(4):518–529.
- Stock, James and Motohiro Yogo. 2005. Testing for Weak Instruments in Linear IV Regression. In *Identification and Inference for Econometric Models*, ed. Donald W.K. Andrews. New York: Cambridge University Press pp. 80–108.
- Tokdemir, Efe and Seden Akcinaroglu. 2016. "Reputation of Terror Groups Dataset: Measuring Popularity of Terror Groups." *Journal of Peace Research* 53(2):268–277.
- Wagstaff, William A. and Danielle F. Jung. 2020. "Competing for Constituents: Trends in Terrorist Service Provision." *Terrorism and Political Violence* 32(2):293–324.