Online Appendix for "Tug of War: The Heterogeneous Effects of Outbidding between Terrorist Groups"

Casey Crisman-Cox* Michael Gibilisco†

Contents

*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: Equilibrium invariant distributions in the numerical example.

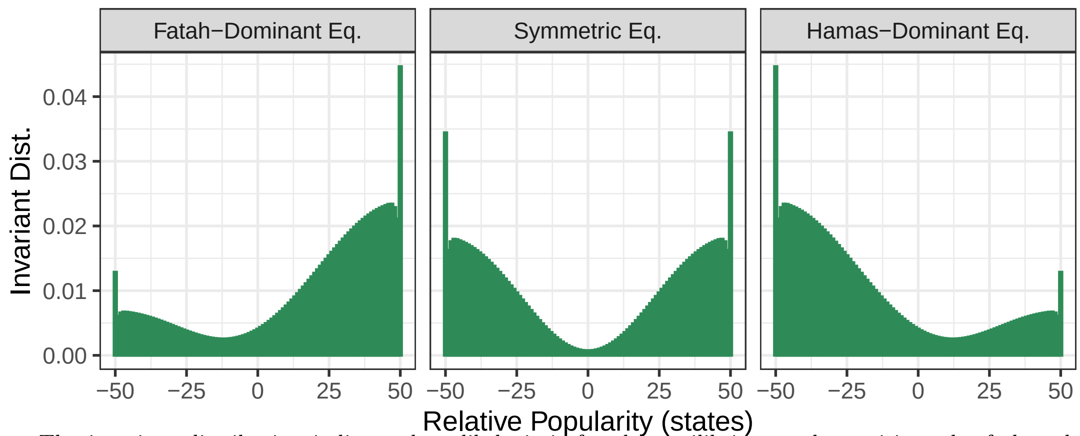

Note: The invariant distribution indicates how likely it is for the equilibrium path to visit each of the relative popularity levels in the long run. The equilibrium path in an asymmetric equilibrium is more likely to visit popularity levels that are favorable to the dominant (more violent) actor. In 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.

Figure A.2: Empirical distribution of the number of GTD attacks in each month.

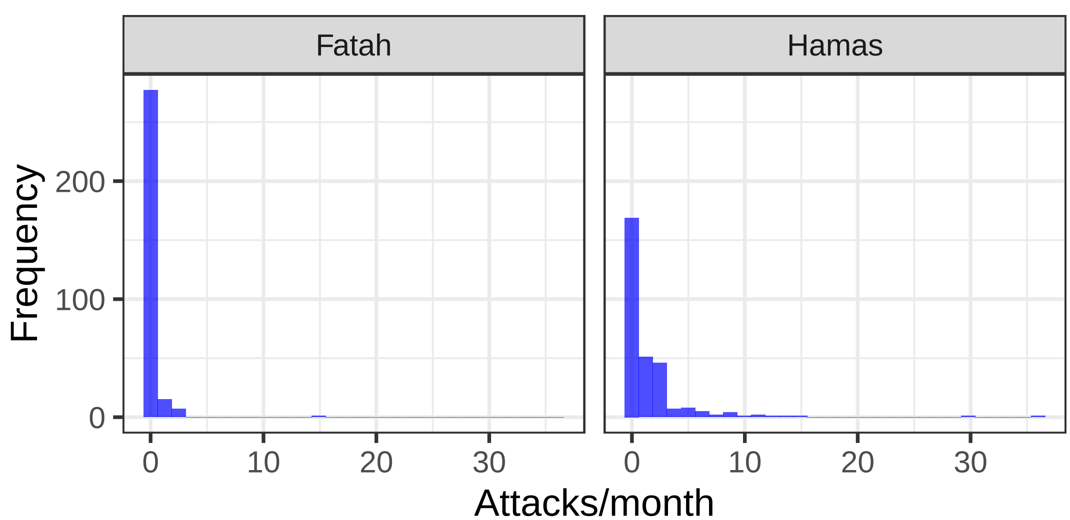

Note: For each month, we count the number of attacks attributed to Fatah (left) and Hamas (right) in the GTD and plot the distribution. For Hamas, the mean is 1.5, median is 0, and range is 0–36. For Fatah, the mean is 0.2, median is 0, and range is 0–15.

Figure A.3: Comparative statics in the numerical example.

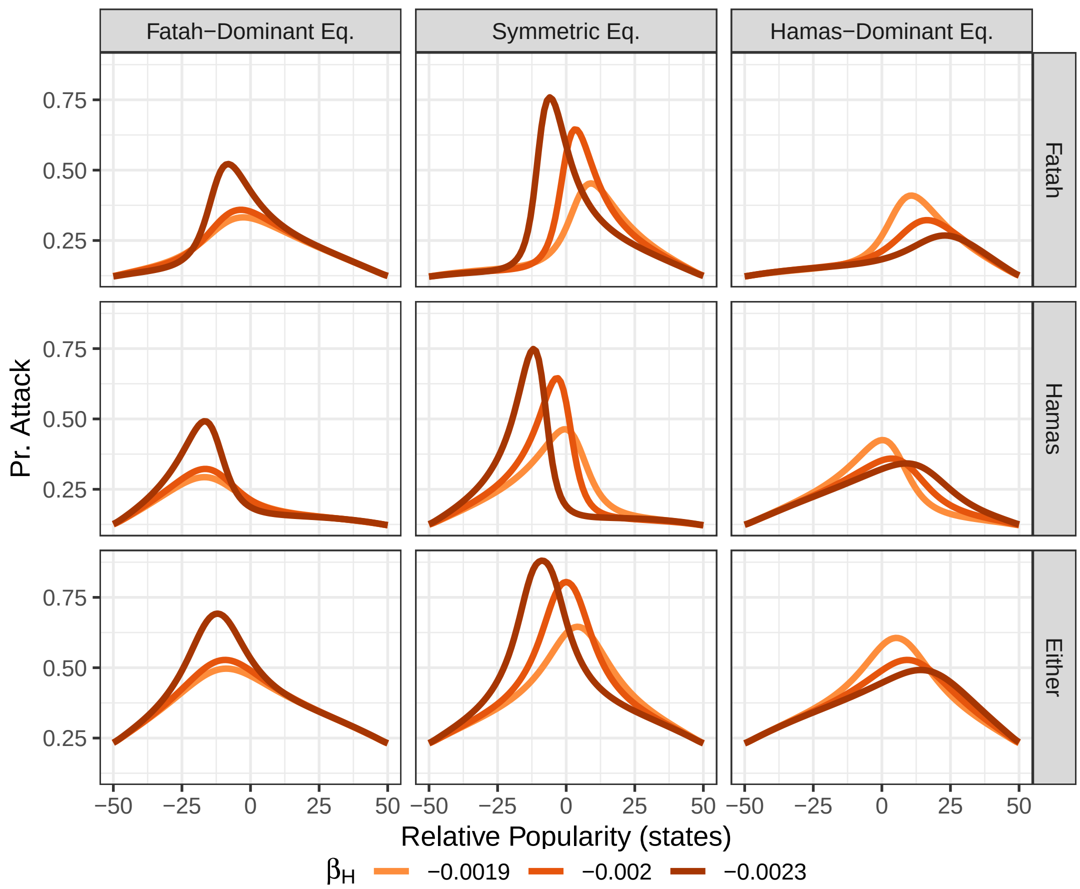

Note: We increase and decrease Hamas's value of popularity, β_H, from the baseline numerical example where $\beta_H = -0.002$. The horizontal axis is the relative popularity levels (smaller values are more favorable to Hamas) and the vertical axis is the probability of an attack. Columns correspond to 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.

Figure A.4: Estimated equilibrium attack probabilities as a function of the state

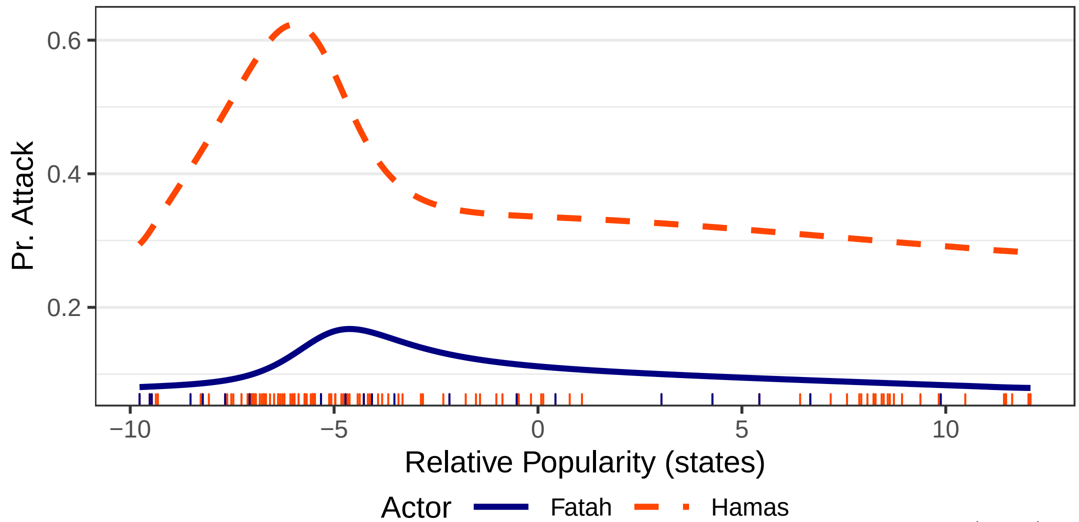

Note: Estimated probability that each group attacks as a function of the popularity level s. Hamas (Fatah) prefers smaller (larger) states. The horizontal axis includes a rug plot of observed attacks. Figure A.4 should be compared to Figure 5 which graphs these same probabilities over time as functions of observed relative popularity level s^t.

Figure A.5: Effects of competition on violence for all states s

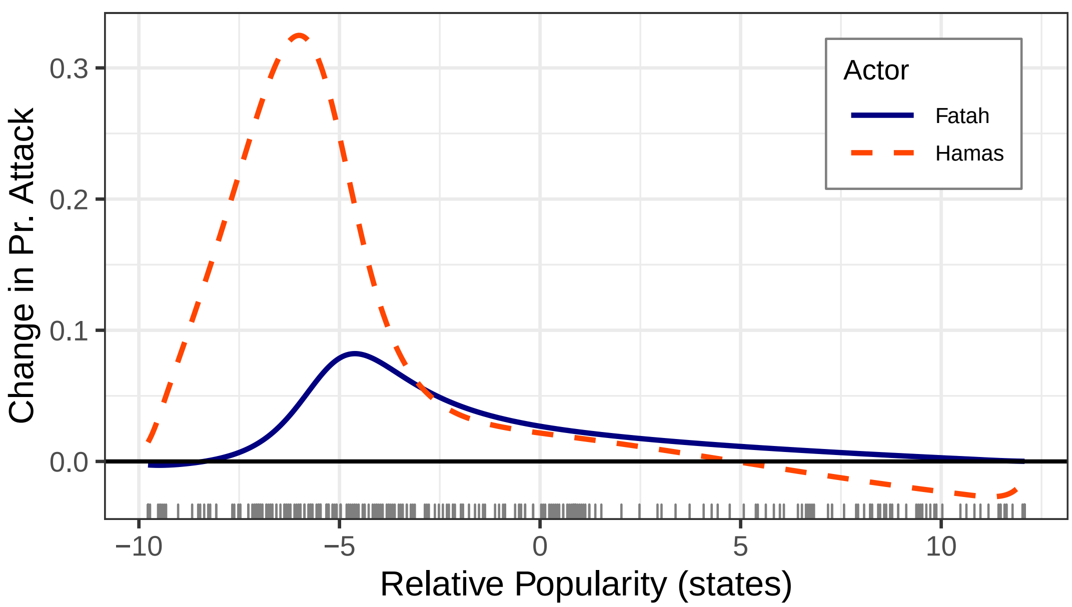

Note: We compare group i's equilibrium probability of terrorism in state s to the probability that would arise if i expects its rival to never use violence, by subtracting the latter from the former. Whereas Figure 6 graphs the difference over time conditional on the observed relative popularity s^t, Figure A.5 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.6: Relationship between terrorism and value of support in observed states.

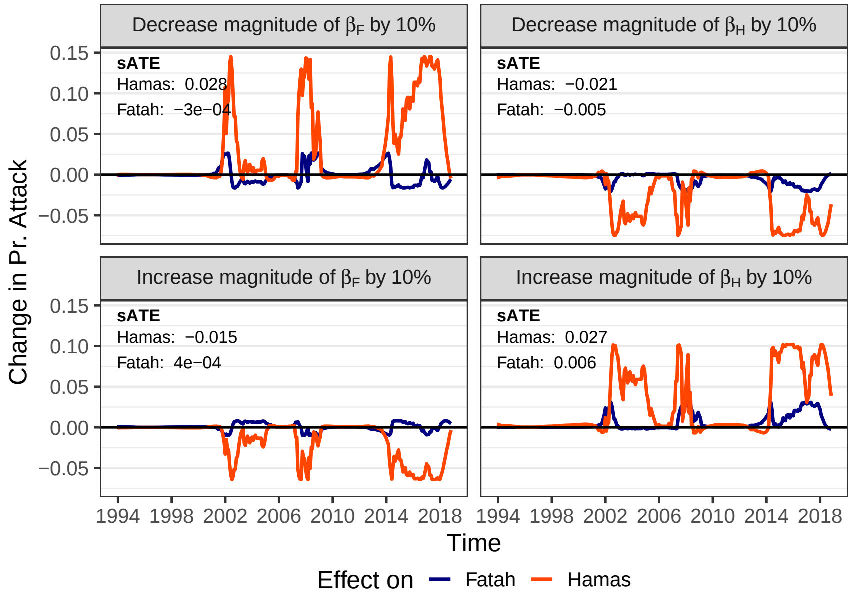

Note: In each panel, we increase and decrease the magnitude of β_i for $i = H, F$ from its estimated value by 10%; all other parameters are held constant at their estimated values. Incentives to compete are greater when the value of support, β_i, is larger in magnitude. The horizontal axis denotes the period/month t. The vertical axis is the difference between equilibrium attack probabilities (Figure 5) 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.7: Relationship between terrorism and cost of attacking in observed states.

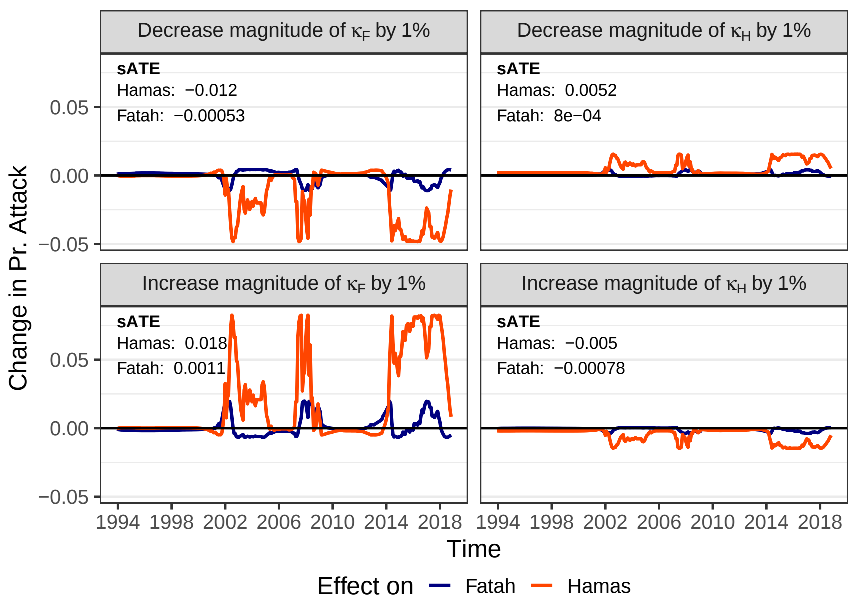

Note: In each panel, we increase and decrease the magnitude of κ_i for $i = H, F$ from its estimated value by 1%; all other parameters are held constant at their estimated values. Incentives to compete are greater when attack costs, κ_i, are closer to zero. The horizontal axis denotes the period/month t. The vertical axis is the difference between equilibrium attack probabilities (Figure 5) 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 More details on the survey data and dynamic factor model

Table B.1 lists the wordings of each survey question we used and the frequency at which it was asked. Table B.2 shows the raw correlations among the survey questions. Within each group, trust and support are highly correlated (0.69 and 0.77 for Hamas and Fatah, respectively). Likewise, support 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 is 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.

Table B.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

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

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

Let z^t denote the vector of z-transformed surveys, where for survey $j = 1, \ldots, 6$, $z_j^t = \frac{y_j^t - \bar{y}_j}{\sqrt{\text{Var}(y_j)}}$. 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

$$z^t = \tilde{s}^t \omega + \xi^t, \tag{B.2}$$

and

$$\tilde{s}^t = \rho \tilde{s}^{t-1} + \alpha_0 + a_H^{t-1} \cdot \alpha_H + a_F^{t-1} \cdot \alpha_F + \eta^t. \tag{B.3}$$

Table B.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	−0.44	0.98	−0.72
Trust in Fatah	−0.19	1.00	−0.26	0.69	−0.54	0.92
Support for Hamas	0.77	−0.26	1.00	−0.57	0.89	−0.83
Support for Fatah	−0.44	0.69	−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

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 states and actions. Note that simply including attacks in the measurement model does not 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.

The estimates of (ω, α) are reported in Table B.3, while the estimates of $\tilde{s}^t$ are presented in Figure 4 (main text). Notice that the variables 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 $\tilde{s}^t$ from the following six models.

1. Main specification (described above)

Table B.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

2. Fix $\rho = 1$. Modify Eq. B.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. B.2 s.t. $\xi^t \sim N(0, \sigma_\xi^2 \mathbf{1})$.

4. Estimated and heteroskedastic variance. Modify Eq. B.2 s.t. $\xi^{t,j} \sim N(0, \sigma_{\xi,j}^2 \mathbf{1})$, where $j = 1, \ldots, 6$ indexes surveys

5. A model with $\alpha = 0$. Modify Eq. B.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.

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 six specifications, each of which produces its own estimate of $\tilde{s}^t$. In Table B.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 $\tilde{s}^t$.

Table B.4: Correlations across measurement model specifications

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Model 1	1.00	1.00	1.00	0.88	1.00	0.99
Model 2	1.00	1.00	0.99	0.87	0.99	0.99
Model 3	1.00	0.99	1.00	0.88	1.00	0.98
Model 4	0.88	0.87	0.88	1.00	0.89	0.91
Model 5	1.00	0.99	1.00	0.89	1.00	0.99
Model 6	0.99	0.99	0.98	0.91	0.99	1.00

C First-stage robustness and transition matrix

In this appendix, we first consider the robustness of the first-stage estimates (Table 1) to additional controls variables (Table C.3) and to changes in how we measure attacks (Table C.4). We then describe how to transform the first-stage estimates into the Markov transition probabilities used in the CMLE to estimate the payoff parameters $\theta = (\beta, \kappa)$.

C.1 Robustness

First, we want to ensure that the relationship between attacks and relative popularity is not driven by alternative factors. In particular, we want to be sure that variables omitted from the first-stage 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 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 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 B.2 and B.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 $\tilde{s}^t$). 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 Jerusalem Media and Communication Centre (JMCC) (N.d) 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
- 3 responses about whether the current peace process is alive, dead, or unclear.

Additionally, we add 4 responses from surveys by Palestinian Center for Policy and Survey Research (PCPSR) (N.d) that record support for armed attacks against

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

Many of these variable are correlated. 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 C.1. All of 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. Additionally, we see that the latent support for violence reaches a minimum during Oslo and a maximum during the second Intifada. 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 pollster they are unemployed to Jerusalem Media and Communication Centre (JMCC) (N.d) pollsters
2. % of respondents telling pollster they are unemployed to Palestinian Center for Policy and Survey Research (PCPSR) (N.d) pollsters
3. Estimated true unemployment by Palestinian Center for Policy and Survey Research (PCPSR) (N.d)
4. Unemployment rate reported in Labor Force Surveys published by the Palestinian Central Bureau of Statistics (PCBS) (N.d)

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

Variable	Est.
% Supporting two-state solution	-0.14
% Supporting a one shared state solution	0.02
% Supporting a one Islamic state solution	0.14
% Saying there is no solution	0.13
% Supporting a peace process	-0.16
% Opposing a peace process	0.16
% Supporting military action against Israel	0.16
% Opposing military action against Israel	-0.16
% Supporting suicide bombings	0.16
% Opposing suicide bombings	-0.16
% Very optimistic about peace	-0.15
% Optimistic about peace	-0.15
% Pessimistic about peace	0.12
% Very pessimistic about peace	0.16
% Strongly support Oslo	-0.11
% Support Oslo	-0.10
% Oppose Oslo	0.13
% Support the Intifada	0.09
% Oppose the Intifada	-0.08
% Who think peace is dead	0.09
% Who think the peace process is stalled	-0.02
% Who think peace is alive	-0.12
% Support armed attacks generally	0.16
% Support armed attacks against civilians	0.15
% Support armed attacks against soldiers	0.12
% Support armed attacks against settlers	0.11

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

Our robustness checks are reported in Table C.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 Models 2 and 3, 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 4 and 5, we consider some political context with an indicator for whether the Second Intifada is ongoing and the time since the last Israeli election. Finally, in Model 6, we control for the

Table C.2: ML estimates for latent unemployment conditions

Variable	Est.
Self reported unemployment rate (JMCC)	0.15
Self reported unemployment rate (PCPSR)	0.16
Estimated unemployment rate (PCPSR)	0.06
Estimated unemployment rate (PCBS)	0.17

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.[1]

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

	Dependent variable:					
	Δ State					
	(1)	(2)	(3)	(4)	(5)	(6)
Hamas Attacks	−0.21	−0.20	−0.23	−0.22	−0.20	−0.34
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.05)
Fatah Attacks	1.00	1.06	1.04	1.04	1.05	1.13
	(0.08)	(0.04)	(0.04)	(0.04)	(0.03)	(0.09)
Δ Lag state	0.34	0.30	0.24	0.20	0.18	0.14
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Δ unemployment		−0.25	−0.13	−0.09	−0.06	−0.06
		(0.09)	(0.08)	(0.05)	(0.06)	(0.08)
Δ support for violence			−0.26	−0.30	−0.34	−0.26
			(0.06)	(0.06)	(0.06)	(0.07)
Second Intifada				−0.12	−0.13	−0.12
				(0.04)	(0.04)	(0.03)
Time since last election					−0.003	−0.01
					(0.001)	(0.001)
Palestinian fatalities						0.005
						(0.01)
Constant	−0.01	−0.01	0.01	0.02	0.08	0.11
	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)
State-attack interactions	No	Yes	Yes	Yes	Yes	Yes
T	298	298	298	298	298	212
adj. R^2	0.720	0.748	0.788	0.800	0.812	0.822
$\hat{\sigma}$	0.183	0.174	0.160	0.155	0.150	0.146

Note: Newey-West standard errors in parenthesis.

[1] Data are only available after the start of the Second Intifada, which explains why Model 6 has fewer observations than other models in the table.

Additionally, we want to be sure that our results are not dependent on the binary coding of attacks. As such we consider four alternative measurements for attacks: counts, binary with fatalities as a control, only fatalities, and fatalities per attacks. Across these four models in Table C.4 (which should be compared to Table 1 in the main text), we see that the main results survive: (i) violence by actor i makes i more relatively popular and (ii) the effects of Fatah's violence are larger in magnitude than the effects of Hamas's violence.[2] This consistency across attacks measures 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.

C.2 Building the transition matrix

To build the transition matrix, we first define the lowest (most Hamas friendly) state s_1 of the discrete state as the 2.5th percentile of the continuous state variable, $\tilde{s}^t$, which we estimated using the dynamic factor analysis. Likewise, the largest state (most Fatah friendly) s_K is set to the 97.5th percentile. States between 1 and K are defined at equally spaced intervals of $2d = 0.05$ (i.e., $s_2 - s_1 = 0.05$). Let $\mu[a, s; \hat{\gamma}]$ be the fitted values from the first model in Table 1 for all possible combinations of action profiles with the discrete states. Plugging these fitted values and the estimated standard deviation $\hat{\sigma}$ (also in Table 1) into Equation 3 produces the estimated transition probabilities. Finally, we discretize the observed states variable by mapping values of the continuous latent variable $\tilde{s}^t$ into the closest value of the discrete state space $\mathcal{S}$, i.e., $s^t = \underset{s_k \in \mathcal{S}}{\operatorname{argmin}} (\tilde{s}^t - s_k)^2$.

D Standard errors and sensitivity analysis

In this appendix, we describe the standard errors reported for the two-step CMLE estimates and consider how sensitive the structural estimates of β and κ 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). Aguirregabiria and Mira (2007) use this result to describe the asympotic distribution of the two-step pseudo-likelihood estimator from Hotz and Miller (1993) and we follow the same approach here. Specifically, let $\theta_2 = (\beta, \kappa, \upsilon)$ be the set of parameters estimated in the second stage, then the two-step correction gives the variance of $\hat{\theta}_2$ as

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

[2]The only circumstance where we fail to reject the null of equally effective groups is in model 4 (fatalities per attack) from Table C.4, but only when $\tilde{s}^t \leq -10$. So if Hamas is *extremely* popular, the two groups might be equally effective at moving their relative popularity level according to this one specification.

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

	Dependent variable:			
	Δ State			
	(1)	(2)	(3)	(4)
Hamas attacks (count)	−0.03			
	(0.01)			
Fatah attacks (count)	0.41			
	(0.05)			
Hamas attacks (binary)		−0.20		
		(0.03)		
Fatah attacks (binary)		1.02		
		(0.04)		
Hamas Fatalities		0.0005	−0.01	
		(0.001)	(0.002)	
Fatah Fatalities		0.01	0.31	
		(0.01)	(0.05)	
Hamas fatalities/attack				−0.02
				(0.01)
Fatah fatalities/attack				0.67
				(0.07)
Second Intifada	−0.17	−0.13	−0.16	−0.20
	(0.04)	(0.04)	(0.04)	(0.05)
Δ Lag unemployment	−0.08	−0.06	−0.06	−0.06
	(0.06)	(0.06)	(0.06)	(0.06)
Δ Lag support for violence	−0.31	−0.34	−0.31	−0.31
	(0.07)	(0.06)	(0.06)	(0.08)
Time since last election	−0.004	−0.003	−0.004	−0.003
	(0.001)	(0.001)	(0.001)	(0.001)
Δ Lag state	0.11	0.18	0.15	0.17
	(0.06)	(0.04)	(0.06)	(0.07)
Constant	0.08	0.08	0.07	0.07
	(0.04)	(0.03)	(0.03)	(0.04)
State-attack interactions	Yes	Yes	Yes	Yes
T	298	298	298	298
adj. R^2	0.508	0.813	0.394	0.431
$\hat\sigma$	0.243	0.150	0.270	0.261

Note: Newey-West standard errors in parenthesis.

as described in Aguirregabiria and Mira (2007, Proposition 1). Here $\hat\Sigma_{\theta_2}$ is the ordinary CMLE covariance matrix described by Silvey (1959) which is

$$\hat\Sigma_{\theta_2} = \begin{bmatrix} H_{\theta_2} L(\hat v|Y) + J_{\theta_2}\mathcal{V}(\hat\theta_2;\hat\gamma)^{\mathrm{T}} J_{\theta_2}\mathcal{V}(\hat\theta_2;\hat\gamma) & -J_{\theta_2}\mathcal{V}(\hat\theta_2;\hat\gamma)^{\mathrm{T}} \\ -J_{\theta_2}\mathcal{V}(\hat\theta_2;\hat\gamma) & \mathbf{0} \end{bmatrix}^{-1},$$

where H_x and J_x respectively denote the Hessian and the Jacobian of a function with respect to x.

The 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

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

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

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

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.[3] The final piece is the first-stage covariance matrix $\hat{\Sigma}_\gamma$, which we construct using a parametric bootstrap.

Figure D.1: Sensitivity of structural estimates to first-stage results

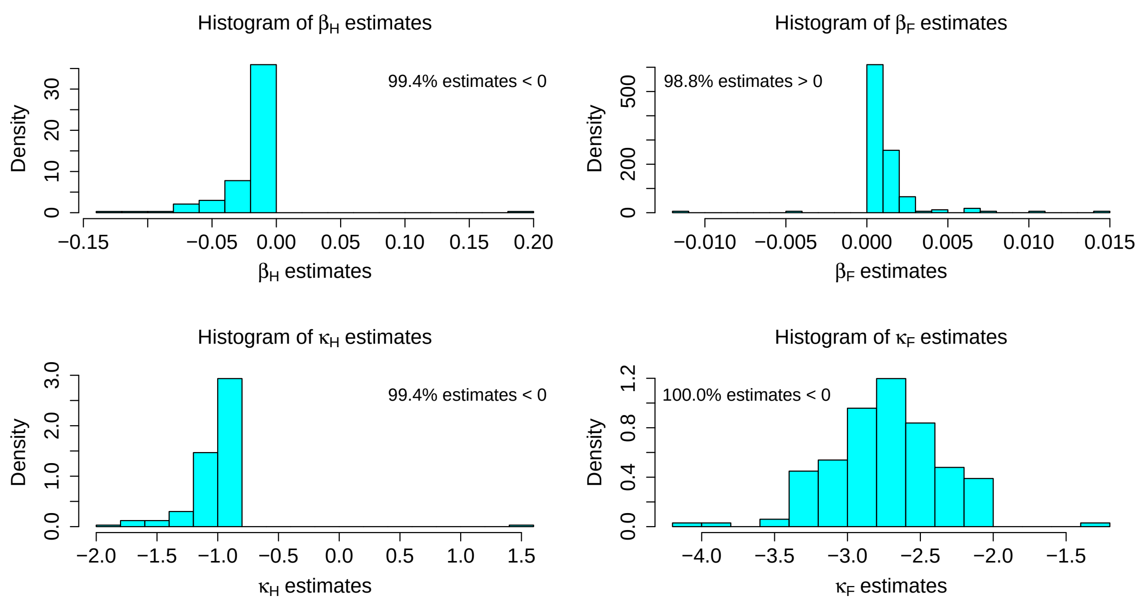

Notes: Hypothesis tests are one-sided z tests at the same levels as reported in Table 2 (Main text).

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 $b = 1, 2, \ldots, B$ we conduct the following exercise:

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

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

2. Re-fit the first-stage model to produce new estimates $\hat{\gamma}_b$. Estimate $\Sigma_{\gamma,b}$ with a parametric bootstrap.

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-stage 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 the results are reported as histograms in Figure D.1.

There are few points of interest in Figure D.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.

E Model fit

We conduct two model fit exercises to weigh the evidence that outbidding theory has some explanatory power in the Palestinian conflict. In both exercises, we compare our fitted model to a "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 and 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 and cannot be part of that because violence does not affect popularity.

In the first exercise, 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

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

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 exercise a little more, by considering an overall measure that aggregates across actors

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

In the second exercise, we consider a likelihood ratio test of the restricted no-competition model against the main model. Here we refit both steps of the model under the null hypothesis that $\gamma_{F,1} = \gamma_{F,2} = \gamma_{H,1} = \gamma_{H,2} = 0$. Note that this hypothesis removes 6 degrees of freedom as the β parameters are unidentified under the null.

Table E.1: Model fit results

	Main model	No Competition
ePCP–Hamas	0.54	0.51
ePCP–Fatah	0.83	0.86
ePCP–Overall	0.46	0.44
LR test (d.f. = 6)	$\chi^2 = 275$ ($p < 0.01$)	

Table E.1 reports the results from both exercises. The first 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 presented here is expected to correctly predict 54% of Hamas' actions and 83% of Fatah's. For the no-competition model these values are 51% and 86% respectively. Note that both 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 no-competition model is slightly preferred. This difference suggests that outbidding plays a larger role in Hamas' overall plans during this period while Fatah may have other concerns that lead it to place less emphasis on competition. This understanding broadly fits with the historical accounts of this period where Hamas is the up-and-coming actor for much of the data. 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. This exercise opens the door for other, competing theories and models of these data that may improve on our outbidding-based approach. Turning to likelihood ratio test, we find strong evidence against the null hypothesis. To put this another way, the competition based model does a better job of explaining the observed actions and state transitions than a model without competition.

Table E.2: Robustness to different time periods

	Full Sample	2014 Agreement	2011 Agreement	2006 Elections	Second Intifada†
β_H	-0.007	-0.007	-0.007	-0.021	-0.007
	(0.004)	(0.005)	(0.004)	(0.013)	(0.004)
β_F	0.0005	0.0006	0.0005	0.0005	0.0005
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0152)
κ_H	-0.95	-0.90	-0.69	-0.74	-0.89
	(0.23)	(0.23)	(0.22)	(0.26)	(0.10)
κ_F	-2.45	-2.47	-2.31	-2.34	-2.89
	(0.28)	(0.31)	(0.28)	(0.26)	(0.23)
T	300	243	208	144	80
LL	-278.09	-230.98	-206.84	-125.27	-71.82

Note: †CMLE did not converge, estimates from nested-pseudo-likelihood (NPL) estimator. Samples begin at Jan. 1994 and end in one month prior to the event listed.

F Robustness to different time spans

In this appendix, we consider four different alternative time frames. 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 three cutpoints:

1. The formation of the Fatah-Hamas unity government in April 2014 (last month is Mar. 2014).

2. The signing of the first Fatah-Hamas unity agreement in 2011 (last month is Apr. 2011).

3. Hamas wins the 2006 legislative elections (last month is Dec. 2005).

4. The start of the Second Intifada in 2001 (last month is Aug. 2000).

The results of these short-T robustness checks are presented in Table E.2 along with a reproduction of the main model (1994-2018) for comparison. Overall we see very few changes across the four models. In general, the estimates are very stable across samples. The CMLE failed to converge in the last sample so we employ an alternative estimation method called the nested-pseudo-likelihood estimator (Aguirregabiria and Mira 2007; Crisman-Cox and Gibilisco 2021).

Table G.1: Discount factors and model fit

δ	Log-Likelihood
0	-284.18
0.9	-280.98
0.925	-281.54
0.95	-282.43
0.975	-283.67
0.99	-280.72
0.999	-278.09
0.9999*	-288.14

Note: * Model failed to converge.

G 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 G.1 shows the log-likelihood 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 G.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 A.4 from the baseline analysis. The graphs in Figure G.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 across for $\delta \in \{0.99, 0.999\}$. Substantively, the predictions in the estimated model are roughly invariant to choosing a discount factor between $[0.99, 0.999]$. These similarities likely arise because discount factors are not identified in dynamic discrete choice models generally (Magnac and Thesmar 2002). Even with suitable exclusion restrictions, the discount factor is not even guaranteed to be point identified (Abbring and Daljord 2020). When δ is not identified, we would expect several values of δ to return essentially identical equilibrium attack probabilities, which is what we see in Figure A.4.

References

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

Figure G.1: Discount factors and equilibrium attack probabilities

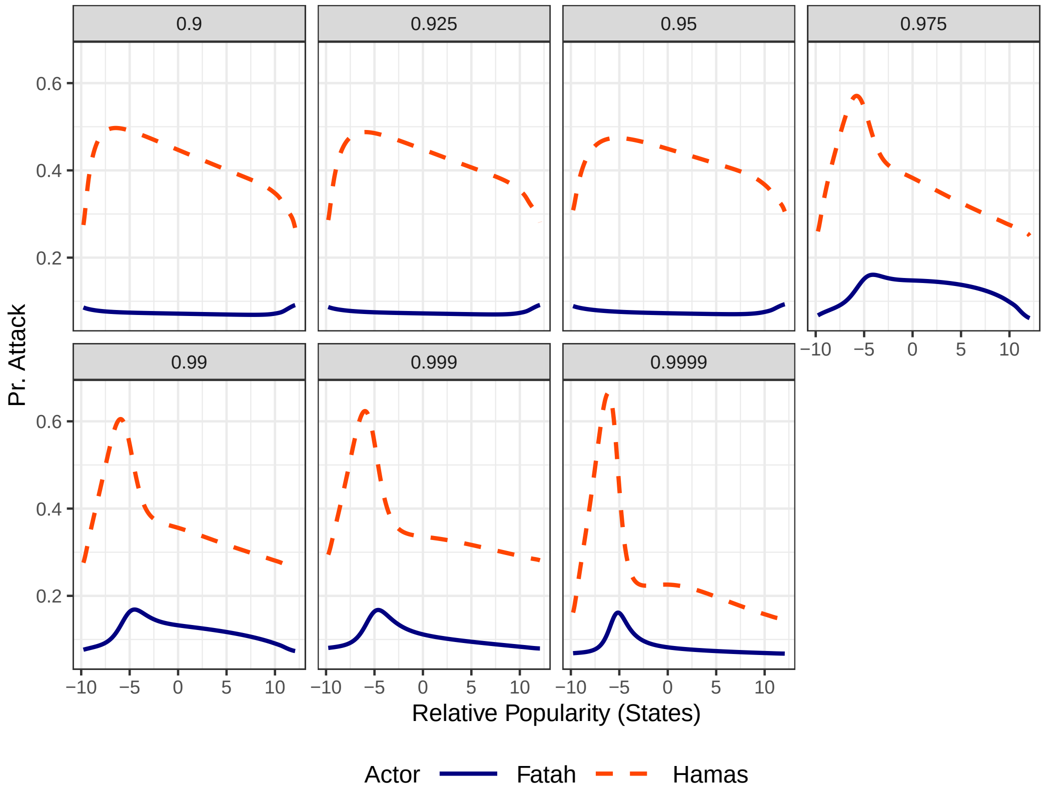

Note: Graphs of the estimated equilibrium attack probabilities at 7 different discount factors. Setting $\delta = 0.999$ (second row, second column) is identical to Figure A.4. The CMLE model failed to converge when $\delta = 0.9999$ as in Table G.1.

Aguirregabiria, Victor and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75(1):1–53.

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.

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.

Jerusalem Media and Communication Centre (JMCC). N.d. "Polls." Accessed November 2019 from http://www.jmcc.org/polls.aspx.

Magnac, Thierry and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70(2):801–816.

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.

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.

Silvey, S. D. 1959. "The Lagrangian Multiplier Test." *The Annals of Mathematical Statistics* 30(2):389–407.