

University of Vermont

UVM ScholarWorks

College of Engineering and Mathematical
Sciences Faculty Publications

College of Engineering and Mathematical
Sciences

2-1-2020

Noncooperative dynamics in election interference

David Rushing Dewhurst
University of Vermont

Christopher M. Danforth
University of Vermont

Peter Sheridan Dodds
University of Vermont

Follow this and additional works at: <https://scholarworks.uvm.edu/cemsfac>



Part of the [Community Health Commons](#), [Human Ecology Commons](#), [Nature and Society Relations Commons](#), [Place and Environment Commons](#), and the [Sustainability Commons](#)

Recommended Citation

Dewhurst DR, Danforth CM, Dodds PS. Noncooperative dynamics in election interference. *Physical Review E*. 2020 Feb 19;101(2):022307.

This Article is brought to you for free and open access by the College of Engineering and Mathematical Sciences at UVM ScholarWorks. It has been accepted for inclusion in College of Engineering and Mathematical Sciences Faculty Publications by an authorized administrator of UVM ScholarWorks. For more information, please contact scholarworks@uvm.edu.

Noncooperative dynamics in election interferenceDavid Rushing Dewhurst,^{*} Christopher M. Danforth,[†] and Peter Sheridan Dodds[‡]*MassMutual Center of Excellence in Complex Systems and Data Science, Computational Story Lab, Vermont Complex Systems Center, and Department of Mathematics and Statistics, University of Vermont, Burlington, Vermont 05405, USA*

(Received 7 August 2019; revised manuscript received 29 October 2019; accepted 20 January 2020; published 19 February 2020)

Foreign power interference in domestic elections is an existential threat to societies. Manifested through myriad methods from war to words, such interference is a timely example of strategic interaction between economic and political agents. We model this interaction between rational game players as a continuous-time differential game, constructing an analytical model of this competition with a variety of payoff structures. All-or-nothing attitudes by only one player regarding the outcome of the game lead to an arms race in which both countries spend increasing amounts on interference and counterinterference operations. We then confront our model with data pertaining to the Russian interference in the 2016 United States presidential election contest. We introduce and estimate a Bayesian structural time-series model of election polls and social media posts by Russian Twitter troll accounts. Our analytical model, while purposefully abstract and simple, adequately captures many temporal characteristics of the election and social media activity. We close with a discussion of our model's shortcomings and suggestions for future research.

DOI: [10.1103/PhysRevE.101.022307](https://doi.org/10.1103/PhysRevE.101.022307)**I. INTRODUCTION**

In democratic and nominally democratic countries, elections are societally and politically crucial events in which power is allocated [1]. In fully democratic countries, elections are the method of legitimate governmental change [2]. One country, labeled “Red,” wishes to influence the outcome of an election in another country, labeled “Blue,” because of the impact that elections in Blue have on Red’s national interest. Such attacks on democracies are not new. It is estimated that the United States (U.S.) and Russia (and its predecessor, the Soviet Union) often interfere in the elections of other nations and have consistently done this since 1946 [3]. Though academic study of this area has increased [4], we are unaware of any formal modeling of noncooperative dynamics in an election interference game. Recent approaches to the study of this phenomenon have focused mainly on the compilation of coarse-grained (e.g., yearly frequency) panels of election interference events and qualitative analysis of this data [5,6] and data-driven studies of the aftereffects and second-order effects of interference operations [7,8]. Attempts to create theoretical models of interference operations are less common. These attempts include qualitative causal models of cyberoperation influence on voter preferences [9] and models of the underlying reasons that a state may wish to interfere in the elections of another [10].

We consider a two-player game in which one country wants to influence a two-candidate, zero-sum election taking place in another country. We think of Red as the foreign

intelligence service of the influencing country and Blue as the domestic intelligence service of the country in which the election is held. Red wants a particular candidate, which we will set to be candidate A without loss of generality, to win the election, while Blue wants the effect of Red’s interference to be minimized. We derive a noncooperative, non-zero-sum differential game to describe this problem and then explore the game numerically. We find that all-or-nothing attitudes by either Red or Blue can lead to arms-race conditions in interference operations. In the event that one party credibly commits to playing a fixed, deterministic strategy, we derive further analytical results.

We then confront our model with data pertaining to the 2016 U.S. presidential election contest, in which Russia interfered [11]. We fit a Bayesian structural time-series model to election polls and social media posts authored by Russian military intelligence-associated troll accounts. We demonstrate that our model, though simple, captures many of the observed and inferred parameters’ dynamics. We close by proposing some theoretical and empirical extensions to our work.

II. THEORY**A. Election interference model**

We consider an election between two candidates with no electoral institutions such as an Electoral College. We assume that the election process at any time $t \in [0, T]$ is represented by a public poll $Z_t \in [0, 1]$. The model is set in continuous time, though when we estimate parameters statistically in Sec. III we move to a discrete-time analog. We hypothesize that the election dynamics take place in a latent space where dynamics are represented by $X_t \in \mathbb{R}$. We will set $X_t < 0$ to be values of the latent poll that favor candidate A and $X_t > 0$

^{*}david.dewhurst@uvm.edu[†]chris.danforth@uvm.edu[‡]peter.dodds@uvm.edu




FIG. 1. The random-walk latent space election model is an accurate approximation to multiple different population candidate preference updates. The latent election process evolves according to $X_{k+1} = X_k + \frac{1}{N} \sum_{1 \leq n \leq N} \xi_{n,k}$. The random variable $\xi_{n,k}$ is voting agent n 's shift toward the left or right of the political spectrum at time k . In the center panel, the solid curve is a draw from the latent election process resulting from the preference updates $\xi_{n,t} \sim B(0.1 \frac{T-t}{T} + 1.5 \frac{t}{T}, 0.1 \frac{T-t}{T} + 1.5 \frac{t}{T})$, where $B(\alpha, \beta)$ is the Beta distribution and we have set $T = 365$. As $t \rightarrow T$, the electorate exhibits increasing resistance to change in their political viewpoints. We display the preference shift distributions at $t = 0$ in panel (a) and at $t = T$ in panel (b). For contrast, the dashed curve is a draw from the latent election process resulting from $\xi_{n,t} \sim B(1.5 \frac{T-t}{T} + 0.1 \frac{t}{T}, 1.5 \frac{T-t}{T} + 0.1 \frac{t}{T})$, which describes an electorate in which agents more often have changing political preferences as $t \rightarrow T$. We show the corresponding preference shift distributions at $t = 0$ in panel (d) and at $t = T$ in panel (e). Despite these preference updates that are, in some sense, opposites of each other, the latent processes X_t are statistically very similar and are both well modeled by the continuum approximation $dX_t = \sigma dW_t$.

that favor candidate B. The latent and observable space are related by $Z_t = \phi(X_t)$, where ϕ is a sigmoidal function which we choose to be $\phi(x) = \frac{1}{1+e^{-x}}$. (Any sigmoidal function that is bounded between zero and one will suffice and lead only to different parameter estimates in the context of statistical estimation.) The actual result of the election—the number of votes that are earned by candidate B—is given by $\phi(X_T)$. The election takes place in a population of N voting agents. Each voting agent updates their preferences over the candidates at each time step t_n by a random variable ξ_{n,t_n} . These random variables satisfy $E_n[\xi_{n,t_k}] = 0$ and $E_n[\xi_{n,t_k}^2] < \infty$ for all t . The increments of the election process are the sample means of the voting agents' preferences at time t . In the absence of interference, the stochastic election model is an unbiased random walk:

$$X_{t_{k+1}} = X_{t_k} + \frac{1}{N} \sum_{1 \leq n \leq N} \xi_{n,t_k} \Delta t_k, \quad (1)$$

where we have put $\Delta t_k = t_{k+1} - t_k$. We display sample realizations of this process for different distributions of ξ_{n,t_k} in Fig. 1. Though one distribution of preference changes has a larger variance than the other, the sample paths of X_k are statistically similar for each since $\frac{1}{N} \sum_n \xi_{n,t_k}$ does not vary much

between the distributions. When N is large we can reasonably approximate this discrete agent electoral process by a Wiener process, $dX_t = \sigma dW_t$, where $\sigma^2 \approx \text{Var}(\frac{1}{N} \sum_{1 \leq n \leq N} \xi_{n,t})$. This limit is valid in the limit of large N .

If the preference change random variables $\xi_{n,k}$ did not satisfy $E_n[\xi_{n,k}] = 0$, then the random-walk approximation to Eq. (1) would not necessarily be valid. For example, if $\{\xi_{n,k}\}_{k \geq 0}$ were a random walk or were trend stationary for each n , then $\{E_n[\xi_{n,k}]\}_{k \geq 0}$ would also respectively be a random walk or trend stationary. A trend-stationary univariate time series is a stochastic process $x_t = f(t) + \varepsilon_t$, where ε_t is a stationary process and $f(\cdot)$ is a deterministic function of time [12,13]. A univariate time series with a unit root is a time series that can be written as an autoregressive process of order p [$AR(p)$], $x_t - \sum_{p'=1}^p \beta_{p'} x_{t-p'} = \varepsilon_t$, such that the polynomial $z^p - \sum_{p'=1}^p \beta_{p'} z^{t-p'}$ has a root on the unit circle when solved over the complex numbers. A random walk is a special case of the $AR(p)$ process with characteristic polynomial given by $z - 1 = 0$, which has the unit root $z = 1$. Trend-stationary and unit-root time series differ fundamentally in that a trend-stationary process subjected to an exogenous shock will eventually revert to its mean function $f(t)$. This is not the case for a stochastic process with a

unit root. A unit root or trend-stationary $\xi_{n,k}$ would model a population in which political preferences were undergoing a shift in population mean rather than just in individual preferences.

However, there do exist cases where the random-walk approximation is valid even when $E_n[\xi_{n,k}] \neq 0$. If the stochastic evolution equation for $E_n[\xi_{n,k}]$ has a stationary colored noise with exponentially decaying covariance function as its solution, then the integral of this noise satisfies a Stratonovich-type equation [14–16]. This equation would be a generalized version of the basic random-walk model considered here, but we will not consider this scenario in the remainder of this work.

We denote the control policies of Red and Blue (the functions by which Red and Blue attempt to influence (or prevent influence on) the election) by $u_R(t)$ and $u_B(t)$. These functions are one-dimensional continuous-time stochastic processes (time series). The term “policy” originates from the fields of economics and reinforcement learning [17–19]. These control policies are abstract variables in the context of our model, but we interpret them as expenditures on interference operations. We assume that Red and Blue can affect the mean trajectory of the election but not its volatility (standard deviation of its increments). We make this assumption because X_t is an approximation to the process described by Eq. (1). As we show in Fig. 1, the variance of the electoral process does not change much even when the voting population’s underlying preference change distributions have differing variance and kurtosis. Under the influence of Red’s and Blue’s control policies, the election dynamics become

$$dX_t = F(u_R(t), u_B(t)) dt + \sigma d\mathcal{W}_t, \quad X_0 = y. \quad (2)$$

The function F captures the mechanism by which Red and Blue affect the mean dynamics of the latent electoral process. We assume that F is at least twice continuously differentiable for convenience. To first-order expansion we have $F(u_R(t), u_B(t)) = a_0 + a_R u_R(t) + a_B u_B(t) + \mathcal{O}(u^2)$, which is most accurate near $u = 0$. We approximate the state equation by

$$dX_t = [u_R(t) + u_B(t)] dt + \sigma d\mathcal{W}_t, \quad X_0 = y, \quad (3)$$

since we have assumed zero endogenous drift and can absorb constants into the definition of the control policies. We will use Eq. (3) as the state equation for the remainder of the paper.

B. Subgame-perfect Nash equilibria

Red and Blue each seek to minimize separate scalar cost functionals of their own control policy and the other agent’s control policy. We will assume that the agents do not incur a running cost from the value of the state variable, although we will revisit this assumption in Sec. IV. The cost functionals are therefore

$$E_{u_R, u_B, X} \left\{ \Phi_R(X_T) + \int_0^T C_R(u_R(t), u_B(t)) dt \right\} \quad (4)$$

and

$$E_{u_R, u_B, X} \left\{ \Phi_B(X_T) + \int_0^T C_B(u_R(t), u_B(t)) dt \right\}. \quad (5)$$

The functions C_R and C_B represent the running cost or benefit of conducting election interference operations. We assume the cost functions have the form

$$C_i(u_R, u_B) = u_i^2 - \lambda_i u_{-i}^2 \quad (6)$$

for $i \in \{R, B\}$. The notation $-i$ indicates the set of all other players. For example, if $i = R$, $-i = B$. This notation originates in the study of noncooperative economic games. The non-negative scalar λ_i parameterizes the utility gained by player i from observing player $-i$ ’s effort. If $\lambda_i > 0$, player i gains utility from player $-i$ ’s expending resources, while if $\lambda_i = 0$, player i has no regard for $-i$ ’s level of effort but only for their own running cost and the final cost. Our assumption that cost accumulates quadratically with magnitude of the control policy is common in optimal control theory [20–22]. We can justify the functional form of Eq. (6) as follows. Suppose that an arbitrary analytic cost function for player i as $\mathcal{C}_i(u_R, u_B) = \mathcal{C}^{(i)}(u_i) - \lambda_i \mathcal{C}^{(-i)}(u_{-i})$. We make the following assumptions:

(i) It is equally costly to for player i to conduct operations that favor candidate A or candidate B. This imposes the constraint that $\mathcal{C}^{(i)}$ and $\mathcal{C}^{(-i)}$ are even functions.

(ii) Player i conducting no interference operations results in player i ’s incurring no direct cost from this choice. In other words, if $u_i(t) = 0$ at some t , then player i does not incur any cost from this.

With these assumptions, the first nonzero term in the Taylor expansion of \mathcal{C}_i is given by Eq. (6).

1. Choice of final conditions

Finding optimal play in noncooperative games often requires solving the game backward through time [23–26]. Therefore, we must define final conditions that specify the cost that Red and Blue incur from the actual election result $\phi(X_T)$. Red and Blue might have different final conditions because of their qualitatively distinct objectives. Since Red wants to influence the outcome of the election in Blue’s country in favor of candidate A, their final cost function Φ_R must satisfy $\Phi_R(x) < \Phi_R(y)$ for all $x < 0$ and $y > 0$. In the final conditions that we are about to present, we also assume that Φ_R is monotonically nondecreasing everywhere. We relax these assumption in Sec. III when we confront this model with election interference-related data. To the extent that this model describes reality, it is probably not true that these restrictive assumptions on the final condition are always satisfied. However, one simple final condition that satisfies these requirements is $\Phi_R(x) = c_0 + c_1 x$, but this allows the unrealistic limiting condition of infinite benefit if candidate A gets 100% of the vote in the election and infinite cost if candidate A gets 0% of the vote. We will also consider two Red final conditions with cost that remains bounded as $x \rightarrow \pm\infty$: one smooth, $\Phi_R(x) = \tanh(x)$, and one discontinuous, $\Phi_R(x) = \Theta(x) - \Theta(-x)$. By $\Theta(\cdot)$ we mean the Heaviside step function.

Blue wants to reduce the overall impact of Red’s interference operations on the electoral process. Since Blue is *a priori* indifferent between the outcomes of the election, it initially seems that $\Phi_B(x) = 0$. However, if $\lambda_B = 0$, then this results in Blue taking no action due to the functional form of Eq. (6). In

other words, if Blue does not gain utility from Red expending resources, then Blue will not try to stop Red from interfering in an election in Blue's country. Hence we believe that Blue cannot be indifferent about the election outcome.

We present three possible final conditions representing Blue's preferences over the election result. Blue might believe that a result was due to Red's interference if X_T is too far from $E_0[X_T] = 0$. An example of a smooth function that represents this belief is $\Phi_B(x) = \frac{1}{2}x^2$. However, this neglects the reality that Red's objective is not to have either candidate A or candidate B win by a large margin but rather to have candidate A win, i.e., have $X_T < 0$. Thus Blue might be unconcerned about larger positive values of the state variable and, modifying the previous function suitably, have $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Alternatively, Blue may accept the result of the election if it does not deviate "too far" from the initial expected value. An example of a discontinuous final condition that represents these preferences is $\Phi_B(x) = \Theta(|x| - \Delta) - \Theta(\Delta - |x|)$, where $\Delta > 0$ is Blue's accepted margin of error.

Though we could propose many other possible final conditions, these example functions demonstrate some possible payoff structures. We include:

(i) "first-order" functions that could result from the Taylor expansion about zero of an arbitrary analytic final condition these functions are linear, in the case of Red's antisymmetric final condition, and quadratic in the case of Blue's smooth symmetric final condition (which is the first nonconstant term in the Taylor expansion of an even analytic function);

(ii) smooth functions that represent bounded preferences over the result of the electoral process and the recognition that Red favors one candidate in particular; and

(iii) discontinuous final conditions that model "all-or-nothing" preferences over the outcome (either candidate A wins or they do not; either Red interferes less than a certain amount or they interfere more).

These functions do not capture some behavior that might exist in real election interference operations. For example, Red's preferences could be as follows: "We would prefer that candidate A wins the election, but if they cannot, then we would like candidate B to win by a landslide so that we can claim the electoral system in Blue's country was rigged against candidate A." These preferences correspond to a final condition with a global minimum at some $x < 0$ but a secondary local minimum at $x \gg 0$. This situation is not modeled by any of the final conditions that we have stated. In Sec. III we relax the assumption that the final conditions are parameterized according to any of the functional forms considered in these sections and instead infer them from observed election and election interference proxy data using the method described in Sec. II B 3.

2. Value functions

Applying the dynamic programming principle [17,18] to Eqs. (3), (4), and (5) leads to a system of coupled Hamilton-Jacobi-Bellman (HJB) equations for the Red and Blue value functions,

$$-\frac{\partial V_R}{\partial t} = \min_{u_R} \left\{ \frac{\partial V_R}{\partial x} [u_R + u_B] + u_R^2 - \lambda_R u_B^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_R}{\partial x^2} \right\} \quad (7)$$

and

$$-\frac{\partial V_B}{\partial t} = \min_{u_B} \left\{ \frac{\partial V_B}{\partial x} [u_R + u_B] + u_B^2 - \lambda_B u_R^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_B}{\partial x^2} \right\}. \quad (8)$$

The dynamic programming principle does not result in an Isaacs equation because the game is not zero-sum and the cost functionals for Red and Blue can have different functional forms. (The Isaacs equation is a nonlinear elliptic or parabolic equation that arises in the study of two-player, zero-sum games in which one player attempts to maximize a functional and the other player attempts to minimize it [27,28].) Performing the minimization with respect to the control variables gives the Nash equilibrium control policies

$$u_R(t) = -\frac{1}{2} \frac{\partial V_R}{\partial x} \Big|_{(t,x)}, \quad (9)$$

$$u_B(t) = -\frac{1}{2} \frac{\partial V_B}{\partial x} \Big|_{(t,x)}, \quad (10)$$

and the exact functional forms of Eqs. (7) and (8),

$$\begin{aligned} -\frac{\partial V_R}{\partial t} = & -\frac{1}{4} \left(\frac{\partial V_R}{\partial x} \right)^2 - \frac{1}{2} \frac{\partial V_R}{\partial x} \frac{\partial V_B}{\partial x} \\ & - \frac{\lambda_R}{4} \left(\frac{\partial V_B}{\partial x} \right)^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_R}{\partial x^2}, \quad V_R(x, T) = \Phi_R(x); \end{aligned} \quad (11)$$

$$\begin{aligned} -\frac{\partial V_B}{\partial t} = & -\frac{1}{4} \left(\frac{\partial V_B}{\partial x} \right)^2 - \frac{1}{2} \frac{\partial V_B}{\partial x} \frac{\partial V_R}{\partial x} \\ & - \frac{\lambda_B}{4} \left(\frac{\partial V_R}{\partial x} \right)^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_B}{\partial x^2}, \quad V_B(x, T) = \Phi_B(x). \end{aligned} \quad (12)$$

When solved over the entirety of state space, solutions to Eqs. (11) and (12) constitute the strategies of a subgame-perfect Nash equilibrium. No matter the action taken by player $-i$ at time t , player i is able to respond with the optimal action at time $t + dt$. This is the (admittedly informal) definition of a subgame-perfect Nash equilibrium in a continuous-time differential game [26]. Given the solution pair $V_R(x, t)$ and $V_B(x, t)$, we can write the distribution of x , u_R , and u_B analytically. Substitution of Eqs. (9) and (10) into Eq. (3) gives $dx = -\frac{1}{2} \left\{ \frac{\partial V_R}{\partial x} \Big|_{(t,x)} + \frac{\partial V_B}{\partial x} \Big|_{(t,x)} \right\} dt + \sigma d\mathcal{W}$. We discretize this equation over N time points to obtain

$$x_{n+1} - x_n + \frac{\Delta t}{2} [V'_{Rn} + V'_{Bn}] - (\Delta t)^{1/2} \sigma w_n - y \delta_{n,0} = 0, \quad (13)$$

with $w_n \sim \mathcal{N}(0, 1)$, $\Delta t = t_{n+1} - t_n$, and $n = 0, \dots, N-1$ and where we have put $V'_{in} \equiv V'_i(x_n, t_n)$. Thus the distribution of an increment of the latent electoral process is

$$p(x_{n+1}|x_n) = \frac{1}{\sqrt{2\pi\sigma^2\Delta t}} e^{-\frac{\Delta t}{2\sigma^2} \left(\frac{x_{n+1}-x_n}{\Delta t} + \frac{1}{2} [V'_{Rn} + V'_{Bn}] - y \frac{\delta_{n,0}}{\Delta t} \right)^2}. \quad (14)$$




FIG. 2. Example value functions corresponding to the system Eqs. (11) and (12). Panels (a) and (b) display $V_R(x, t)$ and $V_B(x, t)$, respectively, for $\lambda_R = \lambda_B = 0$, $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$, and $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$ with $\Delta = 0.1$, while panels (c) and (d) display $V_R(x, t)$ and $V_B(x, t)$, respectively, for $\lambda_R = \lambda_B = 2$, $\Phi_R(x) = 2 \tanh(x)$, and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. For each solution we enforce Neumann no-flux boundary conditions and set $\sigma = 0.6$. The solution is computed on a grid with $x \in [-3, 3]$, setting $dx = 0.025$, and integrating for $N_t = 8,000$ time steps.

Now, using the Markov property of X_t , we have

$$p(x_1, \dots, x_N | x_0) = \prod_{n=0}^{N-1} p(x_{n+1} | x_n) \quad (15)$$

$$= \frac{1}{(2\pi\sigma^2\Delta t)^{N/2}} \exp \left\{ -\frac{1}{2\sigma^2} S(x_1, \dots, x_N) \right\}, \quad (16)$$

where

$$S(x_1, \dots, x_N) = \sum_{n=0}^{N-1} \Delta t \left[\frac{x_{n+1} - x_n}{\Delta t} + \frac{1}{2} [V_{Rn}' + V_{Bn}'] - \frac{y\delta_{n,0}}{\Delta t} \right]^2. \quad (17)$$

Taking $N \rightarrow \infty$ as $N\Delta t = T$ remains constant gives a functional Gaussian distribution,

$$p(x(0 \rightarrow T) | x_0) = \frac{1}{Z} \exp \left\{ -\frac{1}{2\sigma^2} S[x(0 \rightarrow T)] \right\}, \quad (18)$$

with action

$$S[x(0 \rightarrow T)] = \int_0^T \left\{ \frac{dx}{dt} + \frac{1}{2} \left[\frac{\partial V_R}{\partial x} \Big|_{x=x(t)} + \frac{\partial V_B}{\partial x} \Big|_{x=x(t)} - y\delta(t - t_0) \right]^2 \right\} dt \quad (19)$$

and partition function

$$Z = \int_{x(0)}^{x(T)} Dx(0 \rightarrow T) \exp \left\{ -\frac{1}{2\sigma^2} S[x(0 \rightarrow T)] \right\}. \quad (20)$$

We have denoted by $x(s \rightarrow t)$ the actual path followed by the latent state from time s to time t . The measure $Dx(0 \rightarrow T)$ is classical Wiener measure. Since $u_R(0 \rightarrow T)$ and $u_B(0 \rightarrow T)$ are deterministic time-dependent functions of $x(0 \rightarrow T)$, we can find their distributions explicitly using the probability distribution Eq. (16) and the appropriate time-dependent Jacobian transformation. These analytical results are of limited utility because we are unaware of analytical solutions to the

system given in Eqs. (11) and (12), and hence $V'_R(x, t)$ and $V'_B(x, t)$ must be approximated. In Sec. II C we will derive analytical results that are valid when player i announces a credible commitment to a particular control path.

We find the value functions $V_R(x, t)$ and $V_B(x, t)$ numerically through backward iteration, enforcing a Neumann boundary condition at $x = \pm 3$, which corresponds to bounding polling popularity of candidate B from below by $100 \times \phi(-3) = 4.7\%$ and from above by $100 \times \phi(3) = 95.3\%$ [30]. We display example realizations of the value functions for different λ_i and final conditions in Fig. 2. The value functions display diffusive dynamics because the state equation is driven by Gaussian white noise. The value functions also depend crucially on the final condition. When the final conditions are discontinuous (as in the top panels of Fig. 2) the derivatives of the value function reach larger magnitudes and vary more rapidly than when the final conditions are continuous. This has consequences for the game-theoretic interpretation of these results, as we discuss in Sec. II B 4. Figure 2 also demonstrates that the extrema of the value functions are not as large in magnitude when $\lambda_R = \lambda_B = 0$ as when $\lambda_R = \lambda_B = 2$; this is because higher values of λ_i mean that player i derives utility not only from the final outcome of the game but also from causing player $-i$ to expend resources in the game.

Equations (7) and (8) give the closed-loop control policies given the current state X_t and time t , u_R and u_B . We display examples of u_R , u_B and the electoral process Z_t in Fig. 3. For this example, we simulate the game with parameters $\lambda_R = \lambda_B = 2$, $\Phi_R(x) = x$, and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. We plot the control policies in the top panel. The mean control policies $E[u_R]$ and $E[u_B]$ are displayed in thicker curves. For this parameter set, it is optimal for Red to begin play with a larger amount of interference than Blue does and on average decrease the level of interference over time. Throughout the game Blue increases their resistance to Red's interference. Even though Blue resists Red's interference, Red is able to accomplish their objective of causing candidate A to win.

3. Inference and prediction

The solutions to Eqs. (11) and (12) are functions of the final conditions $V_R(x, T) = \Phi_R(x)$ and $V_B(x, T) = \Phi_B(x)$. It is possible to perform both inference and prediction at times $t < T$ even when $\Phi_R(x)$ and $\Phi_B(x)$ are not known. To do this, we assume that the system given by Eqs. (11) and (12) has a unique solution given particular final conditions Φ_R and Φ_B . Though we have numerical evidence to suggest that such solutions do exist and are unique, we have not proved that this is the case. In inference, we want to find the distributions of values of some unobserved parameters of the system. We will suppose that we want to infer Φ_R and Φ_B given the observed paths $x(0 \rightarrow t)$, $u_R(0 \rightarrow t)$, $u_B(0 \rightarrow t)$ with $t < T$. For simplicity we assume that we know all other parameters of Eqs. (11) and (12) with certainty. Then the posterior distribution of Φ_R and Φ_B reads

$$p(\Phi_R, \Phi_B | X_s) \propto p(x(0 \rightarrow t) | \Phi_R, \Phi_B) p(\Phi_R, \Phi_B). \quad (21)$$

The likelihood $p(x(0 \rightarrow t) | \Phi_R, \Phi_B)$ is Gaussian, as shown in Eq. (18), and depends on the time-dependent Jacobian transformations defined implicitly by the solutions of Eqs. (11)




FIG. 3. We display realizations of u_R and u_B in the top panel and paths of the electoral process in the bottom panel. We draw these realizations from the game simulated with parameters $\lambda_R = \lambda_B = 2$, $\Phi_R(x) = x$, and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. For this parameter set, Blue is fighting a losing battle since optimal play by both players results in lower $E[Z_t]$ than for the electoral process without any interference.

and (12). The prior over final conditions $p(\Phi_R, \Phi_B)$ can be set proportional to unity if we want to use a maximum-likelihood approach and not account for our prior beliefs about the form of Φ_R and Φ_B . We can approximate Φ_R and Φ_B with functions parameterized by a finite set of parameters $a_{i,k}$, where $i \in \{R, B\}$ and $k = 0, \dots, K$. The functional prior $p(\Phi_R, \Phi_B)$ is then approximated by the multivariate distribution $p(k_{R,0}, \dots, k_{R,K}, k_{B,0}, \dots, k_{B,K})$. We will take this approach when performing inference in Sec. III.

We can predict future values of $x(t)$, and hence $u_R(t)$ and $u_B(t)$, similarly. Now we want to find the probability of observing $x(t \rightarrow T)$ given observed $x(0 \rightarrow t)$. To do this, we integrate out all possible choices of Φ_R and Φ_B weighted by their posterior likelihood given the observed path $x(0 \rightarrow t)$. The integration is taken with respect to a functional measure, $\mathcal{D}(\Phi_R(x), \Phi_B(x))$. This means that the integration is taken over all possible choices of Φ_R and Φ_B that lie in some particular class of functions [31]. As in the case of inference, we can approximate Φ_R and Φ_B by functions parameterized by a finite set of parameters $a_{i,k}$ and integrate over the $2K$ -dimensional domain of these parameters. In the present work we do not predict any future values of the latent electoral process or control policies.




FIG. 4. Example sweeps over the coupling parameters λ_R and λ_B when Blue’s final condition is set to $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. We vary the coupling parameters over $[0,3]$ and display the resulting standard deviation of the control policies $u_R(x)$ and $u_B(x)$. Panels (a) and (b) represent one coupled system of equations, while panels (c) and (d) represent a coupled system of equations with a different set of final conditions. In panel (a), Red’s value function is set to $\Phi_R(x) = \tanh(x)$, while in panel (b) it is given by $\Phi_R(x) = \Theta(x) - \Theta(-x)$, where $\Theta(\cdot)$ is the Heaviside function. We display a glyph of the corresponding final condition in the upper right corner of each panel. Changing Red’s continuous final condition $\tanh(x)$ to the discontinuous $\Theta(x) - \Theta(-x)$ results in substantially increased variation in the control policies of both players.

4. Dependence of value functions on parameters

We conducted a coarse parameter sweep over λ_R , λ_B , Φ_R , and Φ_B to explore qualitative behavior of this game. We display the results of this parameter sweep for two combinations of final conditions in Figs. 4 and 5. The upper right-hand corner of each panel of the figures displays the final condition of each player. Holding Blue’s final condition of $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$ constant, we compare the means and standard deviations of the Nash equilibrium strategies $u_R(t)$ and $u_B(t)$ across values of the coupling parameters $\lambda_R, \lambda_B \in [0, 3]$ as Red’s final condition changes from $\Phi_R(x) = \tanh(x)$ to $\Phi_R(x) = \Theta(x) - \Theta(-x)$.

For these combinations of final conditions, higher values of the coupling parameters λ_i cause the control policies to have higher variance. This increase in variance is more pronounced when Red’s final condition is discontinuous, which is sensible since in this case $\lim_{t \rightarrow T^-} u_R(x, t) = -\frac{1}{2}\delta(x)$. Appendix contains similar figures for each $3^2 =$

9 combinations of Red example final conditions, $\Phi_R(x) \in \{\tanh(x), \Theta(x) - \Theta(-x), x\}$ and Blue example final conditions, $\Phi_B(x) \in \{\frac{1}{x}x^2, \frac{1}{2}x^2\Theta(-x), \Theta(|x| - \Delta) - \Theta(\Delta - |x|)\}$. We also find that certain combinations of parameters lead to an “arms-race” effect in both players’ control policies. For these parameter combinations, Nash equilibrium strategies entail superexponential growth in the magnitude of each player’s control policy near the end of the game. Figure 6 displays $E[u_R]$ and $E[u_B]$ for some of these parameter combinations, along with the middle 80% credible interval (10th to 90th percentile) of $u_R(t)$ and $u_B(t)$ for each t . A credible interval for the random variable $Y \sim p(y)$ is an interval into which Y falls with a particular probability [32]. For example, the middle 80% credible interval for $Y \sim p(y)$ is the interval (a, b) for which $\int_a^b p(y) dy = 0.8$ and $\int_{-\infty}^a p(y) dy = \int_b^{\infty} p(y) dy = 0.1$.

This growth in the magnitude of each player’s control policy occurs when either player has a discontinuous final




FIG. 5. Example sweeps over the coupling parameters λ_R and λ_B when Blue's final condition is set to $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. We vary the coupling parameters over $[0,3]$ and display the resulting means of the control policies $u_R(x)$ and $u_B(x)$. Panels (a) and (b) represent one coupled system of equations, while panels (c) and (d) represent a coupled system of equations with a different set of final conditions. In panel (a), Red's value function is set to $\Phi_R(x) = \tanh(x)$, while in panel (b) it is given by $\Phi_R(x) = \Theta(x) - \Theta(-x)$. Altering Red's final condition from continuous to discontinuous causes a greater than 100% increase in the maximum value of the mean of Blue's control policy.

condition. Although a discontinuous final condition by player i leads to a greater increase in mean magnitude of player i 's control policy than in player $\neg i$'s, the standard deviation of each player's policy exhibits similar superexponential growth. To the extent that this model reflects reality, this points to a general statement about election interference operations: An all-or-nothing mindset by either Red or Blue about the final outcome of the election leads to an arms race that negatively affects both players. This is a general feature of any strategic interaction to which the model described by Eqs. (3)–(5) applies.

C. Credible commitment

If player $\neg i$ credibly commits to playing a particular strategy $v(t)$ on all of $[0, T]$, then the problem of player i 's finding a subgame-perfect Nash equilibrium strategy profile becomes an easier problem of optimal control. A credible commitment by player $\neg i$ to the strategy $v(t)$ means that

(i) player $\neg i$ tells player i , either directly or indirectly, that player $\neg i$ will follow $v(t)$; and

(ii) player i should rationally believe that player $\neg i$ will actually follow $v(t)$.

An example of a mechanism that makes commitment to a strategy credible is the Soviet Union's "Dead Hand" automated second-strike nuclear response system. This mechanism would launch a full-scale nuclear attack on the United States if it detected a nuclear strike on the Soviet Union [33,34]. The existence of this mechanism made the commitment to the strategy "launch a full-scale nuclear attack on the United States given that any nuclear attack on my country has occurred" credible, even though the potential cost to both parties of executing the strategy was high.

When player $\neg i$ credibly commits to playing $v(t)$, player i 's problem reduces to finding the policy $u(t)$ that minimizes the functional

$$E_{u,x} \left\{ \Phi(X_T) + \int_0^T (u(t)^2 + \lambda v(t)^2) dt \right\}, \quad (22)$$

subject to the modified state equation

$$dx = [u(t) + v(t)]dt + \sigma dW. \quad (23)$$




FIG. 6. In the case of strong coupling (λ_R and $\lambda_B \gg 0$), discontinuous final solutions by either player cause superexponential growth in the magnitude of each player's control policy. Here we set $\lambda_R = \lambda_B = 3$ and integrate three systems, varying only one final condition in each. Panel (a) displays a system with two continuous final conditions: $\Phi_R(x) = \tanh(x)$ and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Panel (b) displays the mean Red and Blue control policies when the Red final condition is changed to $\Phi_R(x) = \Theta(x) - \Theta(-x)$ as the Blue final condition remains equal to $\frac{1}{2}x^2\Theta(-x)$, while panel (c) shows the control policies when $\Phi_B(x) = \Theta(|x| > 1) - \Theta(|x| < 1)$ and $\Phi_R(x) = \tanh(x)$. The shaded regions correspond to the middle 80 percentiles (10th to 90th percentiles) of $u_R(t)$ and $u_B(t)$ for each t . When either player has a discontinuous final condition, the interpercentile range is substantially wider for both players than when both players have continuous final conditions.

Player i 's value function is now given by the solution to the HJB equation

$$-\frac{\partial V}{\partial t} = \min_u \left\{ \frac{\partial V}{\partial x} [u + v] + u^2 - \lambda v^2 + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2} \right\}. \quad (24)$$

Performing the minimization gives the control policy $u(t) = -\frac{1}{2} \frac{\partial V}{\partial x} \Big|_{x=x(t)}$ and the explicit functional form of the HJB equation,

$$-\frac{\partial V}{\partial t} = -\frac{1}{4} \left(\frac{\partial V}{\partial x} \right)^2 + v(t) \frac{\partial V}{\partial x} + \lambda v(t)^2 + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2},$$

$$V(x, T) = \Phi(x). \quad (25)$$

1. Path-integral control

Though nonlinear, this HJB equation can be transformed into a backward Kolmogorov equation (BKE) through a change of variables. The BKE can be solved using path-integral methods [35]. Setting $V(x, t) = -\eta \log \varphi(x, t)$, substituting in Eq. (25), and performing the differentiation, we are able to remove the nonlinearity if and only if $\frac{\eta^2}{4} \frac{1}{\varphi^2} \left(\frac{\partial \varphi}{\partial x} \right)^2 = \frac{\sigma^2 \eta}{2} \frac{1}{\varphi^2} \left(\frac{\partial \varphi}{\partial x} \right)^2$. Setting $\eta = 2\sigma^2$ satisfies this condition. Performing the change of variables, Eq. (25) is now linear and has a time-dependent drift and sink term,

$$\frac{\partial \varphi}{\partial t} = \frac{\lambda}{2\sigma^2} v(t)^2 \varphi(x, t) - v(t) \frac{\partial \varphi}{\partial x} - \frac{\sigma^2}{2} \frac{\partial^2 \varphi}{\partial x^2},$$

$$\varphi(x, T) = \exp \left\{ -\frac{1}{2\sigma^2} \Phi(x) \right\}. \quad (26)$$

Application of the Feynman-Kac formula gives the solution to Eq. (26) as [36]

$$\varphi(x, t) = \exp \left\{ -\frac{\lambda}{2\sigma^2} \int_t^T v(t')^2 dt' \right\} \times E_{Y_t} \left\{ \exp \left[-\frac{1}{2\sigma^2} \Phi(Y_T) \right] \Big| Y_t = x \right\}, \quad (27)$$

where Y_t is defined by

$$dY_t = v(t) dt + \sigma dW_t, \quad Y_0 = x. \quad (28)$$

Using this formalism, we apply path-integral control to estimate the value function for arbitrary $v(t)$. Figure 7 displays example path-integral solutions to Eq. (25) when player $-i$ credibly commits to playing $v(t) = t^2$ for the duration of the game and player i 's final cost function takes the form $\Phi(x) = \Theta(|x| - 1) - \Theta(1 - |x|)$. In this figure we display the approximate value functions $V(x, t)$ along with their corresponding approximate control policies $u(x, t) = -\frac{1}{2} \frac{\partial V(x, t)}{\partial x}$. We calculated these approximations on a grid of $N_x = 500$ linearly spaced $x_n \in [-2, 2]$. Since the approximated $u(x_n, t)$ are noisy stochastic functions we also plot smoothed versions of them in Fig. 7. These smoothed versions are defined by

$$u_s(x_n, t) = \sum_{n'=-k}^k u(x_{n+n'}, t). \quad (29)$$

We set $k = 7$ and hence $u_s(x_n, t)$ are $2k + 1 = 15$ -step moving averages of the more noisy $u(x_n, t)$. As $t \rightarrow T$, $u_s(x_n, t)$ approaches the analytical solution of the control problem at $t = T$, which is given by $u(x, T) = -\frac{1}{2} [\delta(x - 1) - \delta(x + 1)]$.

In the further restricted case where there is a credible commitment by one party to play a constant control policy $v(t) = v$, we can derive further analytical results. Under this assumption, the probability law corresponding with Eq. (28)




FIG. 7. Result of the path-integral Monte Carlo solution method applied to Eq. (25) with the final condition $\Phi(x) = \Theta(|x| > 1) - \Theta(|x| \leq 1)$ and $v(t) = t^2$. Approximate value functions are computed using $N = 10000$ trajectories sampled from Eq. (28) for each point (x, t) . We display approximate value functions in panel (a) for $t \in \{0, 0.75, 1 - dt\}$ and the corresponding approximate control policies $u(x, t)$ in panel (b), along with smoothed versions of $u(x, t)$, which we denote by $u_s(x, t)$, plotted in dashed curves. Panel (c) displays realizations of Y_t , the process generating the measure under which the solution is calculated. The analytical control policy at $t = T$ is given by $u(t) = -\frac{1}{2}[\delta(x - 1) - \delta(x + 1)]$.

is given by

$$u(y, t) = \frac{1}{\sqrt{2\pi\sigma^2 t}} \exp\left\{-\frac{1}{2\sigma^2 t}[(y - x) - vt]^2\right\}, \quad (30)$$

so that the (exponentially transformed) value function reads

$$\begin{aligned} \varphi(x, t) &= E_{u(y, T-t)} \left[\exp\left\{-\frac{\lambda v^2}{2\sigma^2}(T-t)\Phi(Y_T)\right\} \right], \quad (31) \\ &= \frac{\exp\left\{-\frac{\lambda v^2}{2\sigma^2}(T-t)\right\}}{\sqrt{2\pi\sigma^2(T-t)}} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2}\left\{\Phi(y) + \frac{[(y-x) - v(T-t)]^2}{T-t}\right\}\right) dy. \quad (32) \end{aligned}$$

This integral can be evaluated exactly for many $\Phi(y)$ and, for many other final conditions, it can be approximated using the method of Laplace. When $t \rightarrow T$ so that the denominator of the argument of the exponential in Eq. (31) approaches zero, Laplace's approximation to the integral reads

$$\begin{aligned} &\int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2}\left\{\Phi(y) + \frac{[(y-x) - v(T-t)]^2}{T-t}\right\}\right) dy \\ &\approx \sqrt{2\pi\sigma^2(T-t)} \exp\left\{-\frac{1}{2\sigma^2}\Phi[x + (T-t)v]\right\}. \quad (33) \end{aligned}$$

Inverting the transformation φ , the value function is approximated by

$$V(x, t) = \lambda v^2(T-t) + \Phi[x + (T-t)v], \quad (34)$$




FIG. 8. When player $-i$ commits to playing a constant strategy profile $v(t) = v$ for a fixed interval of time, an analytic approximate form for player i 's value function $V(x, t)$ is given by $V(x, t) \approx \lambda v^2(T-t) + \Phi[x + (T-t)v]$, which we derive from a Laplace approximation argument. We show the numerically determined value functions at time $t = 0$ in solid black curves. We display the Laplace approximations at $t = 0$ in dashed black curves. The lighter-hue curves are the value functions at the final time T , i.e., the respective final conditions. Panel (a) demonstrates results for the final condition $\Phi(x) = \frac{1}{2}x^2$, while panel (b) has $\Phi(x) = \tanh(x)$.

and the control policy by

$$u(x, t) = -\frac{1}{2}\Phi'[x + (T-t)v]. \quad (35)$$

We display the results of approximating the value function with Eq. (34) at $t = 0$ in Fig. 8, along with the actual numerically determined value function at both $t = 0$ and, for reference, $t = T$.

2. Dependence on a free parameter

The Laplace-approximated value function may depend on a free parameter a that can be used as a ‘‘control knob’’ to adjust the approximation. For example, player i might use a to tune the sensitivity of the approximation to the electoral process's distance from a dead-heat. Ideally, the approximated control policy should have similar scaling and asymptotic properties as the true control policy. Solving an optimization problem for optimal values of a is one approach to satisfying this desideratum.

As a case study we consider the behavior of the Laplace-approximated value function $V^{(a)}(x, t)$ and its corresponding control policy $u^{(a)}(x, t)$ when we set $\Phi^{(a)}(x) = \tanh(ax)$. We consider this specific example because $\Phi^{(a)}(x) \rightarrow \Theta(x) - \Theta(-x)$ as $a \rightarrow +\infty$. This limit can be the source of complicated behavior in a variety of fields such as piecewise-smooth dynamical systems (both deterministic and stochastic) [37,38], Coulombic friction [39], and evolutionary biology [40].

Figure 9 displays player i 's exponentially transformed value function Eq. (31) with final condition $\Phi_i(x) = \tanh(ax)$. Here player $-i$ credibly commits to playing a constant strategy of $v = 0.01$. As $t \rightarrow T$, larger values of a lead to a sharp boundary between regions of the state space that are costly




FIG. 9. If player $-i$ credibly commits to a strategy of playing a constant strategy with value equal to v for the entire duration of the game, then player i 's (exponentially transformed) value function $\varphi(x, t)$ has an integral representation given by Eq. (31). We display dynamics of $\varphi(x, t)$ with the final condition set to $\Phi(x) = \tanh(ax)$. We calculated this value function for $x \in (-\frac{3}{2}, \frac{3}{2})$ and logarithmically equally spaced values of $a \in [10^{-3}, 10^5]$. For $a < 10^{-1}$, the value function is nearly constant. When $a > 10^1$, $\frac{\partial}{\partial x}\varphi(x, t)$ increases in magnitude near $x = 0$ as $t \rightarrow T$.

for player i (positive values of x) and those that are less costly (negative values of x). This behavior is qualitatively similar to behavior arising from the final condition $\Phi_i(x) = \Theta(x) - \Theta(-x)$. However, we will show that there are significant scaling differences in the control policies resulting

from using $\Phi(x)$ versus $\Phi^{(a)}(x)$. From Eqs. (33) and (34), we approximate the value function by

$$V^{(a)}(x, t) \approx \lambda v^2(T - t) + \tanh\{a[x + v(T - t)]\}, \quad (36)$$

and hence the control policy is approximately

$$u^{(a)}(x, t) \approx -\frac{a}{2} \operatorname{sech}^2\{a[x + v(T - t)]\}, \quad (37)$$

with both expansions increasingly accurate as $t \rightarrow T$. When $\Phi(x) = \Theta(x) - \Theta(-x)$, we can compute the value function analytically:

$$\begin{aligned} & \frac{1}{\sqrt{2\sigma^2(T-t)}} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2} \left\{ \Theta(y) - \Theta(-y) + \frac{[(y-x) - v(T-t)]^2}{T-t} \right\}\right) dy \\ &= \cosh\left(\frac{1}{2\sigma^2}\right) + \sinh\left(\frac{1}{2\sigma^2}\right) \operatorname{erf}\left[-\frac{x + v(T-t)}{\sqrt{2\sigma^2(T-t)}}\right], \end{aligned} \quad (38)$$

where we find that

$$V(x, t) = \lambda v^2(T - t) - 2\sigma^2 \log \left\{ \cosh\left(\frac{1}{2\sigma^2}\right) + \sinh\left(\frac{1}{2\sigma^2}\right) \operatorname{erf}\left[-\frac{x + v(T-t)}{\sqrt{2\sigma^2(T-t)}}\right] \right\} \quad (39)$$

and

$$u(x, t) = -\sqrt{\frac{2\sigma^2}{\pi(T-t)}} \frac{\exp\left\{-\frac{[x+v(T-t)]^2}{2\sigma^2(T-t)}\right\}}{\coth\left(\frac{1}{2\sigma^2}\right) + \operatorname{erf}\left[-\frac{x+v(T-t)}{\sqrt{2\sigma^2(T-t)}}\right]}. \quad (40)$$

The approximate control policy $u^{(a)}(x, t)$ and the limiting control policy have similar negative “bell-like” shapes but also differ in important ways. The true control policy decays

as a Gaussian modulated by the asymmetric function $\operatorname{erf}(\cdot)$. The approximate policy decays logarithmically and hence more slowly than the true control policy. While the approximate




FIG. 10. The solution to Eq. (41) is a superexponentially increasing $a(t)$ parameter in the Laplace method-derived value function $V^{(a)}(x, t) = \tanh(ax)$. We use this value function as an approximation to the exact value function given in Eq. (39). Dashed curves indicate $u(x, t)$, while solid curves indicate $u^{(a)}(x, t)$. Panel (c) displays the same data as the main axis and also includes $u(x, t)$ and $u^{(a)}(x, t)$ at the last simulation time step, $t = 0.9975$, to demonstrate the increasing accuracy of the approximation as $t \rightarrow T$. We display the optimal $a(t)$ in panel (b).

policy is symmetric, the true policy is asymmetric due to the error function term. Using the Laplace approximation results in more control being applied to the electoral process than is optimal. This is because the tails of $u^{(a)}(x, t)$ are heavier than those of $u(x, t)$.

We can maximize the similarity between $u^{(a)}(x, t)$ and $u(x, t)$ by letting the free parameter a be a function of t and solving the functional minimization problem

$$\min_{a(t)} \int_t^T \int_{-\infty}^{\infty} \{u^{[a(t)]}(x, t) - u(x, t)\}^2 dx dt. \quad (41)$$

A stationary point of this problem is given by an $a(t)$ that solves

$$\int_{-\infty}^{\infty} \{u^{[a(t)]}(x, t) - u(x, t)\} \frac{\partial u^{[a(t)]}(x, t)}{\partial a(t)} dx = 0. \quad (42)$$

We are unable to compute this integral analytically on substituting Eqs. (37) and (40). We find the solution to this problem by numerically solving Eq. (42) using the secant method for each of 100 linearly spaced $t \in [0.5, 0.9975]$. We display the optimal $a(t)$, along with the corresponding $u^{[a(t)]}(x, t)$ and true $u(x, t)$ in Fig. 10. We find that the optimal $a(t)$ grows superexponentially as $t \rightarrow T$ and that the accuracy of the approximation increases in this limit. This is expected given that $u^{(a)}(x, t)$ is derived using the Laplace approximation and it is in this limit that the Laplace approximation is valid.

Even with the assumption of credible commitment to a constant control policy v , we can use this theory to approximate the value function in a noncooperative scenario. For arbitrary $v(t)$, expansion about $t + \Delta t$ gives $v(t + \Delta t) = v(t) + v'(t)\Delta t$, leading to an approximate value function iteration over a small time increment Δt ,

$$V(x, t + \Delta t) \approx \lambda v(t)^2(T - t) + \Phi\{x + (T - t)[v(t) + v'(t)\Delta t]\}. \quad (43)$$

The approximation is increasingly accurate as $t \rightarrow T$ and when Δt is small. In application, both of $v(t)$ and $v'(t)$ can be estimated from possibly noisy data on $t' \in [0, t]$.

III. APPLICATION

An example of election interference operations is the Russian military foreign intelligence service (Red team) activity in the 2016 U.S. presidential election contest. Red team attempted to harm one candidate's (Hillary Clinton's) chances of winning and aid another candidate (Donald Trump) [11]. Though Russian foreign intelligence had conducted election interference operations in the past at least once before, in the Ukrainian elections of 2014 [41], the 2015 and 2016 operations were notable in that Red team operatives used the microblogging website Twitter in an attempt to influence the election outcome. When this attack vector was discovered, Twitter shut down accounts associated with Red team activity and all data associated with those accounts was collected and analyzed [42–44]. There has been analysis of the qualitative and statistical effects of these and other election attack vectors (e.g., Facebook advertisement purchases) on election polling and the outcome of the election [45] and on the detection of election influence campaigns more generally [46,47]. However, to the best of our knowledge, there exists no publicly available effort to reverse-engineer the quantitative nature of the control policies used by Russian military intelligence and by U.S. domestic and foreign intelligence agencies.

We first fit a discrete-time formulation of the model described in Sec. II A. We then compare it to theoretical predictions by finding values of free parameters in the theoretical model that best describe the observed data and inferred latent controls. We are faced with two distinct sources of uncertainty in this procedure. First, we cannot observe either Red's or Blue's control policy directly because foreign and domestic




FIG. 11. We approximate the time-series components of the analytical model defined in Sec. II A by a Bayesian structural time-series (BSTS) model. We subsequently confront the BSTS model with 2016 U.S. presidential election data. Observed random variables are denoted by gray-shaded nodes, while latent random variables are represented by unshaded nodes or red ($u_{R,t}$, lighter hue) and blue ($u_{B,t}$, darker hue) nodes. We observe a noisy election poll, denoted by Z_t , and a time series of tweets associated with Russian military intelligence, denoted by $Tweets$. Our objective in this modeling stage is to infer the latent electoral process, denoted by X_t , and the latent control policies.

intelligence agencies shroud their activities in secrecy. Second, each player’s final time payoff structure is also secret and unknown to us. To partially circumvent these issues, we construct a two-stage model. The first stage is a Bayesian structural time-series model, depicted graphically in Fig. 11, through which we are able to infer distributions of discretized analogs of $u_R(t)$, $u_B(t)$, and $x(t)$. Once we have inferred these distributions, we minimize a loss function that compares the means of these distributions to the means of distributions produced by the model described in Sec. II A.

We make the simplifying assumptions about the format of the election that we stated in Sec. I when constructing the discrete-time election model. Namely, we assume that only two candidates contest the election and that the election process is modeled by a simple “candidate A versus candidate B” poll. Though there are methods for forecasting elections that make fewer and less restrictive assumptions than these, such as compartmental infection models [48], prediction markets [49], and more sophisticated Bayesian models [50,51], we construct our statistical model to mimic the underlying election model of Sec. II A. We do this to test the ability of this underlying theoretical model to reproduce inferred control and observed election dynamics.

We can observe neither the Red $u_R(t)$ nor Blue $u_B(t)$ control policies. However, we are able to observe a proxy for u_R : the number of tweets sent by Russian military intelligence-associated accounts in the year leading up to the 2016 election [52]. This dataset contains a total of 2 973 371 tweets from 2848 unique Twitter handles. Of these tweets, a total of 1 107 361 occurred in the year immediately preceding the election (August, 11, 2015 to August 11, 2016). We grouped these tweets by day and used the time series of total number of tweets on each day as an observable from which we could infer u_R . We restricted the time range of the model to begin at the later of the end dates of the Republican National Convention (July 21, 2016) and Democratic National Convention (July 28, 2016). We did this because the later of these dates, July 28, 2016, is the day on which the race was officially between two major party candidates. Of all Russian military intelligence-associated tweets in 2016, 363 131 occurred during the 102 days beginning on July 28, 2016 and ending the day before Election Day. Though the presence of minor party candidates probably played a role in the result of the election, even the most prominent minor parties (Libertarian and Green) received only single-digit support [53,54]. We do not model these minor parties and instead consider only the electoral contest between the two major party candidates. We used the RealClearPolitics poll aggregation as a proxy for the electoral process itself [55], averaging polls that were recorded on the same date and using the earliest date in the date range of the poll if it was conducted over multiple days as the timestamp of that observation. We weighted all polls equally when averaging.

Using these two observed random variables, we fit a Bayesian structural time-series model [56] of the form presented in Fig. 11. We now describe the structure of the model and explain our choices of priors and likelihood functions. In the analytical model, we model the latent control policies $u_R(t)$ and $u_B(t)$ by time- and state-dependent Wiener processes. To see this, recall that the state equation evolves according to a Wiener process and apply Ito’s lemma to the deterministic functions of a random variable $-\frac{1}{2} \frac{\partial V_R}{\partial x'}|_{x'=x_t}$ and $-\frac{1}{2} \frac{\partial V_B}{\partial x'}|_{x'=x_t}$, which define the control policies. A discretized version of the Wiener process is a simple Gaussian random walk. We thus model the latent Red and Blue control policies by Gaussian random walks:

$$p(u_{R,t} | u_{R,t-1}, \mu_R, \sigma) = \mathcal{N}(u_{R,t-1} + \mu_R, \sigma^2), \quad (44)$$

$$p(u_{B,t} | u_{B,t-1}, \mu_B, \sigma) = \mathcal{N}(u_{B,t-1} + \mu_B, \sigma^2). \quad (45)$$

Similarly, we model the latent election process by a discretized version of the state evolution equation Eq. (3):

$$p(X_t | X_{t-1}, u_R, u_B) = \mathcal{N}(X_{t-1} + u_{B,t-1} - u_{R,t-1}, 1). \quad (46)$$

We assume that the latent election model is subject to normal observation error in latent space. Since we chose a logistic function as the link between the latent and real [on (0,1)] election spaces, the likelihood for the observed election process is thus given by a logit-normal distribution. The probability

density function (pdf) of this distribution is

$$p(Z_t|X_t, \sigma_Z) = \sqrt{\frac{1}{2\pi\sigma_Z^2}} \frac{\exp\left\{-\frac{[\text{logit}(Z_t)-X_t]^2}{2\sigma_Z^2}\right\}}{Z_t(1-Z_t)}. \quad (47)$$

Though the number of Russian military intelligence tweets that occur on any given day is obviously a non-negative integer, we chose not to model it this way. A common and simple model for a “count” random variable, such as the tweet time series, is a Poisson distribution with possibly time-dependent rate parameter [57–60]. This model imposes a strong assumption on the variance of the count distribution (namely that the variance and mean are equal) which does not seem realistic in the context of the tweet data. Instead of searching for a discrete count distribution that meets some optimality criterion, we instead normalized the tweet time series to have zero mean and unit variance, making it a continuous random variable rather than a discrete one. We then shifted the time series so that the new time series was equal to zero on the day during our study with the fewest tweets. We then modeled this time series Tweets_t by a normal observation likelihood,

$$p(\text{Tweets}_t|u_{R,t}, \sigma_{\text{Tweets}}) = \mathcal{N}(u_{R,t}, \sigma_{\text{Tweets}}^2). \quad (48)$$

We placed a weakly informative prior, a log-normal distribution, on each standard deviation random variable (σ , σ_Z , σ_{Tweets}), and zero-centered normal priors on each mean random variable (μ_R , μ_B). This model is high dimensional, since the latent time series X , u_R , and u_B are inferred as T -dimensional vectors. In total this model has $3T + 5 = 311$ degrees of freedom.

We display a graphical representation of this model in Fig. 11. We fit this model using the No-U-Turn Sampler algorithm [61], sampling 2000 draws from the model’s posterior distribution from each of two independent Markov chains. We do not include 1000 draws per chain of burn-in in the samples from the posterior. The sampler appeared to converge well based on graphical consideration (i.e., the “eye test”) of draws from the posterior predictive distribution of Z_t and Tweets_t , and, more importantly, because maximum values of Gelman-Rubin statistics [62] for all variables satisfied $R_{\max} < 1.01$ except for that of σ_Z , which had $R_{\max} = 1.07646$. Each of these values is well below the level $R = 1.1$ advocated by Brooks and Gelman [63]. Figure 12 displays draws from the posterior and posterior predictive distribution of this model. Figure 12(a) displays draws of X_t from the posterior distribution, along with $E[X_t]$ and $\text{logit}(Z_t)$, while in Fig. 12(b) we show posterior draws of u_R and u_B , along with $E[u_R]$ and $E[u_B]$ in thick red and blue curves, respectively. In Fig. 12(c), we display Tweets_t and draws from its posterior predictive distribution. On October 6, 2016, Tweets_t exhibited a large spike that is very unlikely under the posterior predictive distribution. This spike likely corresponds with a statement made by the U.S. federal government on this date that officially recognized the Russian government as culpable for hacking the Democratic National Committee computers.

After inferring the latent control policies and electoral process, we searched for the parameter values $\theta = (\lambda_R, \lambda_B, \sigma, \Phi_R, \Phi_B)$ of the theoretical model that best explain the observed data and inferred latent variables. For clarity




FIG. 12. Panel (a) displays the logit of the observed election time series (black curve) $\text{logit}(Z_t)$, along with the posterior distribution of the latent electoral process X_t . Panel (b) displays the mean latent control policies in thick red (lighter hue) and blue (darker hue) curves, along with their posterior distributions. Panel (c) shows the true tweet time series (subject to the normalization described in the main body) along with draws from its posterior predictive distribution. The large spike in the tweet time series that is very unlikely under the posterior predictive distribution corresponds to the day (October 6, 2016) on which the U.S. federal government officially accused Russia of hacking the Democratic National Committee computers.




FIG. 13. We display credible intervals of latent election process X and Red and Blue control policies, u_R and u_B , generated using optimal $\theta = (\lambda_R, \lambda_B, \sigma, a_{0,r}, \dots, a_{K,r}, a_{0,b}, \dots, a_{K,b})$ values. We ran the optimization algorithm with the number of terms of the Legendre expansion of Φ_R and Φ_B set to $K = 10$ and set the variance regularization in the algorithm to $\eta = 0.002$. This resulted in fit parameters of $\lambda_R = 0.849$, $\lambda_B = 0.727$, and $\sigma = 1.509$. Panel (a) displays draws from the latent electoral process under \mathcal{Q} , along with $\text{logit}(Z_t)$, the logit-transformed real polling popularity process. Panel (b) displays draws from the distributions of \hat{u}_R and \hat{u}_B under \mathcal{Q} , while panel (c) displays the inferred final conditions $\Phi_R(x)$ and $\Phi_B(x)$.

in reference, we will refer to the theoretical model as \mathcal{Q} and the Bayesian structural time-series model as \mathcal{M} . We use Legendre polynomials to approximate the final conditions Φ_R and Φ_B , as discussed in Sec. II B 3. Approximating $\Phi_i(x)$ by the finite sum $\sum_{k=0}^K a_{ik} P_k(x)$, \mathcal{Q} 's parameter vector is $\theta = (\lambda_R, \lambda_B, \sigma, a_{0,r}, \dots, a_{K,r}, a_{0,b}, \dots, a_{K,b})$. In contrast with \mathcal{M} , \mathcal{Q} has relatively few degrees of freedom since the assumption of state and policy coevolution via solution of coupled partial differential equations substantially restricts the system's dynamics. In total, \mathcal{Q} has $2K + 3$ free parameters. The smaller the value of K , the less accurate the approximation to the true final conditions will be. Conversely, large K could lead to overparameterization of the model and increases the size of the search space. We thus chose $K = 10$ as a compromise between these two extremes. With $K = 10$, the model has $2K + 3 = 23$ degrees of freedom.

The theoretical model \mathcal{Q} can be viewed as a generative probabilistic function. To find optimal parameter values, we generate $(\hat{u}_R, \hat{u}_B, \hat{X})$ from \mathcal{Q} and minimize a loss function of these generated values and the values inferred by \mathcal{M} . We defined this loss function as

$$L(\theta|\mathcal{Q}) = \sum_{(y, \hat{y})} [\|\mu_y - \mu_{\hat{y}}\|_2^2 + \eta\sigma_{\hat{y}}^2], \quad (49)$$

where $y \in \{u_R, u_B, X\}$ and $\hat{y} \in \{\hat{u}_R, \hat{u}_B, \hat{X}\}$. We have defined the mean and standard deviation under the corresponding distribution by μ and σ , respectively. The ℓ_2 terms in Eq. (49) penalize deviation by \mathcal{Q} from the mean of \mathcal{M} 's inferred posterior distribution. The standard deviation term in Eq. (49) imposes a penalty on dispersion. We minimized $L(\theta|\mathcal{Q})$ using a Gaussian-process Bayesian optimization algorithm. The details of this algorithm are beyond the scope of this work but are readily found in any review paper on the subject [64–66]. Figure 13 displays the result of this optimization procedure for $K = 10$ and $\eta = 0.002$. For this set of hyperparameters, we found coupling parameter values of $\lambda_R = 0.1432$ and $\lambda_B = 1.7847$ and a latent space volatility of $\sigma = 0.7510$. In Fig. 13, we use credible intervals to denote ranges into which our estimates of model parameters fall.

Figure 13(a) displays $\text{logit}(Z_t)$ in a thick black curve and a middle 80% (10% to 90%) credible interval of \hat{X} from \mathcal{Q}

in gray shading. The observed $\text{logit}(Z_t)$ is centered in the credible interval of \hat{X} and hence has a high probability under \mathcal{Q} . In Figure 13(b), we show $E[u_R]$ and $E[u_B]$ in thick red (lighter hue) and blue (darker hue) curves, respectively, along with middle 80% credible intervals of \hat{u}_R and \hat{u}_B . The mean paths of the latent Red and Blue control policies do not lie in the middle 80% credible intervals for approximately the first two weeks after the end of the Democratic National Convention, but do lie in these credible intervals for the remainder of the time until the election. Our model is able to capture the election interference dynamics in the middle range of this time span but is not able to capture the dynamics immediately after the race becomes a two-candidate election. Though the election does officially become a two-candidate contest at that time (notwithstanding our previous comments about third-party candidates), the effects of the Republican and Democratic primaries may take time to dissipate. Our model does not capture the dynamics of noncooperative games in the presence of many candidates. We comment on this finding more in Sec. IV. Finally, we display the inferred final conditions $\Phi_R(x)$ and $\Phi_B(x)$ in Fig. 13(c).

IV. DISCUSSION AND CONCLUSION

We introduce, analyze, and numerically (analytically in simplified cases) solve a simple model of noncooperative strategic election interference. This interference is undertaken by a foreign intelligence service from one country (Red) in an election occurring in another country (Blue). Blue's domestic intelligence service attempts to counter this interference. Though simple, our model is able to provide qualitative insight into the dynamics of such strategic interactions and performs well when fitted to polling and social media data surrounding the 2016 U.S. presidential election contest. We find that all-or-nothing attitudes regarding the outcome of the election interference, even if these attitudes are held by only one player, result in an arms race of spending on interference and counterinterference operations by both players. We then find analytical solutions to player i 's optimal control problem when player $-i$ credibly commits to a strategy $v(t)$. We detail an analytical value function approximation that can be used by




FIG. 14. We demonstrate the lack of qualitative changes in the value functions given by solutions to Eqs. (52) and (53) when compared to the solutions of Eqs. (11) and (12). We plot $V_i(x, 0)$ in solid curves and $V_i(x, T)$ in dashed curves, $i \in \{R, B\}$. We set final conditions here to $\Phi_R(x) = \tanh x$ and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$.

player i even when player $-i$ does not commit to a particular strategy as long as player $-i$'s current strategy and its time derivative can be estimated. We demonstrate the applicability of our model to real election interference scenarios by analyzing the Russian effort to interfere in the 2016 U.S. presidential election through observation of Russian troll account posts on the website Twitter. Using these data, along with aggregate presidential election polling data, we infer the time series of Russian and U.S. control policies and find parameters of our model that best explain these inferred control policies. We show that, for most of the time under consideration (after the Democratic National Convention and before Election Day), our model provides a good explanation for the inferred variables. However, our model does not accurately or precisely capture the interference dynamics immediately after the race becomes a two-candidate race.

There are several areas in which our work could be improved. While our model is justifiable on the grounds of parsimony and acceptable empirical performance on at least one election contest, the kind of assumptions that we make in constructing our modeling framework are unrealistic. Though a pure random-walk model for an election is not without serious precedent [67], an extension of this work could incorporate noninterference-related state dynamics as a generalization of Eq. (3). For example, the state equation could read

$$dx = [\mu_0 + \mu_1 x + u_R(t) + u_B(t)]dt + \sigma d\mathcal{W}. \quad (50)$$

This state equation accounts for simple drift in the election results as a candidate endogenously becomes more or less

popular. It can also account for possible mean-reverting behavior in a hotly contested race. Another extension could introduce state-dependent running costs, particularly in the case of the Red player. Though the action of election interference is nominally intended to cause a particular candidate to win or lose, Red could have other objectives as well, such as undermining the Blue citizens' trust in their electoral process. Red might gain utility from having a candidate lead in polls multiple times when that candidate would not have otherwise done so, even if the candidate does not actually win the election. In the context of our model, this is represented by setting Red's cost functional to be

$$E_{u_R, u_B, X} \left\{ \Phi_R(X_T) + \int_0^T [-\Theta(-X_t) + u_R^2(t) - \lambda_R u_B^2(t)] dt \right\}. \quad (51)$$

Both of these modifications are easy to incorporate into the model and do not change the qualitative nature of Red and Blue's HJB equations since their effects will simply be to introduce an additional drift term [Eq. (50)] or a continuous, nondifferentiable source term [Eq. (51)] into the HJB equations [Eqs. (11) and (12)]. That is, the fundamental nature of these equations as nonlinear parabolic equations coupled through quadratic terms of self- and other-player first spatial derivatives remains unchanged as these modifications to the theory do not introduce any new coupling terms. The solutions to these equations do not demonstrate shock or traveling wave behavior with the addition of the drift or source terms, as we show in Figs. 14 and 15. With the




FIG. 15. We demonstrate the lack of qualitative changes in the value functions given by the solution to Eqs. (54) when compared to the solutions of Eqs. (11) and (12). We plot $V_i(x, 0)$ in solid curves and $V_i(x, T)$ in dashed curves, $i \in \{R, B\}$. We set final conditions here to $\Phi_R(x) = \tanh x$ and $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. We also plot solutions for an augmented HJB equation for Blue in panels (c) and (d). This HJB equation is identical to Eq. (54) except with $R \mapsto B$ and the sign of argument of the discontinuous running cost function reversed.

modification of Eq. (50), the HJB equations become

$$\begin{aligned}
 -\frac{\partial V_R}{\partial t} &= (\mu_0 + \mu_1 x) \frac{\partial V_R}{\partial x} - \frac{1}{4} \left(\frac{\partial V_R}{\partial x} \right)^2 \\
 &\quad - \frac{1}{2} \frac{\partial V_R}{\partial x} \frac{\partial V_B}{\partial x} - \frac{\lambda_R}{4} \left(\frac{\partial V_B}{\partial x} \right)^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_R}{\partial x^2}, \\
 V_R(x, T) &= \Phi_R(x)
 \end{aligned} \tag{52}$$

and

$$\begin{aligned}
 -\frac{\partial V_B}{\partial t} &= (\mu_0 + \mu_1 x) \frac{\partial V_B}{\partial x} - \frac{1}{4} \left(\frac{\partial V_B}{\partial x} \right)^2 \\
 &\quad - \frac{1}{2} \frac{\partial V_B}{\partial x} \frac{\partial V_R}{\partial x} - \frac{\lambda_B}{4} \left(\frac{\partial V_R}{\partial x} \right)^2 + \frac{\sigma^2}{2} \frac{\partial^2 V_B}{\partial x^2}, \\
 V_B(x, T) &= \Phi_B(x).
 \end{aligned} \tag{53}$$

In Fig. 14 we give examples of solutions of Eqs. (52) and (53) at $t = 0$ and $t = T$. These solutions do not display qualitative changes, such as the formation of shock or traveling waves, with the inclusion of nonzero drift terms of the form $\mu_0 + \mu_1 \frac{\partial V_i}{\partial x}$, $i \in \{R, B\}$. With the modification of Eq. (51), Red's HJB equation reads

$$\begin{aligned}
 -\frac{\partial V_R}{\partial t} &= -\frac{1}{4} \left(\frac{\partial V_R}{\partial x} \right)^2 - \frac{1}{2} \frac{\partial V_R}{\partial x} \frac{\partial V_B}{\partial x} \\
 &\quad - \frac{\lambda_R}{4} \left(\frac{\partial V_B}{\partial x} \right)^2 - \Theta(-x) + \frac{\sigma^2}{2} \frac{\partial^2 V_R}{\partial x^2}, \\
 V_R(x, T) &= \Phi_R(x).
 \end{aligned} \tag{54}$$

We plot solutions of Eq. (54) in Fig. 15. These solutions also do not change qualitatively from the solutions to Eqs. (11) and (12) in that there is no shock or traveling wave formation [68]. A more fundamental qualitative change would be to expand the scope of Red's interference to alter the latent volatility of the election process. Red's additional objective might be to increase the uncertainty in polling results.

In addition to theoretical modifications, other work could extend these results to other elections using similarly fine-grained or more granular data. This approach is difficult because there is very little granular public data on election interference [6]. We are able to confront our model to data only because the Russian interference in the 2016 U.S. presidential election was well publicized and because the interference took place at least partially through the mechanism of Twitter, which is a public data source. We were unable to find any other publicly available data at daily (or finer) temporal resolution for any other publicly acknowledged election interference episode.

In the case study of the 2016 U.S. presidential election, our theoretical model accurately captured election interference dynamics from approximately 2 weeks after the Democratic National Convention until Election Day. However, it did not capture the dynamics of election interference accurately or precisely during the first fortnight of time under study. We believe this is because even though the election was then a two-candidate contest, there were additional election state and interference dynamics that we did not model. We believe that

these dynamics arise because the transition from interfering in many candidates' primary campaigns to interfering in only one electoral contest is not immediate. It likely takes time for the foreign intelligence agency to recalibrate their interference strategy. In addition, the foreign intelligence agency may still expend resources on influencing other candidates' supporters, even though those other candidates had been unsuccessful in their quest for inclusion in the general election. Suppose that there are initially N_R "red candidates" (candidates that the foreign intelligence service would like to win the election) and N_B "blue candidates" (candidates that the foreign intelligence service would not like to win the election). Then modeling the transition between the conventions and the general election requires collapsing the state equation from a $N_R + N_B - 2$ -dimensional stochastic differential equation (SDE) to a one-dimensional SDE. The cost functions of Red and Blue would probably also change during this transition, but we are unsure of how to model this change. Because of the dimensionality reduction in the state equation, the coupled HJB equations would change from being PDEs solved in $N_R + N_B - 2$ spatial dimensions to ones solved in one spatial dimension, as now. We did not attempt to model these dynamics, but this could be a useful expansion of our model.

We used two models in our analysis of interference in the 2016 U.S. presidential election. We used a Bayesian structural time-series model to infer the latent random variables u_R , u_B , and X , and then used these inferred values to fit parameters of the theoretical model described in Sec. II B. While the theoretical model does not have many free parameters ($2K + 3 = 23$ degrees of freedom), the structural time-series model does have many free parameters ($3T + 5 = 311$ degrees of freedom). The large number of free parameters of the structural time-series model does not mean that the model is overparameterized. At each t , we observe a number of tweets Tweets_t and a popularity rating for the candidates Z_t . From these we want to infer the distributions of the random variables $u_{R,t}$, $u_{B,t}$, and X_t . Since we want to infer the distribution of each of these three random variables at each of the T time steps, this large number of parameters is expressly necessary. If we observe M identical trials of an election process over T time steps, then the ratio of structural time-series model parameters to observed data points is given by

$$R(T, M) = \frac{3T + 5}{2MT} = \frac{3}{2}M^{-1} \left(1 + \frac{5}{3}T^{-1} \right). \quad (55)$$

With T held constant, $R(T, M) \rightarrow 0$ as M grows, while with M held constant, $R(T, M) \rightarrow \frac{3}{2}$ as T grows large. Since we observe only one draw from the election interference model, $M = 1$ in our case. However, this approach of inferring each random variable's distribution is not tractable when T becomes large since the number of parameters to fit still grows linearly with T .

One way of partially circumventing this problem is to use a variational inference approach combined with amortization of the random variables. Denote the vector of all observed random variables at time t by y_t and the vector of all latent random variables at time t by w_t . Variational inference re-

places the actual posterior distribution with an approximate posterior distribution that has a known normalization constant (thereby eliminating the need for MCMC routines to compute this constant) [69–71]. The parameters of the approximate posterior are found through optimization, which is generally much faster than Monte Carlo sampling. If the joint distribution of $y = (y_1, \dots, y_T)$ and $w = (w_1, \dots, w_T)$ is given by $p(y, w) = \prod_{t=1}^T p(y_t | w_t) p(w_t)$, then the true posterior is given by $p(w | y) = p(y, w) / p(y)$. The approximate (variational) posterior is given by $q_\theta(w) = \prod_{t=1}^T q_{\theta_t}(w_t)$, where $\theta = (\theta_1, \dots, \theta_T)$ is the vector of parameters found through optimization and $q_{\theta_t}(w_t)$ are the probability distributions for each time step. The normalizing constants are known for each $q_{\theta_t}(w_t)$.

Amortization of the random variables means that, instead of finding the optimal value of the entire length- T vector θ , we model the approximate posterior as $q_\psi(w) = \prod_{t=1}^T q(w_t | f_\psi(y_t))$ [70,72]. The vector ψ is the vector of parameters of the (probably nonlinear) function $f_\psi(\cdot)$ and does not scale with T . The function $f_\psi(\cdot)$ models the effect of the time-dependent parameters θ_t of the variational posterior. This amortized variational posterior is also fit using an optimization routine. Since the number of parameters of this model does not scale with time, Eq. (55) for this model becomes

$$R_{\text{amort}}(T, M) = \frac{P}{2MT}, \quad (56)$$

where P is the constant number of parameters (the dimension of ψ) in the amortized model. For fixed M , $R_{\text{amort}}(M, T) \rightarrow 0$ as T becomes large. Another useful extension of our present work would be to reimplement our Bayesian structural time-series model using amortized variational inference. This would also eliminate the problem of choosing the "correct" number of parameters in the Legendre polynomial approximation that we described in Sec. III.

ACKNOWLEDGMENTS

The authors are grateful for financial support from the Massachusetts Mutual Life Insurance Company and are thankful for the truly helpful comments from an anonymous reviewer.

APPENDIX: COUPLING PARAMETER SWEEPS

We conducted parameter sweeps over different values of the coupling parameters λ_R and λ_B for multiple combinations of Red and Blue final conditions. To do this, we integrated Eqs. (11) and (12) for $\lambda_R, \lambda_B \in [0, 3]$ for each of the $3^2 = 9$ combinations of the final conditions

$$\Phi_R(x) \in \{x, 2 \tanh x, \Theta(x) - \Theta(-x)\} \quad (A1)$$




FIG. 16. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to mean of control policy.




FIG. 18. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to mean of control policy.




FIG. 17. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = 2[|\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to mean of control policy.




FIG. 19. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to mean of control policy.




FIG. 20. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to mean of control policy.




FIG. 22. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to mean of control policy.




FIG. 21. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to mean of control policy.




FIG. 23. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to mean of control policy.




FIG. 24. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to mean of control policy.




FIG. 26. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to std of control policy.




FIG. 25. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to std of control policy.




FIG. 27. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = x$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to std of control policy.




FIG. 28. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to std of control policy.




FIG. 30. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to std of control policy.




FIG. 29. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = 2[\Theta(x) - \Theta(-x)]$ and Blue final condition $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to std of control policy.




FIG. 31. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2\Theta(-x)$. Intensity of color corresponds to std of control policy.




FIG. 32. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = 2[\Theta(|x| - 0.1) - \Theta(0.1 - |x|)]$. Intensity of color corresponds to std of control policy.

and

$$\Phi_B(x) \in \left\{ \frac{1}{2}x^2, \frac{1}{2}x^2\Theta(-x), 2[\Theta(|x| - \Delta) - \Theta(\Delta - |x|)] \right\}. \quad (\text{A2})$$

We integrated Eqs. (11) and (12) for $N_t = 8000$ time steps, setting $t_0 = 0$ and $T = 1$ year. We enforced Neumann boundary conditions on the interval $x \in [-3, 3]$ and set the spatial step size to be $dx = 0.025$. After integration, we then drew 1000 paths $u_R^{(n)}$ and $u_B^{(n)}$ from the resulting probability distribution over control policies. [This probability distribution is a time-dependent multivariate Gaussian, since the control policies are deterministic functions of the random variable defined by the Ito stochastic differential equation Eq. (3)]. We then calculated the intertemporal means and standard deviations of these control policies. We draw N control policies $u_R(t)^{(n)}$ and $u_B(t)^{(n)}$ from the theoretical model and then compute

$$\text{mean}(u_i) = \frac{1}{TN} \sum_{n=1}^N \int_0^T u_i^{(n)}(t) dt \quad (\text{A3})$$

- [1] P. S. A. Renaud and F. A. A. M. van Winden, On the importance of elections and ideology for government policy in a multi-party system, in *The Logic of Multiparty Systems* (Springer, Berlin, 1987), pp. 191–207.
 [2] J. Elklit and P. Svensson, What makes elections free and fair? *J. Democr.* **8**, 32 (1997).




FIG. 33. Parameter sweep over coupling parameters λ_R, λ_B with Red final condition $\Phi_R(x) = \tanh(x)$ and Blue final condition $\Phi_B(x) = \frac{1}{2}x^2$. Intensity of color corresponds to std of control policy.

and

$$\text{std}(u_i) = \sqrt{\frac{1}{TN} \sum_{n=1}^N \int_0^T [u_i^{(n)}(t) - \text{mean}(u_i)]^2 dt}, \quad (\text{A4})$$

for $i \in \{R, B\}$. We display the mean control policies in Sec. A 1 and the standard deviations of the control policies in Sec. A 2.

1. Expected value of $u_i(t)$

We display the mean paths of draws from the distributions of control policies generated by solutions to Eqs. (11) and (12) in Figures 16–24.

2. Standard deviation of $u_i(t)$

We display the standard deviation of draws from the distributions of control policies generated by solutions to Eqs. (11) and (12) in Figures 25–33.

- [3] D. H. Levin, When the great power gets a vote: The effects of great power electoral interventions on election results, *Int. Stud. Q.* **60**, 189 (2016).
 [4] S. Shulman and S. Bloom, The legitimacy of foreign intervention in elections: The ukrainian response, *Rev. Int. Stud.* **38**, 445 (2012).

- [5] D. Corstange and N. Marinov, Taking sides in other people's elections: The polarizing effect of foreign intervention, *Am. J. Polit. Sci.* **56**, 655 (2012).
- [6] D. H. Levin, Partisan electoral interventions by the great powers: Introducing the PEIG Dataset, *Conflict Manage. Peace Sci.* **36**, 88 (2019).
- [7] E. D. Borghard and S. W. Loneragan, Confidence building measures for the cyber domain, *Strateg. Stud. Q.* **12**, 10 (2018).
- [8] D. H. Levin, Voting for trouble? Partisan electoral interventions and domestic terrorism, *Terrorism and Political Violence*, **1** (2018).
- [9] I. Hansen and D. J. Lim, Doxing democracy: Influencing elections via cyber voter interference, *Contemp. Polit.* **25**, 150 (2019).
- [10] J. Bubeck, K. Jäger, N. Marinov, and F. Nanni, Why do states intervene in the elections of others? The role of incumbent-opposition divisions, available at SSRN 3435138 (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3435138.
- [11] Read the Mueller Report: Searchable Document and Index, *New York Times*, April 2019, <https://www.nytimes.com/interactive/2019/04/18/us/politics/mueller-report-document.html>.
- [12] C. R. Nelson and C. R. Plosser, Trends and random walks in macroeconomic time series: Some evidence and implications, *J. Monet. Econ.* **10**, 139 (1982).
- [13] D. N. DeJong, J. C. Nankervis, N. Eugene Savin, and C. H. Whiteman, Integration versus trend stationarity in time series, *Econom.: J. Econom. Soc.* **60**, 423 (1992).
- [14] C. W. Gardiner, *Handbook of Stochastic Methods*, Vol. 3 (Springer, Berlin, 1985).
- [15] P. Jung and P. Hänggi, Dynamical systems: A unified colored-noise approximation, *Phys. Rev. A* **35**, 4464 (1987).
- [16] P. Hänggi and P. Jung, Colored noise in dynamical systems, *Adv. Chem. Phys.* **89**, 239 (1995).
- [17] R. Bellman, Dynamic programming and a new formalism in the calculus of variations, *Proc. Natl. Acad. Sci. USA* **40**, 231 (1954).
- [18] R. Bellman, Dynamic programming, *Science* **153**, 34 (1966).
- [19] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [20] H. J. Kappen, An introduction to stochastic control theory, path integrals and reinforcement learning, *AIP Conf. Proc.* **887**, 149 (2007).
- [21] K. J. Åström, *Introduction to Stochastic Control Theory* (Courier Corporation, London, 2012).
- [22] T. T. Georgiou and A. Lindquist, The separation principle in stochastic control, redux, *IEEE Trans. Autom. Control* **58**, 2481 (2013).
- [23] E. Zermelo, Über eine anwendung der mengenlehre auf die theorie des schachspiels, in *Proceedings of the 5th International Congress of Mathematicians*, Vol. 2 (Cambridge University Press, Cambridge, UK, 1913), pp. 501–504.
- [24] J. Nash, Non-cooperative games, *Ann. Math.* **54**, 286 (1951).
- [25] D. Fudenberg and J. Tirole, Noncooperative game theory for industrial organization: An introduction and overview, *Handbook of Industrial Organization* (Elsevier, Amsterdam, Netherlands, 1989), Vol. 1, pp. 259–327.
- [26] A. Mas-Colell, M. D. Whinston, J. R. Green *et al.*, *Microeconomic Theory*, Vol. 1 (Oxford University Press, New York, 1995).
- [27] R. Buckdahn and J. Li, Stochastic differential games and viscosity solutions of Hamilton–Jacobi–Bellman–Isaacs equations, *SIAM J. Contr. Optim.* **47**, 444, (2008).
- [28] T. Pham and J. Zhang, Two person zero-sum game in weak formulation and path dependent Bellman–Isaacs equation, *SIAM J. Contr. Optim.* **52**, 2090 (2014).
- [29] C. K. I. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning*, Vol. 2 (MIT Press, Cambridge, MA, 2006).
- [30] Code to recreate simulations and plots in this paper, or to create new simulations and “what-if” scenarios, is available at <https://gitlab.com/daviddehurst/red-blue-game>.
- [31] There is a large literature on nonparametric functional approximation that we cannot review here. Here is an example of this type of approximation. If $\{x_k\}_{1 \leq k \leq K}$ is any partition of $[a, b]$ and the vector $(\Phi_R(x_k), \Phi_B(x_k))$ were jointly distributed Gaussian on \mathbb{R}^{2K} , then we say that Φ_R and Φ_B are distributed according to a Gaussian process [29]. Though this infinite-dimensional formalism can be useful when deriving theoretical results, in practice we would approximate this process by a finite-dimensional multivariate Gaussian. In Sec. III we actually approximate Φ_R and Φ_B by finite sums of Legendre polynomials and infer the coefficients of these polynomials as our multivariate approximation.
- [32] M.-H. Chen and Q.-M. Shao, Monte Carlo estimation of Bayesian credible and HPD intervals, *J. Comput. Graph. Stat.* **8**, 69 (1999).
- [33] V. E. Yarynich, *C3: Nuclear Command, Control Cooperation* (Center for Defense Information, Washington, DC, 2003).
- [34] A. M. Barrett, False alarms, true dangers? RAND Corporation document PE-191-TSF (2016).
- [35] H. J. Kappen, Path integrals and symmetry breaking for optimal control theory, *J. Stat. Mech.* (2005) P11011.
- [36] M. Kac, On distributions of certain Wiener functionals, *Trans. Am. Math. Soc.* **65**, 1 (1949).
- [37] Y. Chen, A. Baule, H. Touchette, and W. Just, Weak-noise limit of a piecewise-smooth stochastic differential equation, *Phys. Rev. E* **88**, 052103 (2013).
- [38] J. Leifeld, K. Hill, and A. Roberts, Persistence of saddle behavior in the nonsmooth limit of smooth dynamical systems, [arXiv:1504.04671](https://arxiv.org/abs/1504.04671).
- [39] K. Kanazawa, T. G. Sano, T. Sagawa, and H. Hayakawa, Asymptotic derivation of Langevin-like equation with non-Gaussian noise and its analytical solution, *J. Stat. Phys.* **160**, 1294 (2015).
- [40] S. H. Piltz, L. Harhanen, M. A. Porter, and P. K. Maini, Inferring parameters of prey switching in a 1 predator–2 prey plankton system with a linear preference tradeoff, *J. Theor. Biol.* **456**, 108 (2018).
- [41] P. Tanchak, The invisible front: Russia, trolls, and the information war against Ukraine, *Revolution and War in Contemporary Ukraine: The Challenge of Change* (ibidem Press, Stuttgart, Germany, 2017), Vol. 161, p. 253.
- [42] B. C. Boatwright, D. L. Linvill, and P. L. Warren, Troll factories: The internet research agency and state-sponsored agenda building, Resource Centre on Media Freedom in Europe, 2018.
- [43] A. Badawy, E. Ferrara, and K. Lerman, Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign, in *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (IEEE, 2018), pp. 258–265.

- [44] R. L. Boyd, A. Spangher, A. Fourney, B. Nushi, G. Ranade, J. Pennebaker, and E. Horvitz, Characterizing the Internet Research Agency's Social Media Operations During the 2016 US Presidential Election using Linguistic Analyses, 2018.
- [45] D. J. Ruck, N. M. Rice, J. Borycz, and R. A. Bentley, Internet research agency Twitter activity predicted 2016 US election polls, *First Monday* **24** (2019).
- [46] N. G. des Mesnards and T. Zaman, Detecting influence campaigns in social networks using the Ising model, [arXiv:1805.10244](https://arxiv.org/abs/1805.10244).
- [47] J. Im, E. Chandrasekharan, J. Sargent, P. Lighthammer, Taylor Denby, A. Bhargava, L. Hemphill, D. Jurgens, and E. Gilbert, Still out there: Modeling and identifying Russian troll accounts on Twitter, [arXiv:1901.11162](https://arxiv.org/abs/1901.11162).
- [48] A. Volkening, D. F. Linder, M. A. Porter, and G. A. Rempala, Forecasting elections using compartmental models of infections, [arXiv:1811.01831](https://arxiv.org/abs/1811.01831).
- [49] D. Rothschild, Forecasting elections: Comparing prediction markets, polls, and their biases, *Publ. Opin. Q.* **73**, 895 (2009).
- [50] D. A. Linzer, Dynamic Bayesian forecasting of presidential elections in the states, *J. Am. Stat. Assoc.* **108**, 124 (2013).
- [51] W. Wang, D. Rothschild, S. Goel, and A. Gelman, Forecasting elections with non-representative polls, *Int. J. Forecast.* **31**, 980 (2015).
- [52] Data can be downloaded at <https://github.com/fivethirtyeight/russian-troll-tweets/>.
- [53] R. Skibba, Pollsters struggle to explain failures of us presidential forecasts, *Nat. News* **539**, 339 (2016).
- [54] R. Neville-Shepard, Constrained by duality: Third-party master narratives in the 2016 presidential election, *Am. Behav. Sci.* **61**, 414 (2017).
- [55] Data can be downloaded at https://www.realclearpolitics.com/epolls/2016/president/us/general_election_trump_vs_clinton_vs_johnson_vs_stein-5952.html.
- [56] D. Barber, A. T. Cemgil, and S. Chiappa, *Bayesian Time Series Models* (Cambridge University Press, Cambridge, UK, 2011).
- [57] A. K. Erlang, Sandsynlighedsregning og telefonsamtaler, *Nyt tidsskrift Mat.* **20**, 33 (1909).
- [58] G. Rasch, The poisson process as a model for a diversity of behavioral phenomena, in *International Congress of Psychology*, Vol. 2 (International Union of Psychological Science, Montreal, Canada, 1963), p. 2.
- [59] D. R. Cox, Regression models and life-tables, *J. R. Stat. Soc.: Ser. B (Methodol.)* **34**, 187 (1972).
- [60] D. R. Cox and V. Isham, *Point Processes*, Vol. 12 (CRC Press, Boca Raton, FL, 1980).
- [61] M. D. Hoffman and A. Gelman, The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* **15**, 1593 (2014).
- [62] A. Gelman and D. B. Rubin, Inference from iterative simulation using multiple sequences, *Stat. Sci.* **7**, 457 (1992).
- [63] S. P. Brooks and A. Gelman, General methods for monitoring convergence of iterative simulations, *J. Comput. Graph. Stat.* **7**, 434 (1998).
- [64] E. Brochu, V. M. Cora, and N. De Freitas, A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, [arXiv:1012.2599](https://arxiv.org/abs/1012.2599).
- [65] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, Taking the human out of the loop: A review of Bayesian optimization, *Proc. IEEE* **104**, 148 (2015).
- [66] P. I. Frazier, A tutorial on Bayesian optimization, [arXiv:1807.02811](https://arxiv.org/abs/1807.02811).
- [67] N. N. Taleb, Election predictions as martingales: An arbitrage approach, *Quant. Financ.* **18**, 1 (2018).
- [68] Using the HJB integrator `red_blue_pdes.py` included in the code at <https://gitlab.com/daviddewhurst/red-blue-game>, the reader may further investigate the effects of changes in the drift and running cost terms on the qualitative nature of solutions to the coupled HJB system.
- [69] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, Stochastic variational inference, *J. Mach. Learn. Res.* **14**, 1303 (2013).
- [70] D. J. Rezende and S. Mohamed, Variational inference with normalizing flows, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [71] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, Variational inference: A review for statisticians, *J. Am. Stat. Assoc.* **112**, 859 (2017).
- [72] C. Zhang, J. Bütepage, H. Kjellström, and S. Mandt, Advances in variational inference, *IEEE Trans. Pattern Anal. Mach. Intell.* **41**, 2008 (2018).