

Noncooperative dynamics in election interference

David Rushing Dewhurst,¹ Christopher M. Danforth,¹ and Peter Sheridan Dodds¹

¹*Computational Story Lab, Vermont Complex Systems Center,
and Department of Mathematics and Statistics, University of Vermont, Burlington, VT 05405*

Foreign power interference in domestic elections is an age-old, 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. Structures corresponding to all-or-nothing attitudes regarding the effect of the interference operations by only one player lead to an arms race in which both countries spend increasing amounts on interference and counter-interference operations. We then confront our model with data pertaining to the Russian interference in the 2016 United States presidential election contest, introducing and estimating a Bayesian structural time series model of election polls and social media posts by Russian internet trolls. We show that our analytical model, while purposefully abstract and simple, adequately captures many temporal characteristics of the election and social media activity.

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, whom we will label “Red”, may wish to influence or appear to influence the outcome of an election in another country, whom we will label “Blue”, because of the importance or perceived importance of elections in Blue with respect to Red’s interests. Such attacks on democracies are not new; it is estimated that the United States (U.S.) and Russia or 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 much formal modeling of noncooperative dynamics in an election interference game.

Here, we consider a Red - Blue two-player game in which Red wishes to influence a two-candidate, zero-sum election taking place in Blue’s country, as outlined above. In this context, we think of Red and Blue as the respective foreign (Red) and domestic (Blue) intelligence services of the two countries. Red wants a particular candidate (candidate A) to win the election, while Blue wants the effect of Red’s interference to be minimized. We characterize this problem theoretically, deriving a noncooperative, non-zero-sum differential game, 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 particular strategy, we derive further analytical results.

Turning to a recent instance of election interference, we confront our model with the 2016 U.S. presidential election in which Russia conducted interference operations [5]. After fitting a Bayesian structural time series model to election polls and social media posts associated with Russian Internet Research Agency internet trolls, We

show that our model, though simple, is able to adequately capture 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 the case of a simple election between two candidates in a homogeneous environment (e.g., no institutions such as an Electoral College) so that the election process at any time $t \in [0, T]$, a noisy representation of which is a public poll, can be represented by a scalar $Z_t \in [0, 1]$. The actual result of the election is Z_T . We hypothesize that the election dynamics take place in a latent space, where dynamics are represented by $X_t \in \mathbb{R}$. Without loss of generality, we will set $x < 0$ to be values of the latent poll that favor candidate A and $x > 0$ that favor candidate B. The latent and observable space are related by $Z_t = \phi(X_t)$, where ϕ is a smooth sigmoidal function. We will choose $\phi(x) = \frac{1}{1+e^{-x}}$. The election takes place in a population of N voting agents, each of whom updates their preferences over the candidates in the latent space at each time step t_n by a small random variable ξ_{n,t_n} . These random variables satisfy $E_n[\xi_{n,t_k}] = 0$ for all t . The election process’s increments are the sample mean of the realizations of the voting agents’ preferences at time t . In the absence of interference, the stochastic election model is thus very simple—an unbiased random walk, which we write as

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

where $\Delta t = t_{k+1} - t_k$. We display sample realizations of this process for different distributions of ξ_{n,t_k} in Fig. 1. Though one distribution of ξ_{n,t_k} describes the process




FIG. 1. Though simple, the random walk latent space election model is a unified approximation to varied population candidate preference updates. The latent election process evolves according to $X_{k+1} = X_k + \frac{1}{N} \sum_{1 \leq n \leq N} \xi_{n,k}$, where $\xi_{n,k}$ is voting agent n 's shift toward the left (< 0) or right (> 0) 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 $T = 365$. This change in political preference shift distribution describes an electorate with increasing resistance to change in their political viewpoints. We display the preference shift distributions at $t = 0$ ($t = T$) in Panel A (Panel B). For contrast, the dashed curve is a draw from the latent election process resulting from $\xi_{n,t} \sim B(1.5 \frac{T-k}{T} + 0.1 \frac{k}{T}, 1.5 \frac{T-k}{T} + 0.1 \frac{k}{T})$, which describes an electorate for which political viewpoints are in increasing flux. We show the corresponding preference shift distributions at $t = 0$ ($t = T$) in Panel D (Panel E). Despite these preference updates that are, in some sense, opposites of each other, the latent processes X_t are statistically effectively identical and are both well-modeled by the continuum approximation $dX_t = \sigma dW_t$.

of hardening of political preferences and another characterizes preferences that are in increasing flux, the sample paths of X_{t_k} are statistically effectively identical 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 process by the Wiener process, $dX_t = \sigma dW_t$, where $\sigma^2 \simeq \text{Var}\left(\frac{1}{N} \sum_{1 \leq n \leq N} \xi_{n,t}\right)$. 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 control policies are abstract variables in the context of our model but can be interpreted as monetary expenditures on interference

operations.

We will assume that Red and Blue can affect the mean trajectory of the election but not its volatility (variance of its increments). We make this assumption because X_t is an approximation to the process described by Eq. 1 and, as displayed in Fig. 1 and described above, this process's statistical characteristics do not change much even when the voting population's underlying preference change distributions are significantly different. Thus, under the influence of Red's and Blue's control policies, the election dynamics become

$$dX_t = f(u_R(t), u_B(t))dt + \sigma dW_t, \quad X_0 = y. \quad (2)$$

To first order expansion we have $f(u_R(t), u_B(t)) = a_0 + a_R u_R(t) + a_B u_B(t) + \mathcal{O}(u^2)$, so the state equation becomes

$$dX_t = [u_R(t) + u_B(t)]dt + \sigma dW_t, \quad X_0 = y, \quad (3)$$

since we have assumed *a priori* zero 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; for now, we will assume that the agents do not incur a running cost from the value of the state variable. The cost functionals can thus be written

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

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. With the assumption that it is equally costly to conduct operations that favor candidate A or candidate B and the additional requirement that the cost functions be smooth, we take the cost functions to be

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

for $i \in \{R, B\}$. The non-negative scalar λ_i parameterizes the utility gained by player i from observing player $-i$'s effort; if λ_i is high, 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.

Though the running cost functions are equivalent across players in form, the final conditions differ between Red

and Blue 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 $y > 0$; it seems reasonable to further assume that Φ_R is monotonically non-decreasing everywhere. A straightforward approximation is to assume $\Phi_R(x) = c_0 + c_1x + \mathcal{O}(x^2) \sim cx$, since the constant does not affect the relative dynamics. However, this allows the somewhat unrealistic limiting condition of infinite benefit (cost) if candidate A (B) gets 100% of the vote in the election. We will thus also consider two final conditions with bounded extremal cost: one smooth, $\Phi_R(x) \propto \tanh(x)$;

and one discontinuous, $\Phi_R(x) \propto \begin{cases} 1 & \text{if } x \geq 0 \\ -1 & \text{if } x < 0 \end{cases}$. Blue is

attempting to exactly counteract the effects of Red's control policy, hence the form of the state dynamics presented in Eq. 3. Since Blue is *a priori* indifferent between the outcomes of the election, at first glance it appears that the final condition $\Phi_B(x) = 0$ is a reasonable modeling choice. However, for the case $\lambda_B = 0$, this results in Blue taking no action at all in the game 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 it appears likely that Blue must actually have nontrivial preferences over the election outcome.

We present three possible alternatives for a cost function representing Blue's preferences. Blue may simply be suspicious that a result was due to Red's interference if X_T is too far from $E_0[X_T] = 0$, so, again enforcing smoothness, we have $\Phi_B(x) = x^2$. 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 have $\Phi_B(x) \propto \begin{cases} 0 & \text{if } x > 0 \\ x^2 & \text{if } x \leq 0 \end{cases}$.

Alternatively, Blue may accept the result of the election as long as it does not stray "too far" from the initial expected value; we have a discontinuous final condition

$$\Phi_B(x) \propto \begin{cases} 1 & \text{if } |x| > \Delta \\ -1 & \text{if } |x| < \Delta \end{cases}, \text{ where } \Delta \text{ is Blue's accepted}$$

margin of error. A nondenumerable panoply of other final conditions can be hypothesized, but these capture the range of possible payoff structures.

The application of the dynamic programming principle [6, 7] to Eqs. 3, 4, and 5 leads to a system of coupled Hamilton-Jacobi-Bellman equations for the value functions of Red and Blue,

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

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. Performing the minimization with respect to the control variables gives the Nash equilibrium control policies,

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

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

and the exact functional form of Eqs. 7 and 8,

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

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

When solved over the entirety of state space, solutions to Eqs. 11 and 12 constitute the strategies of a subgame-perfect Nash equilibrium because, no matter the action taken by player $-i$ at time t , player i is able to respond with the optimal action at time $t + dt$. Given the solution pair $V_R(x, t)$ and $V_B(x, t)$, the distribution of $Z = (x, u_R, u_B)^T$ can be written analytically. Substitution of Eqs. 9 and 10 into Eq. 3 gives

$dx = -\frac{1}{2} \left\{ \frac{\partial V_R}{\partial x} \Big|_{(t,x)} + \frac{\partial V_B}{\partial x} \Big|_{(t,x)} \right\} dt + \sigma dW$. We discretize the state equation 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 $w_n \sim \mathcal{N}(0, 1)$ 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

$$p(x_{n+1}|x_n) = \frac{1}{\sqrt{2\pi\sigma^2\Delta t}} e^{-\frac{\Delta t}{2\sigma^2} \left(\frac{x_{n+1}-x_n}{\Delta t} + \frac{1}{2}[V'_{Rn} + V'_{Bn}] - y \frac{\delta_{n0}}{\Delta t} \right)^2}. \quad (14)$$

Now, using the Markov property of X_t , we have

$$\begin{aligned} p(x_1, \dots, x_N|x_0) &= \prod_{n=0}^{N-1} p(x_{n+1}|x_n) \\ &= \frac{1}{(2\pi\sigma^2\Delta t)^{N/2}} \exp\{-S(x_1, \dots, x_N)\}, \end{aligned} \quad (15)$$

$$(16)$$

where

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

Taking $N \rightarrow \infty$ as $N\Delta t = T$ remains constant gives a standard Gaussian path integral with an action $S(x(t))$ that incorporates the derivatives of the value functions. Since u_R and u_B are just time-dependent functions of $x(t)$, their distributions can also be found explicitly using the probability distribution Eq. 16 and the appropriate (time-dependent) Jacobian transformation. Unfortunately, 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. However, we will have something to say about analytical solutions presently in the case that player i announces a credible commitment to a particular control path.

In the general case presented above, we find the value functions $V_R(x, t)$ and $V_B(x, t)$ numerically through backward iteration, enforcing a Neumann boundary condition at a large but finite value of the latent state variable. Fig. 2 displays example realizations of the value functions for different λ_i and final conditions. The value functions display diffusive behavior in common due to the game's stochasticity, but also differ qualitatively depending on the effect of the final condition propagating backward in time. When the final conditions are discontinuous, as is the case in the top panels of Fig. 2, the derivatives of the value function assume greater magnitudes and vary more rapidly throughout the game than do the derivatives of the value function when the final conditions are continuous; this is a typical feature of solutions to equations of HJB-type [8] and has consequences for the game-theoretic interpretation of these results, as we discuss below. Fig. 2 also demonstrates that the extrema of the value functions are not as large in absolute 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.

Eqs. 7 and 8 give the closed-loop control policies u_R and u_B respectively given the current state X_t and time t . We display samples of u_R , u_B , and the electoral process Z_t in Fig. 3 to illuminate some of the qualitative properties of this game before considering a more comprehensive sweep over parameters. We simulate the game with parameters

$$\lambda_R = \lambda_B = 2, \Phi_R(x) = x, \text{ and } \Phi_B(x) = \begin{cases} 0 & x > 0 \\ \frac{1}{2}x^2 & x \leq 0 \end{cases}.$$

We plot the control policies in the top panel including the mean control policies $E[u_R]$ and $E[u_B]$, displayed in thicker curves. For this parameter set, it is optimal for Red to begin play with a relatively large amount of interference and, in the mean, decrease the level of interference over time. Conversely, throughout the game Blue increases their resistance to Red's interference. Despite this resistance, the bottom panel reveals that, for this parameter set, Red is able to accomplish their objective of causing candidate A to win: in the mean case, candidate A enjoys a comfortable lead in the election poll by the final time.

To gain a better idea of the qualitative nature of this game for a more varied set of parameters, we conduct a coarse parameter sweep over λ_R , λ_B , Φ_R , and Φ_B . Fig. 4 displays the results of this parameter sweep. We plot Red and Blue Nash equilibrium strategies for each of a set of 81 ($= 3^4$) discrete parameter combinations: $\lambda_R, \lambda_B \in \{0, 1, 2\}$,

$$\Phi_R(x) \in \left\{ x, \tanh(x), \begin{cases} 2 & \text{if } x \geq 0 \\ -2 & \text{if } x < 0 \end{cases} \right\}, \text{ and } \Phi_B(x) \in$$

$$\left\{ \frac{1}{2}x^2, \begin{cases} 0 & \text{if } x > 0 \\ \frac{1}{2}x^2 & \text{if } x \leq 0 \end{cases}, \begin{cases} 2 & \text{if } |x| > \Delta \\ -2 & \text{if } |x| < \Delta \end{cases} \right\}.$$

For each combination, we generated 1000 trials of the election process and averaged the resulting strategies to generate estimates of the mean strategy profile, shown as thick red and blue curves. Higher values of the coupling parameters λ_i cause greater mean magnitude in control policies. Certain combinations of parameters lead to an "arms-race" effect in both players' control policies, wherein the Nash equilibrium strategies entail superexponential growth in the control policies near the end of the game. This occurs when at least one player has a discontinuous final condition. 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 regarding the final outcome of the election leads to an arms race that negatively affects both players. This is a general feature of any situation to which the model described by Eqs. 3–5 applies.




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 $\lambda_R = \lambda_B = 0$, $\Phi_R(x) \propto \begin{cases} 1 & \text{if } x \geq 0 \\ -1 & \text{if } x < 0 \end{cases}$, and $\Phi_B(x) \propto \begin{cases} 1 & \text{if } |x| > \Delta \\ -1 & \text{if } |x| < \Delta \end{cases}$ 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 $\Phi_B(x) \propto \begin{cases} 0 & \text{if } x > 0 \\ x^2 & \text{if } x \leq 0 \end{cases}$. 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 $N_t = 8000$ timesteps.

C. Credible commitment

If player $-i$ credibly commits to playing a particular strategy $v(t)$ on all of $[0, T]$, then the difficult problem of player i 's finding a subgame-perfect Nash equilibrium strategy profile becomes a slightly easier problem of optimal control. Player i now seeks to find the policy $u(t)$ that minimizes the functional

$$E_{u, X} \left\{ \Phi(X_T) + \int_0^T (u(t)^2 + \lambda v(t)^2) dt \right\}, \quad (18)$$

subject to the modified state equation

$$dx = [u(t) + v(t)]dt + \sigma d\mathcal{W}. \quad (19)$$

Following the logic of Eqs. 7 and 9, player i 's value function is now given by the solution to the considerably-simpler HJB equation

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

$$V(x, T) = \Phi(x). \quad (20)$$

Though nonlinear, this HJB equation can be linearized through a change of variables and subsequently be formally solved through path integral methods [9]. Setting $V(x, t) = -\eta \log \varphi(x, t)$, substituting in Eq. 20, and performing the differentiation, we are able to remove the nonlinearity iff $\frac{\eta^2}{4} \frac{1}{\varphi^2} \left(\frac{\partial \varphi}{\partial x} \right)^2 = \frac{\sigma^2 \eta}{2} \frac{1}{\varphi^2} \left(\frac{\partial \varphi}{\partial x} \right)^2$, so we find

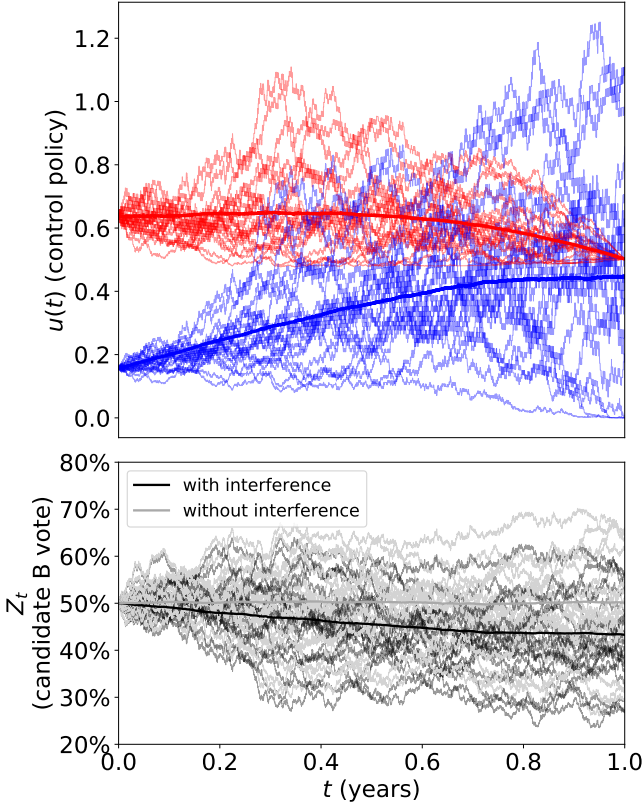


FIG. 3. We display realizations of u_R and u_B in the top panel and paths of the electoral process in the bottom panel. We draw these realizations from the game simulated with parameters $\lambda_R = \lambda_B = 2$, $\Phi_R(x) = x$, and $\Phi_B(x) = \begin{cases} 0 & x > 0 \\ \frac{1}{2}x^2 & x \leq 0 \end{cases}$. For this parameter set, Blue is fighting a losing battle—the bottom panel clearly shows that, even with Blue attempting to stop Red from interfering in the game, optimal play by both players results in a significantly lower $E[Z_t]$ than for the electoral process without any interference.

$\eta = 2\sigma^2$. Eq. 20 becomes a backward Kolmogorov equation with time-dependent drift and sink term,

$$\frac{\partial \varphi}{\partial t} = \frac{\lambda}{2\sigma^2} v(t)^2 \varphi(x, t) - v(t) \frac{\partial \varphi}{\partial x} - \frac{\sigma^2}{2} \frac{\partial^2 \varphi}{\partial x^2}, \quad (21)$$

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

Application of the Feynman-Kac formula gives the solution to Eq. 21 as [10]

$$\varphi(x, t) = \exp \left\{ -\frac{\lambda}{2\sigma^2} \int_t^T v(t')^2 dt' \right\} \times E_{Y_t} \left\{ \exp \left[-\frac{1}{2\sigma^2} \Phi(Y_T) \right] \middle| Y_t = x \right\}, \quad (22)$$

where Y_t is defined by

$$dY_t = v(t) dt + \sigma dW_t, \quad Y_0 = x. \quad (23)$$

Using this formalism, path integral control can be applied to estimate the value function for arbitrary $v(t)$. Fig. 5 displays path integral solutions to Eq. 20 when player $\neg i$ credibly commits to playing $v(t) = t^2$ for the duration of the game and player i 's final cost function takes the form

$$\Phi(x) = \begin{cases} 1 & \text{if } |x| > 1 \\ -1 & \text{if } |x| < 1 \end{cases}. \quad \text{In the further restricted case}$$

where there is a credible commitment by one party to play $v(t) = v$, a constant control policy, we can say more about the nature of solutions. We will also show presently why this constraint is actually not all that restrictive. Under this assumption, the probability law corresponding with Eq. 23 is given by

$$u(y, t) = \frac{1}{\sqrt{2\pi\sigma^2 t}} \exp \left\{ -\frac{1}{2\sigma^2 t} [(y-x) - vt]^2 \right\}, \quad (24)$$

so that the (exponentially-transformed) value function reads

$$\varphi(x, t) = \frac{\exp \left\{ -\frac{\lambda v^2}{2\sigma^2} (T-t) \right\}}{\sqrt{2\pi\sigma^2 (T-t)}} \times \int_{-\infty}^{\infty} \exp \left\{ -\frac{1}{2\sigma^2} \left[\Phi(y) + \frac{((y-x) - v(T-t))^2}{T-t} \right] \right\} dy. \quad (25)$$

This integral can be evaluated exactly for many $\Phi(y)$ and can always be approximated using the method of Laplace. If σ is small, Laplace's approximation to the integral reads

$$\int_{-\infty}^{\infty} \exp \left\{ -\frac{1}{2\sigma^2} \left[\Phi(y) + \frac{((y-x) - v(T-t))^2}{T-t} \right] \right\} dy \simeq \sqrt{2\pi\sigma^2 (T-t)} \exp \left\{ -\frac{1}{2\sigma^2} \Phi(x + (T-t)v) \right\},$$

so that, inverting the transformation above, the value function can be approximated by

$$V(x, t) = \lambda v^2 (T-t) + \Phi(x + (T-t)v), \quad (26)$$

and the control policy by

$$u(t) = -\frac{1}{2} \Phi'(x + (T-t)v). \quad (27)$$

Fig. 6 displays the results of approximating the value function with Eq. 26 at $t = 0$, along with the true (numerically-determined) value function at both $t = 0$ and $t = T$. Even with these seemingly-restrictive assumptions, this theory can be used to provide a model-agnostic method for value function approximation 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 ,

$$V(x, t + \Delta t) \simeq \lambda v(t)^2 (T-t) + \Phi(x + (T-t)[v(t) + v'(t)\Delta t]). \quad (28)$$




FIG. 4. The universe of possible strategies varies qualitatively based on the coupling parameters λ_R and λ_B and the final conditions Φ_R and Φ_B . We obtain the value functions V_R and V_B through numerical solution of Eqs. 11 and 12, and then calculate the realized Nash equilibrium strategies u_R and u_B for 1000 simulations of the election process. We then average these simulations and plot them in thick red and blue curves, representing $E[u_R(t)]$ and $E[u_B(t)]$ respectively. Along with these mean strategies, we plot two arbitrary realizations of u_R and u_B in thinner red and blue curves. For parameter combinations leading to large fluctuations near the final time, we display the final 10% of the curves in inset plots with a linear-log scaled vertical axis and log scaled horizontal axis.

In application, both of $v(t)$ and $v'(t)$ can be estimated from possibly-noisy data on $t' \in [0, t]$ and used in this approximation.

III. APPLICATION

A recent notable example of election interference operations is that of the Russian military foreign intelligence




FIG. 5. Result of the path integral Monte Carlo solution method applied to Eq. 20 with the final condition $\Phi(x) = \begin{cases} 1 & \text{if } |x| > 1 \\ -1 & \text{if } |x| < 1 \end{cases}$ and $v(t) = t^2$. Approximate value functions are computed using $N = 10000$ trajectories sampled from Eq. 23 for each point (x, t) . Approximate value functions are displayed in Panel A for $t \in \{0, 0.75, 1 - dt\}$ and the corresponding approximate control policies in Panel B, along with their smoothed counterparts (15-step moving averages, plotted in dashed curves). Panel C displays realizations of Y_t , the process generating the measure under which the solution is calculated. This method can be advantageous over numerical solution of the nonlinear PDE when the final condition is discontinuous, as here, since in this case Eq. 20 has a solution for all $t \in [0, T]$ only in the sense of distributions. The analytical control vector at $t = T$ is given by $u(t) = -\frac{1}{2}[\delta(x - 1) - \delta(x + 1)]$.

service (GRU)’s and Internet Research Agency (IRA)’s operations in the 2016 U.S. presidential election contest to attempt to harm one candidate’s chances of winning (Hillary Clinton) and aid another candidate (Donald Trump) [5]. Though Russian foreign intelligence had conducted election interference operations in the past at least once before, in the Ukrainian elections of 2014 [11], the 2015–16 operations were notable in that IRA operatives used the microblogging website Twitter in an attempt to influence the election outcome. When this attack vector was discovered, Twitter accounts associated with IRA activity were shut down and all data associated with those accounts was collected and analyzed [12, 13]. However, to the best of the authors’ knowledge, there exists no publicly-available effort to reverse-engineer the exact qualitative nature of the control vectors used by the the IRA—Red team—and by U.S. domestic and foreign intelligence agencies—Blue team.




FIG. 6. 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 $V(x, t) \simeq \lambda v^2(T - t) + \Phi(x + (T - t)v)$. The numerically-determined value functions at time $t = 0$ are shown above in black curves, while the Laplace approximations at $t = 0$ are displayed in dashed curves. The curves of lighter hue are the value functions at the final time T . The top panel demonstrates results for the final condition $\Phi(x) = \frac{1}{2}x^2$, while the bottom panel has $\Phi(x) = \tanh(x)$.

In an effort to perform data-driven simulation of Red-Blue dynamics, we fit a form of the model described in Sec. II A and compare it to qualitative theoretical predictions, finding the free parameters in the model that best describe the observed data and inferred latent controls.

Since the U.S. presidential election system is of nontrivial complexity, owing both to the number of minor party candidates that also compete and also to the unique Electoral College system, we make the simplifying assumptions stated in Sec. I—namely, that only two candidates contest the election and that the election process is modeled by a simple “candidate A versus candidate B” poll. We can observe neither the Red $u_R(t)$ nor Blue $u_B(t)$ control policies, but are able to observe a proxy for u_R , namely, the number of tweets sent by IRA-associated accounts in the year leading up to the 2016 election [14]. This dataset contains a total of 2,973,371 tweets from 2,848 unique Twitter handles. Of these tweets, a total of 1,107,361 occurred in the year immediately preceding the election (08/11/2015 - 08/11/2016). We grouped these tweets by day and used the time series $N_{\text{tweets}}(t) = \text{total number of tweets on day } t$ as an observable from which u_R could be inferred. This time series was subsequently normalized to have (intertemporal) zero mean and unit variance for ease in inference. We used the RealClearPolitics poll aggregation as a proxy for the electoral process itself [15], averaging polls that occurred at the same timestamp 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.

With these two observables, we fit a Bayesian structural

time series model of the form presented in Fig. 7, written as

$$\left\{ \begin{array}{l} \mu_B \sim \text{Normal}(0, 1) \\ \mu_R \sim \text{Normal}(0, 1) \\ \sigma \sim \text{HalfCauchy}(2) \\ u_B \sim \text{GaussianRandomWalk}(\mu_B, \sigma) \\ u_R \sim \text{GaussianRandomWalk}(\mu_R, \sigma) \\ X \sim \text{GaussianRandomWalk}(u_B - u_R, 1) \\ \sigma_{\text{tweet}} \sim \text{HalfCauchy}(0.5) \\ N_{\text{tweets, norm}} \sim \text{Normal}(u_R, \sigma_{\text{tweet}}) \\ \sigma_e \sim \text{HalfCauchy}(0.25) \\ \alpha \sim \text{HalfCauchy}(0.5) \\ e \sim \text{Normal}\left(\frac{1}{1+e^{-\alpha X}}, \sigma_e\right). \end{array} \right. \quad (29)$$

We have written X for the election process in latent space and e for the observed election poll results on $[0, 1]$. We fit the model of Eq. 29 using the No-U-Turn Sampler algorithm [16], sampling 2000 draws from the model’s distribution from each of two independent Markov chains, not including 1000 draws per chain of burn-in. The sampler appeared to converge well based on graphical consideration of draws from the posterior predictive distribution of $N_{\text{tweets, norm}}$ and e , and because Gelman-Rubin statistics for all variables had mean value $E[\hat{R}] < 1.01$ except for that of σ ($E[\hat{R}_\sigma] = 1.0865$). Fig. 8 displays draws from the latent distributions of u_R and u_B , along with draws from the posterior predictive distributions of e and $N_{\text{tweets, norm}}$ overlaid with the respective observed series. Latent control policies display a roughly-convex shape, with higher values in the beginning of the time under study, a steady decrease toward the middle of time, and a more rapid increase toward larger values near $t = 1$.

To gain qualitative insight from the theoretical model proposed above, we need to find the parameter values $\theta = (\lambda_R, \lambda_B, \sigma, \Phi_R, \Phi_B)$ that best explain the observed data and latent control vectors fit using the model of Eq. 29. We use Legendre polynomials to approximate the functions Φ_R and Φ_B , setting $\Phi_i(x) \simeq \sum_{k=0}^K a_{ik} P_k(x)$, so that the actual parameter vector considered is $\theta = (\lambda_R, \lambda_B, \sigma, a_{0,r}, \dots, a_{K,r}, a_{0,b}, \dots, a_{K,b})$. Bayes’s theorem states that

$$p(\theta|Z, M) = \frac{p(Z|\theta, M)p(\theta|M)}{p(Z|M)},$$

where M is the model presented above, and Z is the collection of observed data and unobservable inferred latent variables, $Z = (X, u_R, u_B)$. Though Eq. 16 and subsequent application of time-dependent Jacobian transformations gives us an analytical method of determining the distributions of X , u_R , and u_B , and hence a method for determining $p(\theta|Z, M)$, in practice this result is not much help since we still must numerically solve the PDEs Eqs. 11 and 12 in order to estimate V'_R and V'_B , as well as numerically estimate the appropriate Jacobians.

To avoid at least the numerical Jacobian estimation, we approximate the likelihood function using a nonparametric kernel-based method with kernel function K , so that

$$P(Z|\theta, M) \simeq \frac{1}{N} \sum_{n=1}^N K(Z - Z_n|\theta). \quad (30)$$

The values $Z_n|\theta$ are draws from the distribution of $(X_t, u_R(t), u_B(t))$ generated by the analytical model given a particular value of θ . A maximum *a posteriori* estimate of θ is then given by

$$\theta^* \simeq \arg \max_{\theta} p(\theta|M) \sum_{n=1}^N K(Z - Z_n|\theta). \quad (31)$$

We will not state any *a priori* beliefs over θ and so set $p(\theta|M) \propto 1$, so that an MAP estimate of θ can be found through maximization of $\sum_{n=1}^N K(Z - Z_n|\theta)$. We set $K(x)$ to be a Gaussian kernel and maximize a stochastic estimation of $\log_{10} \sum_{n=1}^N K(Z - Z_n|\theta)$ using $N = 500$ draws from the distribution of paths corresponding with the value functions given by solutions of the system of HJB equations. This maximization was performed using the differential evolution algorithm [17]. Fig. 9 displays draws from the estimated $p(Z|\theta^*, M)$, the latent u_R , u_B , and X inferred in fitting the model of Eq. 29, and the the estimated Φ_R and Φ_B . The coupling parameters are estimated to be $\lambda_R = 0.716$ and $\lambda_B = 1.054$, while the final-time cost functions are approximated in the optimum by $\Phi_R(x) \simeq -1.713P_0(x/2) + 3.269P_1(x/2) - 1.221P_2(x/2) + 0.266P_3(x/2) - 2.884P_4(x/2)$ and $\Phi_B(x) \simeq -2.823P_0(x/2) - 3.993P_1(x/2) + 1.812P_2(x/2) + 0.696P_3(x/2) + 2.239P_4(x/2)$. As hypothesized in Sec. II, Red’s final cost is increasing on most of its spatial domain, while Blue’s final cost is decreasing on most of its spatial domain. Draws of u_R and u_B from $p(Z|\theta^*, M)$ show qualitatively good agreement with u_R and u_B inferred from data over most of the time range under consideration, but do not accurately model the dynamics at the very beginning (approximately first 45 days) or end (approximately last 45 days) of the 2016 election cycle. Draws of X from $p(Z|\theta^*, M)$ show qualitatively good agreement with X inferred from data over nearly all of the time range under consideration except for the first approximately 20 days.

DISCUSSION AND CONCLUSION

We introduce, analyze, and numerically solve (analytically solve in simplified cases) a simple, first-principles model of noncooperative strategic interference by a foreign intelligence service from one country (Red) in an election occurring in another country (Blue) and attempts by Blue’s domestic intelligence service to counter this interference. Though simple, our model is able to provide qualitative insight into the dynamics of




FIG. 7. A discrete-time version of the model discussed in Section II with observable proxies for the latent election process and Red control policy. Latent control vectors by Red and Blue agents can be inferred by noisy observations of two associated time series. Tweets by Twitter “troll” accounts associated with the Russian Internet Research Agency (Red tweets) were grouped by day and the total count measured, while a relative poll (election $\equiv \frac{\text{Clinton}}{\text{Clinton} + \text{Trump}}$) in the 365 days prior to the 2016 U.S. presidential election (which occurred on 08/11/2016) was computed using RealClearPolitics poll aggregation data. The number of tweets was normalized as $N_{\text{tweets, norm}} = \frac{N_{\text{tweets}} - E[N_{\text{tweets}}]}{\text{Std}(N_{\text{tweets}})} - \min \frac{N_{\text{tweets}} - E[N_{\text{tweets}}]}{\text{Std}(N_{\text{tweets}})}$ for stable model fitting.

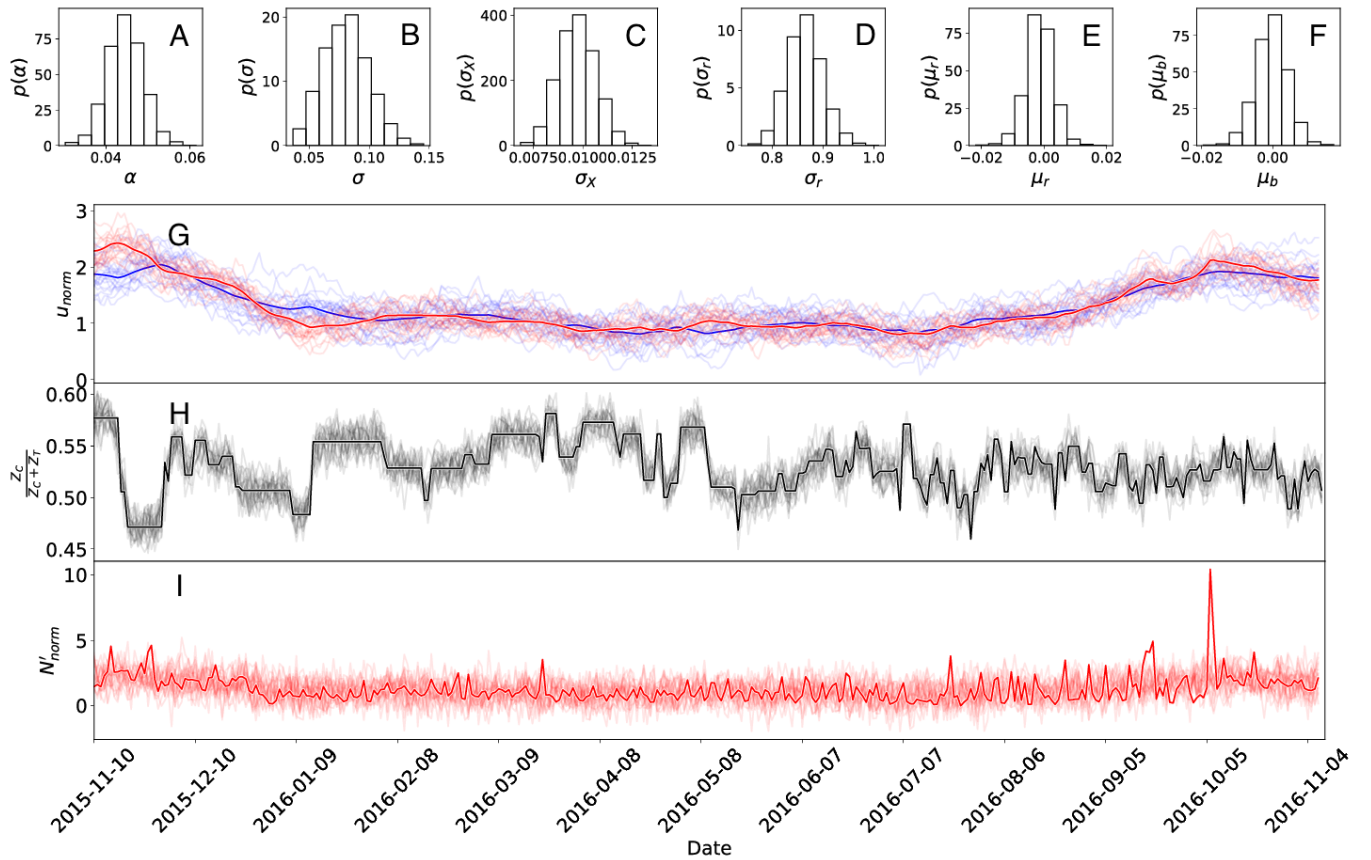


FIG. 8. Fitting of the graphical model displayed in Fig. 7 results in the latent control policies displayed in Panel G with the inferred u_R displayed in red and inferred u_B displayed in blue, and posterior predictive distributions of the observed election process (Panel H) and observed normalized IRA tweet activity (Panel I). Panel A–F display marginal posterior distributions of model parameters α , μ_R , μ_B , σ , σ_e , and σ_{tweets} . As noted above, the number of tweets was normalized as $N_{\text{tweets, norm}} = \frac{N_{\text{tweets}} - E[N_{\text{tweets}}]}{\text{Std}(N_{\text{tweets}})} - \min \frac{N_{\text{tweets}} - E[N_{\text{tweets}}]}{\text{Std}(N_{\text{tweets}})}$ for stable model fitting; this normalization procedure is why posterior predictive samples N_{norm} may have values below zero. The spike in tweets in early October of 2016 co-occurred with the U.S. Director of National Intelligence and Department of Homeland Security’s joint statement accusing Russia of interference in the presidential election process.




FIG. 9. Using a kernel-based approximation to the functional density $p(Z|\theta, M)$, we used the differential evolution algorithm to compute a maximum *a posteriori* estimate of the parameter vector θ . Panel A displays draws of u_R and u_B from the distribution $p(Z|\theta^*, M)$ in thin curves, along with the inferred values of u_R and u_B from the 2016 U.S. presidential election. The model describes the data well except near the beginning and end of the time period under study, in which dynamics other than those modeled appear to be predominant. Panel C displays draws of X from $p(Z|\theta^*, M)$, which also appear to describe the data well. Panel B shows the MAP estimates for Φ_R (red curve) and Φ_B (blue curve).

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 (whether or not it was successful) with no gradation of utility, even if these attitudes are held by only one player, result in an arms race of spending on interference and counter-interference operations by both players. We then find analytical solutions to player i 's optimal control problem in the case where player $-i$ credibly commits to a strategy $v(t)$ and detail an analytical value function approximation that can be used by player i even when player $-i$ does not commit to a particular strategy as long as player $-i$'s current strategy and its 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 Internet Research Agency (IRA) troll posts on the website Twitter. Using this 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 (latent) control policies. We show that, for most of the time under consideration, our model provides a good explanation for the inferred variables.

There are several areas in which our work could be improved. From a theoretical point of view, our model is one of the simplest that can be proposed to model this situation. While from an *a priori* point of view it is derived from first principles and makes a minimum of assumption about the election mechanics, electorate, and cost (equivalently, utility) functions of the respective intelligence agencies and hence is justifiable on the grounds of parsimony and acceptable empirical performance (on at least one election contest), the lack of assumptions that we make is rather unrealistic. Though a pure random walk model for an election is not without serious precedent [18], a prudent extension of this work could incorporate non-interference-related state dynamics as a generalization of Eq. 3, e.g., as

$$dx = [\mu_0 + \mu_1 x + u_R(t) + u_B(t)]dt + \sigma dW. \quad (32)$$

This state equation can account for simple drift in the election results as a candidate endogenously becomes more or less popular or capture possible mean-reverting behavior in a hotly-contested race. Another interesting extension would allow for state-dependent running

costs, particularly in the running cost of the Red player. Though the action of election interference is nominally intended to cause a particular candidate to win or lose, there are often other goals as well, such as undermining the Blue citizens' trust in their electoral process. Thus, Red might gain utility even just from having a particular candidate pull ahead 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 can be represented by setting Red's cost functional to be

$$E_{u_R, u_B, X} \left\{ \Phi_R(X_T) + \int_0^T [-\mathbf{1}_{X_t < 0}(X_t) + u_R^2(t) - \lambda_R u_B^2(t)] dt \right\}. \quad (33)$$

Both of these modifications are relatively easy to incorporate into the model, as demonstrated, and will not qualitatively change the nature of the value functions V_R and V_B since their effects will simply be to introduce an additional drift term (Eq. 32) or a continuous, non-differentiable source term (Eq. 33). A more fundamental qualitative change would be to expand the scope of Red's interference to alter the latent volatility of the election process. For example, the objective of Red's interference operations might be not only to change the drift of the state equation to make it more likely for candidate A to win, but also to increase the uncertainty associated with the election's polling.

In addition to theoretical modifications, other work could simply extend the present results to other elections. The principal difficulty with this approach lies in the inherent difficulty of finding any data at all with which to work. As we note in Sec. III, we are able to confront our model to data only because the Russian interference in the 2016 U.S. presidential election was so well-publicized, and even so we found it necessary to infer the variables in which we were actually interested. Other than this event, we were unable to find any publicly-available data for any other election interference episode that is in the public domain.

ACKNOWLEDGEMENTS

The authors are grateful for insightful conversations with Brian Tivnan and financial support from the Massachusetts Mutual Life Insurance Company.

[1] Paul SA Renaud and Frans AAM van Winden. On the importance of elections and ideology for government policy in a multi-party system. In *The logic of multiparty systems*, pages 191–207. Springer, 1987.

[2] Jorgen Elklit and Palle Svensson. What makes elections free and fair? *Journal of democracy*, 8(3):32–46, 1997.

[3] Dov H Levin. When the great power gets a vote: The effects of great power electoral interventions on election results. *International Studies Quarterly*, 60(2):189–202, 2016.

[4] Stephen Shulman and Stephen Bloom. The legitimacy of foreign intervention in elections: The ukrainian response.

- Review of International Studies*, 38(2):445–471, 2012.
- [5] Read the Mueller Report: Searchable Document and Index. *New York Times*, Apr 2019.
- [6] Richard Bellman. Dynamic programming and a new formalism in the calculus of variations. *Proceedings of the National Academy of Sciences of the United States of America*, 40(4):231, 1954.
- [7] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.
- [8] Michael G Crandall, Hitoshi Ishii, and Pierre-Louis Lions. Users guide to viscosity solutions of second order partial differential equations. *Bulletin of the American mathematical society*, 27(1):1–67, 1992.
- [9] Hilbert J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.
- [10] Mark Kac. On distributions of certain Wiener functionals. *Transactions of the American Mathematical Society*, 65(1):1–13, 1949.
- [11] Peter Tanchak. The Invisible Front: Russia, Trolls, and the Information War against Ukraine. *Revolution and War in Contemporary Ukraine: The Challenge of Change*, 161:253, 2017.
- [12] Brandon C Boatwright, Darren L Linvill, and Patrick L Warren. Troll factories: The internet research agency and state-sponsored agenda building. *Resource Centre on Media Freedom in Europe*, 2018.
- [13] Ryan L Boyd, Alexander Spangher, Adam Fourney, Besmira Nushi, Gireeja Ranade, James Pennebaker, and Eric Horvitz. Characterizing the Internet Research Agencies Social Media Operations During the 2016 US Presidential Election using Linguistic Analyses. 2018.
- [14] Data can be downloaded at <https://github.com/fivethirtyeight/russian-troll-tweets/>.
- [15] Data can be downloaded at https://www.realclearpolitics.com/epolls/2016/president/us/general_election_trump_vs_clinton_vs_johnson_vs_stein-5952.html.
- [16] Matthew D Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [17] Rainer Storn and Kenneth Price. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4):341–359, 1997.
- [18] Nassim Nicholas Taleb. Election predictions as martingales: An arbitrage approach. *Quantitative Finance*, 18(1):1–5, 2018.