Influence Maximization and Equilibrium Strategies in Election Network Games

Anya Zhang, **Andrew Perrault**
Center for Research on Computation and Society, Harvard
anyazhang@college.harvard.edu, aperrault@g.harvard.edu

Abstract

Social media has become an increasingly important political domain in recent years, especially for campaign advertising. In this work, we develop a linear model of advertising influence maximization in two-candidate elections from the viewpoint of a fully-informed social network platform, using several variations on classical DeGroot dynamics to model different features of electoral opinion formation. We consider two types of candidate objectives—margin of victory (maximizing total votes earned) and probability of victory (maximizing probability of earning the majority)—and show key theoretical differences in the corresponding games, including advertising strategies for arbitrarily large networks and the existence of pure Nash equilibria. Finally, we contribute efficient algorithms for computing mixed equilibria in the margin of victory case as well as influence-maximizing best-response algorithms in both cases and show that in practice, as implemented on the Adolescent Health Dataset, they contribute to campaign equality by minimizing the advantage of the higher-spending candidate.

1 Introduction

The impact of various factors on the success of political advertising has been extensively studied in political science [Rothschild, 1978; Ridout and Franz, 2011; Fowler *et al.*, 2016; Wanat, 2020]. A newly relevant question is how the structure of social networks themselves can inform candidates' advertising strategies—or more likely, the strategies of social media giants acting on the behalf of multiple candidates at once. As advertising increasingly shifts away from traditional media and towards digital platforms built on existing friendship networks, and as the voter-specific "micro-targeting" strategies formerly relied on by Google and Facebook come under scrutiny [Zuiderveen Borgesius *et al.*, 2018; McCarthy, 2020], network effects may very well become one of the main ways in which social media companies decide how to spend the millions of advertising dollars entrusted to them. Doing so in an equitable and optimal manner for each candidate is crucial to the health of our democracy.

In this work, we study the problem of competitive influence maximization—or how agents should advertise to a network of nodes in order to maximize their total influence—in a two-candidate election using the DeGroot model. Our contributions are as follows: (1) Proposing a new model of election network games that enables candidates to spend varying amounts per voter, accounts for varying levels of voter receptiveness to each side, and updates using a DeGroot model over finite time horizons, (2) showing key differences in the properties of election games corresponding to the margin of victory and probability of victory objectives and providing best-response (influence-maximizing) algorithms for both, as well as an algorithm for computing equilibria in the former, and (3) demonstrating via simulation how these algorithms can be used as a decision-support system for fully-informed companies tasked with allocating budgets for competing candidates, and showing that equilibrium outcomes lead to a fairer democratic process.

2 Related Work

The problem of influence maximization under the independent cascade and linear threshold models was first introduced by Kempe *et al.* [2003]. Bharathi *et al.* [2007] showed that many results extended to the competitive variant, including algorithmic tractability through exploiting submodularity. Borodin *et al.* [2010] explored similar competitive extensions to the linear threshold model. Also using an independent cascade-based model, Clark and Poovendran [2011] took a game-theoretic approach to competitive influence maximization, while Wilder [2018] introduced the election setting, where each of the players represents a candidate who tries to influence the individuals in the network to support him/her in the election. The independent cascade and linear threshold models provide some analytical and computational tractability through its submodular properties, but they make the assumption that agents are committed to an opinion once "activated" once by a peer and are thereafter removed from influence. While this naturally models one-off purchases and discrete choices, it does not adequately reflect opinions lying along a continuous spectrum that shift slowly over time due to social influences, as argued by Brede *et al.* [2019].

There exists a parallel literature on network opinion dynamics, where nodes update their opinions continuously over time using either Bayesian models or non-Bayesian rules of thumb, the most well-studied of which is the DeGroot model [DeGroot, 1974; Jadbabaie *et al.*, 2011]. The goal of most work in this area is to study the conditions under which (voter) opinions converge and how quickly, such as in the presence of "stubborn" agents or competitors exist [Vial and Subramanian, 2019; Zhao *et al.*, 2014]. As far as we know, ours is the first work to study a finite-time symmetric competitive influence game on a DeGroot network with a focus on learning equilibria for two types of objectives, margin and probability of victory.

3 Election Model

We consider an election with two candidates A and B and n voters on a network, assuming a single round of advertising occurring some amount of time T before an election. Our model thus proceeds in four stages: initial state at time $t = 0$, candidate advertising, voter peer updating, and final vote at $t = T$. We use the notation $\theta_i^{(t)}$ to represent node i's opinion at time t.

1. Initial State: Voters begin with opinions $\boldsymbol{\theta}_i \in [0, 1]$ representing their likelihood of voting for candidate A.

2. Candidate Advertisement: In the advertisement stage, each candidate X's decision variable is the vector $\boldsymbol{S}_X \geq 0$, representing the real number amount they choose to spend on each voter subject to budget constraints $\sum \boldsymbol{S}_X \leq k_X$. Given $\boldsymbol{S}_A$, $\boldsymbol{S}_B$ and additional vector parameters $\boldsymbol{p}_A$ and $\boldsymbol{p}_B$ characterizing the receptiveness of each voter to advertising from each candidate, voters independently update their opinions from $\boldsymbol{\theta}_i$ to $\boldsymbol{\theta}_i^{(0)}$ as defined by the following linear process. We let $\boldsymbol{p}_{X_i}\boldsymbol{S}_{X_i}$ be the overall probability that advertising from X succeeds at reaching voter i, so that both voter receptiveness and candidate expenditures factor into i's updated opinion. To maintain this interpretation, we also require that $\boldsymbol{S}_{X_i} \leq 1/\boldsymbol{p}_{X_i}$. We now define a random variable Y_i, such that if only one candidate successfully advertises to i, Y_i switches from $\boldsymbol{\theta}_i$ to 1 (for A) or 0 (for B). If both or neither successfully advertise, the effects cancel each other out and Y_i remains at the original opinion $\boldsymbol{\theta}_i$. We then set $\boldsymbol{\theta}_i^{(0)} = E(Y_i)$ in order to obtain a deterministic and linear result for the true updated opinion of i:

$$\begin{aligned}\boldsymbol{\theta}_i^{(0)} &= \boldsymbol{\theta}_i + (1 - \boldsymbol{\theta}_i)p_{A_i}\boldsymbol{S}_{A_i} - \boldsymbol{\theta}_i p_{B_i}\boldsymbol{S}_{B_i} \\ &\quad + (2\boldsymbol{\theta}_i - 1)p_{A_i}p_{B_i}\boldsymbol{S}_{A_i}\boldsymbol{S}_{B_i}.\end{aligned}$$

3. Peer Updates: Given post-advertisement opinions at time 0 and the network's fixed trust matrix P (where P is row-stochastic and P_{ij} measures how much i trusts j), the network updates itself for the next T discrete time steps according to the DeGroot model. Each node's opinion at each time step is the trust-weighted average of those of its out-neighbors (including itself) at the previous step, and the final opinion state of the network is given by $\boldsymbol{\theta}^{(T)} = P^T\boldsymbol{\theta}^{(0)}$.

4. Final Vote: From here, nodes cast independent votes as Bernoulli random variables based on their final opinions $\boldsymbol{\theta}_i^{(T)}$, where a vote of 1 goes to A and 0 goes to B. This stochastic voting system, rather than a threshold model, is both the most common in standard election literature and allows for our ultimate distinction between the margin and probability of victory cases. The final number of votes V_A for A is the sum of the Bernoullis, or a Poisson Binomial distribution with parameter $\boldsymbol{\theta}^{(T)}$. The election outcome is determined by simple majority (A wins iff $V_A > n/2$).

4 Margin and Probability of Victory Games

Having fully defined our election model, we proceed to define two types of election network games based on objectives commonly studied in election literature: the Margin of Victory (MOV) game and the Probability of Victory (POV) game. Both are defined by agents $\{A, B\}$, action sets $\{\boldsymbol{S}_A, \boldsymbol{S}_B\}$, and constant parameters $\{n, P, \theta, T, k_A, k_B, p_A, p_B\}$.

Margin of Victory. In the MOV game, the objective of candidate A (B) is to maximize (minimize) A's expected margin of victory $\mathbb{E}(V_A - V_B)$, as in elections based on proportional representation. By the Poisson Binomial expectation, payoffs are given by $MOV_A = \mathbb{E}(V_A - V_B) = 2\mathbb{E}(V_A) - n = 2\sum_{i=0}^{n}\boldsymbol{\theta}_i^{(T)} - n$ and $MOV_B = \mathbb{E}(V_B - V_A) = -MOV_A$. Dropping unnecessary constants, the MOV objective for candidate A is simply maximizing $E(V_A)$.

Probability of Victory. In the POV game, the objective of candidate A (B) is to maximize (minimize) A's probability of earning the majority $P(V_A > n/2)$ regardless of by how large a margin, as in winner-take-all elections such as the US Presidential election. Using the Poisson Binomial CDF, payoffs are given by

$$\begin{aligned}POV_A &= P(V_A > n/2) \\ &= 1 - \sum_{\ell=0}^{n/2}\sum_{S' \in F_\ell}\left[\prod_{i \in S'}\boldsymbol{\theta}_i^{(T)}\prod_{j \in S'^c}(1 - \boldsymbol{\theta}_j^{(T)})\right]\end{aligned}$$

and $POV_B = P(V_A < n/2) = 1 - POV_A$, where F_ℓ is the set of subsets of our n nodes of size ℓ. Computing POV_A even once, not to mention optimization, requires performing a sum that is exponential in the number of voters. Thus, we rely instead on the Normal approximation provided by Hong [2013], assuming n is large enough that the central limit theorem applies:

$$POV_A = P(V_A > n/2) \approx 1 - \Phi\left(\frac{\frac{n+1}{2} - \mu}{\sigma}\right), \quad (1)$$

where μ is the expected value and σ the standard deviation of V_A, defined as $\mu = \sum_{i=0}^{n}\boldsymbol{\theta}_i^{(T)}$, $\sigma = \sqrt{\sum_i \boldsymbol{\theta}_i^{(T)}(1 - \boldsymbol{\theta}_i^{(T)})} = \sqrt{\mu - \sum_i\left(\boldsymbol{\theta}_i^{(T)}\right)^2}$, and Φ is the standard Normal CDF. By the monotonicity of Φ, maximizing Equation (1) is in turn equivalent to minimizing $\frac{\frac{n+1}{2} - \mu}{\sigma}$.

Observe that both games are constant-sum with infinite strategy sets and continuous payoffs.

Proposition 1. *Both the MOV and POV games have at least one mixed strategy Nash equilibrium.*

Proof. Because strategy sets are convex and compact, Glicksberg's existence theorem holds [Fudenberg and Levine, 1998]. However, the same is not true of pure equilibria. □

Proposition 2. *The existence of pure Nash equilibria is guaranteed in MOV games, but not POV games.*

Proof. The first part follows from Debreu-Fan-Glicksberg Theorem [Fudenberg and Levine, 1998], because the strategy set is convex and compact, and the utility function MOV_A is both continuous and quasiconcave (linear) in $\boldsymbol{S}_A$. To show the second part, we present a counterexample. Consider the following 3-node network with a doubly-stochastic trust matrix, such that every node has equal influence, where:

$$\boldsymbol{\theta} = \mathbf{p}_A = \mathbf{p}_B = \begin{bmatrix} 1 & 0 & 1 \end{bmatrix}.$$

Let A and B have identical budgets, $k_A = k_B = 1$. We claim that there are no pure strategy equilibria for the POV game by reduction to Matching Pennies: if A and B advertise to the same voter, A wins, and if they advertise to different voters, B wins. Similar reasoning applies to any possible split of budgets between the two. (Observe, however, that $[0.5, 0, 0.5]$ is a pure equilibrium for the corresponding MOV game.) $\square$

Previous work conjectures that POV and MOV amount to the same thing for large electorates by the Central Limit Theorem [Hinich, 1977; Coughlin and Nitzan, 1981]. However, we show that even for arbitrarily large networks, objectives for the two types of games may differ, illustrating the need for us to study them separately:

Theorem 1. *The margin and probability of victory objectives may differ for arbitrarily large networks.*

Proof. The main idea is that variance plays a key role in POV games as can be seen from Equation (1). In particular, we can construct arbitrarily large networks in which a losing candidate has a 0 chance of winning the POV game unless he deviates from his MOV strategy: For any large N, consider a network with two components that are entirely disconnected. Let the first component be of size $s = \lceil \frac{N+1}{2} \rceil$, and set the opinions of each node in this component to 0. Let the second component be of size $N - s$, where the nodes can have arbitrary opinions. If we then pick $\boldsymbol{\alpha}$ and $\mathbf{p}_A$ such that the maximum value of $u_A(\mathbf{X}_B)$ in the first component is still less than the minimum $u_A(\mathbf{X}_B)$ value in the second component, candidate A will never have an incentive to advertise to a node in component 1 from an expected value perspective as long there remain nodes to be convinced in component 2. However, without advertising to at least one member of the first component (increasing variance), candidate A has a zero probability of winning the election. We may similarly construct examples in which a leading candidate can increase his probability of winning to 100% by prioritizing getting $\lceil \frac{N+1}{2} \rceil$ nodes to an opinion of exactly 1 (reducing variance) instead of increasing expected votes overall. $\square$

As we discuss in the following section, this means in practice that in MOV-based elections, candidates' best strategies follow predictable patterns of advertising to the voters most susceptible to their messaging. On the other hand, in POV-based elections, both should generally focus on the strongest supporters of the predicted leading candidate, with the leader benefiting from shoring up their support and the underdog from moving as many as possible back towards the middle.

5 Maximizing Influence

For each election network game, we present best-response algorithms *MOV-Oracle* and *POV-Oracle* for maximizing a candidate's influence given the actions of the other.

MOV-Oracle. In the MOV case, rewriting the objective based on the classic DeGroot model and linearity:

$$\begin{aligned} E(V_A) &= \sum_{i=0}^{n} (P^T \boldsymbol{\theta}^{(0)})_i = \sum_{i=0}^{n} \sum_{j=0}^{n} P_{ij}^T \boldsymbol{\theta}_j^{(0)} = \sum_{j=0}^{n} \boldsymbol{\theta}_j^{(0)} \sum_{i=0}^{n} P_{ij}^T \\ &= \boldsymbol{S}_A \cdot u_A(\boldsymbol{S}_B) + \boldsymbol{\alpha} \cdot \mathbf{c}, \end{aligned}$$

defining the influence coefficient vector $\boldsymbol{\alpha} = \sum_{i=0}^{n} P_{ij}^T$, constant vector $\mathbf{c} = \boldsymbol{\theta} - \boldsymbol{\theta} \circ p_B \circ \boldsymbol{S}_B$, and marginal payoff vector $u_A(\boldsymbol{S}_B) = \boldsymbol{\alpha} \circ \left[(\vec{1} - \boldsymbol{\theta}) \circ p_A + (2\boldsymbol{\theta} - \vec{1}) \circ p_A \circ p_B \circ \boldsymbol{S}_B \right]$. Thus, *MOV-Oracle*$_A$ is simply the following linear program:

$$\max_{\boldsymbol{S}_A} \quad \boldsymbol{S}_A \cdot u_A(\boldsymbol{S}_B), \quad \text{s.t.} \quad \sum_i \boldsymbol{S}_{A_i} \leq k_A$$

This outputs A's best response given the (potentially mixed, due to linearity) $\boldsymbol{S}_B$ strategy of her opponent. Our election model thus allows the MOV influence maximization problem to be solved in quadratic time.

POV-Oracle. We define *POV–Oracle*$_A$ to minimize

$$\frac{\frac{n+1}{2} - \mu}{\sigma} = \frac{\frac{n+1}{2} - \mu}{\sqrt{\mu - \sum_i (\boldsymbol{\theta}_i^{(T)})^2}} \tag{2}$$

POV–Oracle$_A$ accepts candidate B's pure strategy $\boldsymbol{S}_B$ and generates A's best response by iterating over fixed guesses for μ and either minimizing or maximizing variance based on how μ compares to $(n+1)/2$. In particular, if we have some fixed $\mu \leq \frac{n+1}{2}$ (A is losing in expectation), (2) is maximized by minimizing $\sum_i (\boldsymbol{\theta}_i^{(T)})^2$ subject to our constraint on μ:

$$\min_{\boldsymbol{S}_B} \quad f_{\text{obj}} = \sum_i (\boldsymbol{\theta}_i^{(T)})^2, \quad \text{s.t.} \quad \sum_i \boldsymbol{\theta}_i^{(T)} = \mu$$

This is a convex problem with linear constraints, and is thus solvable using quadratic programming as long as our budget is sufficient for μ to be feasible. Putting our constraints in terms of the original decision variables yields

$$\begin{aligned} \min_{\boldsymbol{S}_A} \quad & \sum_i \left(\sum_{j=0}^{n} P_{ij}^T \boldsymbol{\theta}_j^{(0)} \right)^2, \\ \text{s.t.} \quad & \boldsymbol{S}_A \cdot u_A(\boldsymbol{S}_B) = \mu - \boldsymbol{\alpha} \cdot \mathbf{c} \\ & \boldsymbol{S}_{A_i} p_{A_i} \leq 1 \quad \forall i, \quad \sum_i \boldsymbol{S}_{A_i} \leq k_A. \end{aligned}$$

If instead $\mu > \frac{n+1}{2}$ (A is winning in expectation), (2) is maximized by maximizing $\sum_i (\boldsymbol{\theta}_i^{(T)})^2$. The form of the optimization is the same as before, with the direction reversed, and can still be solved (albeit more slowly) with nonconvex optimization methods, such as those offered by Gurobi 9.

Thus, *POV-Oracle* finds a candidate's POV influence-maximizing strategy given the pure strategy of his opponent. Unlike *MOV-Oracle*, it cannot handle mixed opponent strategies as input because of the nonlinearity of Φ.

6 Computing Equilibria

Due to the linear structure of the MOV objective, we may efficiently find equilibria of the MOV game using the *Follow the Perturbed Leader (FTPL)* algorithm of Kalai and Vempala (Alg. 1). At each iteration, each player chooses the action that would have maximized their expected payoff against the uniform distribution over all previous opponent actions, which is equivalent by linearity to best-responding to their opponent's expected action using *MOV-Oracle*. Because of FTPL's no-regret guarantee for online linear optimization, neither player can gain significantly by deviating from their historical actions after a sufficient number of iterations.

Theorem 2. *Uniform distributions over* $\{\boldsymbol{S}_B^r\}$ *and* $\{\boldsymbol{S}_B^r\}$ *form an* ϵ*-equilibrium with* $\frac{4n^2 \max(k_a, k_b)}{\epsilon^2}$ *iterations of FTPL.*

Algorithm 1: FTPL(ϵ)

Initialize $\mathbf{X}_A^0$ and $\mathbf{X}_B^0$ to all zeros
for $r = 0 \ldots R$ **do**
  Draw ρ_a, ρ_b uniformly at random from $[0, \frac{1}{\epsilon}]^m$
  $\boldsymbol{S}_A^r = MO_A(\sum_{s=1}^{r-1} u_A(\boldsymbol{S}_B^s) + \frac{\rho_a}{r-1}, k_a)$
  $\boldsymbol{S}_B^r = MO_B(\sum_{s=1}^{r-1} u_A(\boldsymbol{S}_A^s) + \frac{\rho_b}{r-1}, k_b)$
end
return means of $\{\mathbf{X}_A^r\}$ and $\{\mathbf{X}_B^r\}$

---

7 Experiments

Our experiments use the National Longitudinal Study of Adolescent to Adult Health Dataset (AHD) [Harris, 2009]. This dataset is convenient for several reasons, namely including information on both strength of relationships and demographics, which we can use as a proxy for political opinion, and being conducted in a high school setting with dense networks.

Follow the Perturbed Leader (MOV). We run the FTPL algorithm for a 32-node network, fixing the total budget at 300 and changing its division between the two candidates to minimize the effects of diminishing returns. Defining single-round best-response (SBR) to be the algorithm where each candidate maximizes their own influence using *MOV-Oracle* against an opponent budget of 0 (essentially, optimization with no knowledge of the opponent's strategy), we then plot A's expected margin of victory using FTPL against that using SBR. To compare the relative equity of the two algorithms, we borrow the concept of the *Budget Multiplier* from Goyal et. al. [2019], which measures the extent to which the wealthier player's final payoff exceeds their initial share of the budget, or $\frac{MOV_X}{MOV_Y} \times \frac{k_Y}{k_X}$ in our case, where $k_X > k_Y$. In the resulting graph (Fig. 1), we can clearly see that FTPL gives a slight advantage as compared to SBR to the candidate with a lower budget, which translates to a strictly lower Budget Multiplier for FTPL than for SBR except for when the division of budgets is exactly equal. In other words, computing equilibrium strategies using FTPL serves to level the playing field. We also see that FTPL is far more efficient than the upper bound given by Theorem 2, converging in only about a dozen iterations even for $n = 64$, $k = 300$, $\epsilon = 0.1$.

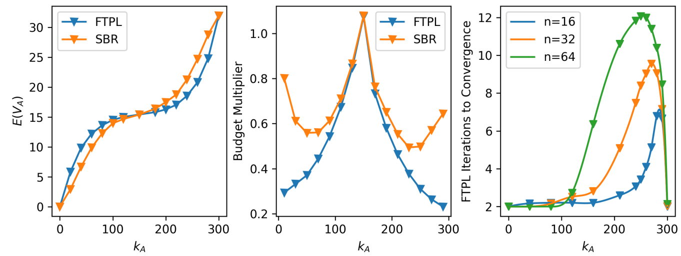

Figure 1: Left and Middle: Candidate A's utility and Budget Multiplier using FTPL and SBR for various divisions of total budget 300 between k_A and k_B averaged across randomly chosen subsets of $n = 32$ nodes of the first 50 AHD networks, with $\epsilon = 0.1$, $T = 10$, and initial opinions correlated with sex. Using FTPL always leads to a weakly lower Budget Multiplier than SBR. Right: Varying both n and k, observe that FTPL converges quickly in practice.

Alternating Best Response (POV). To find pure Nash equilibria in the probability of victory game, we devise an alternating best response algorithm that restarts randomly as soon as a cycle is detected and stops when an ϵ-equilibrium is found. This algorithm performs well for our test of up to $n = 50$ nodes, converging a majority of the time in less than 10 iterations and only occasionally taking long detours as a result of cycles (Fig. 2). However, as the number of nodes increases, the optimization problem required for each iteration also takes longer to complete, especially in the nonconvex case when the candidate of interest is predicted to win.

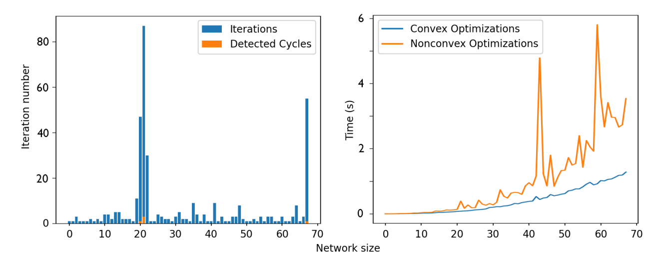

Figure 2: Results of alternating best response applied to AHD network 1 with $\epsilon = .01$, budgets $k_A = 20$ and $k_B = 30$, $T = 5$, and initial opinions correlated with sex. The image on the left shows the number of iterations required for the algorithm to converge, and the image on the right shows average runtime per optimization, of which each iteration was allowed at most 40. We see that, for the most part, alternating response converges in very few iterations, and that long convergence periods are related to the presence of cycles.

Conclusion. Our results show that the properties of social networks make them a tractable alternative to individual-focused strategies for electoral advertising, especially in proportional representation elections. We hope that, in particular, the existence of efficiently computable equilibria in election games and their dampening effect on imbalances in campaign spending argues this model is worthy of further investigation.

References

[Bharathi *et al.*, 2007] Shishir Bharathi, David Kempe, and Mahyar Salek. Competitive influence maximization in social networks. In *WINE*, 2007.

[Borodin *et al.*, 2010] Allan Borodin, Yuval Filmus, and Joel Oren. Threshold models for competitive influence in social networks. *Lecture Notes in Computer Science*, page 539–550, 2010.

[Brede *et al.*, 2019] Markus Brede, Valerio Restocchi, and Sebastian Stein. Effects of time horizons on influence maximization in the voter dynamics. *Journal of Complex Networks*, 7(3):445–468, 2019.

[Clark and Poovendran, 2011] Andrew Clark and Radha Poovendran. Maximizing influence in competitive environments: a game-theoretic approach. In *International Conference on Decision and Game Theory for Security*, pages 151–162. Springer, 2011.

[Coughlin and Nitzan, 1981] Peter Coughlin and Shmuel Nitzan. Electoral outcomes with probabilistic voting and nash social welfare maxima. *Journal of Public Economics*, 15(1):113–121, 1981.

[DeGroot, 1974] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69, 1974.

[Fowler *et al.*, 2016] Erika Franklin Fowler, Travis N Ridout, and Michael M Franz. Political advertising in 2016: The presidential election as outlier? In *The Forum*, volume 14, pages 445–469. De Gruyter, 2016.

[Fudenberg and Levine, 1998] Drew Fudenberg and David K Levine. *The Theory of Learning in Games*, volume 2. MIT press, 1998.

[Goyal *et al.*, 2019] Sanjeev Goyal, Hoda Heidari, and Michael Kearns. Competitive contagion in networks. *Games and Economic Behavior*, 113:58–79, Jan 2019.

[Harris, 2009] Kathleen Mullan Harris. *The National Longitudinal Study of Adolescent to Adult Health (Add Health)*. Chapel Hill, NC: Carolina Population Center, University of North Carolina at Chapel Hill, 2009.

[Hinich, 1977] Melvin J. Hinich. Equilibrium in spatial voting: The median voter result is an artifact. *Journal of Economic Theory*, 16(2):208–219, Dec 1977.

[Hong, 2013] Yili Hong. On computing the distribution function for the poisson binomial distribution. *Computational Statistics & Data Analysis*, 59:41–51, 2013.

[Jadbabaie *et al.*, 2011] Ali Jadbabaie, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi. Non-Bayesian social learning. *Games Econ. Behav.*, 76:210–225, 2011.

[Kalai and Vempala, 2005] Adam Tauman Kalai and Santosh S. Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71:291–307, 2005.

[Kempe *et al.*, 2003] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. *Theory of Computing*, 11:105–147, 2003.

[McCarthy, 2020] Justin McCarthy. In U.S., most oppose micro-targeting in online political ads. *Gallup*, Mar 2020.

[Ridout and Franz, 2011] Travis N Ridout and Michael M Franz. *The persuasive power of campaign advertising*. Temple University Press, 2011.

[Rothschild, 1978] Michael L Rothschild. Political advertising: A neglected policy issue in marketing. *Journal of Marketing Research*, 15(1):58–71, 1978.

[Vial and Subramanian, 2019] Daniel Vial and Vijay Subramanian. Local non-Bayesian social learning with stubborn agents. *arXiv*, 2019.

[Wanat, 2020] John Wanat. Political broadcast advertising and primary election voting. *Journal of Broadcasting*, 2020.

[Wilder, 2018] Bryan Wilder. Controlling elections through social influence. *AAMAS*, 2018.

[Zhao *et al.*, 2014] Jiuhua Zhao, Qipeng Liu, and Xiaofan Wang. Competitive dynamics on complex networks. *Scientific Reports*, 4(1), Jul 2014.

[Zuiderveen Borgesius *et al.*, 2018] Frederik Zuiderveen Borgesius, Judith Möller, Sanne Kruikemeier, Ronan Ó Fathaigh, Kristina Irion, Tom Dobber, Balazs Bodo, and Claes H de Vreese. Online political microtargeting: promises and threats for democracy. *Utrecht Law Review*, 14(1):82–96, 2018.