

The Complexity of Strategic Behavior in Primary Elections

Colin Cleveland
King’s College London
London, United Kingdom
colin.cleveland@kcl.ac.uk

Bart de Keijzer
King’s College London
London, United Kingdom
bart.de_keijzer@kcl.ac.uk

Maria Polukarov
King’s College London
London, United Kingdom
maria.polukarov@kcl.ac.uk

ABSTRACT

We study the computational complexity of strategic behaviour in primary elections. Unlike direct voting systems, primaries introduce a multi-stage process in which voters first influence intra-party nominees before a general election determines the final winner. While previous work has evaluated primaries via welfare distortion, we instead examine their game-theoretic properties. We formalise a model of primaries under first-past-the-post with fixed tie-breaking and analyse voters’ strategic behaviour. We show that determining whether a pure Nash equilibrium exists is Σ_2^P -complete, computing a best response is NP-complete, and deciding the existence of subgame-perfect equilibria in sequential primaries is PSPACE-complete. These results reveal that primaries fundamentally increase the computational difficulty of strategic reasoning, situating them as a rich source of complexity-theoretic challenges within computational social choice.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory; Representations of games and their complexity; Problems, reductions and completeness.**

KEYWORDS

Computational Social Choice; Algorithmic Game Theory; Strategic voting; Primary elections; PSPACE-completeness

ACM Reference Format:

Colin Cleveland, Bart de Keijzer, and Maria Polukarov. 2026. The Complexity of Strategic Behavior in Primary Elections. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/RXOJ5992>

1 INTRODUCTION

Democracy is often understood at the level of *electorate* → *government*, yet an equally important question arises within political parties: how should parties select their own candidates? Historically, in many systems party elites simply chose nominees, with primaries serving, at best, as informal opinion polls. The 1968 U.S. Democratic convention, where Hubert Humphrey obtained the nomination without entering any primary, triggered widespread protests and prompted the institutionalisation of primary elections in the United States [8, 22]. Similar reforms have since spread worldwide: in some countries (e.g. the United States, Chile, Argentina)

primaries are regulated by the state [6, 29], while in others (e.g. France, Italy) they remain intra-party procedures run by party organizations [3, 17, 28]. Even in the United Kingdom, where primaries are not formally adopted, major parties increasingly allow supporters to participate in leadership selection [24].

Primary elections appear in several institutional variants—*open*, *closed*, *semi-open*, or hybrid systems such as *poll-weighted primaries*, in which opinion surveys influence results. These hybrids are especially common in newer democratic countries such as Taiwan and South Korea, where weak party membership structures leads parties to integrate public polling into nomination procedures [12, 16, 25].

Primaries are not only decision procedures but also communication channels between parties and their supporters. Yet, they invite strategic behaviour. For instance, in Colorado’s 2022 Republican primary, thousands of Democrats reportedly re-registered as unaffiliated voters to oppose Representative Lauren Boebert, ultimately without affecting the outcome [14]. Such “crossover” behaviour illustrates that additional choice can create incentives for manipulation and unintended losses.

From the party’s perspective, primaries may fail to maximize general-election performance, as polarizing candidates can prevail. From the voter’s perspective, they offer opportunities to reshape competition, but at non-trivial computational and informational costs. These observations motivate a formal analysis of the strategic structure of primaries.

Our contribution. We study primary elections through the lens of computational social choice and algorithmic game theory. Specifically, we:

- Formalize a two-stage election model in which voters’ primary and general-election choices constitute explicit strategies (Section 3.1);
- Characterize the computational complexity of computing best responses, verifying, and deciding the existence of equilibria (Section 4);
- Extend the framework to sequential multi-stage primaries, showing how temporal conditioning increases strategic complexity (Section 5).

Our results show that multi-stage primaries amplify the reasoning burden faced by strategic agents: best-response computation is NP-complete, equilibrium existence rises to Σ_2^P -completeness, and sequential conditioning yields PSPACE-completeness. Beyond elections, these results connect to broader questions in multi-agent reasoning—how decomposing a collective decision into stages expands the logical depth of equilibrium analysis within the polynomial hierarchy.

Complexity overview. We use the complexity classes NP, Σ_2^P , and PSPACE. NP consists of decision problems of the form $\exists x \varphi(x)$ with φ computable in polynomial time. The class Σ_2^P consists of



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/RXOJ5992>

problems of the form $\exists x \forall y \varphi(x, y)$, with φ computable in polynomial time. PSPACE consists of decision problems decidable in polynomial space; equivalently, it consists of problems expressible as $Q_1 x_1 \cdots Q_p(n) x_{p(n)} \varphi(x_1, \dots, x_{p(n)})$ with polynomially many alternating quantifiers, as in quantified Boolean formulas (QBF). Increasing the number of electoral stages corresponds to increasing quantifier alternation depth.

2 RELATED WORK

Most research in computational social choice has focused on *direct elections*, where all candidates compete simultaneously. Borodin et al. [4] provide the first formal analysis of *primary elections* in this domain, comparing multi-stage nomination processes with direct elections under various rules. Their results demonstrate that primaries can diverge substantially from direct elections both in theory and in empirical distortion analyses. Our work extends this direction by examining the *strategic and computational* dimensions of voter behaviour in primaries rather than their welfare implications.

A complementary line of work studies *strategic candidacy*, where candidates decide whether to enter or withdraw. This literature, initiated by Dutta et al. [9, 10] and later developed by Brill and Conitzer [5] and Polukarov et al. [23], characterises equilibria over candidate participation decisions. Related work on *electoral control*—for example, adding or deleting candidates—originates with Bartholdi et al. [2]. Harrenstein et al. [11] unify strategic candidacy and spatial competition in a Hotelling–Downs framework where parties nominate candidates, showing complexity bounds for equilibrium existence and links to facility-location and Voronoi games. More recently, Schlotter and Cechlárová [27] analyse the *Possible President* problem under Condorcet-consistent rules, exploring computational and parameterized complexity when nominees depend on rival parties’ choices. In contrast to these *candidate-centric* perspectives, our focus is on the strategic computation of *voters* within primary systems.

Another relevant stream concerns *iterative voting*, where voters repeatedly adjust their ballots in response to others. Meir et al. [20] initiated the study of convergence to equilibrium under plurality, with subsequent extensions by Reijngoud and Endriss [26], Obraztsova et al. [21], and Lev and Rosenschein [18]. Surveys by Meir [19] and Airiau et al. [1] synthesise this literature, which typically assumes a single-stage election and investigates behavioural convergence. Our analysis differs in treating voters as computationally bounded agents: we show that computing a best response is NP-complete once primaries precede the general election, and that equilibrium existence climbs to Σ_2^P - or PSPACE-completeness as sequential conditioning is introduced. This places primary elections within a broader conversation about the complexity of equilibrium reasoning in multi-agent systems.

Finally, beyond computational studies, the design and consequences of primary elections have been widely examined in the social sciences. Kenig [13] surveys candidate-selection methods across parliamentary democracies, while Cohen et al. [8] and Sides and Vavreck [30] analyse elite influence and representativeness in U.S. primaries. These empirical insights complement the computational models developed in AI and game theory, together offering a

more complete view of how institutional design shapes strategic reasoning in multi-agent collectives.

3 PRELIMINARIES

For $n \in \mathbb{N}$, let $[n] = \{1, 2, \dots, n\}$. Let $V = [n]$ be the *set of voters* and $A = [m]$ the *set of politicians*.¹ We write v_i for voter $i \in V$ and a_j for politician $j \in A$.

Let $\Pi = [p]$ denote the *set of parties*. Each party $k \in \Pi$ has a disjoint set of affiliated politicians $A_k \subseteq A$, so that $A_i \cap A_j = \emptyset$ whenever $i \neq j$. Hence every politician belongs to exactly one party.

After all primaries conclude, each party advances a single nominee to the *general election (GE)*. The *set of finalists* is thus a profile

$$c = (c_1, \dots, c_p) \in C := A_1 \times A_2 \times \cdots \times A_p,$$

where $c_k \in A_k$ is party k ’s nominee.

Each voter i has a *cardinal utility function* $u_i : A \rightarrow \mathbb{Q}$, where $u_i(a_j)$ is the utility from politician a_j winning the GE. Because A is finite, we identify u_i with its vector representation

$$(u_{i,1}, u_{i,2}, \dots, u_{i,m}), \quad u_{i,j} = u_i(a_j).$$

The *utility profile* $[u]$ collects all voters’ preferences and fully determines payoffs for every possible GE winner.

3.1 Voting Framework

We model the election as a multiple-stage process: party primaries followed by the GE. Each party $k \in \Pi$ holds its primary in a fixed order $(1, 2, \dots, p)$. At each stage, voters cast their primary ballots simultaneously, and a unique nominee is declared immediately. Unless otherwise stated, both primaries and the GE use the *first-past-the-post (FPTP)* rule: each voter casts one ballot, and the candidate with the most votes wins. Ties are broken according to a fixed, predetermined ordering of candidates (a *fixed-tie breaking list*).

Voter strategies. A voter’s behaviour is described by a strategy

$$s_i = \{b_i, g_i\},$$

where:

- $b_i = (b_{i,1}, \dots, b_{i,p})$ is the **primary ballot vector**. Each coordinate $b_{i,k} \in A_k \cup \{0\}$ records either a vote for a politician in A_k or abstention (0).
- $g_i : C \rightarrow A \cup \{0\}$ is the **general-election decision function**. For each finalist profile $c \in C$, it specifies the GE choice (or abstention) of voter i .

This reduced form captures only the ex-post behaviour relevant to electoral outcomes, rather than a full extensive-form plan. A fully extensive description would need to specify actions for every possible primary path and finalist set—requiring $\Theta(m^{pn})$ space (see supplement material in [7]). The reduced form therefore trades completeness for analytical tractability.

Registration. Institutions vary in whether participation requires prior *party registration*. In a *closed* primary, each voter affiliates with one party before the season begins; if i registers with party k , only $b_{i,k}$ may be non-zero. In an *open* primary, affiliation is not fixed in advance: at each stage k , voter i decides whether to cast $b_{i,k}$, subject to the constraint that no voter participates twice in

¹In a direct election, A coincides with the candidate set.

the same contest. Thus b_i has the same domain in both regimes but different interpretations—static commitment versus contingent choice. Administrative details (deadlines, membership rolls, etc.) are abstracted away.

Participation regimes. A second institutional choice concerns how many primaries a voter may enter:

- **Single-primary participation:** At most one coordinate of b_i is non-zero.
- **Multiple-primary participation:** Each coordinate may be filled independently, $b_i \in \prod_{k=1}^p (A_k \cup \{0\})$.

We adopt the multiple-participation regime as the general case; it subsumes the single-primary setting.

Access model for g_i . The general-election component g_i may be exponentially large: since $|C| = \prod_{k=1}^p |A_k| = \Theta(m^p)$, an explicit table for g_i would require exponential space. We therefore assume a standard RAM model in which g_i is queried as an oracle with cost $\tau = \tau(m, p)$. Our reductions use gadgets admitting constant-time queries ($\tau = O(1)$). We also allow *succinct encodings*, where $g_i(c)$ is computed on demand. A canonical example is the *listed-preference rule* $g_i(\cdot; L)$, where L is an ordered list over $A \cup \{0\}$: given c , the voter selects the first element of L appearing in c (or abstains if 0 precedes all finalists). This encoding uses only $O(m)$ space and supports $O(1)$ queries under the RAM model.

Collecting all voters’ strategies yields the profile $s = (s_1, \dots, s_n)$. The institutional regime then determines the *final winner* $\mathcal{W}(s) \in A$. Throughout, voters are viewed as polynomial-time agents operating under these oracle assumptions.

3.2 Voter Cost and Utility

Casting a GE ballot incurs a participation cost $\kappa_i > 0$, and participation in party k ’s primary incurs $\kappa_i^k > 0$. Given a strategy s_i and the strategies of others $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$, the *total utility* of voter v_i equals the utility derived from the eventual GE winner minus the participation costs incurred in the GE and any primaries. Formally,

$$U_i(s_i, s_{-i}) = u_i(\mathcal{W}(s_i, s_{-i})) - \mathbf{1}[g_i \neq 0] \kappa_i - \sum_{k \in [p]} \mathbf{1}[b_{i,k} \neq 0] \kappa_i^k,$$

where $\mathcal{W}(s_i, s_{-i})$ denotes the GE winner. Voters act to maximise their expected U_i given beliefs about others’ strategies.

Toy Example: Two Parties with Single-Primary Participation. Consider $\Pi = \{1, 2\}$ with $A_1 = \{a, b\}$ and $A_2 = \{c, d\}$, and assume alphabetical tie-breaking. Politicians are positioned at $[10, 6, 4, 0]$ on a line, with four voters located at $[0, 4, 6, 10]$ and no participation costs. Utilities follow $u_i(a_j) = -|x_i - y_j|$ with x_i and y_j the location of voter i and candidate j ². Figure 1 illustrates the setting.

Each voter’s strategy is $s_i = \{b_i, g_i\}$, where g_i is expressed as a ranked list:

$$\begin{aligned} s_1 &= \{[0, d], [d, c, b, a]\}, & s_2 &= \{[0, c], [c, b, d, a]\}, \\ s_3 &= \{[b, 0], [b, c, a, d]\}, & s_4 &= \{[a, 0], [a, b, c, d]\}. \end{aligned}$$

The finalists are $C = (a, c)$, and the GE winner is c , supported by v_1, v_2, v_3 . Hence $[u_1(c), u_2(c), u_3(c), u_4(c)] = [-4, 0, -2, -6]$.

²Utilities may be shifted by an additive constant without affecting comparisons




Figure 1: Politician and voter positions in the toy example.

Two strategic votings improve v_4 ’s utility:

- *Hostile crossover.* v_4 changes b_4 in $s_4 = (b_4, g_4)$ to $[0, d]$. Then d defeats c in Party 2’s primary (2:1) and b defeats a in Party 1’s (1:0), giving $C^* = (b, d)$. In the GE, b beats d (3:1), yielding $U_4(s_4^*, s_{-4}) = u_4(b) = -4 > -6$.
- *Honest coordination.* v_4 instead votes $b_4 = [b, 0]$, giving $C^* = (b, c)$. The GE is a 2:2 tie resolved in favour of b , again improving v_4 ’s utility to -4 .

These examples show that sincere primary voting need not maximise individual utility. Other manipulations, such as *honest crossover*, where a voter supports their favourite politician in a rival party’s primary—may also shift the finalist set toward a more desirable overall outcome.

4 COMPUTATIONAL QUESTIONS OVER PRIMARIES

The framework above gives rise to two central decision problems.

Question 1 (Best Response).

Instance: Voter set V , utility profile $[u]$, politician sets A_1, \dots, A_p , tie-breaking list \mathcal{L} , and strategy profile s .

Task: For a given voter i , compute a strategy s_i^* that maximizes $U_i(s_i^*, s_{-i})$.

Question 2 (Equilibrium).

Instance: V , $[u]$, and politician sets A_1, \dots, A_p .

Task: Decide whether there exists a strategy profile s^* such that, for every voter i and every alternative strategy s'_i ,

$$U_i(s^*) \geq U_i(s'_i, s_{-i}^*).$$

For clarity, we begin with the simplest timing structure in which all parties hold their primaries simultaneously. This reduces the election to two stages: primaries followed by the GE, while retaining the essential strategic components. In Section 5, we relax this assumption and consider the sequential case, where primaries unfold in turn across parties.

4.1 Best-Response Search Space and Running Time

A naïve way to compute the best response of voter i is to enumerate all possible strategies. Each strategy consists of (i) a primary ballot for every party and (ii) a general-election (GE) decision rule $g_i : C \rightarrow A \cup \{0\}$ that *selects among finalists*. This yields the upper bound

$$\left(\prod_{k=1}^p (1 + |A_k|) \right) \cdot (|A| + 1)^{|C|},$$

since each party ballot offers $|A_k|+1$ options (including abstention), and a GE rule must specify an action for each finalist profile, where $|C| = \prod_k m_k = O(m^p)$.

However, when computing a best response, we never need to enumerate the entire g_i . Once the finalists $c \in C$ are known, voter i simply chooses greedily among $\{c_1, \dots, c_p, 0\}$ —the available finalists and abstention—thereby reducing the GE decision to $m+1$ evaluations. Hence the dominant factor in the search space comes from enumerating primary ballots³.

To analyse the running time, we decompose the computation into two components.

Baseline procedure.

- (1) *Pre-computation.* For each party k , compute primary tallies from s_{-i} , requiring $O(np)$ time by scanning all other voters once.
- (2) *Enumeration of i 's primary actions.* For every combination $b_i \in \prod_k (A_k \cup \{0\})$:
 - (a) Update tallies locally to reflect b_i , determine each party's nominee under FPTP (with fixed tie-breaking), and obtain the finalist tuple $c_x \in C$. This takes $O(mp)$ per b_i , using running arg-maxes for each party.
 - (b) *Evaluate the GE outcome.* For each $i' \neq i$, query $g_{i'}(c_x)$ (time τ per query) and increment that choice's tally. This step costs $O(n\tau + p)$ per b_i (p accounts for the final arg-max over GE tallies). Voter i then greedily picks her GE action from $\{c_{x,1}, \dots, c_{x,p}, 0\}$ and computes U_i .

Aggregating these terms yields the following bound.

PROPOSITION 1 (BEST-RESPONSE RUNNING TIME). *Let τ be the time to evaluate a single voter's GE decision $g_i(c)$ for a given finalist profile c . For a fixed voter i , the enumeration algorithm over $b_i \in \prod_{k=1}^p (A_k \cup \{0\})$ runs in*

$$O(np + (m+1)^p (mp + n\tau))$$

For fixed p and constant τ ,

$$T_{BR} = O(n + m^p (mp + n)) = O(m^{p+1} + n m^p),$$

so best response is computable in time polynomial in n and m .

Explicit vs. succinct GE behaviour. The size of the finalist set is $|C| = \prod_k |A_k| = O(m^p)$. Under an *explicit* representation, each $g_i : C \rightarrow A \cup \{0\}$ is stored as a table of $\Theta(m^p)$ entries with $O(1)$ query time, implying that the input already encodes $\Theta(nm^p)$ information. Since $T_{BR} = O((nm^p)^2)$, the best response remains polynomial in input length even if p grows with m .

By contrast, if g_j is given *succinctly* (e.g., the listed-preference presentation), the input size drops to $O(nm)$. The exponential factor $(m+1)^p$ in the action space then re-emerges, making complexity sensitive to the representation gap.

REMARK 1 (POLYNOMIAL TIME VERSUS INPUT SIZE). *For fixed p and succinctly evaluable GE rules, best response is polynomial in n, m . With explicit g , the input size is already exponential in p , so best response remains polynomial in input length. Only when g is succinct do we encounter a genuine complexity increase.*

³When the number of candidates, parties, or primary rounds is small, the question of optimal strategy may be encoded as Presburger arithmetic formulas and thus, may not need enumeration [15]

Special regimes. Two simplified variants are worth noting:

- (1) *Single-Participation primaries.* If each voter can participate in at most one primary, the action space reduces to $O(m)$, and best response runs in $O(m^2 + mn)$.
- (2) *Fixed GE behaviour.* If a voter's GE choice is predetermined (e.g., always supporting a fixed party's nominee), the GE stage adds no computational burden, and the runtime remains as stated in the proposition.

Towards NP-completeness. When GE behaviours are given succinctly, a natural question arises: does a polynomial-time algorithm for best response exist? We show that the corresponding decision version is NP-complete.

DEFINITION 1 (BEST-RESPONSE-AT-LEAST- U (BR- $\geq U$)). *Under FPTP rule, fix $m = |A|$, $p = |\Pi|$, and a profile s_{-i} of all voters except i^* . The decision problem BR- $\geq U$ is defined as follows:*

Instance: An election with parties $\Pi = \{1, \dots, p\}$ and politician sets A_k , a fixed tie-breaking list \mathcal{L} , other voters' strategies s_{-i^*} , and succinct GE functions $g_j : C \rightarrow A \cup \{0\}$ (query-able in $O(1)$). A utility threshold $U \in \mathbb{Q}$.
Question: Does there exist a strategy s_{i^*} such that $U_{i^*}(s_{i^*}, s_{-i^*}) \geq U$?

THEOREM 2 (NP-COMPLETENESS OF BR- $\geq U$). *When the number of parties p is part of the input and GE behaviours are given succinctly, BR- $\geq U$ is NP-complete with fixed tie-breaking, even with zero participation costs.*

PROOF. Membership. A certificate consists of i^* 's strategy. Verification simulates the primaries, evaluates the finalists, queries each $g_j(c)$, tallies GE votes, and computes U_{i^*} —all in polynomial time under the RAM model.

Hardness (from 3-SAT). Given a Boolean formula $\Phi(x_1, \dots, x_p)$, we construct an election instance where the primary stage encodes variable assignments and the GE tallies evaluate clauses.

Politicians (anchors as a fixed party). For each variable x_k create a party with politicians $\{x_k, \neg x_k\}$. Add two *dummy* parties $A_\top = \{x_\top\}$ and $A_\perp = \{x_\perp\}$ that skip meaningful primaries and always advance their unique nominees to the GE (they contribute only a constant factor to $|C|$ and do not affect the analysis). Thus each finalist profile has the form $c = (x_\top, x_\perp, c_1, \dots, c_p)$.

Primaries. Only i^* votes in the variable parties' primaries. Deterministic tie-breaking makes her ballot pivotal, so she freely selects nominees $\{c_1, \dots, c_p\}$ encoding an assignment σ .

GE voters. Two fixed blocs of size $Q = q + 1$ vote for x_\top and x_\perp , respectively. For each clause $C_t = (\ell_{t1} \vee \ell_{t2} \vee \ell_{t3})$, add one clause voter with list-form GE rule

$$g_t(c; L_t), \quad L_t = [\ell_{t1}, \ell_{t2}, \ell_{t3}, x_\perp, 0],$$

which selects the first finalist in L_t (succinct, $O(1)$ query). The distinguished voter i^* votes for x_\top .

Correctness and size. If Φ is satisfiable, i^* picks nominees according to a satisfying assignment, every clause voter supports some literal finalist, and x_\top beats x_\perp by one vote ($q+2$ vs. $q+1$). If Φ is unsatisfiable, some clause is falsified and its voter supports x_\perp , which

then ties or beats x_{\top} ; fixed tie-breaking favours x_{\perp} . The construction uses $O(p)$ politicians and $O(q)$ voters, hence is polynomial.

Thus $BR \geq U$ is NP-hard; combined with membership, it is NP-complete. \square

With $Q = q + 1$ fixed supporters each for x_{\top} and x_{\perp} , and only q clause voters in total, any non-anchor finalist can obtain at most q votes—strictly fewer than $q + 2$. Hence the final winner is always either x_{\top} or x_{\perp} .

The reduction can be further simplified by removing the fixed blocs. Define each clause voter v_t so that if at least one literal ℓ_{ti} of clause C_t appears among the finalists, v_t abstains; otherwise v_t votes for x_{\perp} . This version still uses $O(1)$ -query rules and needs only one fixed supporter for x_{\top} , but it is less behaviourally natural (a kind of “protest abstention”), so we retain the bloc-based construction.

Thus, in our formulation, best-response computation remains tractable only when numbers of primaries or explicit GE rules are fixed.

4.2 Pure Equilibrium: Verification, Existence, and Running Time

As in § 4.1, we assume GE behaviour is given succinctly, so each voter’s strategy is $s_i = (b_{i,1}, \dots, b_{i,p}, g_i)$ with $b_{i,k} \in A_k \cup \{0\}$ and $g_i : C \rightarrow A \cup \{0\}$ queryable in $O(1)$ time.

We study two problems: NE-VERIFY—checking whether a given profile is a pure Nash equilibrium — and NE-EXIST — deciding whether one exists.

DEFINITION 2 (EQUILIBRIUM VERIFICATION (NE-VERIFY)). *Under FPTP, given parties $\Pi = \{1, \dots, p\}$ with politician sets A_k ($|A_k| \geq 1$), a fixed tie-breaking list \mathcal{L} , utilities $[u]$, succinct GE rules $(g_j)_j$, and a strategy profile $s = (s_i)_{i \in V}$, decide whether for all voters i and all alternative strategies s'_i ,*

$$U_i(s) \geq U_i(s'_i, s_{-i}).$$

Trivial equilibria. Before analysing equilibrium existence, we exclude outcome-constant profiles where everyone unanimously supports the same politician across stages; the winner is fixed and no unilateral deviation can change it. We therefore focus on a small subset of strategic voters:

$$V = V_{\text{env}} \cup V_{\text{strat}}.$$

where V_{env} are fixed environment voters and V_{strat} are those whose decisions affect the GE outcome.

DEFINITION 3 (EQUILIBRIUM EXISTENCE (NE-EXIST)). *Given the same model, decide whether there exists a profile $s^* = (s_i^*)_{i \in V}$ such that, for every $i \in V_{\text{strat}}$ and every unilateral deviation s'_i ,*

$$U_i(s^*) \geq U_i(s'_i, s_{-i}^*).$$

4.2.1 Verification. The link between NE-VERIFY and $BR \geq U$ is direct: $BR \geq U$ asks whether there exists a profitable deviation, whereas NE-VERIFY asks whether no such deviation exists for any voter. A certificate for “not NE” is a pair (i, s'_i) ; verifying it recomputes the finalists under (s'_i, s_{-i}) , evaluates i ’s new GE action $y = g'_i(c)$ (a constant choice suffices), queries all $g_j(c)$ for $j \neq i$, tallies, and compares U_i . All these operations are polynomial, so “not NE” lies in NP and hence NE-VERIFY is in coNP.

THEOREM 3 (coNP-COMPLETENESS OF NE-VERIFY). *Given a strategy profile s , deciding whether s is a pure Nash equilibrium is coNP-complete under FPTP with fixed tie-breaking, even with zero participation costs and succinct GE evaluation.*

PROOF. Membership. A witness for “not NE” is (i, s'_i) ; verifying it takes polynomial time, so NE-VERIFY \in coNP.

Hardness. Reduce from UNSAT using the same construction as in Theorem 2. Fix s to select all “false” nominees. If Φ is satisfiable, a pivotal voter can deviate to a satisfying assignment and improve—so s is not an NE. If Φ is unsatisfiable, no unilateral deviation can yield a satisfying assignment, so no voter can improve and s is an NE. Thus $\Phi \in \text{UNSAT} \iff s \in \text{NE-VERIFY}$, giving coNP-hardness. \square

4.2.2 Existence. NE-EXIST has the quantifier form

$$\exists s \forall i \neg(\exists s'_i \text{ profitable}),$$

where the inner predicate (“is there a profitable deviation?”) is in NP; hence NE-EXIST $\in \Sigma_2^P$. A standard gadget-based reduction from 2-QBF (i.e., deciding $\exists X \forall Y \varphi(X, Y)$)—using closed primaries so that v_{odd} controls X and v_{even} controls Y , and the same clause/rescue voter gadgets with succinct GE functions—yields Σ_2^P -hardness. Before the formal proof, we show why pure equilibria may fail to exist.

EXAMPLE 4.1. No pure NE with closed primaries. *Consider FPTP with two parties, $A_{\text{odd}} = \{a_1, a_3\}$ and $A_{\text{even}} = \{a_2, a_4\}$. Four environment voters v_1, \dots, v_4 have cyclic GE lists (a_1, a_2, a_3, a_4) , (a_2, a_3, a_4, a_1) , (a_3, a_4, a_1, a_2) , (a_4, a_1, a_2, a_3) , and two primary voters:*

$$v_{\text{odd}} : a_1 \succ a_3 \succ a_2 \sim a_4, \quad v_{\text{even}} : a_2 \sim a_4 \succ a_1 \sim a_3.$$

Say no cross-party participation ($\kappa_{\text{even}}^{\text{odd}} = \kappa_{\text{odd}}^{\text{even}} = \infty$), so each primary has a single strategic voter. The induced 2×2 game of nominees is

	a_2	a_4
a_1	odd wins	even wins
a_3	even wins	odd wins

which is the matching-pennies payoff structure: no pure Nash equilibrium exists (the unique mixed equilibrium is uniform).

Even without participation costs, pure equilibria may fail to exist: cyclic tie-breaking across primaries ensures that some voter can always deviate profitably (see supplement material in [7] for the full construction).

PROPOSITION 4. *Under FPTP with fixed tie-breaking, primary elections may admit no pure-strategy Nash equilibrium.*

PROOF. Example 4.1 forms a deviation cycle, so at least one voter can always improve unilaterally; hence no pure NE exists. We also show an example in the supplement material [7] forms a deviation cycle without participation cost. \square

THEOREM 5 (Σ_2^P -COMPLETENESS OF NE EXISTENCE). *Under FPTP, deciding whether a pure Nash equilibrium exists is Σ_2^P -complete with fixed tie-breaking, even when GE behaviour is finalist-based (list actions).*

PROOF. Membership. The problem has the form $\exists s \forall i \neg(\exists s'_i \text{ profitable})$. Given s and (i, s'_i) we can recompute finalists, evaluate

list-based GE actions in $O(1)$ per voter, tally, and compare utilities in polynomial time, hence $\text{NE-EXIST} \in \Sigma_2^P$.

Hardness (from 2-QBF). Let $\Phi(X, Y) = \exists X \forall Y : \varphi(X, Y)$ be a quantified 3-CNF. We build a primary-election game G_Φ (FPTP, fixed tie-breaking) that admits a pure NE iff Φ is true.

Parties and players. Each variable $z \in X \cup Y$ corresponds to a party with nominees $\{z, \neg z\}$. Add two top-level parties $A_{\text{odd}} = \{a_1, a_3\}$ and $A_{\text{even}} = \{a_2, a_4\}$ with tie-breaking $a_1 \succ a_2 \succ a_3 \succ a_4$. Two strategic voters act: v_{odd} controls all X -parties and chooses $a_o \in \{a_1, a_3\}$, v_{even} controls all Y -parties and $a_e \in \{a_2, a_4\}$. Both care only about whether their top-level party wins the GE.

Background and clauses. Give A_{odd} a baseline margin +1 (e.g., $r+1$ fixed odd supporters vs. r even). For each clause $C_t = (\ell_{t1} \vee \ell_{t2} \vee \ell_{t3})$ add four clause voters v_{t1}, \dots, v_{t4} with lists

$$L_{ti} = (\ell_{t1}, \ell_{t2}, \ell_{t3}, a_i, a_{i+1}, a_{i+2}, a_{i+3}) \pmod{4}.$$

Each votes for the first finalist in L_{ti} . If C_t is satisfied, all four vote for literals; if falsified, they fallback cyclically to $\{a_1, a_2, a_3, a_4\}$, producing a 3:1 split that favours one top-level party depending on (a_o, a_e) (as in Example 4.1).

The construction uses $O(|X|+|Y|+r)$ politicians and voters, all with succinct GE rules, so its size is polynomial.

(\Rightarrow) If Φ is true, let x^* satisfy $\forall Y : \varphi(x^*, Y)$. In profile $S^* = (x^*, a_1; y, a_e)$ for any (y, a_e) , every clause is satisfied, so all clause voters support literals and the baseline +1 margin lets A_{odd} win. Neither v_{even} (who always loses) nor v_{odd} (already winning) can improve, so S^* is a pure NE.

(\Leftarrow) If Φ is false, then for every x there exists y falsifying at least one clause. Each falsified clause triggers a 3:1 fallback favouring one party; switching the top-level nominee (a_o or a_e) reverses that advantage. Hence, after any $(x, a_o; y, a_e)$ the loser can deviate profitably, and deviations alternate indefinitely—no pure NE exists.

Thus G_Φ has a pure NE iff Φ is true, establishing Σ_2^P -hardness. Combined with membership, NE-EXIST is Σ_2^P -complete. \square

Takeaway. Equilibrium analysis in primary elections spans the second level of the polynomial hierarchy. Verification of a proposed profile lies in coNP , as disproving stability only requires exhibiting a profitable deviation; deciding existence introduces an additional existential quantifier over profiles, yielding Σ_2^P -completeness. While best responses are tractable in isolation, the search for mutually stable outcomes is theoretically intractable.

5 COMPLEXITY UNDER SEQUENTIAL PRIMARIES: COMMITMENT VS. CONDITIONING

Our Σ_2^P result already assumes that voters' full strategies are evaluated *ex post*: each voter commits in advance to a complete plan across all stages, and equilibrium verification checks whether any unilateral deviation is profitable.

The picture changes once primaries hold sequentially, and later-stage voters may *condition* their actions on earlier outcomes, as in subgame-perfect play. In Theorem 5, we used \exists for the designer-controlled variables (v_{odd}) and \forall for the adversarial ones (v_{even}) to obtain the Σ_2^P -complete result. Extending this reasoning, each sequential stage introduces an additional quantifier block of the same

type. Thus a process with p stages behaves like a $\Sigma_{O(p)}^P$ problem, and when p is unbounded, the alternation depth is unbounded, suggesting PSPACE-completeness by analogy with quantified Boolean formulas (QBF). We formalize this correspondence below.

Sequential primaries (model recap). There are p parties indexed by $k \in [p]$, each with a nonempty politician set A_k , and let $C := A_1 \times \dots \times A_p$ denote the set of nominee tuples. Primaries occur in a fixed order $1, 2, \dots, p$. At stage $t \in [p]$, a designated subset of strategic voters $V_{\text{strat}} \subseteq V$ acts, while all others follow fixed behaviours. Within each stage, moves are simultaneous. After resolving the party- k primary (under FPTP with fixed tie-breaking), its nominee $c_k \in A_k$ is fixed, and the process proceeds to stage $k+1$. After all stages, the general-election (GE) outcome is computed from $c = (c_1, \dots, c_p) \in C$ via a polynomial-time evaluable rule $\mathcal{W} : C \rightarrow A$ (e.g., FPTP with fixed tie-breaking).

Information and histories. Let $C_{<k} := A_1 \times \dots \times A_{k-1}$ denote the space of nominee prefixes. Before acting at stage k , all players observe the realized prefix $h_k \in C_{<k}$, i.e., the sequence of winners (c_1, \dots, c_{k-1}) .⁴ A ballot for party k is an element of $A_k \cup \{0\}$. For voter i , a (behavioural) strategy at party k is a function

$$f_{i,k}^p : C_{<k} \rightarrow A_k \cup \{0\}.$$

Hence a *history-dependent strategy* for voter i is the tuple

$$s_i = (f_{i,1}^p, \dots, f_{i,p}^p, g_i),$$

where g_i represents the GE behaviour (e.g., a map $g_i : C \rightarrow A \cup \{0\}$ or a ranked-list policy, depending on the GE model). A full strategy profile is $s = (s_i)_{i \in V}$.

Given s , the induced nominee tuple $c(s) \in C$ is generated by sequentially applying FPTP with fixed tie-breaking at each stage using ballots $\{f_{i,k}^p(h_k)\}_{i \in V_k}$, where h_k is the realized prefix. The GE winner is $\mathcal{W}(c(s))$.




Figure 2: A partial game tree of two voters in a sequential primary. Nodes labelled “1’s” and “2’s | a_j ” represent party 1’s primary and party 2’s primary conditional on candidate a_j winning party 1’s primary, respectively. At each node, both voters simultaneously cast ballots (a_i, a_i') . Leaf nodes $c = \{c_1, c_2\}$ denote the general-election finalists.

The tree in Figure 2 illustrates the extensive form of a two-voter sequential primary. Stage 1 ($1's$) and stage 2 ($2's$) correspond to the primaries of party 1 and party 2, respectively, where nodes labelled “2’s | a_j ” indicate that candidate a_j won party 1’s primary. At each

⁴If only partial or noisy information is available (e.g., observing only c_{k-1}), replace $C_{<k}$ by the corresponding information set \mathcal{H}_k .

node, both voters simultaneously cast ballots $(a_i, a_{i'})$, and each branch represents a possible nominee outcome leading to the GE.

Note that the branching factor at each node of party k 's primary is $(m_k + 1)^n$, reflecting all possible combinations of n voters' simultaneous ballots over m_k politicians (including abstention). Each node can thus be indexed by an n -tuple $(a_{k1}, a_{k2}, \dots, a_{kn})$ representing the voters' ballots. Although the number of nodes grows exponentially in n , the full tree can still be traversed using polynomial space.

Strategic scope and representation. We assume that only a small subset of voters are strategic, while all others are *environmental voters* whose behaviour remains fixed as (b_i, g_i) defined in § 3.1. If each function $f_{i,k}^p$ were represented explicitly, its table size would grow as $\Theta(\prod_{\ell < k} |A_\ell|)$, i.e., $O(m^p)$ per voter, making even the input exponentially large in the number of parties. To avoid this blow-up, we retain environmental voters in their simple form and restrict sequential reasoning to the designated strategic subset.

Instead of the best responses for a single voter,⁵ we focus on *dominant strategies and equilibria among the strategic voters* to ensure that the input size remains polynomially bounded.

Equilibrium notion. We study pure subgame-perfect equilibria (SPNE) in the induced extensive-form game with observed prefixes; that is, strategy profiles s that form a Nash equilibrium in every subgame rooted at any history $h_k \in C_{<k}$.

DEFINITION 4 (SEQUENTIAL DOMINANT STRATEGY AT LEAST U (SEQDS- $\geq U$)). *Given an instance as above, a designated player i^* , and a threshold $U \in \mathbb{Q}$, decide whether i^* has a strategy s_{i^*} that guarantees $U_{i^*}(s_{i^*}, s_{-i^*}) \geq U$ for all admissible multi-stage strategies s_{-i^*} of her opponents.*

THEOREM 6 (SEQDS- $\geq U$ IS PSPACE-COMPLETE). *Under FPTP with fixed tie-breaking and polynomial-time GE evaluation, if the number of stages p is part of the input (polynomially bounded horizon), then SEQDS- $\geq U$ is PSPACE-complete. For constant p , the problem lies in $\Sigma_{O(p)}^P$.*

PROOF. Membership. The extensive form has depth p . At a stage k , the joint action set has size $(m_k)^n$ (cf. Fig. 2 and the discussion: each node can be indexed by the n -tuple of ballots (a_{k1}, \dots, a_{kn})), so the branching factor may be exponential in n , but depth-first evaluation needs only polynomial space: we store one path of length p plus polynomial-time computable summaries at each node. Utilities at leaves are obtained by computing the GE outcome and checking $U_{i^*} \geq U$, both in polynomial time by assumption. Hence the alternating (max–min) evaluation of the game tree for the dominance predicate (“does i^* guarantee utility $\geq U$ against all opponents?”) is decidable in PSPACE via standard DFS with backtracking.

Hardness (from QBF). Let $\Phi = Q_1 X_1 Q_2 X_2 \dots Q_p X_p : \varphi(X)$ be a quantified 3-CNF formula. We construct a sequential-primaries game with p stages encoding Φ .

Politicians. For each variable X_k , create party k with nominees $\{x_k, \neg x_k\}$, and add two dummy parties $A_\top = \{x_\top\}$ and $A_\perp = \{x_\perp\}$.

⁵If all voters followed the simpler (b_i, g_i) form, the game effectively collapses to the simultaneous case. Otherwise, one would need to impose additional succinctness restrictions on $f_{i,k}^p$, which would be ad hoc and fall outside our present scope.

Control. If $Q_k = \exists$, stage k is controlled by i^* (tie-breaking ensures her ballot decides); if $Q_k = \forall$, by an adversary v_a . Thus the nominee tuple $c = (c_1, \dots, c_p)$ encodes a full assignment to X .

GE voters. Two fixed blocs B_\top, B_\perp (each of size $q+1$) always vote x_\top and x_\perp , respectively. For each clause $C_t = (\ell_{t1} \vee \ell_{t2} \vee \ell_{t3})$, add a clause voter with policy

$$g_t(c) = \text{first of } [\ell_{t1}, \ell_{t2}, \ell_{t3}, x_\perp, 0] \text{ appearing among finalists.}$$

If C_t is falsified under c , this voter supports x_\perp ; otherwise, a satisfied literal gains her vote.

Utilities. Let $U := u_{i^*}(x_\top)$, with $u_{i^*}(x_\top) > u_{i^*}(x_\perp) > u_{i^*}$ (others), so $U_{i^*} \geq U$ iff x_\top wins the GE.

Tally. x_\top has baseline $q+2$ votes (B_\top plus i^*); x_\perp starts with $q+1$ (B_\perp) and gains one per falsified clause. With tie-breaking in favour of x_\perp , x_\top wins iff all clauses are satisfied.

Correctness. At each stage: - if $Q_k = \exists$, i^* chooses $c_k \in \{x_k, \neg x_k\}$ (OR-node); - if $Q_k = \forall$, the adversary chooses (AND-node). At the leaf, x_\top wins iff φ is satisfied by the induced assignment. Thus i^* can guarantee $U_{i^*} \geq U$ iff Φ is true.

The construction uses $O(p+q)$ voters and politicians and is polynomial in $|\Phi|$. Hence SEQDS- $\geq U$ is PSPACE-hard. Combined with membership, it is PSPACE-complete. For fixed p , the alternation depth is $O(p)$, yielding the stated $\Sigma_{O(p)}^P$ bound. \square

The sequential dominance framework already captures the essential source of hardness: each stage introduces one quantifier alternation, mirroring the structure of a quantified Boolean formula. Equilibrium reasoning, however, does not fundamentally simplify this picture. In a turn-based extensive form, verifying subgame-perfection simply requires backward induction over the same alternating stages, while hardness stems from the same quantifier nesting. Hence the complexity of equilibrium existence grows in exactly the same way as dominance verification, culminating in the following theorem.

We reuse the *environment/clause voters with cyclic fallback lists* and the two top-level parties $A_{\text{odd}} = \{a_1, a_3\}$, $A_{\text{even}} = \{a_2, a_4\}$ with fixed tie-breaking $a_1 \succ a_2 \succ a_3 \succ a_4$ exactly as in Theorem 5. No new voter types are introduced; we only stage the same construction.

THEOREM 7 (SPNE EXISTENCE REUSING THE NE GADGET). *Under FPTP with fixed tie-breaking and list-based GE behaviour, deciding whether a pure subgame-perfect equilibrium (SPNE) exists is $\Sigma_{O(p)}^P$ -complete for fixed number of stages p , and PSPACE-complete when p is part of the input.*

PROOF. Membership. The same depth-first evaluation argument (cf. Fig. 2 and Theorem 6) applies. Given a polynomially bounded horizon and polynomial-time evaluable GE outcomes, verifying subgame-perfect optimality along the sequential tree requires only polynomial space. Hence SPNE-EXIST is in PSPACE, and for constant p the induced alternation yields membership in $\Sigma_{O(p)}^P$.

Hardness (from QBF). Let $\Psi = \exists X_1 \forall Y_1 \dots \exists X_k \forall Y_k : \varphi(X, Y)$. Create $p = 2k$ stages in order $X_1, Y_1, \dots, X_k, Y_k$. For each variable $z \in X \cup Y$, make a party with nominees $\{z, \neg z\}$. At stage X_i , v_{odd} chooses x_i or $\neg x_i$; at stage Y_i , v_{even} chooses y_i or $\neg y_i$. Keep the same background as in Theorem 5: (i) A_{odd} has baseline GE margin $+1$; (ii) for each clause $C_t = (\ell_{t1} \vee \ell_{t2} \vee \ell_{t3})$, include clause voters

v_{t1}, \dots, v_{t4} with lists $L_{ti} = (\ell_{t1}, \ell_{t2}, \ell_{t3}, a_i, a_{i+1}, a_{i+2}, a_{i+3})$ (indices mod 4), who select the first finalist on their list.

Correctness. (\Rightarrow) If Ψ is true, there exists a *history-contingent plan* for the X -stages: for each i , a rule that maps the previously chosen (Y_1, \dots, Y_{i-1}) to a choice of x_i or $\neg x_i$ such that φ holds for all Y -choices. Let v_{odd} follow this plan. Then along any sequence of Y -moves, every clause is satisfied, clause voters never fall back, and the baseline +1 elects A_{odd} at the GE. Hence neither v_{odd} (already winning) nor v_{even} (cannot induce a fallback) can profitably deviate in any subgame; backward induction yields a pure SPNE.

(\Leftarrow) If Ψ is false, then against any X -stage choices there exists a first Y_i whose choice falsifies some clause. At that subgame, the corresponding clause voters fall back cyclically, yielding the 3:1 split that—together with the baseline—flips the GE winner to A_{even} (as in Thm. 5). Thus v_{even} has a profitable deviation, contradicting subgame-perfection. Therefore no pure SPNE exists.

Hence a pure SPNE exists iff Ψ is true. The instance size is polynomial. For fixed $p = 2k$ this gives $\Sigma_{O(p)}^P$ -hardness; with unbounded p , PSPACE-hardness. Combined with membership, the theorem follows. \square

Summary. The sequential model generalizes the simultaneous setting by introducing stagewise conditioning. When all primaries are resolved simultaneously—or sequentially but with ex-ante commitment to complete strategies, before any primary outcomes are observed—dominance and equilibrium verification involve only a single existential–universal alternation and thus remain within Σ_2^P . Allowing voters to condition on earlier outcomes adds one quantifier block per stage, giving $\Sigma_{O(p)}^P$ complexity for fixed p , and PSPACE-completeness once p is part of the input. Together with Theorem 5 and Theorem 6, this yields a unified view: sequential conditioning acts as a quantifier-expanding mechanism linking the polynomial hierarchy and PSPACE.

6 CONCLUSION

We presented the first formal characterization of the computational complexity of primary elections under both simultaneous and sequential settings. In the ex-ante (commitment) model, equilibrium and dominance verification involve one existential–universal alternation, yielding Σ_2^P -completeness. Allowing voters to condition on earlier outcomes across p stages introduces additional quantifier blocks, giving $\Sigma_{O(p)}^P$ complexity for fixed p and PSPACE-completeness once the number of stages becomes part of the input.

Beyond these classifications, the staged primary process emerges as a *quantifier-expanding mechanism* that bridges the polynomial hierarchy and PSPACE. This unifies several equilibrium notions — dominant strategies, Nash equilibria, and subgame-perfect reasoning — under a common structural lens, showing how temporal sequencing amplifies computational difficulty even under simple voting rules. From a multi-agent perspective, this connects election design to the broader challenge of reasoning about sequential coordination among boundedly rational agents.

Our results suggest a form of *complexity shield*: certain strategically undesirable behaviours are computationally hard to implement in the worst case. At the same time, this hardness depends on parameters such as the number of parties and primary stages,

which may be small in real elections and thus reduce complexity. Conversely, real voters often face richer information environments—such as observing primary tallies, forming beliefs about others’ behaviour, or dealing with timing uncertainty, which can further increase strategic complexity. These issues are beyond the scope of this work but point to promising directions for future research.

7 FUTURE WORK

Our results only begin to capture the strategic and institutional richness of primary elections. Several directions remain open.

Richer strategic and institutional models. Extending the framework to mixed strategies would illuminate probabilistic reasoning and equilibrium selection in competitive settings. Further, while our hardness results apply generally, exploring restricted domains, such as closed or open primaries, may uncover tractable subclasses. Alternative voting mechanisms, including run-off and ranked-choice systems, could lead to qualitatively different complexity boundaries. Some countries even integrate polling data into nomination outcomes, suggesting hybrid models that couple social choice with statistical inference.

Information, dynamics, and candidate behaviour. Our sequential model assumes perfect information and static preferences. Relaxing these assumptions to incorporate noisy or incomplete information, belief updates, or adaptive learning would link primary elections to models of bounded rationality and dynamic games. Another promising direction is to introduce candidate positioning: when politicians choose spatial platforms strategically, primaries may reshape equilibrium positions and polarization patterns.

AI and simulation-based exploration. Agent-based simulations offer a natural complement to our theoretical results, enabling empirical study of strategic adaptation under bounded rationality. Combining such simulations with reinforcement learning or evolutionary dynamics could reveal when complex equilibria are approximable in practice, linking computational hardness with observable collective behaviour.

Overall, this work takes a first step toward a unified theory of primaries within computational social choice, clarifying how sequential collective decisions reshape both strategic reasoning and the computational landscape of multi-agent systems.

ACKNOWLEDGEMENTS

BdK was supported by EPSRC grant EP/X021696/1. MP was supported by EPSRC grant EP/X038351/1. The authors thank the anonymous reviewers for constructive feedback and helpful suggestions.

REFERENCES

- [1] Stéphane Airiau, Umberto Grandi, and Filippo Studzinski Perotto. 2017. Learning Agents for Iterative Voting. In *Algorithmic Decision Theory*, Jörg Rothe (Ed.). Springer International Publishing, Cham, 139–152.
- [2] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 3 (01 Jul 1989), 227–241. <https://doi.org/10.1007/BF00295861>
- [3] Luca Bernardi. 2013. New Adventures in Membership: Primary Elections and the Italian Case. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2305097. SSRN Working Paper, 20 pp.
- [4] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. 2024. Primarily about Primaries. *Artificial Intelligence* 329 (2024), 104095.

- [5] Markus Brill and Vincent Conitzer. 2015. Strategic Voting and Strategic Candidacy. *Proceedings of the AAAI Conference on Artificial Intelligence* 29, 1 (Feb. 2015), 69–84. <https://doi.org/10.1609/aaai.v29i1.9330>
- [6] John M. Carey and Peter M. Siavelis. 2003. Primaries and Politics in Chile. *Latin American Politics and Society* 45, 3 (2003), 1–25. <https://doi.org/10.1111/j.1548-2456.2003.tb00247.x>
- [7] Colin Cleveland, Bart de Keijzer, and Maria Polukarov. 2026. The Complexity of Strategic Behavior in Primary Elections. arXiv:2602.10290 [cs.GT] <https://arxiv.org/abs/2602.10290>
- [8] Marty Cohen, David Karol, Hans Noel, and John Zaller. 2008. *The Party Decides: Presidential Nominations Before and After Reform*. University of Chicago Press, 1427 E. 60th Street Chicago, IL 60637 USA.
- [9] Bhaskar Dutta, Matthew O. Jackson, and Michel Le Breton. 2004. Equilibrium agenda formation. *Social Choice and Welfare* 23, 1 (2004), 21–57. <http://www.jstor.org/stable/41106617>
- [10] Bhaskar Dutta, Matthew O. Jackson, and Michel Le Breton. 2001. Strategic Candidacy and Voting Procedures. *Econometrica* 69, 4 (2001), 1013–1037. <https://doi.org/10.1111/1468-0262.00228> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00228>
- [11] Paul Harrenstein, Grzegorz Lisowski, Ramanujan Sridharan, and Paolo Turrini. 2021. A Hotelling-Downs Framework for Party Nominees. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (Virtual Event, United Kingdom) (AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 593–601.
- [12] Yumi Hwang. 2017. Candidate Selection in South Korea. In *Pathways to Power: Political Recruitment and Candidate Selection in Asia*, Pippa Norris and Bernard Grofman (Eds.). Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP.
- [13] Ofer Kenig. 2009. Classifying Party Leaders' Selection Methods in Parliamentary Democracies. *Journal of Elections, Public Opinion & Parties* 19, 4 (2009), 433–447.
- [14] Andrew Kenney. 2022. Thousands of Democrats are changing their voter registration in Lauren Boebert's district ahead of the primary. Colorado Public Radio (online). <https://www.cpr.org/2022/06/07/lauren-boebert-2022-elections-democrats-unaffiliated-voters/> Accessed: 2025-10-02.
- [15] Martin Koutecký and Nimrod Talmon. 2021. Multi-Party Campaigning. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (May 2021), 5506–5513. <https://doi.org/10.1609/aaai.v35i6.16693>
- [16] Sook Jong Lee. 2022. South Korea's 2022 Presidential Election: Polarization and Change. East Asia Institute Report. <https://www.eai.or.kr/new/en/pub/view.asp?intSeq=10845> Policy report, available online.
- [17] Rémi Lefebvre. 2017. Les primaires ouvertes en France: un nouvel instrument de sélection des candidats présidentiels. *Revue française de science politique* 67, 6 (2017), 1167–1187. <https://doi.org/10.3917/rfsp.676.1167> Analyzes the introduction of open primaries in the Parti Socialiste and Les Républicains.
- [18] Omer Lev and Jeffrey S. Rosenschein. 2012. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (Valencia, Spain) (AAMAS '12)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 611–618.
- [19] Reshef Meir. 2017. Iterative Voting. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, COST Association Avenue du Boulevard – Bolwerklaan 21 1210 Brussels | Belgium, Chapter 4, 69–86. <http://procaccia.info/wp-content/uploads/2019/04/TrendsComputationalSocialChoice.pdf>
- [20] Reshef Meir, Maria Polukarov, Jeffrey Rosenschein, and Nicholas Jennings. 2010. Convergence to Equilibria in Plurality Voting. *Proceedings of the AAAI Conference on Artificial Intelligence* 24, 1 (Jul. 2010), 823–828. <https://doi.org/10.1609/aaai.v24i1.7624>
- [21] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. 2015. On the convergence of iterative voting: how restrictive should restricted dynamics be?. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (Austin, Texas) (AAAI'15)*. AAAI Press, 601 Pennsylvania Ave, NW Suite 900 Washington, DC 20004 USA, 993–999.
- [22] Nelson W. Polsby. 1983. *Consequences of Party Reform*. Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP.
- [23] Maria Polukarov, Svetlana Obraztsova, Zinovi Rabinovich, Alexander Kruglyi, and Nicholas R. Jennings. 2015. Convergence to Equilibria in Strategic Candidacy. In *International Joint Conference on Artificial Intelligence (25/07/15 - 31/07/15)*. AAAI Press / International Joint Conferences on Artificial Intelligence, Palo Alto, California USA, 624–630. <https://eprints.soton.ac.uk/376940/>
- [24] Thomas Quinn. 2024. Parliamentarians versus party members? Leadership selection systems in the British Conservative and Labour parties. *The British Journal of Politics and International Relations* 26, 2 (2024), 569–596. <https://doi.org/10.1177/13691481231191915> arXiv:<https://doi.org/10.1177/13691481231191915>
- [25] Jacob Austin Reidhead. 2020. *Comparing Intraparty Politics in South Korea and Taiwan*. Ph.D. Dissertation. Stanford University. <https://purl.stanford.edu/st454xr1842>
- [26] Annemieke Reijngoud and Ulle Endriss. 2012. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (Valencia, Spain) (AAMAS '12)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 635–644.
- [27] Ildikó Schlotter and Katarína Cechlárová. 2025. Candidate Nomination for Condorcet-consistent Voting Rules. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (Detroit, MI, USA) (AAMAS '25)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1858–1866.
- [28] Antonella Seddone and Fulvio Venturino. 2013. Bringing Voters Back in Leader Selection: the open primaries of the Italian Democratic Party. *Modern Italy* 18, 3 (2013), 303–318. <https://doi.org/10.1080/13532944.2013.801675>
- [29] Peter M. Siavelis. 2009. Chile: The End of Exceptionalism? In *Pathways to Power: Political Recruitment and Candidate Selection in Latin America*, Peter M. Siavelis and Scott Morgenstern (Eds.). Penn State University Press, 8 University Dr., USB 1 Suite C University Park, PA 16802-1003.
- [30] John Sides and Lynn Vavreck. 2020. On the Representativeness of Primary Electorates. *British Journal of Political Science* 50, 2 (2020), 677–685. <https://doi.org/10.1017/S000712341700062X>