Prospect-Theoretic Voting: Strategic Implications and Computational Limits

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

Background: In practice, elections deviate from Median Voter Theorem: candidates polarise rather than converge, and core supporters sometimes threaten abstention when platforms shift. To explain this divergence, we incorporate psychological decision making into spatial competition. By embedding prospect-theoretic voters (loss aversion, reference dependent) into the Hotelling-Downs model, we study how behavioural turnout reshape strategic positioning and equilibrium structure.

Objectives and Research Questions: This work addresses three main questions: 1. How do loss-averse, reference-dependent voters decide between turnout and abstention under a regret-minimising rule? 2. How does such behavioural turnout affect candidate positioning? Does median convergence survive, and when do extreme or asymmetric strategies arise? 3. What is the computational complexity of finding optimal positions or equilibria in this behaviourally enriched spatial model?

Methods: We incorporate prospect-theoretic voters with a discrimination threshold into the one-dimensional Hotelling–Downs framework. We derive structural results on voter support regions and equilibrium properties, including non-existence results. Algorithmically, we design geometric sweep procedures for best responses and equilibrium detection in the discrete case. We also prove NP-completeness and #P-hardness results for optimal positioning under continuous input representations.

Results: The analysis shows that behavioural turnout alters spatial competition. Structurally, voter support regions can be non-convex, median converge can fail, and pure (even mixed) Nash equilibria may not exist. Algorithmically, while best responses are tractable in the discrete case via geometric sweeps, optimal positioning becomes hard under continuous representations (NP- or #P-hard). Thus, incorporating prospect-theoretic voters changes equilibrium predictions and raises barriers for candidates.

Conclusions: The paper shows that incorporating behavioural decision-making into spatial voting models can fundamentally reshape both strategic outcomes and computational feasibility. Prospect-theoretic turnout breaks classical median convergence, may eliminate equilibria, and can induce polarisation even in symmetric electorates. At the same time, behavioural refinements significantly increase the computational difficulty of optimal candidate positioning.

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; Prospect Theory; NP-completeness; #P-hardness

ACM Reference Format:
Colin Cleveland, Bart de Keijzer, and Maria Polukarov. 2026. Prospect-Theoretic Voting: Strategic Implications and Computational Limits. In *Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Paphos, Cyprus, May 2026*, IFAAMAS, 16 pages. https://doi.org/10.65109/

1 INTRODUCTION

In 1957, Anthony Downs [7] adapted the *linear-city* framework of Harold Hotelling [13], to analyse political competition in a one-dimensional policy space. The resulting *Hotelling–Downs model* offers a simple yet powerful prediction: when two rational candidates compete for votes and voters hold single-peaked preferences, both candidates converge to the position of the median voter. This *Median Voter Theorem (MVT)* has become one of the most influential results in political economy, providing a tractable explanation for policy moderation and electoral convergence.

Although foundational, the model's prediction that politicians become indistinguishable already depends on strong assumptions in Downs' original formulation, where all voters participate and hold symmetric preferences around their ideal points. Empirical research has tested this prediction in real elections and found it holds only under certain conditions: e.g., women's enfranchisement shifted the U.S. median voter leftward and increased public-health spending [22], while rising income inequality moved the median relative to the mean and boosted demand for redistribution [21]. Similar convergence patterns appear in Swiss referenda [30], which suggests that convergence predicted by the Median Voter Theorem explains some, but not all, sources of policy change.

Yet, even these regularities leave important deviations unexplained. U.S. senators often vote in line with party ideology rather than the policy preferences of their electoral constituencies [19]; candidate identity (e.g., gender or caste) can outweigh spatial proximity in policy space, pulling outcomes away from the median

Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Armstrong, Curry, Hosseini, Mattei, Tsang, Wąs (Chairs), May 2026, Paphos, Cyprus. © 2026 Copyright held by the owner/author(s). https://doi.org/10.65109/

voter [4]; and polarised campaigns increasingly mobilise ideologically extreme voters while moderates abstain, rather than converge toward the centre [14]. Such asymmetries in voter participation, where ideologically extreme voters turn out at higher rates than moderates , and perception biases lie outside the models that assume full turnout and symmetric voter preferences, under which candidate positions converge to the median.

To explain why candidates may optimally adopt positions away from the median, we extend the Hotelling–Downs model by incorporating reference-dependent, loss-averse voter preferences that make turnout sensitive to relative gains and losses. These behavioural assumptions are motivated by *Prospect Theory* [15].

Prospect Theory identifies three key ingredients: (i) *reference dependence*, where outcomes are judged relative to a status quo (the *reference point*); (ii) *loss aversion*, where losses loom larger than equal gains; and (iii) *probability weighting*. In this work, we adopt the first two ingredients to define voter utility. A voter's utility is defined relative to the *incumbent*, that serves as a reference point, and is distorted by loss aversion. Formal details are given in Section 3 (see the appendix for extensions).

We embed this paradigm in the one-dimensional Hotelling–Downs model, where candidates pick positions in a single-dimensional space (i.e., a line) to maximise the support they receive from the voters. Unlike the classical setting of full turnout, where candidates converge to the median, our model endogenises participation. We allow for behavioural abstention, which disrupts standard equilibrium logic. We ask the following questions:

(1) How do voters choose between turnout and abstention under loss-averse, reference-dependent preferences?
(2) How should candidates position themselves when median convergence no longer holds?
(3) What equilibria and optimal strategies arise, and is it possible to compute them efficiently?

We answer these quesions in Sections 4–5 for the case of two candidates, with a sketch of the multi-candidate case given in the appendix.

Summary of our contributions. Our results, focusing on two-candidate case, draw on behavioural decision theory [12, 33], campaign strategy [27, 34], and algorithmic game theory [31], and are listed below.

- **Model.** We extend the Hotelling–Downs model to include prospect-theoretic voters with reference points, loss aversion, and abstention thresholds (Section 3).
- **Algorithms (Discrete Setting).** For a population of n discrete voters, we compute an $\tilde{O}(n)$ candidate's best-response (Algorithm 1), an $\tilde{O}(n^4)$ pure-Nash search (Algorithm 2), and a recursive boundary tracker for multiple candidates (see the appendix).
- **Complexity (Continuous Setting).** When voters are represented by density functions, we show that candidate optimization becomes NP-hard (Theorem 5) and #P-hard (Theorem 6). Additionally, we show that Nash equilibria do not generally exist in this setting.

Our findings reveal how prospect-theoretic reasoning reshapes strategic positioning and raises the algorithmic *tractability* bar in Hotelling–Downs elections.

2 RELATED WORK

Our study links the canonical theories of Hotelling–Downs spatial competition and prospect theory to three research strands: behavioural voting, strategic candidate positioning, and algorithmic analysis of spatial games.

Voter behaviour. Building on the prospect-theoretic foundations laid out in the Introduction, we survey work that adapts its core elements: "reference dependence", "loss aversion", and "probability weighting" to voting behaviour. Vis [33] assesses the empirical evidence for loss-averse preferences; Herrmann et al. [12] model distorted probability weighting in turnout; and Quattrone and Tversky [23] analyse the notion of relative candidate comparisons. Distinct frameworks such as regret theory [2, 20] and omission bias [25] also predict abstention through aversion to negative outcomes, offering parallel explanations to loss aversion. We combine a prospect-theoretic value function with a minimax-regret rule, capturing loss aversion within an ambiguity-averse decision process. Heuristic-based models add yet another layer: Fairstein et al. [9] propose the "Attainability-Utility" heuristic, where voters employ a cognitive shortcut by weighing utility directly against perceived popularity, illustrating bounded rationality in political choice.

Strategic Candidate Positioning. Positioning can reshape choice without moving voters in the issue space. Wu et al. [34] show that making a designated candidate win by modifying voters' perceived locations is NP-complete. Schlotter et al. [27] study how changing approval thresholds affects outcomes, while Riker's *heresthetic* thesis [24] highlights the power of agenda control and framing. Our results show that even without such interventions, prospect-theoretic voter behaviour alone can drive candidates away from the classical median.

Computational and equilibrium analysis. Classical spatial competition [7] and its refinements (e.g., Aragones and Palfrey [1] who incorporate abstention into probabilistic spatial voting) provide our geometric starting point. Harrenstein et al. [11] show that discrete Hotelling–Downs games with primaries may lack Nash equilibria and that checking whether one exists is NP-complete. Tian et al. [31] extend prospect-based reasoning to Markov games, identifying further computational barriers. We complement this line of research by proving that best-response computation is NP-complete when the cumulative distribution function (that represents the voter distribution) is explicitly described, and #P-hard when the voter distribution is instead described as a probability density function.

Foundational models. Hotelling–Downs model remains the baseline for one-dimensional spatial competition [7]. Its many variants now cover abstention [1], primaries [11], algorithms [3], rationalisable sets [32], popularity shocks [10], and limited attraction [28].

Prospect-theoretic voting continues to expand through probability weighting [12], loss-based thresholds [6], bandwagons [18], and empirical value-function estimates [5]. By embedding prospect-driven abstention directly into the Hotelling–Downs model, our

work connects these strands and charts the resulting computational landscape.

3 PRELIMINARIES

We model elections on the real line $\mathbb{R}$ that represents a one-dimensional policy space. The set of strategic candidates is $C = \{c_1, \ldots, c_m\}$. There is, furthermore, an *incumbent* $c_0 \in \mathbb{R}$, which is a non-strategic candidate that serves as the voters' prospect-theoretic reference point. All candidates, except the incumbent, are considered as strategic players in a game where they must choose their location on the line.[1]

Throughout the main text we focus on the two-candidate case ($m = 2$); the multi-candidate extension ($m > 2$) is deferred to the appendix.

The electorate will be represented either in a *discrete* way, as a finite set $V = \{v_1, \ldots, v_n\}$ where each $v_i \in \mathbb{R}$ is a point on the line that denotes the ideal location of voter i, or in a *continuous* way, by a probability density function (PDF) $f : \mathbb{R} \to \mathbb{R}_{\geq 0}$, with cumulative distribution function (CDF) F, such that for $x \in \mathbb{R}$, the quantity $f(x)$ denotes the density of the voter population, with the total voter population mass assumed to be normalised to 1.

Voter Utility and Behaviour. We assume that all voters perceive the distance between two points $x, y \in \mathbb{R}$ identically, using the standard metric $d(x, y) = |x - y|$. While we demonstrate below that voters may evaluate the same candidate differently, this heterogeneity arises from individual preferences rather than any distortion in the perception of spatial distance.

In the discrete variant of the model, a voter's *improvement from x relative to the incumbent* c_0 is defined as

$$\Delta(v_i, x) := 1 - \frac{d(v_i, x)}{d(v_i, c_0)},$$

which is positive when x lies closer to v_i than c_0.

This relative improvement is passed through a continuous, monotonically increasing *value function* μ_i. This function satisfies $\mu_i(0) = 0$ and is concave on $[0, \infty)$ and convex on $(-\infty, 0]$. The utility is defined as

$$u_i(x) := \mu_i(\Delta(v_i, x)). \tag{1}$$

We further assume that a voter perceives a strict preference between two candidates only when the utility gap surpasses a *discrimination threshold* $\kappa_i > 0$. That is, a voter supports c_1 over c_2 if and only if

$$u_i(c_1) - \kappa_i > u_i(c_2).$$

The full specification of voting behaviour, including abstention conditions and multi-candidate scenarios, is detailed in the appendix. Note that removing this behavioural condition would collapse the model to the classical Hotelling–Downs framework, where all voters vote for the closest candidate. The threshold κ_i represents the minimal utility advantage needed to motivate turnout, and can be interpreted as a cognitive or motivational effort needed to vote.

In the continuous variant of the model, a voter at position $v \in \mathbb{R}$ has a measurable value function $\mu(\cdot)$ and derives utility

$$u(v, x) = \mu(\Delta(v, x)) = \mu\left(1 - \frac{d(v, x)}{d(v, c_0)}\right).$$

[1] If one needs to model the incumbent as a strategic player as well, we consider them as the player c_1, with c_0 still serving as the voters' prospect-theoretic reference point.

We assume that the threshold is a uniform parameter $\kappa > 0$, which is common to all voters. Furthermore, we assume that the distance d and the value function μ are *Borel measurable* with respect to the voter's location v. This ensures that sets such as $\{v : u(v, c_1) - u(v, c_2) > \kappa\}$ are measurable, and that integrals representing vote shares are well-defined.

Candidate Goals. Candidates are assumed to choose their position to maximise an objective function under which receiving a higher support is never worse than receiving a lower support. However, there are multiple natural options for a voter's objective function, and the choice of objective strongly depends on the context and strategic motivation.

Let $\mathbf{c} = (c_1, \ldots, c_m) \in \mathbb{R}^m$ denote a profile of candidate positions, where each $c_i \in \mathbb{R}$ is the position of candidate i. We also denote $c_{-i} = (c_1, \ldots, c_{i-1}, c_{i+1}, \ldots, c_m)$ as the profile of all candidate positions without candidate i. Given a voter population V (discrete or continuous), let $\text{score}(c_i; c_{-i})$ denote the number of voters (in the discrete variant) or the mass of voters (in the continuous variant) that support candidate i against other candidates. We consider three objectives:

$$\begin{aligned}
&\text{(share)} && \max_{c_i \in \mathbb{R}} \text{score}(c_i; c_{-i}), \\
&\text{(victory)} && \max_{c_i \in \mathbb{R}} \mathbf{1}\left[\text{score}(c_i; c_{-i}) > \max_{j \neq i} \text{score}(c_i; c_{-i})\right], \\
&\text{(margin)} && \max_{c_i \in \mathbb{R}} \text{vm}(c_i) := \text{score}(c_i; c_{-i}) - \max_{j \neq i} \text{score}(c_j; c_{-j}).
\end{aligned}$$

The second objective corresponds to a first-past-the-post electoral rule where the candidate that receive the highest support wins the election. The third objective refines the second one by rewarding the margin of victory vm, not merely the event of winning. Importantly, both the victory and margin objectives rely on strict comparisons of vote counts. Thus, tie-breaking is strategically irrelevant in our model.

When the candidate or voter identity is unambiguous, we abuse notation and write c and v for their respective locations. In the later examples, if we assume identical κ and μ across all voters, we omit subscripts accordingly.

Takeaway. This model extends the classical Hotelling–Downs framework by introducing three behavioural refinements in the voter's decision process. Two refinements stem from Prospect Theory: *(PT1) Reference Dependence* and *(PT2) Value Distortion*, and one stems from relaxing the full-turnout assumption: *(A1) Discrimination Threshold.*

- **(PT1) Reference Dependence:** voters evaluate alternatives relative to the incumbent c_0, which serves as a status-quo reference point.
- **(PT2) Value Distortion:** a prospect-theoretic value function μ captures asymmetric perception of gains and losses over a voter's distance to a candidate.
- **(A1) Discrimination Threshold:** a positive discrimination threshold κ introduces abstention behaviour, which is in contrast to the universal participation assumption of the classical model.

Together, these features generate non-convex voting regions and (as we shall see in the sequel) will cause median convergence to not occur in general.

4 BEHAVIOURAL AND STRUCTURAL DEVIATIONS FROM MEDIAN CONVERGENCE

We now analyse the two-candidate case, with positions denoted by ℓ and r (initially $\ell < r$). Without loss of generality, throughout the paper we assume that the incumbent is fixed at $c_0 = 0$. The labels ℓ, r track specific candidates and are fixed; the candidates are not relabelled if they swap positions.

The behavioural refinements introduced in Section 3 have two immediate consequences. First, introducing the Discrimination Threshold (A1) makes the *gap* between candidate positions matter: a voter may abstain unless the perceived utility difference exceeds the threshold κ_i.

Second, introducing Reference Dependence (PT1) and Value Distortion (PT2) alters how voters value spatial distance. Utility differences are now evaluated relative to the incumbent c_0 and transformed by a non-linear value function μ. As a result, some voters may perceive two candidates as "equally distant" even when they are not, while others may exaggerate the difference.

In what follows, we show that (i) the classical MVT fails under prospect-theoretic decision rules, and (ii) this valuation distortion cannot be removed by redefining the locations of the voters in the policy space.

4.1 Breakdown of the Median Voter Theorem

Under the classical MVT, two rational candidates that face single-peaked voters with full turnout rate converge to the median [7].

In the following example, we show that when voters evaluate candidates by reference dependence (PT1) utility function with discrimination threshold (A1), which leads to possible abstentions, such convergence may fail.

Example 1 (Non-existence of a pure-strategy Nash equilibrium). *There are five voters located at*

$$v_1 = v_2 = -1, \quad v_3 = 0, \quad v_4 = v_5 = 2,$$

with discrimination threshold

$$\kappa_1 = \kappa_2 = 1 - \mu(\Delta(-1, 2)) > 0, \qquad \kappa_3 = \kappa_4 = \kappa_5 = 0,$$

and two candidates $\ell < r$ with $c_0 = 0$. Since μ is unspecified, we do not assign numerical values to κ_1 and κ_2. The choice $1 - \mu(\Delta(-1, 2))$ makes voters v_1 and v_2 exactly indifferent between abstention and voting whenever both candidates lie in the interval $[-1, 2]$.

Since all voters satisfy $v_i \in [-1, 2]$, for any candidate $c \in \{\ell, r\}$ with $c < -1$ we have $|v_i - c| \geq |v_i - (-1)|$, and therefore $\Delta(v_i, c) \leq \Delta(v_i, -1)$ for all i . As each value function μ_i is monotonically increasing, this implies $u_i(c) = \mu(\Delta(v_i, c)) \leq u_i(-1)$ for every voter. Hence, moving a candidate further left than -1 cannot increase that candidate's utility for any voter, and thus cannot generate additional support. We therefore restrict attention to $\ell, r \geq -1$ in this example. The validity of this restriction—together with the resolution of the resulting weak-dominance issue—is formalised in Proposition 7 (see the appendix).

Voters v_1 and v_2 vote ℓ only if[2]

$$\mu(\Delta(v_1, \ell)) - \mu(\Delta(v_1, r)) \geq 1 - \mu(\Delta(-1, 2)).$$

When $\ell = v_1 = -1$, this reduces to

$$\mu(\Delta(v_1, 2)) \geq \mu(\Delta(v_1, r)).$$

Since μ is monotonically increasing, this is equivalent to $\Delta(v_1, 2) \geq \Delta(v_1, r)$, i.e. $|v_1 - 2| \leq |v_1 - r|$, which holds exactly when $|r + 1| \geq 3$, that is, when $r \geq 2$ or $r \leq -4$.

The latter case can be ruled out here, since we restrict attention to $r \geq -1$. Accordingly, v_1 and v_2 may support ℓ only when $r \geq 2$ if $\ell \in [-1, 2]$.

Voters v_3, v_4, v_5 simply follow proximity voting and break ties in favour of r. Candidates maximise the vote margin vm.

Case 1: $r < 2$. *Voters v_1, v_2 abstain, while v_4, v_5 support r. Thus r receives at least two votes, which is sufficient to guarantee victory regardless of the behaviour of v_3. In this case, ℓ has an incentive to move into the interval $(r, 2)$; in response, r prefers moving further right towards 2.*

Case 2: $r \geq 2$. *Candidate ℓ can move to -1, which activates voters v_1 and v_2, who then support ℓ. Voter v_3 is closer to -1 than to r, while v_4 and v_5 continue to support r. Hence ℓ wins by* 3:2. *Thus, r prefers moving back to some $r' < 2$, returning to Case 1.*

Best responses cycle and no fixed point exists. Therefore, no pure-strategy Nash equilibrium exists.

Implications. This example shows that the median is no longer a dominant strategy under prospect-theoretic preferences, and that pure-strategy equilibria may fail to exist. In particular, prospect-theoretic abstention can generate cycles of profitable deviations, and make pure-strategy Nash equilibria impossible.

In Proposition 1 (proof in the appendix), we extend this observation to mixed strategies.

Proposition 1 (No Nash equilibrium, even in mixed strategies). *In the setting of Example 1,* no *mixed-strategy profile (σ_ℓ, σ_r) constitutes a Nash equilibrium.*

4.2 Necessity of the Value Function

In our model, voters correctly perceive the left–right ordering of candidates on the policy line, and policy distances themselves are not distorted. However, the value function μ is non-linear in distance. This raises a natural question: can candidate positions be re-embedded on a transformed line so that value becomes linear in distance, while preserving the same left–right ordering? This relativity is captured by the value function μ, which maps spatial distances into perceived utility. A natural question then arises: can the model be "flattened" by removing μ, i.e., by redefining spatial locations so that utility is measured linearly?

For a single voter v, this is possible. Recall that $u(v, x) = \mu(1 - (d(v, x)/d(v, c_0)))$, where $d(v, x) = |v - x|$. By setting the reference candidate $\bar{c}_0 = 0$ and normalising the voter's position to $\bar{v} = 1$, we can map any candidate position x to a transformed location

$$\bar{x} = \begin{cases} 1 + |1 - u(v, x)|, & \text{if } x > v, \\ 1 - |1 - u(v, x)|, & \text{if } x < v. \end{cases}$$

[2]For simplicity, we use a weak inequality ($\geq$) for activation, unlike the strict form used in the Preliminaries.

This results in a identity value representation $\mu(\Delta(v,x)) = \mu(\frac{|v-x|}{|v|}) = \frac{|v-x|}{|v|}$ in the transformed space.

However, there is no transformation that can jointly linearise the values of multiple voters. As a candidate moves, each voter experiences changes in their utility, filtered through their own (possibly non-linear) value function μ_i, and differing curvatures of these functions prevent a single spatial mapping from preserving utility orderings across all voters.

As shown below, this impossibility persists even when all voters share the same monotonic value function.

Lemma 2. *Let d be the standard distance metric, and c_0 the incumbent. Fix an instance with two voters i, j with distinct ideal points $v_i \neq v_j \neq c_0$. Each voter evaluates an alternative x by the utility function defined in* (1)*, where the common value function μ is strictly increasing and strictly concave on the gain side, with $\mu(0) = 0$ and $\mu(1) = 1$.*

Then, there does not exist a function $g : \mathbb{R} \to \mathbb{R}$ such that for both voters $k \in \{i, j\}$,

$$u_k(x) = 1 - \left| \frac{g(x) - g(v_k)}{g(v_k) - g(c_0)} \right|.$$

Since the impossibility already occurs with two voters, it must also occur with more than two. The proof (see the appendix) shows that this impossibility arises purely from non-linearity, not from heterogeneity: it holds even when all voters share the same value function $\mu_i = \mu$.

5 BEST RESPONSES AND EQUILIBRIA

We now study strategic interaction between the two candidates. Each voter evaluates the two positions through their prospect-theoretic utilities and supports the candidate that offers a sufficiently larger perceived gain after accounting for their discrimination threshold κ. To analyse best responses and (non-)existence of pure-strategy equilibria, we formalise this comparison using a surplus function.

A voter v prefers candidate x to y exactly when

$$u(v,x) - u(v,y) > \kappa \quad \Longleftrightarrow \quad \mu(\Delta(v,x)) - \mu(\Delta(v,y)) > \kappa.$$

We therefore introduce the general surplus function

$$S(v;x,y,\kappa) := \mu\left(1 - \frac{d(v,x)}{d(v,c_0)}\right) - \mu\left(1 - \frac{d(v,y)}{d(v,c_0)}\right) - \kappa, \qquad (2)$$

which measures the net perceived improvement from voting for x instead of y. Without loss of generality, in the best-response analysis we focus on candidate ℓ. Accordingly, we treat r as fixed and use the shorthand

$$S^{\ell}(v) := S(v;\ell,r,\kappa). \qquad (3)$$

We analyse two objectives for candidate ℓ: the *share* objective, which maximises the vote count $\text{score}(\ell;r)$, and the *margin* objective, which maximises the vote margin $\text{vm}(\ell) = \text{score}(\ell;r) - \text{score}(r;\ell)$.

5.1 Discrete Model

In the discrete setting with voter set $V = \{v_1, \ldots, v_n\}$, the support for ℓ is

$$\rho(\ell,r) := \{v_i \in V : S(v_i;\ell,r,\kappa) > 0\},$$

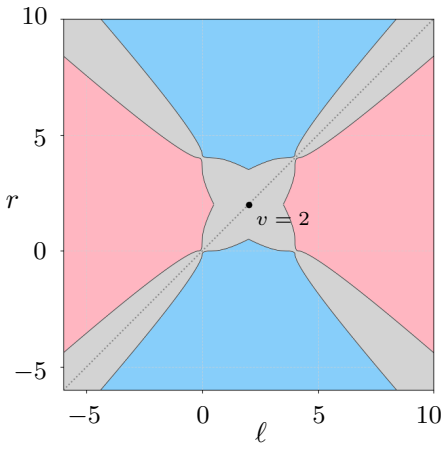

Figure 1: Each point (ℓ, r) shows candidate positions for a voter at $v = 2$ with threshold $\kappa = 0.5$ and value function $\mu = \sqrt{\Delta(v,x)}$ for gains and $-2\sqrt{-\Delta(v,x)}$ for losses. Blue, red, and grey areas indicate support for ℓ, r, and abstention.

and the score of ℓ is simply $\#\rho(\ell,r)$.

Candidate ℓ receives a voter's support exactly when ℓ lies in the region where that voter obtains positive surplus against r. For a voter at location $v \in \mathbb{R}$, define the *activation regions*

$$\gamma_v(r) = \{\ell \in \mathbb{R} : S(v;\ell,r,\kappa) > 0\},$$
$$\gamma'_v(r) = \{\ell \in \mathbb{R} : S(v;r,\ell,\kappa) > 0\},$$

that describe, respectively, where v would vote for ℓ or for r. The complement of $\gamma_v(r) \cup \gamma'_v(r)$ corresponds to abstention.

A key feature is that these regions admit a simple structure: each $\gamma_v(r)$ is either empty or a single open interval, while each $\gamma'_v(r)$ is always the union of two open intervals. Formally:

Lemma 3. *When $v \neq 0$, the set $\gamma_v(r)$ is a single open interval. In contrast, $\gamma'_v(r)$ is the union of two open intervals. The endpoints of both regions are functions of r.*

The full proof is in the appendix.

Consequently, we may write

$$\gamma_v(r) = (\underline{\gamma}_v(r),\ \overline{\gamma}_v(r)),\ \gamma'_v(r) = (-\infty, \underline{\gamma}'_v(r)) \cup (\overline{\gamma}'_v(r), \infty).$$

The four endpoints $\underline{\gamma}_v(r), \overline{\gamma}_v(r), \underline{\gamma}'_v(r), \overline{\gamma}'_v(r)$ are functions of the opponent position r and represent the values of ℓ at which the voter's surplus against r changes sign. When plotted endpoints in the (ℓ, r)-plane, they form the *boundary curves* between regions of distinct voting behaviour for voter v. An illustration of these regions is shown in Figure 1.

Fix an opponent position r. For each voter v_i, the endpoints $\underline{\gamma}_{v_i}(r), \overline{\gamma}_{v_i}(r), \underline{\gamma}'_{v_i}(r), \overline{\gamma}'_{v_i}(r)$ mark where v_i changes behaviour as ℓ varies. Collecting and sorting these breakpoints over all voters partitions the ℓ-axis into $O(n)$ intervals. Within each interval, every voter's action is fixed, and hence the vote margin $\text{vm}(\ell;r)$ is constant. Algorithm 1 sweeps these intervals to identify a vote-margin–maximising best response for ℓ against r.

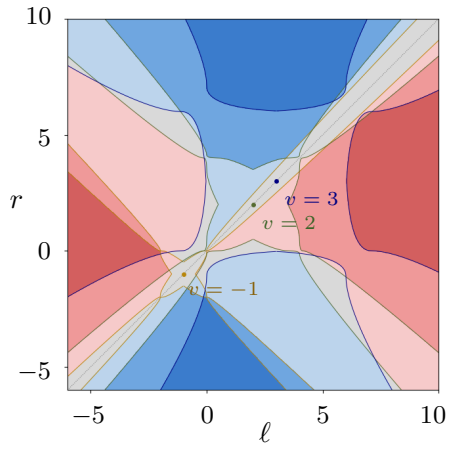

Figure 2: Victory margin with three loss-averse voters: Each point (ℓ, r) represents candidate positions for $(v_i, \kappa_i) = \{(-1, 0.3), (2, 0.5), (3, 1.2)\}$ with the same value function as in Fig. 1. Darker Blue (resp., darker red) regions indicate where ℓ (resp., r) wins by a larger margin. Curves show each voter's switching boundaries, coloured according to their labels.

The share objective can be handled within the same framework by modifying the aggregation rule: negative contributions (corresponding to votes for the opponent) are treated as zero, so only the inner threshold region contributes to the score. The underlying threshold structure is unchanged.

5.1.1 Equilibria. Since vote-margin optimisation forms a two-player zero-sum game, any Nash equilibrium satisfies $\mathsf{vm}(\ell) = \mathsf{vm}(r) = 0$.

Each voter's threshold functions $\gamma_i(r)$ and $\gamma_i'(r)$ partition the line into five zones with local vote contributions following the pattern $(-1, 0, +1, 0, -1)$. Here, the numbers indicate v_i's contribution to $\mathsf{vm}(\ell)$: +1 indicates a vote for ℓ, −1 indicates a vote for r, and 0 indicates abstention. When $\underline{\gamma}_i(r) > \overline{\gamma}_i(r)$, the central +1 zone vanishes, and the pattern collapses to $(-1, 0, -1)$: the voter always prefers r or abstains, thus contributing non-positively to $\mathsf{vm}(\ell)$.

While a single voter's contribution to $\mathsf{vm}(\ell)$ is quasiconcave, different voters induce different patterns, and the global vote margin vm is the aggregate of these local effects. As a result, vm is generally not quasiconcave. See Figure 2 for a visual example of such overlapping patterns.

To find a Nash equilibrium, we scan left to right over the domain of ℓ, recording intersections between boundary curves $\gamma_i(r)$ and $\gamma_i'(r)$, and tracking changes in vote margin $\mathsf{vm}(l)$. A symmetric scan for r returns $\mathsf{vm}(r)$. By comparing maximal margins across intervals and testing for overlaps, we can identify pure-strategy Nash equilibria.

When the value function μ is non-linear, the boundary curves induced by different voters may overlap in a highly complex manner. In particular, without additional regularity assumptions, the number of intersections among these curves can be unbounded. Since our algorithms work by enumerating and testing regions bounded by these curves, tractability hinges on keeping their combinatorial complexity under control. To this end, we restrict attention to instances in which each curve intersects any other curve only $O(1)$ times. With n voters, this implies that the (ℓ, r)-plane is partitioned into $\tilde{O}(n^2)$ regions.[3]

The algorithm for computing a mini-max strategy for a single candidate runs in $\tilde{O}(n^2)$ time (Algorithm 2; see the appendix). Since $\mathsf{vm}(\ell, r) = -\mathsf{vm}(r, \ell)$, it suffices to analyse one direction.

Suppose the solution set for r consists of curved regions of the form

$$\{(\ell, r) \mid r \in (p_i, p_j), \ell \in [\gamma_k(r), \gamma_k'(r)]\},$$

where $r \in (p_i, p_j)$ and the best response ℓ lies between γ_k and γ_k'. A symmetric construction gives the regions for ℓ. Equilibrium candidates are those where these regions overlap. Assuming each γ-type curve intersects any other only $O(1)$ times, the number of curved regions is $\tilde{O}(n^2)$. If pairwise region intersection can be tested in $O(1)$ time (e.g., under monotonicity), equilibrium detection runs in $\tilde{O}(n^4)$ time. However, as Example 1 shows, even mixed equilibria may fail to exist. This motivates future work on approximate equilibria or alternative solution concepts.

5.2 Continuous Model

In the continuous model, the support for ℓ is a measurable set

$$\rho(\ell, r) := \{v \in \mathbb{R} : S^\ell(v) > 0\}. \tag{4}$$

Given a bounded voter density function f, the score of ℓ is therefore

$$\mathsf{score}(\ell; r) = \int_{\rho(\ell,r)} f(v)\mathrm{d}v. \tag{5}$$

We assume μ is measurable in v (as discussed in Preliminaries, the paragraph about μ's measurability), which ensures that both $\rho(\ell, r)$ and the above integral are well-defined.

To compute $\mathsf{score}(\ell; r)$, let us illustrate how we could proceed if $\rho(\ell, r)$ would be guaranteed to always form a single interval: $\rho(\ell, r) = (a(\ell), b(\ell))$. In this case,

$$\mathsf{score}(\ell; r) = F(b(\ell)) - F(a(\ell)),$$

where F is the CDF of f. Thus, the score depends only on the three functions F, $a(\ell)$, and $b(\ell)$, and we would need to evaluate F at boundary points $a(\ell)$ and $b(\ell)$. Maximising $\mathsf{score}(\ell)$ then reduces to tracking how the two boundary points $a(\ell)$ and $b(\ell)$ move and how $F(a(\ell)) - F(b(\ell))$ varies. While it is not immediately clear how to turn the latter into a precise algorithm, having $\rho(\ell, r)$ be a single interval seemingly makes this problem a lot more approachable than if $\rho(\ell, r)$ were an arbitrary unstructured support region.

Thus, a meaningful question to ask now is whether $\rho(\ell, r)$ is guaranteed to be a single interval. Unfortunately, the answer turns out to be negative. If two voters at locations v_1, v_3 support ℓ, a voter v_2 between them may abstain if $S^\ell(v_2) < 0$. Thus, $v_1, v_3 \in \rho(\ell, r)$ but not v_2.

Example 2 (Non-convexity of the Support Region). *Let $\mu(x) = \sqrt{x}$ for gains (i.e., on $[0, \infty)$) and $-2\sqrt{-x}$ for losses (i.e., on $(-\infty, 0]$). Set $\ell = 2$, $r = 5$, and $\kappa = 3$. Then:*

$$S^\ell(0.2) \approx 0.93, \quad S^\ell(0.9) \approx -0.17, \quad S^\ell(1) \approx 0.46.$$

[3] If any two curves intersect only $O(1)$ times, the total number of distinct regions formed is $\tilde{O}(n^2)$ [8].

Voters at 0.2 *and* 1 *would vote for* ℓ*, but voters at* 0.9 *would abstain. The support region is thus non-convex. We visualise the support region in Figure 3.*

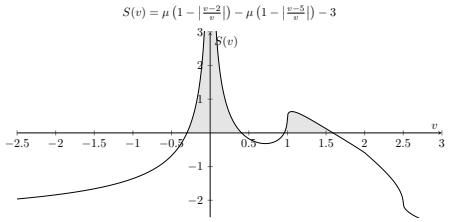

Figure 3: Example of support region $\rho(\ell, r)$ that is not convex.

Structure of $\rho(\ell, r)$. The function $S^\ell(v)$ is continuous on $\mathbb{R} \setminus \{0\}$, with at most a single discontinuity at $v = 0$. Since $v \to 0, S^\ell(v)$ exists ($\pm\infty$), the set $\rho(\ell, r) := \{v \in \mathbb{R} : S^\ell(v) > 0\}$ is open and can be written as

$$\rho(\ell, r) = \bigcup_{i=1}^{\infty} (a_i, b_i),$$

a countable union of disjoint open intervals. It is measurable and differs from an open set only by a measure-zero point.[4]

In Example 2, the resulting ρ consists of only a few intervals. In full generality, ρ may consist of infinitely many connected components, as shown in Theorem 4 (full proof in the appendix).

Theorem 4 (Infinite–Interval Theorem). *For every ℓ, r with $\ell < r$, a constant $\kappa > 0$, and $S^\ell(v)$ as defined by* (3) *and $v < 0$, there exists a value function $\mu : (-\infty, 0] \to \mathbb{R}$ such that the set $\rho(\ell, r)$ as defined by* (4) *is a countable union of disjoint open intervals (hence infinitely many components) that yet satisfies for the Lebesgue measure λ that $\lambda(\rho) < \infty$.*

In the continuous setting, candidate ℓ's score becomes:

$$\text{score}(\ell) = \sum_{i=1}^{\infty} (F(b_i) - F(a_i)).$$

Although this conclusion involves potentially infinite intervals, the $\lambda(\rho(\ell, r))$ is finite.[5] Since f is bounded, the score in the continuous model can be approximated to arbitrary precision by integrating over a finite union of intervals.

5.2.1 Optimal Location and Equilibrium. Unlike in the discrete model, vote-share maximisation remains NP-hard in the continuous model, even if the electorate is given as a polynomial-time computable CDF (i.e. the integral of the density f).

Theorem 5 (NP-completeness). *Let F be a CDF, represented by a polynomial-time Turing machine M that evaluates F at a given input point up to n bits, n given in unary.*[6] *Then, the following decision problem is* NP*-complete:*

> ***Instance:*** *Turing machine M representing F, and a rational threshold n.*
> ***Question:*** *is there a $\ell \in \mathbb{R}$ s.t. score$(\ell) \geq n$, where $\mu(v, \cdot)$ is linear on $\mathbb{R}$ for all v?*

The proof is given in the appendix. If the PDF f is provided, instead of the CDF, then the optimal location for ℓ becomes a #P-hard problem (proof in the appendix). We borrow the result of Kawamura [16, Section 4.3.2], which proved that integration of polynomial-time computable functions is #P-hard.

Theorem 6 (#P-hardness). *Let f be represented by a polynomial-time Turing machine M that evaluates F at a given input point up to n bits, n given in unary. Then, the following decision problem is* #P*-hard:*

> ***Instance:*** *Turing machine M representing f, rational number a.*
> ***Question:*** *Is there an $\ell \in \mathbb{Q}$ such that score$(\ell) \geq a$, where $\mu(v, \cdot)$ is linear on $\mathbb{R}$ for all v?*

While these results show that the continuous model is computationally hard in general, this does not mean it is always inferior in practice. As an idealisation of the discrete setting, continuous models often simplify analysis and enable compact representations. However, computational tractability depends not only on voter behaviour but also on how it is encoded. The above result shows that continuous formulations, though elegant, do not inherently offer computational advantages and should be adopted with care.

6 CONCLUSION

Classical models of electoral competition often assume full turnout and rational, utility-maximising behaviour. In this paper, we asked what happens when voters deviate from these assumptions systematically, as prescribed through Prospect Theory. By incorporating reference dependence, loss aversion, and the possibility of abstention based on perceived regret, we showed how familiar results like the Median Voter Theorem can unravel. A candidate's strategy choice needs to be reconsidered with care: The candidate's optimal position depends not only on the distribution of preferences but also on how voters interpret and filter their choices.

On the technical side, we analysed how difficult it is for candidates to compute best responses to each other in this setting, and present some important initial algorithmic results on this behaviourally enriched voting model. For discrete sets of voters, the task is tractable, but for more general input formats, it becomes provably computationally hard.

Several aspects remain unresolved. It would be worthwhile to investigate follow-up questions around equilibrium existence under more general conditions, multidimensional preferences, and the strategic implications of candidate asymmetries. Whether this model can explain empirical voting patterns, is also an open and intriguing question.

REFERENCES

[1] Enriqueta Aragones and Thomas Palfrey. 2002. Mixed Strategy Equilibrium in a Downsian Model with a Favored Candidate. *Journal of Economic Theory* 103, 1 (2002), 131–161.

[2] David E. Bell. 1982. Regret in decision making under uncertainty. *Operations Research* 30, 5 (1982), 961–981.

[4] See Royden and Fitzpatrick [26, Prop. 8, Prop. 9 and Prop. 22, Sec. 1.4];

[5] As $v \to \pm\infty$, both $\Delta(v, \ell)$ and $\Delta(v, r)$ approach zero: this follows since $\Delta(v, \ell) = 1 - \frac{d(v,\ell)}{d(v,c_0)}$, and for any fixed $\ell \neq c_0$, both numerator and denominator tend to infinity, but $\frac{d(v,\ell)}{d(v,c_0)} \to 1$, so the ratio tends to 1 , so $S^\ell(v) \to -\kappa$. Hence $S^\ell(v) > 0$ holds only within a bounded region.

[6] See Ko and Friedman [17]: F is in the class PF.

[3] Umang Bhaskar and Soumyajit Pyne. 2024. Equilibrium Computation in the Hotelling-Downs Model of Spatial Competition. arXiv:2412.12523 [cs.GT] https://arxiv.org/abs/2412.12523

[4] Raghabendra Chattopadhyay and Esther Duflo. 2004. Women as Policy Makers: Evidence from a Randomized Policy Experiment in India. *Econometrica* 72, 5 (2004), 1409–1443.

[5] Aleksandra Conevska and Can Mutlu. 2025. When Do Voters Stop Caring? Estimating the Shape of Voter Utility Function. arXiv:2501.03196 [econ.GN] https://arxiv.org/abs/2501.03196

[6] Kohei Daido and Tomoya Tajika. 2022. Impact of Information Concerning the Popularity of Candidates on Loss-Averse Voters' Abstention. *Economic Theory Bulletin* 10, 1 (2022), 41–51. https://doi.org/10.1007/s40505-021-00214-y

[7] Anthony Downs. 1957. *An Economic Theory of Democracy.* Harper and Row, Harper & Row, Publishers, Incorporated, 49 East 33rd Street, New York, N. Y. 10016.

[8] Herbert Edelsbrunner, Leonidas Guibas, János Pach, Richard Pollack, Raimund Seidel, and Micha Sharir. 1992. Arrangements of curves in the plane—topology, combinatorics, and algorithms. *Theoretical Computer Science* 92, 2 (1992), 319–336. https://doi.org/10.1016/0304-3975(92)90319-B

[9] Roy Fairstein, Adam Lauz, Reshef Meir, and Kobi Gal. 2019. Modeling People's Voting Behavior with Poll Information. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (Montreal QC, Canada) *(AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1422–1430.

[10] Gaëtan Fournier and Marc Schröder. 2024. Popularity in location games. arXiv:2401.07732 [cs.GT] https://arxiv.org/abs/2401.07732

[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] Oliver Herrmann, Richard Jong-A-Pin, and Lambert Schoonbeek. 2019. A Prospect-Theory Model of Voter Turnout. *Journal of Economic Behavior and Organization* 168 (2019), 362–373. https://doi.org/10.1016/j.jebo.2019.10.012

[13] Harold Hotelling. 1929. Stability in Competition. *The Economic Journal* 39, 153 (1929), 41–57. http://www.jstor.org/stable/2224214

[14] Matthew I. Jones, Antonio D. Sirianni, and Feng Fu. 2022. Polarization, Abstention, and the Median Voter Theorem. *Humanities and Social Sciences Communications* 9, 1 (2022), 43. Examines how polarization shifts campaign strategies toward bases, often excluding moderates.

[15] Daniel Kahneman and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47, 2 (1979), 263–291.

[16] Akitoshi Kawamura. 2011. *Computational Complexity in Analysis and Geometry.* Ph.D. Dissertation. University of Toronto.

[17] Ker-I. Ko and Harvey Friedman. 1982. Computational complexity of real functions. *Theoretical Computer Science* 20, 3 (1982), 323–352.

[18] Anastasia Leontiou, Georgios Manalis, and Dimitrios Xefteris. 2023. Bandwagons in costly elections: The role of loss aversion. *Journal of Economic Behavior & Organization* 209 (2023), 471–490. https://doi.org/10.1016/j.jebo.2023.03.011

[19] Steven D. Levitt. 1996. How Do Senators Vote? Disentangling the Role of Voter Preferences, Party Affiliation, and Senator Ideology. *American Economic Review* 86, 3 (1996), 425–441.

[20] Graham Loomes and Robert Sugden. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal* 92, 368 (1982), 805–824.

[21] Branko Milanovic. 2000. The Median Voter Hypothesis, Income Inequality, and Income Redistribution: An Empirical Test with the Required Data. *European Journal of Political Economy* 16, 3 (2000), 367–410.

[22] Grant Miller. 2008. Women's Suffrage, Political Responsiveness, and Child Survival in American History. *The Quarterly Journal of Economics* 123, 3 (2008), 1287–1327.

[23] George A. Quattrone and Amos Tversky. 1988. Contrasting Rational and Psychological Analyses of Political Choice. *American Political Science Review* 82, 3 (1988), 719–736.

[24] William H. Riker. 1986. *The Art of Political Manipulation.* Yale University Press, 302 Temple Street, New Haven, CT 06511-8909, USA.

[25] Ilana Ritov and Jonathan Baron. 1992. Omission bias and the preference for the status quo. *Journal of Risk and Uncertainty* 5, 1 (1992), 49–61.

[26] Halsey L. Royden and Patrick M. Fitzpatrick. 2010. *Real Analysis* (4th ed.). Pearson, Boston.

[27] Ildiko Schlotter, Piotr Faliszewski, and Edith Elkind. 2017. Campaign Management Under approval-Driven Voting Rules. *Algorithmica* 77, 1 (01 Jan 2017), 84–115. https://doi.org/10.1007/s00453-015-0064-0

[28] Weiran Shen and Zihe Wang. 2017. Hotelling-Downs Model with Limited Attraction. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems* (São Paulo, Brazil) *(AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 660–668.

[29] Maurice Sion. 1958. On general minimax theorems. *Pacific J. Math.* 8, 1 (1958), 171–176.

[30] David Stadelmann, Marco Portmann, and Reiner Eichenberger. 2012. Evaluating the Median Voter Model's Explanatory Power. *Economics Letters* 114, 3 (2012), 312–314.

[31] Ran Tian, Liting Sun, and Masayoshi Tomizuka. 2021. Bounded Risk-Sensitive Markov Games: Forward Policy Design and Inverse Reward Learning with Iterative Reasoning and Cumulative Prospect Theory. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 7 (May 2021), 6011–6020. https://doi.org/10.1609/aaai.v35i7.16750

[32] Job van Sloun. 2023. Rationalizable behaviour in the Hotelling-Downs model of spatial competition. *Theory and Decision* 95, 4 (2023), 359–379.

[33] Barbara Vis. 2011. Prospect Theory and Political Decision Making. *Political Studies Review* 9, 3 (2011), 334–343. https://doi.org/10.1111/j.1478-9302.2011.00238.x arXiv:https://doi.org/10.1111/j.1478-9302.2011.00238.x

[34] Junlin Wu, Andrew Estornell, Lecheng Kong, and Yevgeniy Vorobeychik. 2022. Manipulating Elections by Changing Voter Perceptions. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 557–563. https://doi.org/10.24963/ijcai.2022/79 Main Track.

A APPENDIX: DECISION MODEL OF VOTERS IN MULTI-CANDIDATE SITUATION

In elections with more than two candidates, voter behaviour becomes substantially more complex than the simple left–right comparison assumed in the Hotelling–Downs model. A voter may need to evaluate trade-offs across multiple plausible winners, consider strategic or psychological factors, or account for the possibility of abstaining.

To clarify the behavioural assumptions underlying our multi-candidate extensions, we summarise three decision policies commonly used in the literature: *baseline truthful voting*, *regret minimisation*, and *subjective expected utility*, each capturing a different way in which voters may process utility differences and uncertainty.

Baseline Truthful Voting. Voters are assumed to vote for their utility-maximising candidate if it's strictly better than all others *and* exceeds the abstention threshold:

$$u_i(c_k) > u_i(c_j)\ \forall j \neq k, \quad \text{and} \quad u_i(c_k) > \kappa_i.$$

This model does not account for the *magnitude* of differences: even a marginally better candidate will be chosen, as long as their utility exceeds the threshold.

Regret-Minimising. This model borrow the idea from knowledge theory.

We model voters as regret-averse agents. Each considers hypothetical "worlds" $\omega \in \Omega$, where their action affects the election outcome. Let $S = \{\alpha, c_1, \ldots, c_m\}$ be the set of actions, with α denoting abstention. A world ω maps each action $s \in S$ to a resulting winner.

The regret of choosing action s in world ω is:

$$\text{Regret}_i(\omega, s) = \max_{s' \in S} U_i(\omega, s') - U_i(\omega, s),$$

where $U_i(\omega, s)$ is the utility voter i derives from s in ω, accounting for voting cost κ_i.

Since the voter does not know whether they are pivotal, they adopt a minimax regret strategy:

$$\min_{s \in S} \max_{\omega \in \Omega} \text{Regret}_i(\omega, s). \tag{6}$$

This captures how voters may abstain when the potential regret from backing the "wrong" winner outweighs the perceived benefit, reflecting both ambiguity aversion and responsibility avoidance.

For example, let $\omega_{j,k}$ denote a world in which abstaining causes candidate c_k to win, but voting for c_j would cause c_j to win. If $u_i(c_j) - u_i(c_k) > \kappa_i$, abstaining creates more regret than voting, so the voter participates. Conversely, if utility differences are small, abstention may minimise regret by avoiding both the cost and the burden of influencing an unfavourable outcome.

In particular, when two preferred candidates are close and the voter is unsure who will win, the act of voting for one may lead to greater regret than abstaining, despite both being better than the least-favourite.

Subjective Expected Utility. In the *Subjective Expected Utility (SEU)* model, voters hold a belief distribution $P_i(\omega)$ over the states of the world and maximise the expected utility:

$$s_i^* \in \arg\max_{s \in S} \sum_{\omega \in \Omega} P_i(\omega) U_i(\omega, s).$$

Prospect Theory modifies this by applying a probability weighting function π_i to obtain:

$$V_i(s) = \sum_{\omega \in \Omega} \pi_i\left(P_i(\omega)\right) U_i(\omega, s),$$

capturing how voters overweight rare events and underweight likely ones.

B APPENDIX: A FORMAL PROOF OF NO EQUILIBRIUM OUTSIDE VOTERS' SUPPORT IN EXAMPLE 1

In Example 1, we argued informally that no candidate has an incentive to locate outside the interval $[-1, 2]$, which coincides with the support of the voter distribution. We now formalise this observation by showing that no pure-strategy equilibrium $(\ell^\star, r^\star)$ can place either candidate strictly outside this interval when candidates maximise vote margin.

Proposition 7. *In the setting of Example 1, under the margin-maximisation objective, there is no pure-strategy equilibrium $(\ell^\star, r^\star)$ with $\ell^\star \leq r^\star$ such that either $\ell^\star < -1$ or $r^\star > 2$.*

Proof. We consider all cases in which at least one candidate lies outside the voter support $[-1, 2]$, and show that in each case a profitable deviation exists.

Case 1: $\ell^\star < -1$ and $r^\star < 2$. In this configuration, voters v_4 and v_5 are strictly closer to $r^\star$ than to $\ell^\star$, while voters v_1 and v_2 are strictly closer to $\ell^\star$ than to $r^\star$ but remain inactive due to their discrimination thresholds. Candidate r can deviate to any $r' \in (-1, 1)$, thereby attracting v_3 while retaining v_4 and v_5. This guarantees r at least three votes. Candidate ℓ receives no votes in this outcome.

However, ℓ can then deviate to any position $\ell' \in (r', 2)$, which makes v_4 and v_5 strictly closer to ℓ' than to r'. This deviation gives ℓ two votes instead of zero, strictly increasing her vote margin. Therefore, $(\ell^\star, r^\star)$ cannot be an equilibrium.

Case 2: $r^\star > 2$. Since $r^\star$ lies to the right of all voters, voters v_1 and v_2 may become active and support $\ell^\star$ whenever

$$\mu(\Delta(v_1, \ell^\star)) > \kappa_1 + \mu(\Delta(v_1, r^\star)).$$

In particular, ℓ can choose a position $\ell^\star$ with $|\ell^\star| < |r^\star|$, ensuring that v_3 also prefers $\ell^\star$. Thus, ℓ can secure at least three votes.

We claim that r has a profitable deviation. If $\ell^\star = -1$, let $r' = 0$; otherwise, let $r' = 2$. By the definition of κ_1 in Example 1, we have $\kappa_1 + \mu(\Delta(v_1, 0)) > 1$, while $\mu(\Delta(v_1, -1)) \leq 1$. Moreover, whenever $\ell^\star \neq -1$, $\Delta(v_1, \ell^\star) < 1$ and hence $\mu(\Delta(v_1, \ell^\star)) < 1$, whereas $\kappa_1 + \mu(\Delta(v_1, 2)) \geq 1$ by assumption. Therefore, in either situation,

$$\kappa_1 + \mu(\Delta(v_1, r')) > \mu(\Delta(v_1, \ell^\star)),$$

so voters v_1 and v_2 do not support $\ell^\star$ against r' (they either abstain or switch to r'). Hence r strictly improves his vote margin, contradicting that $(\ell^\star, r^\star)$ is an equilibrium. Consequently, no equilibrium can satisfy $r^\star > 2$.

Since $r^\star > 2$ is impossible, any equilibrium has $r^\star \leq 2$; and given $r^\star < 2$, the configuration $\ell^\star < -1$ is impossible by the deviation argument in Case 1, so no equilibrium places a candidate outside $[-1, 2]$.

□

C APPENDIX: PROOF OF THE NON-EXISTENCE OF MIXED-NASH EQUILIBRIUM IN EXAMPLE 1

Proof. Let σ_r be an arbitrary mixed strategy of the right-hand candidate (player r); denote its distribution function by $F_r(x) = \Pr_{r\sim\sigma_r}[r \le x]$. We proceed in three steps.

Step 1: σ_r puts no mass on $[2,\infty)$. If $\Pr[r \ge 2] > 0$, the left-hand candidate (player ℓ) can deviate to the pure strategy $\ell = -1$. Whenever $r \ge 2$ the two leftmost voters (v_1, v_2) activate and support ℓ, voter v_3 is closer to -1 than to r, and r keeps only v_4, v_5; hence ℓ wins 3:2 with strictly positive probability. Because the maximal payoff in the game is 1, this deviation raises ℓ's expected payoff, contradicting equilibrium. Thus $\Pr[r \ge 2] = 0$.

Step 2: $\sup\operatorname{supp}\sigma_r = 2$. Assume instead that $\bar r := \sup\operatorname{supp}\sigma_r < 2$. Fix $\varepsilon \in (0, 2-\bar r)$ and let $\ell^\varepsilon := \bar r + \varepsilon$ (note $\ell^\varepsilon < 2$). Then for *every* realisation $r \le \bar r$ the distance order at v_4, v_5 is $|2-\ell^\varepsilon| < |2-r|$, so voters v_4, v_5 both choose ℓ. Voter v_3 favours the closer of (ℓ^ε, r), i.e. r, and the two leftmost voters abstain because $r < 2$. Hence ℓ wins 2:1 with probability 1, again contradicting equilibrium. Therefore $\sup\operatorname{supp}\sigma_r = 2$.

Step 3: A profitable deviation for ℓ when $\sup = 2$. Because σ_r has no mass at 2 (by Step 1) but its support reaches all the way up to 2 (by Step 2), for every $n \in \mathbb{N}$ the interval $I_n := [2-\frac{1}{n},\, 2)$ carries strictly positive probability: $F_r(2) - F_r(2-\frac{1}{n}) = 1 - F_r(2-\frac{1}{n}) > 0$.

Let $\ell_n := 2 - \frac{1}{n}$. If ℓ deviates to the pure strategy ℓ_n, they beats the right-hand candidate whenever $r < \ell_n$, because then v_4, v_5 are closer to ℓ_n than to r (whereas the other three voters behave exactly as in Step 2). Hence the winning probability of ℓ under this deviation is

$$\Pr_{r\sim\sigma_r}[r < \ell_n] = F_r\big(2-\tfrac{1}{n}\big).$$

But F_r is right-continuous and $F_r(2) = 1$, so $F_r(2-\frac{1}{n}) \longrightarrow 1$ as $n \to \infty$. Consequently ℓ can secure a payoff arbitrarily close to 1 by choosing n large enough.

Since in any putative equilibrium players must be indifferent among the strategies in the support of their mixes, the original expected payoff to ℓ cannot already be 1; hence there exists an n for which the deviation to ℓ_n is strictly profitable. This contradicts the assumption that (σ_ℓ, σ_r) was an equilibrium.

□

D APPENDIX: FULL PROOF OF LEMMA 2 (UNFLATTEN LEMMA)

Proof. Assume, for contradiction, that such a flattening mapping exists. Without loss of generality, let $\bar c_0 = 0$. Pick two voters: v_1 and $v_{\frac12}$ with $\bar v_1 = 1$ and $v^-_{\frac12} = \frac12$. For simplicity, let $e_1 := d(v_1, c_0)$ and $e_{1/2} := d(v_{\frac12}, c_0)$.

Step 1: View from v_1. By the flattening assumption,

$$u_1(v_{\frac12}) = 1 - \left|\frac{v^-_{\frac12} - \bar v_1}{\bar v_1 - \bar c_0}\right| = 1 - \left|\frac{\frac12 - 1}{1}\right| = \frac12.$$

So,

$$\mu_1\left(1 - \frac{d(v_1, v_{\frac12})}{e_1}\right) = \frac12.$$

Since μ_1 is concave at $\Delta > 0$ and normalized, we must have

$$1 - \frac{d(v_1, v_{\frac12})}{e_1} < \frac12 \quad\Rightarrow\quad d(v_1, v_{\frac12}) > \frac12 e_1.$$

Step 2: View from $v_{\frac12}$. Similarly,

$$u_{\frac12}(v_1) = 1 - \left|\frac{\bar v_1 - v^-_{\frac12}}{v^-_{\frac12} - \bar c_0}\right| = 1 - \left|\frac{1-\frac12}{\frac12}\right| = 0,$$

so

$$\mu_{\frac12}\left(1 - \frac{d(v_{\frac12}, v_1)}{e_{1/2}}\right) = 0.$$

As $\mu_{\frac12}$ is strictly increasing with $\mu_{\frac12}(0) = 0$, we conclude that $d(v_{\frac12}, v_1) = e_{1/2}$.

Step 3: Contradiction from additivity. By the additivity of d and the fact that v_1, $v_{\frac12}$, and c_0 lie on a line, we have

$$e_1 = d(v_1, c_0) = d(v_1, v_{\frac12}) + d(v_{\frac12}, c_0) = d(v_1, v_{\frac12}) + e_{1/2}.$$

Using $d(v_1, v_{\frac12}) = d(v_{\frac12}, v_1) = e_{1/2}$, this implies $e_1 = 2e_{1/2}$, or equivalently,

$$d(v_1, v_{\frac12}) = \frac12 e_1.$$

But this contradicts the inequality derived in Step 1, where we showed $d(v_1, v_{\frac12}) > \frac12 e_1$.

Conclusion: Therefore, our assumption of the existence of a universal flattening mapping must be false. □

E FULL PROOF OF THEOREM 4 (INFINITE INTERVAL THEOREM)

Proof. **Step 0. Change of variables.** For every voter $v < 0$ define

$$x := \Delta(v, r) = 1 - \frac{|v-r|}{|v|}, \qquad y := \Delta(v, \ell) = 1 - \frac{|v-\ell|}{|v|}.$$

With $\ell = 1, r = 2$ we have $y = \frac12 x$ on $v < 0$. Hence, for $x < 0$,

$$S(x) = \mu(x) - \mu(\tfrac12 x) - \kappa.$$

Step 1. Build μ on $[-1, 0]$ with decreasing slopes.

Cut $[-1, 0]$ into the intervals

$$I_n := \Big(\frac{1-2^{n+1}}{2^{n+1}},\ -\frac{1-2^n}{2^n}\Big], \qquad n = 0, 1, 2, \dots$$

so $I_0 = (-\frac12, 0]$, $I_1 = (-\frac34, -\frac12]$, $I_2 = (-\frac78, -\frac34]$, etc. Their lengths sum to $\frac12$.

Assign the tangent

$$t_n := \frac12 + \frac{1}{2^n} \quad (t_0 = 1,\ t_{n+1} < t_n).$$

Define the value function by

$$\mu(0) := 0, \quad \mu(x) := \mu(\inf I_n) + t_n\big(x - \inf I_n\big) \quad (x \in I_n).$$

Because slopes decrease, μ is increasing and convex. The series $\sum_{n\ge0} |I_n| t_n = \sum_{n\ge0} 2^{-n}\big(\frac12 + 2^{-n}\big)$ converges, so the limit $\mu(-1) = \frac{13}{12}$ exists.

Step 2. Copy–shrink blocks left of -1.

Bundle five consecutive intervals into $J_n := I_{5n+1} \cup \cdots \cup I_{5n+5}$ $(n \ge 0)$, and let $\alpha_0^n < \cdots < \alpha_5^n$ be their endpoints.

Define two affine interpolants

$$g_1 :\ (\alpha_0^n, \mu(\alpha_0^n)) \to (\alpha_2^n, \mu(\alpha_2^n)),$$

$$g_2 : (\alpha_3^n, \mu(\alpha_3^n)) \to (\alpha_5^n, \mu(\alpha_5^n)).$$

Both slopes lie in $[\frac{1}{2}, 1]$.

For $x \le -1$ satisfying $x/2 \in J_n$ set

$$\mu(x) := \max\{g_1(x/2), g_2(x/2)\} + \big(\mu(-\tfrac{1}{2}) - \mu(-1)\big).$$

The maximum of two affine maps is convex, so μ remains increasing and convex on $(-\infty, 0]$.

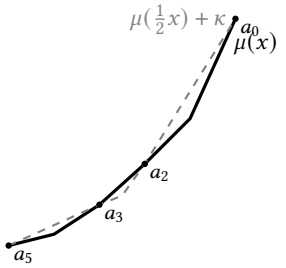

Figure 4: Inside a block J_n the two graphs meet at a_5, a_3, a_2, a_0, creating two disjoint intervals with $S(x) < 0$.

Step 3. Surplus sign pattern and measure bound.

Let $\kappa := \mu(-\frac{1}{2}) - \mu(-1) > 0$. Inside each block J_n, $\mu(x)$ and $\mu(\frac{1}{2}x) + \kappa$ intersect at $\alpha_0^n, \alpha_2^n, \alpha_3^n, \alpha_5^n$, giving

$$S(x) \begin{cases} < 0 & x \in (\alpha_0^n, \alpha_2^n) \cup (\alpha_3^n, \alpha_5^n), \\ > 0 & x \in (\alpha_2^n, \alpha_3^n). \end{cases}$$

Thus each J_n contributes two positive-surplus sub-intervals. Because $|J_n| = 2^{-5n-1}$,

$$\lambda(\rho(\ell, r)) \;\le\; 2 \sum_{n \ge 0} |J_n| \;=\; 2 \sum_{n \ge 0} 2^{-5n-1} \;<\; \infty,$$

while the countably many blocks ensure ρ has infinitely many components.

□

F THE IDENTITY VALUE FUNCTION

In this section, we specialise to the identity value function. Since the full model introduces several new behavioural components at once, analysing the linear case provides a clean baseline: it isolates (PT1) reference dependence and (A1) participation cost while removing curvature-induced non-linearities (PT2). This allows the structure of the voting region to be characterised more transparently before returning to the general case.

If the value function μ is the identity function, then the surplus condition simplifies to: [7]

$$S'(v_i) := S^\ell(v_i) \cdot |v_i| = |v_i - r| - |v_i - \ell| - \kappa_i |v_i|.$$

A voter supports ℓ if $S^\ell(v_i) > 0$, which corresponds to $S'(v_i) > 0$ under this transformation. Unlike the original surplus $S^\ell(v_i)$, which diverges at $v_i = 0$, the simplified form $S'(v_i)$ remains finite and well-defined. At $v_i = 0$, we get $S'(0) = |r| - |\ell|$, so the model predicts support for ℓ precisely when $|r| - |\ell| > 0$, consistent with the limiting behaviour of the original surplus.

[7] We use $S'(v_i)$ to denote the surplus under the identity value function. The superscript ℓ is omitted since we fix the comparison against candidate ℓ.

Candidate positions fall into three configurations:

$$\text{(I) } 0 < \ell < r, \quad \text{(II) } \ell < 0 < r, \quad \text{(III) } \ell < r < 0.$$

For each configuration, the function $S'(v)$ is piecewise linear with breakpoints at $\{\ell, 0, r\}$, dividing the real line into four affine segments, denoted (S'_1, S'_2, S'_3, S'_4). Of the twelve possible orderings of the breakpoints (three configurations, each with four segment placements), only nine lead to non-trivial support regions. When $v > \max\{r, 0\}$, we always have $S'(v) < 0$, so the outermost segment S'_4 does not contribute. This is because $|v - r| < |v - \ell|$, so the net surplus becomes negative even before subtracting the cost term $\kappa|v|$.

For $v \le \min\{\ell, 0\}$, the function $S'(v)$ increases. Since $S'(v)$ is piecewise affine with breakpoints at $\{\ell, 0, r\}$, we define its *slope pattern* as the sequence of monotonicities over the four resulting segments. The global slope pattern thus takes the form

$$(\nearrow, ?, ?, \searrow),$$

where the inner slopes depend on ℓ, r, κ. A nonempty solution set can form multiple regions only if the pattern

$$(\nearrow, \searrow, \nearrow, \searrow)$$

occurs. Even in that case, since $S'_4(v) < 0$ strictly, any second local maximum is cut off. Therefore, if $S'(v) > 0$ holds at all, it defines a *single open interval.*

F.0.1 Dynamic Implications in Vote-Share Optimisation. Assume $\rho(\ell, r) = (\underline{\rho}, \overline{\rho})$. We analyse the location of $(\underline{\rho}, \overline{\rho})$ and the length of the interval $\rho(\ell, r)$ based on the relative positions of $(\ell, 0, r)$, as categorised in Table 1 and Table 2 (Appendix G).

Assuming a uniform voter distribution over $(-M, M)$, i.e., $f(v) = 1$, we have:

$$\text{score}(\ell) = \min(\overline{\rho}, M) - \max(\underline{\rho}, -M).$$

As $\ell \to r$ from the far left and $\kappa < 1$, both endpoints increase, but the interval length shrinks. Consequently, score(ℓ) decreases, and the optimal strategy is to move ℓ away from the centre.

If $1 \le \kappa < 2$, a tipping point emerges at $\ell = -r/(\kappa - 1)$, beyond which $\rho(\ell, r)$ becomes empty. Once ℓ decreases past this threshold (i.e., moves further away from the origin), the region $\rho(\ell, r)$ reappears but continues to shrink. Hence, optimality still favours a more extreme position.

For $\kappa \ge 2$, the vote share function score(ℓ) becomes unimodal: it increases up to a tipping point and then declines. The maximum is attained precisely at $\ell = \frac{-r}{\kappa - 1}$, which balances the directional advantage with cost sensitivity. Unlike in the low-κ regime, pushing ℓ further from the centre eventually backfires due to the growing cost penalty.

Now assume the voter distribution has a differentiable CDF F. Then,

$$\text{score}(\ell) = F(\overline{\rho}) - F(\underline{\rho}) \tag{7}$$

$$\frac{\partial \text{score}}{\partial l} = F'(\overline{\rho}) \cdot \overline{\rho}' - F'(\underline{\rho}) \cdot \underline{\rho}'.$$

If an interior optimum exists with $-r < \ell < 0$, a necessary condition of this location is:

$$\frac{F'(\overline{\rho}(\ell))}{F'(\underline{\rho}(\ell))} = \frac{2 + \kappa}{\kappa}.$$

Thus, higher central density F' mitigates polarisation even when κ is small.

F.0.2 Refined Victory Model. In the two-candidate model, the vote margin is defined as $\mathsf{vm}(\ell, r) = \mathrm{score}(\ell, r) - \mathrm{score}(r, \ell)$. Assume r is fixed; then both $\mathrm{score}(\ell, r)$ and $\mathrm{score}(r, \ell)$ depend on ℓ. Let $\rho(\ell, r) = (\underline{\rho}^{\ell}, \overline{\rho}^{\ell})$ and $\rho(r, \ell) = (\underline{\rho}^{r}, \overline{\rho}^{r})$.

Under a uniform voter distribution f, when all boundary points of $\rho(\ell, r)$ and $\rho(r, \ell)$ lie within the support of f, we have:

$$\mathsf{vm}(l, r) = (\overline{\rho}^{\ell} - \underline{\rho}^{\ell}) - (\overline{\rho}^{r} - \underline{\rho}^{r}).$$

As $\ell \to r$ with $\kappa < 1$, $\mathsf{vm}(\ell, r)$ increases except within the interval $(r(1-\kappa), r)$, where $\mathrm{score}(r, \ell) = 0$. Thus, a margin-maximising candidate prefers to move beyond the centre.

Now suppose r also seeks to maximise the margin. When $0 < \ell < r(1-\kappa)$, we obtain:

$$\mathsf{vm}(r, \ell) = \frac{(\kappa+1)r + 3\ell}{(\kappa-2)(\kappa+1)},$$

and the first-order condition with respect to r is negative. Hence, both candidates converge to the same location, and the resulting vote margin is zero, resulting in a trivial equilibrium.

More generally, let $\mathrm{score}(\ell, r) = F(\overline{\rho}(\ell, r)) - F(\underline{\rho}(\ell, r))$, where we now treat the previously defined endpoints $\underline{\rho}^{\ell}, \overline{\rho}^{\ell}$ as functions of both ℓ and r. The ρ-bounds are piecewise linear, and F is differentiable.

Then $\mathsf{vm}(\ell, r)$ is bounded, continuous, and piecewise differentiable over $[-M, M]^2$. The domain admits a finite partition such that vm is smooth on each region. Interior equilibria satisfy the conditions:

$$\frac{\partial \mathsf{vm}}{\partial \ell} = 0, \quad \frac{\partial \mathsf{vm}}{\partial r} = 0.$$

A pure-strategy equilibrium occurs when:

$$\mathsf{vm}(\ell^*, r) \le 0 \quad \text{and} \quad \mathsf{vm}(\ell, r^*) \ge 0 \quad \forall \ell, r.$$

Although Sion's minimax theorem [29] does not apply due to non-convexity, the antisymmetry $\mathsf{vm}(\ell, r) = -\mathsf{vm}(r, \ell)$ and the identity $\mathsf{vm}(\ell, \ell) = 0$ make the diagonal $\ell = r$ a natural candidate for equilibrium. Whether this holds depends on the gradient directions and possible discontinuities, which could be analysed region by region.

G APPENDIX: TABLE OF CANDIDATES' ρ UNDER LINEARISED μ

To characterise the solution set, we examine where the segments $S_1'(v), S_2'(v), S_3'(v)$ intersect zero. Define:

- $\Lambda := r - \ell > 0$ — distance between the candidates,
- $\xi := \ell + r$ — sum of the anchors,
- $v_1 := -\dfrac{\Lambda}{\kappa}$ — root of $S_1'(v) = 0$, if it exists,
- $v_2 := \begin{cases} -v_1, & \ell > 0, \\ \dfrac{\xi}{2-\kappa}, & \text{otherwise} \end{cases}$ — root of $S_2'(v) = 0$, if it exists,
- $v_3 := \dfrac{\xi}{2+\kappa}$ — root of $S_3'(v) = 0$, if $\xi > 0$.

Table 1 summarises the shape of $\rho(\ell, r)$ under different relative locations of $(\ell, 0, r)$ and κ

Remarks.

- The inequality has no solution if κ lies outside the listed range — define this null case as (0).
- All intervals are open; endpoints are excluded due to the strict inequality.
- When $\ell < 0 < r$, we have $\Lambda = |r| + |\ell|$, so
$$\frac{\Lambda}{|\ell|} = 1 + \frac{r}{|\ell|} > 1, \quad \text{and} \quad \frac{\Lambda}{|\ell|} = 2 \text{ when } \ell = -r.$$
This ratio diverges as $|\ell| \to 0$, suggesting that for $\kappa < 1$, voters at $v = \ell$ will vote.
- When $\ell > 0$, $\frac{\Lambda}{|\ell|} = \frac{r}{\ell} - 1 \to 0$ as $\ell \to r$.

Suppose ℓ moves from far left to r, the region will progress according to the κ. The breakdown is:

- $\kappa < 1$: the region progresses as: (3) → (4) → (1) → (2).
- $1 \le \kappa < 2$: Similar to $\kappa < 1$, except that ℓ may pass through a null zone ((0)) before entering (3).
- $\kappa \ge 2$: (0) → (5) → (4) → (1) → (2),

Similarly to Table 1, we may observe the dynamic geometric change of $\rho(r, \ell)$ as ℓ move, which is presented in Table 2:

- As $\ell \to r$ with $\kappa < 1$, $\rho(r, \ell)$ moves through: (b) → (d) → (a) → (0). The full joint evolution is: ((3), (b)) → ((4), (d)) → ((1), (a)) → ((1), (0)) → ((2), (0)).

H APPENDIX: THEOREMS AND ALGORITHMS FOR TWO CANDIDATES CASE

H.1 Full proof of Lemma 3

Proof. We fix v, r, and $\kappa > 0$, and analyse the shape of $\gamma_r(v)$ and $\gamma_r'(v)$ as subsets of $\ell \in \mathbb{R}$.

Analysis of $\gamma_r(v)$. By definition,

$$\gamma_r(v) = \{\, \ell : \mu(\Delta(v, \ell)) > \mu(\Delta(v, r)) + \kappa \,\}.$$

Let $C := \mu(\Delta(v, r)) + \kappa$, which is a constant. Since μ is strictly increasing and invertible on its domain,

$$\gamma_r(v) = \{\, l : \Delta(v, \ell) > \mu^{-1}(C) \,\} = \{\, l : \Delta(v, \ell) > C' \,\},$$

where $C' := \mu^{-1}(C)$. Recall that $\Delta(v, \ell) = 1 - \left|1 - \frac{\ell}{v}\right|$, defined for $v \neq 0$.

Solving the inequality $1 - \left|1 - \frac{\ell}{v}\right| > C'$, we obtain:

$$\left|1 - \frac{\ell}{v}\right| < 1 - C' \quad \Longleftrightarrow \quad -(1 - C') < 1 - \frac{\ell}{v} < (1 - C').$$

Rewriting this yields:

$$(1 - (1 - C'))v < \ell < (1 + (1 - C'))v,$$

i.e.

$$C''v < \ell < (2 - C'')v$$

when $v > 0$ (and flip the inequality if $v < 0$), where $C'' := 1 - C'$. Since $C' \in (0, 1)$, we have $C'' \in (0, 1)$, and thus $\gamma_r(v)$ is a single open interval (or empty, if the inequality has no solution).

Since $C' = \mu^{-1}\big(\mu(\Delta(v, r)) + \kappa\big)$, a function of r, and C'' is a function of C' we know the endpoints $C''v$ and $(2 - C'')v$ are also functions of r.

Geometry	Valid κ range	shape of $\rho(\ell, r)$	Size of $\rho(\ell, r)$.
$0 < \ell < r$	$0 < \kappa < \frac{\Lambda}{\lvert\ell\rvert}$ $\kappa \geq \frac{\Lambda}{\lvert\ell\rvert}$	$v \in (v_1,\ v_3)$ (1) $v \in (v_1,\ v_2)$ (2)	$\frac{(2+2\kappa)r-2\ell}{(2\kappa+2)\kappa}$ $\frac{2(r-\ell)}{\kappa}$
$\ell < 0 < r$ and $\lvert\ell\rvert \geq r$ ($\xi \leq 0$)	$0 < \kappa < \frac{\Lambda}{\lvert\ell\rvert}$	$v \in (v_1,\ v_2)$ (3)	$\frac{2r+(2\kappa-2)\ell}{(2-\kappa)\kappa}$
$\ell < 0 < r$ and $\lvert\ell\rvert < r$ ($\xi > 0$)	$0 < \kappa \leq \frac{\Lambda}{\lvert\ell\rvert}$ $\kappa > \frac{\Lambda}{\lvert\ell\rvert}$	$v \in (v_1,\ v_3)$ (4) $v \in (v_2,\ v_3)$ (5)	$\frac{(2+2\kappa)r-2\ell}{(2\kappa+2)\kappa}$ $\frac{-2\kappa(r+\ell)}{(2+\kappa)(2-\kappa)}$
$\ell < r < 0$	$0 < \kappa < \frac{\Lambda}{\lvert\ell\rvert}$	$v \in (v_1,\ v_2)$ (6)	$\frac{2r+(2\kappa-2)\ell}{(2-\kappa)\kappa}$

Table 1: Summary of geometric cases under linearised utility function: for each configuration of candidate positions, valid values of κ, the shape of the voting region $\rho^C(l)$, and its size are shown.

Geometry	Valid κ range	Shape of $\rho(r, \ell)$	Size of $\rho(r, \ell)$
$0 < \ell < r$	$0 < \kappa < 1 - \frac{\ell}{r}$	$v \in (v_2,\ -v_1)$ (a)	$\frac{(2-2\kappa)r-2\ell}{(2-\kappa)\kappa}$
$\ell < 0 < r$ and $\lvert\ell\rvert \geq r$ ($\xi \leq 0$)	$0 < \kappa \leq 1 + \lvert\frac{\ell}{r}\rvert$ $\kappa > 1 + \lvert\frac{\ell}{r}\rvert$	$v \in (v_3,\ -v_1)$ (b) $v \in (v_3,\ v_2)$ (c)	$\frac{2r-(2+2\kappa)\ell}{(2\kappa+2)\kappa}$ $\frac{2\kappa(r+\ell)}{(2+\kappa)(2-\kappa)}$
$\ell < 0 < r$ and $\lvert\ell\rvert < r$ ($\xi > 0$)	$0 < \kappa < 1 + \lvert\frac{\ell}{r}\rvert$	$v \in (v_2,\ -v_1)$ (d)	$\frac{(2-2\kappa)r-2\ell}{(2-\kappa)\kappa}$
$\ell < r < 0$	$0 < \kappa < \frac{\ell}{r} - 1$ $\kappa \geq \frac{\ell}{r} - 1$	$v \in (v_3,\ -v_1)$ (e) $v \in (-v_2,\ -v_1)$ (f)	$\frac{2r-(2+2\kappa)\ell}{(2\kappa+2)\kappa}$ $\frac{2(r-\ell)}{\kappa}$

Table 2: Analogous to Table 1, for each candidate configuration, we list valid ranges for κ, the form of $\rho(r, \ell)$ with respect to ℓ, and its size.

Analysis of $\gamma'_r(v)$. Similarly, $\gamma'_r(v) = \{\, \ell : \mu(\Delta(v, r)) > \mu(\Delta(v, \ell)) + \kappa \,\}$. Define $\hat{C} := \mu(\Delta(v, r)) - \kappa$, and let $\hat{C}' := \mu^{-1}(\hat{C})$. Then

$$\gamma'_r(v) = \{\, \ell : \Delta(v, \ell) < \hat{C}' \,\}.$$

Again, $\Delta(v, \ell) = 1 - \left|1 - \frac{\ell}{v}\right|$. Solving the inequality:

$$1 - \left|1 - \frac{\ell}{v}\right| < \hat{C}' \quad \Longleftrightarrow \quad \left|1 - \frac{\ell}{v}\right| > 1 - \hat{C}'.$$

When $v > 0$ (flip the inequality if $v < 0$), this yields:

$$1 - \frac{\ell}{v} < -(1 - \hat{C}') \quad \text{or} \quad 1 - \frac{\ell}{v} > (1 - \hat{C}'),$$

i.e.,

$$\frac{\ell}{v} > 1 + (1 - \hat{C}') = 2 - \hat{C}' \quad \text{or} \quad \frac{\ell}{v} < 1 - (1 - \hat{C}') = \hat{C}'.$$

Hence:

$$\gamma'_r(v) = \{\, \ell : \ell < \hat{C}'v \,\} \cup \{\, \ell : \ell > (2 - \hat{C}')v \,\},$$

which is always the union of two open intervals when $\hat{C}' \in (0, 1)$. This holds because $\mu(\Delta(v, r)) < 1, \kappa > 0$, so $\hat{C} < 1$, implying $\hat{C}' < 1$.

Similarly, because $\hat{C} = \mu^{-1}\big(\mu(\Delta(v, r)) - \kappa\big)$ is a function of r, we know $\hat{C}'v$ and $(2 - \hat{C}')v$ are also functions of r.

□

H.2 Algorithms to Find Best Response and Equilibria

In Algorithm 1 and Algorithm 2, We assume it takes $O(1)$ time to find the $\gamma_a(r)$, which is the function of boundaries to some γ_i γ_i'. For two different curves γ_a, γ_b, we also assume it takes $O(1)$ times to find all the intersection. That is, we also assume the amount of intersection of two curves are finite and bounded.

The core idea of Algorithm 1 is to compute the best response of candidate ℓ for a fixed opponent position r. Given the value function and reference-dependent structure, each voter defines a preference region on the real line where their vote shifts between candidates or abstention. Since both v and r are fixed, the boundaries $\gamma_v(r)$ and $\gamma_v'(r)$ for each voter are determined.

Each voter contributes at most two transition points—switches in their vote direction—leading to a total of $O(n)$ critical points for candidate ℓ. Sorting these boundaries allows us to scan the line and update the vote margin incrementally. Because the model is one-dimensional and transitions occur at isolated thresholds, the number of segments with constant score is linear. Thus, the algorithm maintains a $\tilde{O}(n)$ complexity to find score $\mathsf{vm}(l)$ across intervals and returns the position ℓ^* with the highest vote margin as we only need to sort $O(n)$ breakpoints and perform a linear sweep.

Algorithm 1 Best response of candidate ℓ given fixed r

Input: Candidate location r, voter set $\{(v_i, \kappa_i, \mu_i)\}_{i=1}^n$
Output: Location $l^\star$ maximising vote margin $\mathsf{vm}(\ell)$

Initialise empty list of breakpoints P
for each voter v_i **do**
  Compute endpoints $\underline{\gamma}_i(r), \overline{\gamma}_i(r), \underline{\gamma}_i'(r), \overline{\gamma}_i'(r)$
  Add all four points to P with tags: oz, uz, oz', uz'
end for
Sort P by increasing position
Initialise current margin $M \leftarrow$ total contribution in region $(-\infty, P[1])$
Initialise $M_{\max} \leftarrow M$, $\ell^\star \leftarrow$ midpoint in first interval
for each breakpoint $p_i \in P$ in order **do**
  if p_i is oz or oz' **then**
    $M \leftarrow M + 1$
  else if p_i is uz or uz' **then**
    $M \leftarrow M - 1$
  end if
  if $M > M_{\max}$ **then**
    $M_{\max} \leftarrow M$, $\ell^\star \leftarrow$ midpoint in (p_i, p_{i+1})
  end if
end for
return $\ell^\star$

Algorithm 2 can be viewed as a higher-level procedure built on top of Algorithm 1. While Algorithm 1 computes the best response for a fixed opponent position, Algorithm 2 searches for a pure-strategy Nash equilibrium by identifying mutual best responses. For each fixed r, we use Algorithm 1 to determine the interval(s) of optimal ℓ, where the vote margin $\mathsf{vm}(\ell)$ is maximised. Repeating this for all r partitions the domain into curved regions of the form $\{(\ell, r) \mid r \in (p_i, p_j), \ell \in [\gamma_k(r), \gamma_k'(r)]\}$, representing best-response sets. A symmetric process constructs best-response regions for r against fixed ℓ.

A pure-strategy equilibrium occurs when these regions intersect. Assuming each pair of boundary curves intersects $O(1)$ times, the total number of such regions is $\tilde{O}(n^2)$. Testing all pairs of regions for overlap yields a worst-case complexity of $\tilde{O}(n^4)$, which remains tractable for moderate n. Although the curves involved can be complex in general, empirically reasonable assumptions on the value function μ (e.g., monotonicity, bounded curvature) ensure that the geometric structure remains manageable.

Algorithm 2 Min-Max Strategy Sweep over r

Input: A collection of curves $\overline{\gamma}_i(r), \underline{\gamma}_i(r), \overline{\gamma}_i'(r), \underline{\gamma}_i'(r)$ for all $i = 1, \ldots, m$
Output: Intervals (p_k, p_{k+1}) that minimise $\max_\ell \mathsf{vm}(\ell, r)$

Identify all pairwise intersection points among the $\overline{\gamma}$ and $\underline{\gamma}$ curves
  Total number of curves: m, number of intersections: $n = O(m^2)$
Sort all intersection points in increasing order: $p_1 < p_2 < \cdots < p_n$
Define $p_0 := p_1 - \varepsilon$ for small $\varepsilon > 0$
Compute the value of all curves at $r = p_0$
Sort curves vertically by their values at p_0 to initialise ordering
Initialise Table:
  Columns = intervals (p_k, p_{k+1}) for $k = 0, \ldots, n-1$
  Rows = ordered list of curves at each interval
  Between each pair of adjacent curves, define a voter at position ℓ
  Compute $\mathsf{vm}(\ell, r)$ at representative $r \in (p_0, p_1)$
for $k = 1$ to n **do**
  Let p_k be the point where a pair of curves (c_1, c_2) intersect
  Create new column for interval (p_k, p_{k+1})
  Update curve order by swapping c_1 and c_2
  for each voter position l between two adjacent curves **do**
    Let $\mathsf{vm}_{\text{prev}}(\ell, r)$ be from previous interval
    if $(c_1, c_2) = (\underline{\gamma}_i, \overline{\gamma}_i)$ **(disappearance of γ_i) then**
      $\mathsf{vm}(\ell, r) \leftarrow \mathsf{vm}_{\text{prev}}(\ell, r) - 1$
    else if $(c_1, c_2) = (\overline{\gamma}_i, \underline{\gamma}_i)$ **(re-emergence of γ_i) then**
      $\mathsf{vm}(\ell, r) \leftarrow \mathsf{vm}_{\text{prev}}(\ell, r) + 1$
    else
      Apply ± 1 adjustment to $\mathsf{vm}(\ell, r)$ depending on: Whether $\overline{\gamma}'$ moves above $\overline{\gamma}$ or $\underline{\gamma}$ above $\underline{\gamma}'$
    end if
  end for
end for
For each interval (p_k, p_{k+1}), compute $\max_\ell \mathsf{vm}(\ell, r)$
Identify all intervals minimising this maximum
return Set of optimal intervals (p_k, p_{k+1}) and corresponding min-max value

I FULL PROOF OF THEOREM 5 (NP-COMPLETENESS)

Proof. **Membership in NP.** Given a candidate location ℓ (with polynomial bit-length), we can evaluate the score via the arithmetic circuit for F, so the inequality can be verified in polynomial time.

NP-**hardness via Partition.** Let $S = \{a_1, \ldots, a_n\} \subset \mathbb{N}$ with total $T := \sum a_i$, and set $M := T$.

We define a population density function f on the binary grid of step $2^{-(n+1)}$ inside $(-1.5, 1.5)$. For any $x = 0.x_1x_2 \ldots x_n$ (binary), define:

$$f(x) = \left(\frac{M}{2}\right)^2 - \left(\sum a_i x_i - \frac{M}{2}\right)^2.$$

All other points have zero density. The maximum value $f(x) = M^2/4$ occurs if and only if $\sum a_i x_i = M/2$, i.e., a perfect partition.

Let $F(v) = \sum_{x \le v} f(x)$ be the CDF. The circuit construction for F and a smoothed version are detailed in the later subsections.

We shift the distribution: define $g(x) := f(x + 1.5)$, $G(x) := F(x + 1.5)$. Let $r := 2^{-(n+2)}$, $\kappa := 1$, and $n' := M^2/4$. For any $\ell < -r$, the window $(\ell - r, \ell + r)$ captures exactly one grid point $x^\star$, giving:

$$G(\ell + r) - G(\ell - r) = g(x^\star).$$

This quantity equals n' if and only if the bitstring encoded by $x^\star$ represents a perfect partition of S. Thus, we reduce Partition to our decision problem in polynomial time.

Conclusion. The problem is both in NP and NP-hard, hence NP-complete. □

I.1 Arithmetic Circuit for the CDF

To compute the CDF $F(v) = \sum_{x \le v} f(x)$, we treat $x = 0.x_1x_2 \cdots x_n$ as a binary string of n bits. The function $f(x)$ is a degree-2 polynomial:

$$f(x) = \left(\frac{M}{2}\right)^2 - \left(\sum_{i=1}^{n} a_i x_i - \frac{M}{2}\right)^2.$$

Expanding the square, we obtain:

$$f(x) = \sum_{i,j} a_i a_j x_i x_j - M \sum_i a_i x_i + \frac{M^2}{4}.$$

To evaluate the CDF $F(v) = \sum_{x \le v} f(x)$, we perform an n-fold nested symbolic summation over the bits $x_1, \ldots, x_n$, conditioned on the input prefix $v = 0.v_1v_2 \ldots v_n$. At each level k, the summation range for x_k depends on whether the current prefix $x_1x_2 \ldots x_{k-1}$ is lexicographically less than, equal to, or greater than that of v.

Let $\pi_k \in \{0, 1\}$ be the admissible upper limit for x_k given the prefix comparison. We apply the following symbolic substitutions during summation:

$$\sum_{x_k=0}^{\pi_k} x_k = \pi_k, \qquad \sum_{x_k=0}^{\pi_k} x_k^2 = \pi_k,$$

$$\sum_{x_k=0}^{\pi_k} x_k x_j = \begin{cases} x_j \cdot \pi_k & \text{if } j < k, \\ 0 & \text{if } j > k. \end{cases}$$

Each substitution is linear in the number of monomials present. Since $f(x)$ begins with $O(n^2)$ monomials, and each of the n summation layers multiplies this by at most $O(n)$, the final symbolic form of $F(v)$ contains $O(n^3)$ monomials.

Each resulting monomial is of the form $c \cdot v_{i_1} \cdots v_{i_m}$ for $m \le 3$, where the coefficient c is a product of the a_i and constants bounded by $O(M^3)$. Every monomial can be implemented by a constant-sized sub-circuit (two multiplies and one addition). Therefore, the full circuit has size $O(n^4)$ and all coefficients have bit-length $O(\log M)$, which is polynomial in the input size.

I.2 Continuous Approximation of the CDF

To replace discrete point masses by continuous density, define $\delta := 2^{-(n+4)}$, and let $\mathcal{G} := \{x^\star = k\delta \mid k \in \mathbb{Z}\}$ be the grid. Define a "top-hat" approximation:

$$f_{\text{cont}}(x) := \begin{cases} \frac{f(x^\star)}{2\delta}, & \text{if } |x - x^\star| \le \delta, \\ 0, & \text{otherwise,} \end{cases}.$$

with $x^\star := \delta \left\lfloor \frac{x}{\delta} + \frac{1}{2} \right\rfloor$

This spreads each point mass into a short plateau. The corresponding CDF becomes:

$$F_{\text{cont}}(v) = F(v) + \mathbf{1}_{|v - x^\star(v)| \le \delta} \cdot \frac{f(x^\star)}{2\delta}(v - x^\star + \delta),$$

which differs from F only in a localised linear segment. The construction preserves all asymptotic bounds: computing $x^\star(v)$ takes constant time, and the additional term requires only a constant-sized arithmetic circuit gadget.

J DETAILED PROOF OF THEOREM 6 (#P-HARDNESS)

Proof. According to Kawamura [16, Section 4.3.2], integration of polynomial-time computable functions is #P-hard. We use this fact and prove that the problem in Theorem 6 is #P hard.

Given $g : [0, 1] \to \mathbb{R}$ (polynomial-time computable), write $g = g_+ - g_-$ with $g_\pm(x) = \max\{\pm g(x), 0\}$. Let $B := \sup_x |g(x)| \le 2^{p(|g|)}$ for some poly p. Embed $g_\pm$ into the voting model by choosing $\bar{\ell} < -r$ so the window $(\bar{\ell} - r, \bar{\ell} + r)$ equals $[0, 1]$, and ask the oracle twice:

$$(f, g_+,\ n_+ := t,\ \varepsilon) \quad (f, g_-,\ n_- := t,\ \varepsilon), \qquad \varepsilon := 2^{-(n+2)}.$$

A binary search with $O(n)$ queries yields $\tilde{I}_\pm$ such that $|\int_0^1 g_\pm - \tilde{I}_\pm| \le \varepsilon$. Set $I^* := \tilde{I}_+ - \tilde{I}_-$; then $|I^* - \int_0^1 g| \le 2\varepsilon < 2^{-n}$. Hence $|\int_0^1 g - q| < 2^{-n}$ iff $|I^* - q| < 2^{-n}$, giving a polynomial-time Turing reduction. Therefore, the decision problem is #P-hard. □

K EXTENSION TO MULTIPLE CANDIDATES

While this paper focuses on two-candidate settings, the underlying decision model naturally generalises to elections with more candidates. We briefly sketch the structural and computational aspects in the discrete case.

K.1 Best Response under Victory Margin Objective

Consider a new candidate $\hat{c}$ entering a race with existing candidates $c_1, c_2, \ldots, c_m$. Under the minimax Regret-inspired strategy, each voter v_i identifies their top two candidates c_1^i, c_2^i—those closest to their position.

The decision to vote for $\hat{c}$ depends on comparisons against these favourites:

- If $u_i(\hat{c}) > u_i(c_1^i) + \kappa_i$, then v_i switches to $\hat{c}$.

- If $u_i(\hat{c}) < u_i(c_1^i)$ but $u_i(c_1^i) < u_i(\hat{c}) + \kappa_i$, then v_i abstains.
- If v_i initially supports c_1^i and $u_i(\hat{c}) < u_i(c_2^i)$, then $\hat{c}$ has no effect on their behaviour.

At last, we consider whether $\hat{c}$ could indirectly activate an abstaining voter to support some $c_k \neq \hat{c}$. This seems unlikely: abstention occurs when all candidates fall below the threshold, and introducing $\hat{c}$ does not elevate others above it.

Computing Victory Margins. As in the two-candidate case, we define threshold curves $\gamma_i(\hat{c})$ and $\gamma_i'(\hat{c})$ to describe voter behaviour:

- $\gamma_i(\hat{c})$: region of voters who vote for $\hat{c}$,
- $\gamma_i'(\hat{c})$: region of voters who vote for c_1^i (which is not $\hat{c}$)
- Otherwise, abstain.

However, to compute the vm, we must consider how $\hat{c}$'s entry affects the standing of the **current winner** c_w. Let us define the possible triplets:

- If v_i switches from voting for c_w to $\hat{c}$: gain of 2 (from opponent to self),
- If v_i switches from abstaining or from non-winning $c_k \neq c_w$ to $\hat{c}$: gain of 1,
- Otherwise: no effect.

This approach enables computing a **best response** location for candidate $\hat{c}$, based on the impact on net margin. Assuming fixed voter positions and known κ_i, this process appears tractable.

K.2 Existence of Equilibrium

Classical spatial models often lack equilibria with more than two candidates. In contrast, our prospect-theoretic framework can admit pure-strategy Nash equilibria even in high-candidate settings, due to the discontinuous structure of voter response under loss aversion.

For example, let n voters be located at positions $1, 2, 4, \ldots, 2^n$, each with loss aversion threshold $\kappa_i = 0.49$. Suppose n candidates are positioned exactly at the voter locations. Then no candidate has incentive to deviate: moving away would reduce their support, since no voter would switch to a candidate more than κ_i worse than their current best. Hence, this profile forms a Nash equilibrium.

This suggests that the key challenge is not existence but characterisation: how to describe or compute such equilibria efficiently as the candidate profile space grows in dimension. While exhaustive search is infeasible for large n or m, one promising direction is to exploit the geometric structure of threshold boundaries to guide search or prune the space.

Let $\mathbf{c} := (c_1, c_2, \ldots, c_m) \in \mathbb{R}^m$ denote the vector of candidate positions.

For each voter v_i and each candidate c_j, define a utility dominance region:

$$\gamma_{i,j} := \left\{ \mathbf{c} \in \mathbb{R}^m \;\middle|\; u_i(c_j) > u_i(c_k) + \kappa_i \quad \forall k \neq j \right\},$$

where $\kappa_i \geq 0$ is a loss aversion threshold.

Each voter induces up to m such regions, and the equilibrium structure depends on their intersections.

We then define a voting outcome function:

$$\mathsf{H} : \mathbb{R}^m \to \mathbb{N}^m, \quad \mathsf{H}(\mathbf{c}) := (\mathsf{score}(c_1), \ldots, \mathsf{score}(c_m)),$$

This setup works for any candidate goal—vote share, margin, or just winning—since all of them depend only on $\mathsf{H}(\mathbf{c})$.

The regions $\gamma_{i,j}$ partition $\mathbb{R}^m$ into combinatorially distinct zones. If each voter's preference boundaries intersect in a bounded and finite way, the total number of such regions is at most $O(n^m)$.

This reduces the equilibrium problem to identifying the distinct regions induced by the $\gamma_{i,j}$, evaluating outcome scores and candidate utilities within each, and selecting profiles $\mathbf{c}$ that constitute mutual best responses.

Recursive Tracking over Candidate Profile Space. We study how candidate positions influence outcomes by recursively partitioning the profile space $\mathbb{R}^m$ using preference boundaries $\gamma_{i,j}$. This generalises topological sweep from 2D to higher dimensions, where each $\gamma_{i,j}$ forms a codimension-1 manifold.

Base case ($m = 2$). Let $\mathbf{c} = (l, r) \in \mathbb{R}^2$. Each voter defines up to four boundary curves in the (l, r)-plane, partitioning it into behavioural zones. Sweeping r from left to right, we update intersections to track vote margins and detect equilibria (Section 5).

Recursive step ($m \to m + 1$). Assume boundaries $\mathcal{B}_1, \ldots, \mathcal{B}_s \subseteq \mathbb{R}^m$ have been identified for fixed $(c_1, \ldots, c_m)$. Introducing a new candidate $c_{m+1} \in \mathbb{R}$ lifts each region $\gamma_{i,j} \subseteq \mathbb{R}^{m+1}$, extending the partition structure. Fixing $c_{m+1} = z$ yields a slice $\mathbb{R}^m \times \{z\}$, where the boundaries become $\mathcal{B}_1(z), \ldots, \mathcal{B}_s(z) \subseteq \mathbb{R}^m$.

The algorithm recursively sweeps z from left to right, tracking combinatorial changes in intersection structure. At each step, we: (1) compute features at $z = z_0$; (2) update intersections as z varies; (3) detect event points where the structure changes; and (4) preserve the invariant $\mathsf{H}(\mathbf{c})$ between events. The sweep terminates as $z \to \infty$.

For example, when $m = 2$, fixing c_2 gives curves $\vartheta_i(c_1)$; when $m = 3$, fixing c_3 yields curves defined by intersections $\mathcal{B}_i \times \mathcal{B}_j \subseteq (c_1, c_2)$. In general, each level tracks pairwise combinations of features from the previous level.

Complexity. If each boundary consists of s components, the number of tracked features grows as:

$$s, \quad s^2, \quad s^4, \quad \ldots, \quad s^{2^{m-1}}.$$

This doubly exponential growth limits scalability but captures rich combinatorial structure.

This framework enables detection of constant-score boundaries, identification of mutual best response zones, and reduction of equilibrium search to a finite—though large—set of combinatorial types.

Computational Hardness Hypothesis. We hypothesise that deciding whether a Nash equilibrium exists in the multi-candidate setting is NP-hard in the number of candidates m.

The motivation is the combinatorial explosion induced by pairwise preference boundaries $\gamma_{i,j}$. Each candidate profile defines exponentially many behavioural regions, and verifying mutual best responses across them is computationally intensive. Even with linear utilities, fixed thresholds, and discretised positions, the search space grows as s^{2^m} for grid size s, making naive enumeration infeasible. A formal hardness proof would likely require a reduction from a known NP-complete problem in geometric configuration or preference aggregation.