

An Enriched Model of Strategic Voting under Uncertainty

Henri Surugue¹ and Sébastien Destercke²

¹ LIP6 - Nukkai, Sorbonne University, Paris, France

`henri.surugue@lip6.fr`

² CNRS - Heudiasyc, UTC, Compiègne, France

`sebastien.destercke@cnrs.fr`

Abstract. We present a new strategic voting model where we use uncertainty representation to model preferences. Specifically, we use probability sets as uncertainty representations, together with lower and upper expected utility gains to take strategic decisions. Focusing on belief functions in particular, we demonstrate that this very expressive model includes in one sweep many existing models based on probabilities, sets or incomplete preferences. Additionally, we generalize several well-known convergence results from the literature to this broader representational setting. Furthermore, we illustrate how this model can capture more realistic scenarios for practical applications but also raises theoretical challenges.

Keywords: Preference uncertainty · Belief functions · Strategic voting

1 Introduction

Strategic voting [20] arises when voters have an incentive to misreport their true preferences to obtain a more desirable result. The Gibbard-Satterthwaite theorem [29] shows that no reasonable voting rule is immune to strategic manipulation. This makes it desirable to formalise the notion of strategic votes and investigate their consequences, first in a single-step, single-voter setting and then in multi-steps, iterative settings. This can be done in various ways: formulating the problem as a strategic game [23] or as a Bayesian decision problem [26, 12]. These approaches raise the question of how voters act strategically as well as the role of polls, and show that understanding how information impacts voters is crucial [9, 28].

Several reasons justify the need for an uncertainty-based approach to polls, stemming either from the polling process itself or from the voters themselves. Indeed, building polls, even when done honestly, comes with uncertainty because you cannot ask everyone's preferences (for dishonest polls, see [25]). Voters' uncertainty can come from unsure preferences or votes, especially if the elections are still far in the future. Whatever the uncertainty sources, they translate into uncertainty about the scores when considering the Plurality rule, as these are the

only relevant pieces of information for strategic behavior under this rule. Models of uncertain polls either consider a distribution over possible score vectors [26], or uncertain neighborhood of scores [19]. Uncertain voters, on their side, are commonly modeled by the means of incomplete preferences [3, 8, 31]. Within such models, strategic vote changes have been studied either as a single-step change for a voter [26, 3] or as a multi-step iterative procedure [19, 31, 17].

Our contributions in this paper are the following:

- We consider modeling voting uncertainties by sets of probabilities, a quite expressive model that includes in one sweep set-based and probabilistic approaches and allows voters to better express their uncertainty. Section 2 introduces the representation, and provides several examples showing how its expressive power can help in poll or voter uncertainty modeling;
- In Section 3, we introduce our decision-theoretic part, providing several decision rules including known ones [26, 3], that we can revisit and interpret within our framework. Such a higher expressiveness could also be useful in descriptive studies, as more voter behaviours can be described by the model;
- Finally, our study culminates in Section 4 that considers iterative, multi-step frameworks, for which we study convergence results in the style of Meir [21].

2 Uncertainty model on votes

We start by introducing the elements of our newly proposed model. As the model itself is one of the main contribution of the paper, we will describe and illustrate it in detail.

2.1 Modeling Plurality Elections from the Perspective of Strategic Voting

Let N be a set of voters where $N = \{1, \dots, n\}$, and M be a set of candidates where $M = \{1, \dots, m\}$. As we study strategic voting, we assume that $m > 2$, as voting rules are susceptible to manipulation only in those cases [29]. We assume that each voter $i \in N$ has some true underlying preferences, possibly imperfectly known, over candidates represented by a linear order $\succ_i$ over candidates. Let Π^m be the set of all possible preference orders for m candidates. In this paper, the winner of an election is determined by the Plurality voting rule where ties are broken lexicographically. Let $b_i \in M$ denote the ballot of voter i and $b \in M^n$ denote the ballot profile. The winner under Plurality of the ballot profile b is $\mathcal{W}(b) \in \arg \max_{x \in M} s_x(b)$, where $s_x(b) := |\{i \in N : b_i = x\}|$ and a lexicographic tie-breaking order, denoted by $\triangleright$, is used if necessary. Given a number n of voters, we will denote by $\mathcal{S} = \{s : \sum_{x \in M} s_x(b) = n\}$ the space of possible score vectors. By abuse of notation, we sometimes directly write $\mathcal{W}(s)$ to refer to the winner of a score vector s , and more generally use a ballot profile and its corresponding score interchangeably.

While strategic voting is unavoidable [29], a voter can only strategize if she has some information about the election. In plurality, this is relatively straightforward, as knowing the current scores is sufficient to determine whether strategic voting is beneficial. However, the broadcast score is usually subject to uncertainty, as polling institutes can only survey a subset of voters, and voters themselves may be uncertain about their preferences before the election. Two models have met consensus in the literature to handle this problem.

2.2 Representing uncertain votes and polls

Existing representations Existing uncertainty representations typically focus on sets and probabilities. The main ones encountered in strategic voting literature are:

- sets $S \subseteq \mathcal{S}$ of possible scores, as for instance in the model of Meir [19]. This is arguably one of the simplest representation of uncertainty.
- probabilities $p(s)$ of scores over $\mathcal{S}$, which Myerson and Weber [26] derived from observed votes through a multinomial model. In our case, it suffices to retain that it produces a probability over scores.
- Conitzer et al. [3] builds on the model of Konczak et al. [15] using the notions of possible and necessary winners, considering that a voter uncertainty is given by a partial order over the M candidates, provided in the form of pairwise comparisons. For plurality rules, this is equivalent to considering the set of maximal elements of the partial order as the potential vote ballot.

While our approach can cope with any subset of $\mathcal{S}$, in further examples, we will often adopt an interval-based notation for sets of scores, where an interval $[\underline{s}_x, \bar{s}^x]$ represents the range of a candidate score. This is a natural assumption that allows the neighborhood to be connected.

Example 1. If $n = 5$ and $|M| = 3$ with $M = \{a, b, c\}$, the set-valued score $S = ([2, 3], [1, 3], [0, 2])$ means that the score of b will be between 1 (one voter will vote b for sure) and 3 (up to three voters may vote b). It may result from the following 5 imprecise ballots: $\{a\}$, $\{a\}$, $\{a, b, c\}$, $\{b\}$, $\{b, c\}$, with the last resulting, for example, from the partial preference $b \succ_i a$.

Recall that whenever we use this interval representation, the sum of the scores $s = (s_1, \dots, s_m)$ picked within $[\underline{s}_i, \bar{s}^i]$ still sums up to n , as it is a subset of $\mathcal{S}$.

Probability sets In our proposal, uncertainty on votes is given by *convex probability set* $\mathcal{P}$ with probabilities defined over the space $\mathcal{S}$, either derived from poll results or voters uncertainty. The next example describe a very simple situation that this model can capture but that sets or probabilities alone cannot.

Example 2. Assume we have two voters and two candidates³ $\{a, b\}$. Each voter state that she is more willing to vote for a rather than b , while being unsure. This

³ Other examples will consider more than 2.

can be translated by $\mathcal{P}_i = \{p_i \mid p_i(\{a\}) \geq p_i(\{b\})\}$ for voter i . This information cannot be faithfully captured by a set nor by a precise probability.

Assuming voters are independent, these uncertain votes then induce a set of possible probabilities over $\mathcal{S}$, with $P(\{(2, 0)\}) = p_1(\{a\})p_2(\{a\})$, $P(\{(1, 1)\}) = p_1(\{a\})p_2(\{b\}) + p_1(\{b\})p_2(\{a\})$, $P(\{(0, 2)\}) = p_1(\{b\})p_2(\{b\})$. Given the constraints on $\mathcal{P}_i$, one can then deduce that $P(\{(2, 0)\}) \geq P(\{(0, 2)\})$ and $P(\{(1, 1)\}) \geq P(\{(0, 2)\})$, hence that a will most likely be elected, but not whether $P(\{(2, 0)\}) \geq P(\{(1, 1)\})$, meaning that the most probable score is not known for sure.

This set can be interpreted as a collection of plausible probability distributions. This very generic model includes the previously cited one:

- a **set** S of scores, either given [19] or resulting from incomplete preferences [3] is captured by the set $\mathcal{P}_S = \{P : P(S) = 1\}$. In particular, if S is a neighborhood of the true score s with respect to a distance d that we leave unspecified at this stage, then we recover Meir’s [19] representation.
- a **probability** p over $\mathcal{S}$ simply amounts to considering $\mathcal{P} = \{p\}$, i.e., the set reduces to a singleton. This means that any framework resulting in such a probability can be easily captured by our setting, as is the case for Myerson [26]

Manipulating generic sets of probabilities, whether one directly gives a set $\mathcal{P}$ on $\mathcal{S}$ or builds it from individual voter uncertainties $\mathcal{P}_i$, gives us a very expressive framework to describe uncertainties, but is not very computationally friendly nor easy to represent. This is why we detail some specific models of interest [7], focusing on belief functions and their special cases.

Remark 1. Belief functions offer practical advantages, but some natural assessments need the full language of convex sets of probabilities, such as qualitative comparisons [24]. Example 2 is of this kind.

Belief functions A belief function over $\mathcal{S}$ consists in defining a positive mass function $\mathcal{M} : \mathcal{S} \rightarrow [0, 1]$ that sums to one, i.e., $\sum_{S \subseteq \mathcal{S}, S \neq \emptyset} \mathcal{M}(S) = 1$. From such a mass $\mathcal{M}$, we define two bounds over events $A \subseteq \mathcal{S}$ that are defined as

$$\underline{P}(A) = \sum_{S \subseteq A} \mathcal{M}(S), \quad \overline{P}(A) = \sum_{S \cap A \neq \emptyset} \mathcal{M}(S) \quad (1)$$

and from which one can define a corresponding set of possible probabilities

$$\mathcal{P}_{\mathcal{M}} = \{P : \forall A, \underline{P}(A) \leq P(A) \leq \overline{P}(A)\} \quad (2)$$

Belief functions are still quite expressive, and notably include probabilities and sets as special cases: probabilities correspond to masses bearing only on singletons, while a set $\mathcal{S}$ correspond to the mass $m(\mathcal{S}) = 1$.

Example 3. Consider a belief function over the set of possible scores $\mathcal{S}$. For example, one can define $\mathcal{M}(\{(1, 1, 1), (0, 2, 1)\}) = \mathcal{M}([0, 1], [1, 2], 1) = 1$. This mass function corresponds to a set and describes two voters that are certain to vote for b and c , while one voter is hesitating between a and b .

Remark 2. We can also represent each voter's ballot using a mass function $\mathcal{M}_i$ for voter i defined over M , as we did to recover Conitzer's model [3]. From these, one can build a joint mass function over M^n (hence over $\mathcal{S}$): if E_i are subsets of M having positive masses for voter i , the joint mass given over $E_1 \times \dots \times E_n$ is simply $\mathcal{M}(E_1 \times \dots \times E_n) = \prod_{i=1}^n \mathcal{M}_i(E_i)$, which corresponds to assuming independence between agents [30]. As each imprecise ballot can be mapped to a set of possible score vectors, we assume in the rest models on the score space, that is, mass functions $\mathcal{M} : 2^{\mathcal{S}} \rightarrow [0, 1]$.

A key concept in belief functions that we will use is the *pignistic probability* p^* of $\mathcal{M}$, which selects a single probability from $\mathcal{P}_{\mathcal{M}}$. It is defined as $p^*(s) = \sum_{S \ni s} \frac{\mathcal{M}(S)}{|S|}$, and corresponds to the Shapley value of $\underline{P}(A)$ viewed as a game, also reflecting a Laplacian assumption within each focal set S . It is also a vertex gravity center of $\mathcal{P}_{\mathcal{M}}$, hence a meaningful precise representative of $\mathcal{P}_{\mathcal{M}}$, if one must be selected. Also, if the mass $\mathcal{M}$ is a probability p , then $p^* = p$.

As mentioned earlier, belief functions are quite useful to model uncertain polls [16], especially if we allow voters to express their opinion as sets of possible candidates, just as in the work of Conitzer et al. [3] once incomplete preferences are transformed into possible ballots.

Example 4. Consider the case of single-peaked preferences [2] where candidates $\{a, b, c, d, e\}$ are ordered with a left-right axis. One could then weaken the result of a poll by considering that each voter could also vote for the nearest candidates on each side (votes for a become $\{a, b\}$, etc.). Imagine that an initial poll of 100 voters gives the initial score $(10, 30, 10, 30, 20)$ and 10 voters to $\{a, b\}$, leading to the belief function $\mathcal{M}([0, 40], [0, 50], [0, 70], [0, 60], [0, 50]) = 1$. It is clear that the resulting information is quite weak, hence it would make sense to use a less severe weakening process, e.g., assuming that a voter i initially voting for b becomes $\mathcal{M}_i(\{b\}) = \alpha$, $\mathcal{M}_i(\{a, b, c\}) = 1 - \alpha$

Before stepping to decision, we will now present and illustrate two particular cases of belief functions of interest here, in particular for the convergence results of Section 4.

Necessity measures A necessity measure is a belief function with mass on nested elements, i.e., $m(S_i) \neq 0$ and $m(S_j) \neq 0$ only if $S_i \subseteq S_j$ or $S_j \subseteq S_i$.

Example 5. Consider a voter i providing the following information about $\{a, b, c\}$: "I will not vote for c , and hesitate between a and b , with higher chances to vote for a ". A model for such an opinion is $\mathcal{M}^i(\{a\}) = 0.5$ and $\mathcal{M}^i(\{a, b\}) = 0.5$, meaning that $[\underline{P}(\{a\}), \overline{P}(\{a\})] = [0.5, 1]$ and $[\underline{P}(\{b\}), \overline{P}(\{b\})] = [0, 0.5]$.

Example 6. Consider candidates $\{a, b, c\}$ and the initial poll $s^* = (0, 2, 1)$ as a result on 3 voters. It is natural to consider that s^* is the most plausible result, but that nearby results are plausible as well, albeit less as they are further away from s^* . Building neighborhoods S_r^* by considering r the maximal number of voters changing their votes (so that the sum of votes remains the

same), we get $S_1^* = \{s^*\} \cup \{(0, 2, 1)(1, 1, 1), (0, 1, 2), (1, 2, 0)\}$ and $S_2^* = S_1^* \cup \{(2, 0, 1)(1, 0, 2), (0, 0, 3), (2, 1, 0)\}$. One can then consider the mass $\mathcal{M}(\{s^*\}) = \alpha$, $\mathcal{M}(S_1^*) = \beta$ and $\mathcal{M}(S_2^*) = 1 - \alpha - \beta$ with $\alpha \geq \beta \geq (1 - \alpha - \beta)$ to denote a decreasing trust of capturing the true results as we get further away from s^* .

These examples illustrate the kind of model we will consider to extend Meir et al.'s results [19] convergence results in Section 4.

Inner measures An inner measure is a belief function [5, Ch. 2.], where masses are given on a partition of the space, that is $m(S_i) \neq 0$ and $m(S_j) \neq 0$ if $S_i \cap S_j = \emptyset$ and $\cup_{S, m(S) > 0} S = \mathcal{S}$ (or a subset of interest).

While necessity measures are well-suited to equip set-based approaches with gradual notions of enlargement (due to the nestedness), inner measures are well-suited to extend probabilistic models, as they rely on a partition. They may therefore provide further insights about probabilistic models. This is reflected in Figure 1.




Fig. 1: Overview of Our Uncertainty Model and Its Special Cases

2.3 Relations to existing models

Meir's model: [19] defines a neighborhood of scores with respect to a distance d : $S_r(s) = \{s' \in \mathcal{S} \mid d(s, s') \leq r\}$, d can be for instance some ℓ_p ([19] takes the ℓ_1 distance). However, Meir [19] requires one to fix r and restrict uncertainty to a set. Our models can circumvent this limit by considering multiple r values and by decreasing the plausibility of successive sets, as in Example 6.

Example 7. Consider three candidates $\{a, b, c\}$, $n = 30$ the number of voters and the broadcast score being $s = (10, 9, 11)$. We will define beliefs on the neighborhood of scores directly on $\mathcal{S}$, $\mathcal{M}(S_1) = 0.5$, $\mathcal{M}(S_2) = 0.3$ and $\mathcal{M}(S_3) = 0.2$ with $S_1 \subset S_2 \subset S_3$, hence this particular structure is also a necessity measure.

This belief function, which is a necessity measure, models a certainty decreasing with respect to the distance to the broadcast poll s . They are ideal candidate to extend Meir's approach [19].

We may also model the same idea with a more probabilistic flavour. In this case, one can consider the partition generated by the S_i sets, and associate a probability mass to each member of this partition, generating a so-called inner measure [5, Ch. 2.] on the initial space $\mathcal{S}$. An interest of this model is that it is a probabilistic one, albeit defined on an algebra whose atoms are non-singleton sets of scores.

Example 8. Consider three candidates $\{a, b, c\}$, $n = 30$ the number of voters and the broadcast score being $s = (10, 9, 11)$. We will define uncertain information directly on $\mathcal{S}$, $\mathcal{M}(S_1) = 0.5$, $\mathcal{M}(S_2 \setminus S_1) = 0.3$ and $\mathcal{M}(S_3 \setminus S_2) = 0.2$.

Conitzer's model Incomplete preferences of voters are naturally embedded in our approach, as they give rise to imprecise ballots, which can be mapped to a belief function. The next example illustrates this.

Example 9. Consider three candidates $\{a, b, c\}$, $n = 3$ the number of voters and the following preferences given by voters: $b \succ_1 c \succ_1 a$; $c \succ_2 a \succ_2 b$; $a \succ_3 c$ where voter 3 is sure that she prefers a to c , but nothing else. In all completions $a \succ c \succ b$, $b \succ a \succ c$, or $a \succ b \succ c$, c is never ranked first, whereas the top-ranked candidate may be either a or b . This is naturally modelled by the belief function over $\mathcal{S}$ defined as $\mathcal{M}(\{(1, 1, 1), (0, 2, 1)\}) = \mathcal{M}([0, 1], [1, 2], 1) = 1$.

3 Single-step decision model

We now introduce our decision model, illustrate it and link it with the previous single-step manipulation works we mentioned previously.

3.1 The decision model

From a probability set $\mathcal{P}$ on $\mathcal{S}$ and some utility $u : \mathcal{S} \rightarrow \mathbb{R}$, we define the lower expectation

$$\underline{\mathbb{E}}_{\mathcal{P}}(u) = \inf_{P \in \mathcal{P}} \mathbb{E}_P(u) \quad (3)$$

where $\mathbb{E}_P$ is the standard expectation operator. Upper expectation $\bar{\mathbb{E}}_{\mathcal{P}}(u)$ can be defined likewise, taking a sup over $\mathcal{P}$. Lower and upper expectations form the basis of decision making with imprecise probabilities, and we will use them to model strategic decisions and manipulations. These lower and upper expectations should be interpreted from the voter's perspective as the worst and best expected scenarios regarding the voting situation, based on their perception of uncertainty and their individual preferences.

In the case of belief functions, Equation (3) as well as the upper expectation can be easily solved by adopting the following formula using $\mathcal{M}$, i.e.,

$$\underline{\mathbb{E}}_{\mathcal{M}}(u) = \sum_{S \subseteq \mathcal{S}} \mathcal{M}(S) \inf_{x \in S} u(x), \quad \bar{\mathbb{E}}_{\mathcal{M}}(u) = \sum_{S \subseteq \mathcal{S}} \mathcal{M}(S) \sup_{x \in S} u(x). \quad (4)$$

Let us now encode the strategic behavior of voters under uncertainty. To do so, let us define a function $st(s, a_i, a'_i) : \mathcal{S} \times M^2 \rightarrow \mathcal{S}$ that sends back an updated score if voter i moves his vote from a_i to a'_i . To simplify reading, we will use the shorthand $s_{a_i \rightarrow a'_i} := st(s, a_i, a'_i)$ for the update of score s after a strategic move. Then, let $u_i : M \rightarrow \mathbb{R}$ be for each voter i the associated utility that reflects how satisfied the voter is with a winner in M . We will stick to this simple definition depending only on M , as it is sufficient for our purpose. Also note that by Debreu et al. [4], we can always build a utility function from the strict linear orders $\succ_i$. Therefore, to understand how voters may change their vote, we need to evaluate the benefit of moving from a_i to a'_i that we denote

$$u_i(a'_i | a_i, s), \quad (5)$$

that is the utility of moving to a'_i , given the initial state of affairs s and the current vote a_i . In particular, if we are certain about our current score vector s , it suffices for (5) to be positive for the voter to make a strategic move, that is

$$a'_i \succeq a_i \text{ iff } u_i(a'_i | a_i, s) \geq 0$$

where $\succeq$ denotes here that action of voting for a'_i is preferable to voting for a_i .

However, in our model our agents receive a probability set $\mathcal{P}$ over $\mathcal{S}$, and we will define four decision criteria (denoted DC) under $\mathcal{P}$.

- The first criterion that we will consider is a pessimistic criterion to decide whether an action is better than another as follows:

$$a'_i \succ_{pess} a_i \text{ iff } \mathbb{E}_{\mathcal{P}}(u_i(a'_i | a_i, \cdot)) \geq 0 \text{ and } \overline{\mathbb{E}}_{\mathcal{P}}(u_i(a'_i | a_i, \cdot)) > 0$$

under uncertainty $\mathcal{P}$. This corresponds to the maximin of Gilboa et al. [10].

- A pignistic criterion to decide whether an action is better than another:

$$a'_i \succeq_{pig} a_i \text{ iff } \mathbb{E}_{p^*}(u_i(a'_i | a_i, \cdot)) \geq 0,$$

where p^* is the pignistic probability. It includes probabilistic decision as a special case.

- A mixture criteria to decide whether an action is better than another:

$$a'_i \succeq_{\alpha, p^*} a_i \text{ iff } \alpha \cdot \mathbb{E}_{\mathcal{M}}(u_i(a'_i | a_i, \cdot)) + (1 - \alpha) \cdot \mathbb{E}_{p^*}(u_i(a'_i | a_i, \cdot)) \geq 0$$

with $\alpha \in [0, 1]$. α can be seen as a level of completeness of the obtained preferences between possible moves, as it allows one to go from a rather partial order ($\alpha = 0$) to a complete one ($\alpha = 1$).

- Another mixture of criteria:

$$a'_i \succeq_{H(\alpha)} a_i \text{ iff } \alpha \cdot \mathbb{E}_{\mathcal{M}}(u_i(a'_i | a_i, \cdot)) + (1 - \alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i | a_i, \cdot)) \geq 0$$

with $\alpha \in [0, 1]$. This corresponds to the well-known Hurwicz criterion [13] that intends to balance between optimism ($\mathbb{E}$) and pessimism ($\overline{\mathbb{E}}$), and that has been justified within a belief function setting by Denoeux et al [6].

Note that for the three last rules, the strict relation $\succ$ corresponds to strict positive values on the right side. Those decision rules are akin to commonly used decision rules within the imprecise probabilistic setting [32].

To sum up, we describe an election under plurality with a lexicographic tie breaking and strategic voting under uncertainties with the tuple $(N, M, \succ, \triangleright, \mathcal{P}, (u_i)_{i \in N}, DC)$. Let us give an example putting our framework at work, where the considered uncertainty is neither a set nor a probability.

Example 10. Consider candidates $\{a, b, c\}$ and $n = 3$, with voter 1 as in Example 5, while the second and the third voters have certain ballots b and c . In this case, the belief function over $\mathcal{S}$ is defined as $\mathcal{M}(\{(1, 1, 1)\}) = 0.5$, $\mathcal{M}([0, 1], [1, 2], 1) = 0.5$. Let us now illustrate one of our decision rules. Consider for instance that the preference of the second voter is given by $b \succ_2 c \succ_2 a$, and a Hurwicz criterion with $\alpha = \frac{1}{3}$, meaning that the weight of optimism is larger. We will consider the following utilities

$$\forall i \in N, u_i(a'_i | a_i, s) = \begin{cases} 1 & \text{if } \mathcal{W}(s_{a_i \rightarrow a'_i}) \succ_i \mathcal{W}(s), \\ 0 & \text{if } \mathcal{W}(s_{a_i \rightarrow a'_i}) \sim_i \mathcal{W}(s), \\ -1 & \text{otherwise.} \end{cases} \quad (6)$$

We want to evaluate the deviation from b to c for voter 2, where a deviation simply means a change of ballot. Then, using Eq. (4), we compute

$$\underline{\mathbb{E}}_{\mathcal{M}}(u_2(a'_2 | a_2, \cdot)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0; \quad \overline{\mathbb{E}}_{\mathcal{M}}(u_2(a'_2 | a_2, \cdot)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1.$$

For the precise set $\{(1, 1, 1)\}$ with mass 0.5, the move makes c elected instead of a , and the utility is a constant 1 (there is no uncertainty on set $\{(1, 1, 1)\}$). For the set $([0, 1], [1, 2], 1)$ that also has mass 0.5, the move is still beneficial for $(1, 1, 1)$, but would bring negative utility for $(0, 2, 1)$, as c would be elected instead of b , a clear given $b \succ_2 c \succ_2 a$. Thus, we get:

$$\frac{1}{3} \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_2(a'_2 | a_2, \cdot)) + (1 - \frac{1}{3}) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_2(a'_2 | a_2, \cdot)) = \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3} \geq 0,$$

concluding that voter 2 will do the deviation from b to c with such a criterion.

This results in a rich uncertainty modeling and decision framework, in which one can easily separate what comes from the uncertainty model, and what comes from the decision rule. Let us now see how our approaches connect with the existing literature.

3.2 Relations with Existing Models

We show that the models from Meir et al. [23, 19], Myerson et al. [26] and Conitzer et al. [3] are specific case of our approach, demonstrating its expressiveness.

For the first, we consider the model $S_r(s)$ mentioned in Section 2.3. By abuse of notation, we will use S_r when it is clear from the context. Meir [23] first worked on the special case $r = 0$, meaning the poll information is complete for plurality and certain. We retrieve Meir's model [23] by considering voters $i \in N$, $\mathcal{M}_i(S_0) = 1$ and the particular utility functions:

$$\forall i \in N, u_i(a'_i | a_i, s) = \begin{cases} 1 & \text{if } a'_i = \mathcal{W}(s_{a_i \rightarrow a'_i}) \succ_i \mathcal{W}(s), \\ 0 & \text{if } a'_i \neq \mathcal{W}(s_{a_i \rightarrow a'_i}) \succ_i \mathcal{W}(s) \text{ or } \mathcal{W}(s_{a_i \rightarrow a'_i}) \sim_i \mathcal{W}(s), \\ -1 & \text{otherwise,} \end{cases}$$

and the pessimistic decision rule $\succeq_{\text{pess}}$. Recall that in [19], convergence is guaranteed and we call these strategic moves "direct best response" since all strategic moves where a voter deviates are to the new winner.

Meir [22] then extended this framework to add uncertainty by considering strictly positive value $r_i > 0$ for voter i . This comes down in our case to take, $\forall i \in N$, $\mathcal{M}_i(S_{r_i}) = 1$, the utility functions given by Equation (6) and again the pessimistic decision rule $\succeq_{\text{pess}}$ to recover Meir [22] proposal.

For the probabilistic model of Myerson et al. [26], they derive from observed votes a weight vector $q = (q_1, \dots, q_m)$, and use a multinomial distribution to obtain the probability over score vectors $\mathbb{P}(s)$. The utilities can be chosen arbitrarily provided they satisfy the voter's preferences $u_i(a'_i | a_i, s) \geq 0$ iff $\mathcal{W}(s_{a_i \rightarrow a'_i}) \succ_i \mathcal{W}(s)$. Since in the probabilistic case all decision rules we have described in Section 3.1 reduce to a classical expectation operator, there is no need to specify a specific behaviour.

The last model we want to encompass is the one from [3] with incomplete preferences. As we already showed in Section 2, the incompleteness on each voter's preference can easily be summarized as an uncertainty set on scores that we call S and where $\mathcal{M}(S) = 1$. It remains to choose the decision rule with respect to this set. Conitzer et al. [3] adopt a pessimistic decision criterion, which amounts to use the utilities of (6) and the pessimistic decision rule $\succeq_{\text{pess}}$.

Now that we have a fully fledged decision framework allowing us to model whether or not a voter will proceed to a strategic move or manipulation, we study the multi-step case, in particular showing that classical convergence results can be reinterpreted and extended to specific instantiations of our framework.

4 Multi-step Manipulations

In this section we generalize existing results on convergence from Meir [19] by adding quantitative uncertainty about the score as in Example 7 and Example 8. Let us recall that a neighborhood of scores S_r with respect to a distance d is defined as in Meir's [19]: $S_r(s) = \{s' \in \mathcal{S} \mid d(s, s') \leq r\}$, where d can be $\ell_1, \ell_\infty, \dots$ as in Meir's work [19]. Let us consider the ℓ_1 distance, which is easy to interpret in terms of adding or removing votes, without necessarily preserving the total number of votes, which may vary by one. This behavior can be easily explained by considering a voter who chooses to abstain. For each voter i , we

will let r_i be the support of their uncertainty, so that S_{r_i} is the biggest set of voter i 's uncertainty. We now formalize the notion of equilibrium in this model through the following definition.

Definition 1. *An equilibrium is a situation where no voter has an incentive to deviate from its ballot, i.e. $\forall i \in N, \nexists a'_i$ such that $a'_i \succ_{DC} a_i$. We assume voters will only make a strategic move if their preferences are strict.*

4.1 Extending Meir Framework

Let us consider a belief function for voter i , $\mathcal{M}^i : \mathcal{S} \rightarrow [0, 1]$ that sums to one, i.e., $\sum_{S \subseteq S_{r_i}, S \neq \emptyset} m^i(S) = 1$. This notation for belief functions indexed by voters should not be confused with that of Remark 2, where the question was to model each ballot independently, while here we consider that the scores come from a unique poll, but that each voter may be more or less skeptical about its result. We will consider two particular cases of interest. The first is the case of nested sets, i.e. $\forall i \in N, \forall k \geq 1, m^i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$. Indeed, we might want to describe the fact that the belief mass depends on the distance to the true score. Second, we look at the case of *partitioned* belief function, i.e. $\forall i \in N, \forall k \geq 2, m^i(S_k \setminus S_{k-1}) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$ and $\forall i \in N, m^i(S_1) = \beta_1^i$. In both cases, we will assume decreasing $(\beta_k^i)_{1 \leq k \leq r_i}$, meaning that our evidence decreases as we get further from the observed s . For the rest of this section, we consider utilities given by Equation (6).

Theorem 1. *Voters considering uncertainty given by a nested or a partitioned decreasing belief function, and making strategic decisions according to either pessimistic ($\succeq_{pess}$) or mixed ($\succeq_{\alpha, p^*}, \succeq_{H(\alpha)}$) decision rules with α large enough, will converge to an equilibrium.*

Proof. At first, we will take $\forall i \in N, \forall k \geq 1, \mathcal{M}(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$. We consider a feasible move $a_i \rightarrow a'_i$. We remark that the hypothesis that $\mathbb{E}_{\mathcal{M}}(u_i(a'_i|a_i, s)) \geq 0$ implies that either:

- $\forall s \in S_{r_i}, u_i(a'_i|a_i, s) \neq -1$,
- $\exists \tilde{r} \in [0, r_i], \forall s \in S_{\tilde{r}}, u_i(a'_i|a_i, s) = 1$

since, if some neighborhoods have a negative utility, this must be compensated by a positive contribution, which happens only if all moves within a neighborhood are positive. However, the second case is impossible, for the reason that if we allow a voter to transfer its vote to another candidate ($r = 1$), then there is a situation for which the voter is not pivotal, and therefore $\exists s \in S_1, u_i(a'_i|a_i, s) = 0$, showing that the second case never happens. Therefore, if $\mathbb{E}_{\mathcal{M}}(u_i(a'_i|a_i, s)) \geq 0$, this means that $\forall s \in S_{r_i}, u_i(a'_i|a_i, s) \in \{0, 1\}$, with at least one s giving the null value.

Second, note that we cannot have $\mathbb{E}_{\mathcal{M}}(u_i(a'_i|a_i, s)) > 0$ if $\forall s \in S_{r_i}, u_i(a'_i|a_i, s) = 0$, and there must be a situation for which making this strategic move is a local dominance move, meaning it verifies Theorem 4 from Meir [19].

We do exactly the same reasoning for the second type of belief functions, i.e., $\forall i \in N, \forall k \geq 2, \mathcal{M}_i(S_k \setminus S_{k-1}) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$ and $\forall i \in N, \mathcal{M}_i(S_1) = \beta_1^i$. For the mixed decision, one can prove that the decision behaves as the pessimistic one if α is large enough.

We now want to go a step further by showing the convergence can hold even with a less, but still pessimistic behavior, namely the Hurwicz criterion with $\alpha > \frac{1}{2}$. This allows us to consider negative outcomes in the uncertainty neighborhood of scores. Let us consider belief functions as follows: $\forall i \in N, \forall k \geq 1, m^i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$

Theorem 2. *Voters considering a decreasing belief function around the true score and making strategic votes according to a sufficiently pessimistic Hurwicz criterion (i.e. $\alpha > \frac{1}{2}$) will converge to an equilibrium.*

Proof (sketch). Hurwicz can be rewritten as $\sum_{i=1}^{r_i} \mathcal{M}(S_i)[\alpha \inf_{s \in S}(u_i(a'_i|a_i, s)) + (1 - \alpha) \cdot \sup_{s \in S}(u_i(a'_i|a_i, s))]$ and we show that if $\alpha > \frac{1}{2}$, such an equation is positive only if we are in averaging over situations that are similar from section VIII in Meir [18] which tells us that the convergence holds for any level of uncertainty r_i and any starting point (even non-truthful states) for local dominance strategic behaviors.

4.2 The Case of Pignistic Probability

Let us now consider the case of a pignistic probability on a neighborhood of scores S_{r_i} , which is equivalent to having a uniform distribution $\mathcal{U}$ over all single scores within S_{r_i} for any voter i . If S_{r_i} is our single set with positive mass, the corresponding strategic behaviour is equivalent having the following criterion in the strict case:

$$a'_i \succ_{pig} a_i \text{ iff } \text{card}(s \in S_{r_i} \mid u_i(a'_i|a_i, s) = 1) - \text{card}(s \in S_{r_i} \mid u_i(a'_i|a_i, s) = -1) > 0$$

where card is the cardinality of the set.

Proposition 1. *With a uniform pignistic criterion on a neighborhood of scores of size 1 (with respect to the ℓ_1 distance), convergence is not guaranteed.*

Proof (sketch). If we consider the preferences

$$\begin{aligned} a \succ_1 b \succ_1 c \succ_1 d ; & c \succ_2 a \succ_2 b \succ_2 d ; & c \succ_3 a \succ_3 b \succ_3 d \\ c \succ_4 d \succ_4 a \succ_4 b ; & a \succ_5 c \succ_5 d \succ_5 b ; & d \succ_6 b \succ_6 a \succ_6 c \\ c \succ_7 b \succ_7 d \succ_7 a ; & d \succ_8 b \succ_8 a \succ_8 c ; & b \succ_9 d \succ_9 c \succ_9 a \\ & b \succ_{10} d \succ_{10} c \succ_{10} a \end{aligned}$$

of those ten voters over four candidates, we can show that Voter 8 will move from d to a with a cardinal difference of 1, voter 2 will move from c to a with a cardinal difference of 2, voter 8 will move from a to d with a cardinal difference of 1 and finally voter 2 will move from a to c with a cardinal difference of 3. This creates a cycle, which prevents convergence.

This example is to be put in perspective with our result of Theorem 1 about partitioned belief functions, which indicates that with the right (coarse) algebra, it is possible to consider a probability measure and a corresponding decision rule such that convergence holds, meaning that the counter-example is not so much about having a probabilistic model itself than about the voters being too optimistic in their movement, hinting also at the fact that "optimistic" decision rules are unlikely to lead to convergence of voting behaviors. Note that the distance here is the ℓ_1 distance, and that the number of voters does not sum to n . However, this can easily be interpreted as meaning that one voter does not participate, and is therefore still very reasonable. Moreover, similar examples can be found for other distances, including distances for which the number of voters sums to n .

5 Conclusion and Future Works

In this paper, we provide new tools to model uncertainty in strategic voting. We think that our work provides substantially new views on preference modeling in voting theory, and in particular in strategic voting models. In particular, we have shown that they capture standard settings in a unifying framework, and allow one to represent new uncertainty scenarios that were not captured by previous models. We have illustrated, through many examples, that this model can account for uncertainty arising both from polls and from the voters themselves. Moreover, we establish new convergence results in our framework that generalize existing ones, and we provide a key counterexample that points out the limits of these convergence results.

This work opens up several promising research directions, both in strategic voting models and in preference modeling in voting theory more broadly. A first important question is how incorporating a more realistic model of uncertainty may influence election outcomes compared to classical iterative voting models [21]. Another avenue for future research is to design empirical studies aimed at identifying the parameters of our model, in order to represent the uncertainty present in both polls and voters' preferences more accurately in real-world elections. We also believe that the descriptive and expressive power of the framework is an advantage: it can formally capture more behaviors than other frameworks, and can ensure smooth transitions between set and probabilistic models. A challenge will however be to propose parametric models where the number of parameters does not explode. Suppose we are able to do so, it would allow us to reframe classical questions in iterative voting, such as whether equilibria can be computed efficiently [27], how an external agent might manipulate the outcome [1], or whether it increases social welfare [14]. One might also ask how such a model could be adapted to other voting rules. Finally, we believe this modeling can offer a new perspective on other voting problems, such as possible and necessary winners or elicitation of voters' preferences, by capturing uncertainty in a more quantitative flavor.

Acknowledgement

This work is supported under the France 2030 program by the ANR-23-IACL-0007 grant (AI Cluster PostGenIA).

Bibliography

- [1] Baumeister, D., Selker, A.K., Wilczynski, A.: Manipulation of opinion polls to influence iterative elections. In: Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020). pp. 132–140 (2020)
- [2] Black, D.: The theory of committees and elections. Cambridge University Press (1958)
- [3] Conitzer, V., Walsh, T., Xia, L.: Dominating manipulations in voting with partial information. In: Proceedings of the 25th AAAI conference on artificial intelligence (AAAI 2011). vol. 25, pp. 638–643 (2011)
- [4] Debreu, G., et al.: Representation of a preference ordering by a numerical function. *Decision processes* **3**, 159–165 (1954)
- [5] Denneberg, D.: Non-additive measure and integral, vol. 27. Springer Science & Business Media (2013)
- [6] Denoeux, T., Shenoy, P.P.: An interval-valued utility theory for decision making with dempster-shafer belief functions. *International Journal of Approximate Reasoning* **124**, 194–216 (2020)
- [7] Destercke, S., Dubois, D.: Special cases. *Introduction to Imprecise Probabilities* (chapter 4), 79–91 (2014)
- [8] Dey, P., Misra, N., Narahari, Y.: Complexity of manipulation with partial information in voting. *Theoretical Computer Science* **726**, 78–99 (2018)
- [9] Endriss, U., Obraztsova, S., Polukarov, M., Rosenschein, J.S.: Strategic voting with incomplete information. *AAAI Press/International Joint Conferences on Artificial Intelligence* (2016)
- [10] Gilboa, I., Schmeidler, D.: Maxmin expected utility with non-unique prior. *Journal of mathematical economics* **18**(2), 141–153 (1989)
- [11] Grabisch, M.: Decision under risk and uncertainty. In: *Set Functions, Games and Capacities in Decision Making*, pp. 281–323. Springer (2016)
- [12] Hazon, N., Aumann, Y., Kraus, S., Wooldridge, M.: On the evaluation of election outcomes under uncertainty. *Artificial Intelligence* **189**, 1–18 (2012)
- [13] Hurwicz, L.: Optimality criteria for decision making under ignorance. Tech. rep., Cowles Commission discussion paper, statistics (1951)
- [14] Kavner, J., Xia, L.: Strategic behavior is bliss: iterative voting improves social welfare. *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)* **34**, 19021–19032 (2021)
- [15] Konczak, K., Lang, J.: Voting procedures with incomplete preferences. In: *Proceedings of the Multidisciplinary Workshop on Advances in Preference Handling (IJCAI 2005)*. vol. 20, p. 12 (2005)
- [16] Kreiss, D., Augustin, T.: Undecided voters as set-valued information—towards forecasts under epistemic imprecision. In: *Proceedings of the 14th International Conference of Scalable Uncertainty Management (SUM 2020)*. pp. 242–250. Springer (2020)

- [17] Lev, O., Meir, R., Obraztsova, S., Polukarov, M.: Heuristic voting as ordinal dominance strategies. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI 2019). vol. 33, pp. 2077–2084 (2019)
- [18] Meir, R.: Plurality voting under uncertainty. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015). vol. 29 (2015)
- [19] Meir, R.: Iterative voting. In: Endriss, U. (ed.) Trends in Computational Social Choice, chap. 4, pp. 69–86. AI Access (2017)
- [20] Meir, R.: Strategic voting. Synthesis lectures on artificial intelligence and machine learning **13**(1), 1–167 (2018)
- [21] Meir, R.: Strategic voting. Springer Nature (2022)
- [22] Meir, R., Lev, O., Rosenschein, J.S.: A local-dominance theory of voting equilibria. In: Proceedings of the fifteenth ACM conference on Economics and computation (ACM 2014). pp. 313–330 (2014)
- [23] Meir, R., Polukarov, M., Rosenschein, J., Jennings, N.: Convergence to equilibria in plurality voting. In: Proceedings of the 24th AAAI conference on artificial intelligence (AAAI 2010). vol. 24, pp. 823–828 (2010)
- [24] Miranda, E., Destercke, S.: Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology* **64**, 44–57 (2015)
- [25] Mousseau, V., Surugue, H., Wilczynski, A.: Do we care about poll manipulation in political elections? In: Proceedings of the 27th European Conference on Artificial Intelligence (ECAI 2024) (2024)
- [26] Myerson, R.B., Weber, R.J.: A theory of voting equilibria. *American Political science review* **87**(1), 102–114 (1993)
- [27] Rabinovich, Z., Obraztsova, S., Lev, O., Markakis, E., Rosenschein, J.: Analysis of equilibria in iterative voting schemes. In: Proceedings of the 29th AAAI conference on artificial intelligence (AAAI 2015). vol. 29 (2015)
- [28] Reijngoud, A., Endriss, U.: Voter response to iterated poll information. In: Proceeding of the 4th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012). vol. 4, p. 8 (2012)
- [29] Satterthwaite, M.A.: Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* **10**(2), 187–217 (1975)
- [30] Smets, P., Kennes, R.: The concept of distinct evidence. In: Proceedings of the Information Processing and Management of Uncertainty (IPMU 1992). pp. 789–794 (1992)
- [31] Terzopoulou, Z., Terzopoulos, P., Endriss, U.: Iterative voting with partial preferences. *Artificial Intelligence* **332**, 104133 (2024)
- [32] Troffaes, M.C.: Decision making under uncertainty using imprecise probabilities. *International journal of approximate reasoning* **45**(1), 17–29 (2007)

Technical Appendix

A Complete proofs

Proposition 1. *With a uniform pignistic criterion on a neighborhood of scores of size 1 (with respect to the ℓ_1 distance), convergence is not guaranteed.*

Proof. Here is the counter example with a neighborhood of scores of size 1. Let us take the following profile:

$$\begin{aligned} a \succ_1 b \succ_1 c \succ_1 d ; c \succ_2 a \succ_2 b \succ_2 d ; c \succ_3 a \succ_3 b \succ_3 d \\ c \succ_4 d \succ_4 a \succ_4 b ; a \succ_5 c \succ_5 d \succ_5 b ; d \succ_6 b \succ_6 a \succ_6 c \\ c \succ_7 b \succ_7 d \succ_7 a ; d \succ_8 b \succ_8 a \succ_8 c ; b \succ_9 d \succ_9 c \succ_9 a \\ b \succ_{10} d \succ_{10} c \succ_{10} a \end{aligned}$$

Voter 8 will move from d to a with a cardinal difference of 1. For clarity, let us detail the computation of this first move: the original score is $s = (2, 2, 3, 3)$, so

$$\begin{aligned} S_1 = \{ & (2, 2, 3, 3), (1, 2, 3, 3), (2, 1, 3, 3), (2, 2, 2, 3), (2, 2, 3, 2), \\ & (3, 2, 3, 3), (2, 3, 3, 3), (2, 2, 4, 3), (2, 2, 3, 4) \} \end{aligned}$$

that corresponds to add/remove one vote. When moving from d to a , Voter 8 is improved in the first state $(2, 2, 3, 3)$ as it becomes $(3, 2, 3, 2)$, and a is elected instead of c , which is better from voter 8's perspective. Voter 8 is also improved for states $(2, 1, 3, 3)$, $(2, 2, 3, 2)$ and $(2, 3, 3, 3)$, and deteriorated in states $(2, 2, 2, 3)$, $(2, 3, 3, 3)$ and $(2, 2, 3, 4)$. Other states are not impacted by this move. Then voter 2 will move from c to a with a cardinal difference of 2, voter 8 will move from a to d with a cardinal difference of 1 and finally voter 2 will move from a to c with a cardinal difference of 3. This creates a cycle, which prevents convergence.

Theorem 2. *Voters considering a decreasing belief function around the true score and making strategic votes according to a sufficiently pessimistic Hurwicz criterion (i.e. $\alpha > \frac{1}{2}$) will converge to an equilibrium.*

Proof. We know that the weights β_k^i of our belief functions defined as $\forall i \in N, \forall k \geq 1, \mathcal{M}_i(S_k) = \beta_k^i$ are decreasing.

We consider a feasible move $a_i \rightarrow a_{i'}$. Using the fact that the lower (an upper) expectation is positive homogeneous, i.e., $\alpha \mathbb{E}(f) = \mathbb{E}(\alpha f)$, we get

$$\begin{aligned} & \alpha \cdot \mathbb{E}_{\mathcal{M}}(u_i(a'_i | a_i, s)) + (1 - \alpha) \cdot \mathbb{E}_{\mathcal{M}}(u_i(a'_i | a_i, s)) \\ &= \sum_{i=1}^{r_i} \mathcal{M}(S_i) [\alpha \inf_{s \in S} (u_i(a'_i | a_i, s)) + (1 - \alpha) \cdot \sup_{s \in S} (u_i(a'_i | a_i, s))] \end{aligned} \quad (7)$$

Let us denote by

$$\tilde{u}_i(S, a_i, a'_i) = \alpha \inf_{s \in S} (u_i(a'_i | a_i, s)) + (1 - \alpha) \cdot \sup_{s \in S} (u_i(a'_i | a_i, s))$$

the term associated to subset S , meaning that Equation (7) can be rewritten $\sum_{i=1}^{r_i} \mathcal{M}(S_i) \tilde{u}_i(S, a_i, a'_i)$, and that the strategic move is decided by a weighted average of $\tilde{u}_i$ values.

We distinguish six possible cases:

- Case A: $\exists s \in S, u_i(a'_i | a_i, s) = 1$ and $\forall s \in S, u_i(a'_i | a_i, s) \geq 0$
- Case B: $\exists s \in S, u_i(a'_i | a_i, s) = -1$ and $\forall s \in S, u_i(a'_i | a_i, s) \leq 0$
- Case C: $\exists s \in S, u_i(a'_i | a_i, s) = -1$ and $\exists s \in S, u_i(a'_i | a_i, s) = 1$
- Case D: $\forall s \in S, u_i(a'_i | a_i, s) = 0$
- Case A.1: $\forall s \in S, u_i(a'_i | a_i, s) = 1$
- Case B.1: $\forall s \in S, u_i(a'_i | a_i, s) = -1$

However, cases A.1 and B.1 can never happen for a non-singleton S , for the same reasons as the ones advocated in the proof of Theorem 1. We then get

$$\tilde{u}_i(S, a_i, a'_i) = \begin{cases} 1 - \alpha & \text{Case A,} \\ -\alpha & \text{Case B,} \\ 1 - 2 \cdot \alpha & \text{Case C,} \\ 0 & \text{Case D} \end{cases}$$

It is clear that $\alpha > \frac{1}{2}$ implies that $\tilde{u}_i(S, a_i, a'_i) \geq 0$ in cases A only.

Therefore, if

$$\alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i | a_i, s)) + (1 - \alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i | a_i, s)) \geq 0$$

then

$$\exists r'_i \leq r_i \text{ such that } \exists s \in S_{r'_i}, u_i(a'_i | a_i, s) = 1 \text{ and } \forall s \in S_{r'_i}, u_i(a'_i | a_i, s) \geq 0$$

Which is Case A. The last equivalence comes from the fact that only case A can lead to a positive $\tilde{u}_i(S, a_i, a'_i)$, and that the criterion is an average of such $\tilde{u}_i$ values. Therefore, there must exist $r'_i \leq r_i$ such that $\tilde{u}_i(S_{r'_i}, a_i, a'_i) = 1 - \alpha$. If $R_i^+ = \{j : |\tilde{u}_i(S_j, a_i, a'_i) = 1 - \alpha|\}$ is the set of neighborhood indices in case A, we need for $\sum_{j \in R_i^+} \beta_j^i$ to be large enough for the decision criterion to be positive. In other words, we can accept a move with some case C only if there exists a local dominance move in a smaller neighborhood that receives enough evidence. Using the result from section VIII in Meir [18] which tells us that the convergence holds for any level of uncertainty r_i and any starting point (even non-truthful states) for local dominance strategic behaviors, then we can do a bijection of our strategic moves and get the convergence also.

B A word on complexity

Let us briefly situate our model with respect to the computational complexity of computing a manipulation. Prior work has already investigated related questions, and it is useful to distinguish two broad cases. For plurality, when one must evaluate each scenario induced by uncertainty, for instance, under incomplete preferences, this essentially amounts to reasoning over all completions. This is closely related to the possible winner task and is $\#P$ -hard [3]. In particular, under criterion that aggregate across completions (e.g., computing an expectation under a pignistic rule), one cannot avoid considering the full set of completions. By contrast, when the decision criterion is simpler, such as a pessimistic (worst-case) criterion, plurality requires only limited information; it suffices to identify the winner in the relevant worst-case scenario. In that case, the problem can be rewritten as a flow problem as in Conitzer et al. [3] which can be solved in polynomial time. We may also mention other decision criterion that remain simple enough to be computed in polynomial-time [8], such as the optimistic criterion (i.e., the preferred candidate wins in at least one completion) and opportunistic manipulation (i.e., the preferred candidate wins in every completion in which some manipulation is feasible).

The same dichotomy can be applied to our approaches: if a decision rule requires enumerating all possible completions of a vote, we will end up with $\#P$ -hard problems. This is the case for decision rules involving the pignistic probability, for instance. On the other hand, and provided the belief function mass $\mathcal{M}$ is positive only on a polynomial number of sets, rules that only require identifying one completion remains polynomial.

In our general model, including these cases as particular cases, we then know that under any criteria that needs to compute every possible scenario then the decision will be also $\#P$ -hard to decide, or worse if we consider quite generic/complex models. However, the case of a pessimistic decision criterion remains solvable in polynomial time as long as the number of sets receiving positive mass (in the case of belief functions) remains polynomial, essentially because it does not require computing probabilities or weights. Thanks to this observation, a standard maximum-flow approach [3] can be applied to obtain a polynomial-time algorithm for manipulation with the pessimistic criterion.

This also opens up an interesting avenue of research: given that in the probabilistic case all decision rules amount to classical expected value, could we weaken the initial probabilistic information, e.g., by considering an inner measure on a coarse partition of linear orders, and take a tractable decision rule (pessimistic and optimistic) in order to approximate intractable probabilistic manipulation rules? To be more precise, if computing an expectation $\mathbb{E}(\cdot)$ for a probability is intractable, can we find a set $\mathcal{P}$ containing this probability such that computing $\underline{\mathbb{E}}(\cdot), \bar{\mathbb{E}}(\cdot)$ is tractable, and where $\underline{\mathbb{E}}(\cdot), \bar{\mathbb{E}}(\cdot)$ are reasonably close to each other? This would let us reach a decision in some cases, and in others, quantify the remaining uncertainty. Note that such a question is made possible only by the fact that our models unify set-based with probabilistic-based decision (respectively known as decision under uncertainty, and decision under risk [11]) in a single

framework, thereby allowing one to easily transfer notions used in one setting to the other setting.