An Enriched Model of Strategic Voting under Uncertainty

Henri Surugue[a,*] **and Sébastien Destercke**[b]

[a]MICS, CentraleSupélec, Université Paris-Saclay, France
[b]Université de technologie de Compiègne, CNRS, Heudiasyc, Compiègne, France

Abstract. We present a new model of strategic voting where uncertainty is modelled by probability sets and the decisions are taken according to lower and upper expected utility gains. Focusing on the particular case of belief functions, we show that this generic model encompasses in one sweep previous models that considered either sets or probabilities to model uncertainty, and enables us to generalize some well-known convergence results from the literature. We also discuss the case of uncertain voters, and emphasize the challenges that come with considering richer and therefore more realistic models of uncertainty.

1 Introduction

Strategic voting [15] arises when voters have an incentive to misreport their true preferences to obtain a more desirable result. The Gibbard-Satterthwaite theorem [8, 24] shows that no reasonable voting rule is immune to strategic manipulation. One approach to circumvent this impossibility result is to allow for strategic manipulation and investigate its consequences, for example analyzing its strategic moves as a strategic game [17] or using a Bayesian approach [21, 10]. These approaches raise the question of how voters act strategically and the role of polls. Therefore, understanding how information impacts voter opinions is crucial [7, 23].

In particular, this information is subject to uncertainty for many reasons. First of all, building polls creates uncertainty by construction since you have to query a subset of voters and make some inferences about the remaining voters. This is the assumption of the two main models in the literature, where authors have considered both a probability around received polls [21] for the first model, and the use of an uncertain neighborhood of scores [14] in the second model. However, the first model makes all strategic moves in one-shot, without voters being able to change their votes afterwards, which is hard to justify. The second model is iterative but with a more qualitative behavior, meaning that it considers any score in the neighborhood of scores with the same importance. In this paper, our goal will be to generalize these models by introducing some new concepts that both include the existing literature and give a framework to capture new insights.

Specifically, we will be interested in giving a generalization of Meir's models by including some quantitative aspect, namely introducing the idea that some scores are less important than others, if they are far from the broadcast poll for example. In this context, we will give convergence results of our iterative voting process. However, our model allows us to conceptualize new ideas as an uncertainty about voter's themselves, similar to Kreiss and Augustin but for iterative voting. For example, our model capture situations where voters have either uncertain ballots or incomplete preferences as in [1, 6], which is often the case in practice. We will show how to model such behaviors and explain the difficulties that remain to implement this in practice.

The rest of the paper is organized as follows: Section 2 describes in detail our new model, as it is a central contribution of the paper. Section 3 then motivates it, first by showing how it includes and generalizes previous frameworks, second by providing practical examples where one would need to consider generalized representations of uncertainty in order to have an accurate model. Section 4 provides convergence results for our newly introduced model, showing that while being more general, it preserves some key properties discussed in the literature. Finally, Section 5 provides a discussion rather than a conclusion, as we think that our contribution opens up many avenues of research that would bring strategic voting closer to real-life behaviour of voters, hopefully offering a higher descriptive power for such models.

2 The Model

We start by introducing the elements of the newly proposed model. As the model itself is one of the main contribution of the paper, we will describe it in details, before illustrating some of its aspects.

2.1 A Voting Situation Equipped with an Uncertainty Model

Let N be a set of voters where $N = \{1, \dots, n\}$, and M be a set of candidates where $M = \{1, \dots, m\}$. Since the goal is to study strategic voting, we naturally assume that $m > 2$ (by the Gibbard-Satterthwaite theorem, voting rules are susceptible to manipulation only when there are more than two candidates). We assume that each voter $i \in N$ has some underlying preferences 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. 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 ballot profile b from which b_i is excluded is denoted by b_{-i}. The winner under Plurality of the ballot profile b is $\mathcal{W}_P(b) \in \arg\max_{x \in M} s_x(b)$, where

* Corresponding Author. Email: henri.surugue@centralesupelec.fr.

$s_x(b) := |\{i \in N : b_i = x\}|$ and a lexicographic tie-breaking order, denoted by $\rhd$, is used if necessary.

By abuse of notation, we sometimes directly write $\mathcal{W}_P(s)$ to refer to the winner of a score vector s, and more generally use a ballot profile and its corresponding score interchangeably. Let b^T denote the truthful ballot profile, i.e., $b_i^T \succ_i x$ for every candidate $x \in M \setminus \{b_i^T\}$ and voter $i \in N$, and s^T denote the candidates' scores in b^T.

We equip our election with an uncertainty model on scores. We consider that the voter receives a convex probability set $\mathcal{P}$ over the set of possible scores[1] $\mathcal{S} = \{s : \sum_{x \in M} s_x(b) = n\}$. Two common ways to build $\mathcal{P}$ are (1) to consider an uncertainty model around an original score s and (2) by having uncertainty about each ballot, assumed independent, thus building a global $\mathcal{P}$ over M^n through each individual $\mathcal{P}_i$ (the uncertainty of voter i). Note that since a ballot profile maps to a single score vector, a set defined over M^n can be unambiguously mapped to a set $\mathcal{P}$ over $\mathcal{S}$.

From a probability set $\mathcal{P}$ on $\mathcal{S}$, one can then define, for any function $f : \mathcal{S} \to \mathbb{R}$ (typically a utility of some kind), a lower expectation

$$\underline{\mathbb{E}}_{\mathcal{P}}(f) = \inf_{P \in \mathcal{P}} \mathbb{E}_P(f) \tag{1}$$

where $\mathbb{E}_P$ is the standard, linear expectation operator. Upper expectation $\overline{\mathbb{E}}_{\mathcal{P}}(f)$ can be defined like-wise, taking a sup. Of particular interest are lower and upper expectations over identity functions of events A, that corresponds to lower and upper probabilities over such events, and that we will denote $\underline{P}(A), \overline{P}(A)$. We will also use the following notations for sets of scores: an interval for a given candidate simply means that the score of this candidate can vary within these bounds. For instance, if $n = 5$ and $|M| = 3$, the set-valued score $E = ([2, 3], [1, 3], [0, 2])$ means that the score of b will be between 1 (one voter will vote for b for sure) and 3.

Probability sets $\mathcal{P}$ are very generic uncertainty models that include as special cases [5] quite a number of existing uncertainty representations. We detail some of those that will be of particular interest in this paper.

The specific case of sets A first special case is the one of sets, where we know that $s \in E$ for some subset $E \subseteq \mathcal{S}$. The corresponding probability set is defined as

$$\mathcal{P}_E = \{P : P(E) = 1\},$$

summarised by the constraint $\underline{P}(E) = 1$. In such a case, Equation (1) has for solution $\underline{\mathbb{E}}_{\mathcal{P}_E}(f) = \inf_{x \in E} f(x)$.

The specific case of probability Probability sets obviously generalize probabilities by construction. Therefore, if we take $\underline{P}(A) = \overline{P}(A)$ for all events A, we come back to the very well-known probability setting, in which $\mathcal{P}$ reduces to a singleton.

The specific case of belief functions A belief function over M^n (or $\mathcal{S}$) consists in defining a positive mass function $\mathcal{M} : 2^{M^n} \to [0, 1]$ that sums up to one, i.e., $\sum_{E \subseteq M^n, E \neq \emptyset} m(E) = 1$. From such a mass $\mathcal{M}$ can then be defined two bounds over events $A \subseteq M^n$ that are defined as

$$\underline{P}(A) = \sum_{E \subseteq A} m(E) \tag{2}$$

$$\overline{P}(A) = \sum_{E \cap A \neq \emptyset} m(E) \tag{3}$$

[1] Possible scores are sufficient if we work under plurality rule, otherwise we would have to consider probability sets on the spaces of vote profiles.

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

$$\mathcal{P}_{\mathcal{M}} = \{P : \forall A, \underline{P}(A) \leq P(A) \leq \overline{P}(A)\} \tag{4}$$

Should we collect a mass function $\mathcal{M}_i$ per voter defined over M, $m_i(E)$ would then express the mass of evidence we have that the voter ballot will be in E. From those, one can easily build a joint mass function over M^n (hence over $\mathcal{S}$): if E_i are some subsets of M receiving positive masses for voter i, then the joint mass given over $E_1 \times \ldots \times E_n$ is simply $\mathcal{M}(E_1 \times \ldots \times E_n) = \prod_{i=1}^N \mathcal{M}_i(E_i)$, which is well-suited to denote independence between agents [25]. Each imprecise ballot can then be mapped to a corresponding set of possible score vectors. Belief functions, while strictly less expressive than generic lower probabilities $\underline{P}(A)$, are still a quite expressive model, as they include probabilities and sets as special cases: probabilities correspond to masses bearing only on singletons, while a set E correspond to the mass $m(E) = 1$.

An advantage of belief functions is that Equation (1) as well as the upper expectation can be easily solved by adopting the following formula using $\mathcal{M}$, i.e.,

$$\underline{\mathbb{E}}_{\mathcal{M}}(f) = \sum_{E \subseteq M^n} m(E) \inf_{x \in E} f(x),$$

$$\overline{\mathbb{E}}_{\mathcal{M}}(f) = \sum_{E \subseteq M^n} m(E) \sup_{x \in E} f(x).$$

Note that since M^n is finite the sup (resp. inf) is a max (resp. min).

Of interest within the literature devoted to belief functions is the notion of pignistic probability p^* of $\mathcal{M}$, that allows one to go from $\mathcal{P}_{\mathcal{M}}$ to a single probability within it. This pignistic probability taking the Shapley value of $\underline{P}(A)$ seen as a game and corresponds to

$$p^*(s) = \sum_{E, s \in E} \frac{m(E)}{|E|}.$$

Therefore, we can also see that taking the pignistic probability corresponds to considering a Laplacian assumption within each set E receiving positive evidence. Note also that if the mass $\mathcal{M}$ is equivalent to a probability p, then $p^* = p$, meaning that when one starts with a probability as uncertainty model, the pignistic criterion is just the standard expected utility for this probability.

Finally and before stepping to the decision part of our model, we will mention two particular cases of belief functions that are of interest to us: the first is the case of necessity measures, where the masses are given to nested elements, that is $m(E_i) \neq 0$ and $m(E_j) \neq 0$ iff $E_i \subset E_j$ or $E_i \subset E_j$, and the case of inner measures [3, Ch. 2.], where masses are given on a partition of the space, that is $m(E_i) \neq 0$ and $m(E_j) \neq 0$ if $E_i \cap E_j = 0$ and $\cup_{E, m(E) > 0} E = \mathcal{S}$ (or a subset of interest).

2.2 *Some Strategic Decisions*

In this model, we want to encode the strategic behavior of voters under uncertainty. To do so, let us define a function $st(s, a_i, a_i') : \mathcal{S} \times M^2 \to \mathcal{S}$ that sends back an updated score if voter i moves his vote from a_i to a_i' in its votes. Then, let $u_i : M \to \mathbb{R}$ be for each voter i the associated utility often reflects how satisfied the voter is with the current winner $\mathcal{W}_p(s)$. This utility function could be more complex, e.g., be defined on M^2 to denote pairwise preferences over candidates and represent how much some candidates are preferred to

other ones, however we will therefore stick to this definition depending only on M. Moreover, even though we have strict linear orders $\succ_i$ for each voter, we can always build a utility functions by Debreu et al. [2]. Therefore, a simple way to compute the benefit of moving from a_i to a_i' is to evaluate

$$u_i(st(s, a_i, a_i')) - u_i(s), \tag{5}$$

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

$$a_i' \succeq a_i \text{ iff } u_i(st(s, a_i, a_i')) - u_i(s) \geq 0$$

However, in our model this is not the case since agents receive a probability set $\mathcal{P}$. Assuming that we have some probability set $\mathcal{P}$ defined over $\mathcal{S}$, we will define four decision criteria (named DC) under uncertainty.

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

$$a_i' \succ_{pess} a_i \text{ iff } \underline{\mathbb{E}}_{\mathcal{P}}(u_i(st(S, a_i, a_i')) - u_i(S)) \geq 0$$

$$\text{and } \overline{\mathbb{E}}_{\mathcal{P}}(u_i(st(S, a_i, a_i')) - u_i(S)) > 0$$

  where S denotes uncertain scores whose knowledge is modelled by $\mathcal{P}$. In essence, this corresponds to the maximmin criterion put forward by Gilboa and Schmeidler [9].
- A pignistic criteria to decide whether an action is better than another as follows:

$$a_i' \succeq_{pig} a_i \text{ iff } \mathbb{E}_{p^*}(u_i(st(S, a_i, a_i')) - u_i(S)) \geq 0,$$

  where p^* is the pignistic probability. Note that it includes standard probabilistic decision as a special case.
- A mixture of criteria to decide whether an action is better than another as follows:

$$a_i' \succeq_{\alpha, p^*} a_i \text{ iff}$$

$$\text{iff } \alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S))$$

$$+(1-\alpha) \cdot \mathbb{E}_{p^*}(u_i(st(S, a_i, a_i')) - u_i(S)) \geq 0$$

  for some $\alpha \in [0, 1]$. Note that α can be seen as a level of completeness of the obtained, as it allows one to go from a very partial order ($\alpha = 0$) to a complete one ($\alpha = 1$).
- Another mixture of criteria to decide whether an action is better than another as follows:

$$a_i' \succeq_{H(\alpha)} a_i \text{ iff}$$

$$\text{iff } \alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S))$$

$$+(1-\alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) \geq 0$$

  with $\alpha \in [0, 1]$. In essence, this corresponds to the well-known Hurwicz criterion [11] that intends to balance between optimism ($\overline{\mathbb{E}}$) and pessimism ($\underline{\mathbb{E}}$), and that has been justified within a belief setting by Denoeux and Shenoy [4].

Note that for the three last rules, the strict relation $\succ$ corresponds to having strict positive values on the right side. Those decision rules mirror some of the commonly used decision rules within the imprecise probabilistic setting [26], and we hope to leverage other commonly used rules in future works.

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 $a_i' \succ a_i$ (we assume voters will only make a strategic move if their preferences are strict). To sum up, we can fully describe an election under plurality with a lexicographic tie breaking and strategic voting under uncertainties with the following tuple $(N, M, \succ, \rhd, \mathcal{P}, (u_i)_{i \in N}, DC)$.

3 Motivations of the Model

Our model is enriched because it includes the existing models from the literature but also provides a deeper understanding of the uncertainty in iterative voting in Plurality elections.

3.1 Relations with Existing Models

One of strength of our model is its power to unify two theories that were first considered as different, namely the one from Meir [14], Meir et al. [17] and the one from Myerson and Weber [21].

For the first one, we need to define a neighborhood of scores as Meir [14] with respect to a distance $d : S_r(s) = \{s' \in \mathcal{S} \mid d(s, s') \leqslant r\}$, d can be ℓ_1, ℓ_∞ or another one. By abuse of notation, we will use S_r when it is clear from the context. The seminal work of Meir et al. [17] is working on the special case $r = 0$, meaning the poll information is complete for plurality and certain. This corresponds to considering $\forall i \in N, m_i(S_0) = 1$, particular utilities functions:

$$\forall i \in N, u_i(st(s, a_i, a_i')) - u_i(s)$$

$$= \begin{cases} 1 & \text{if } st(s, a_i, a_i') \succ_i s \text{ and } \mathcal{W}_P(st(s, a_i, a_i')) = a_i', \\ 0 & \text{if } st(s, a_i, a_i') \succ_i s \text{ and } \mathcal{W}_P(st(s, a_i, a_i')) \neq a_i', \\ 0 & \text{if } st(s, a_i, a_i') \sim_i s, \\ -1 & \text{else,} \end{cases}$$

and the pessimistic decision rule $\succeq_{pess}$ to have Meir et al.'s model as a particular case. Recall that in Meir [14], the convergence is guaranteed and we call these strategic moves "direct best response" since all strategic moves where a voter deviates is to the new winner.

Meir et al. [18] then extended this framework to add uncertainty by considering some strictly positive value $r_i > 0$ for voter i. This comes down in our model to take, $\forall i \in N, m_i(S_{r_i}) = 1$, particular utilities functions

$$\forall i \in N, u_i(st(s, a_i, a_i')) - u_i(s) = \begin{cases} 1 & \text{if } st(s, a_i, a_i') \succ_i s, \\ 0 & \text{if } st(s, a_i, a_i') \sim_i s, \\ -1 & \text{else.} \end{cases}$$

and again the pessimistic decision rule $\succeq_{pess}$ to to have Meir's model as a particular case.

For the probabilistic model of Myerson and Weber [21], we define $\mathcal{P}$ as a probability $\mathbb{P}(s) \sim M(s, n)$ where s is the observed score and M a multinomial distribution. The utilities can be chosen arbitrarily provided it verifies the voter preferences:

$$u_i(st(s, a_i, a_i')) \geq u_i(s) \text{ iff } st(s, a_i, a_i') \succ_i s$$

Then, we need to remark that in Myerson and Weber [21] all voters are acting simultaneously and the decision is that each voter maximizes its expected utilities. The computation can be greatly simplified by only looking at the pivotal state and because we look only at plurality:

$$\mathbb{E}_{\mathbb{P}}(u_i(st(S, a_i, a_i')) - u_i(S))$$

$$= \sum_{a_i \in M} \mathbb{P}(a_i, a_i')[u_i(st(S, a_i, a_i')) - u_i(S)]$$

where $\mathbb{P}(a_i, a_i')$ is the probability to be pivotal by doing the action from a_i to a_i'.

We suggest two ways to build a richer model of uncertainty about the poll, either from the polling institute or from the voters themselves. Building polls, even when done honestly, comes with uncertainty because you cannot ask everyone's preferences (see the case of dishonest polls [20]). We can also interpret uncertainty from voters as they may be uncertain or have incomplete knowledge of their preferences. The next section presents present some examples to illustrate these ideas.

3.2 Examples

Our model not only includes other models from the literature but is also able to capture other situations. First of all, one of the weakness of the first model from Meir [14] is that any score of the neighborhood of scores is considered identically. Ideally, we would like to affect weights relative to the distance to the broadcast score. Let us detail that idea in the two following examples.

Example 1. *Let $|M| = 3$, $n = 30$ 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 $\mathcal{M}(S_3) = 0.2$.*

This model describes exactly what we expect with a belief function that models a certainty decreasing with respect to the distance to the broadcast poll s. This corresponds to building possibility distributions, that among imprecise probabilistic models are the most direct extensions of sets. They are therefore the ideal candidate to extend Meir [14]'s approach.

However, we may want to have a model 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 [3, Ch. 2.] on any event of the initial space $\mathcal{S}$. An intrinsic interest of this model is that it remains a probabilistic one, even if defined on an algebra where the atoms are sets of scores.

Example 2. *Let $|M| = 3$, $n = 30$ 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 $\mathcal{M}(S_3 \setminus S_2) = 0.2$.*

Moreover, we can also collect the lack of knowledge or incomplete information about voter's preferences. In fact, in most elections we usually have undecided voters especially far from the election day. Let us develop that idea in the two followings examples. In the following example, voter 1 is undecided between candidate a and candidate b, voter 2 and 3 are certain.

Example 3. *Let $|M| = 3$ and $n = 3$, $m_1(\{a, b\}) = 1$, $m_2(\{b\}) = 1$ and $m_3(\{c\}) = 1$. The joint belief function derived from this information over $\mathcal{S}$ is $\mathcal{M}(\{(1, 1, 1), (0, 2, 1)\}) = \mathcal{M}(([0, 1], [1, 2], 1)\}) = 1$.*

Now, we move to a last example that shows our model is able to capture more complex uncertainty. Indeed, we might have a voter (i.e. voter 1) that is able to tell she is not voting for c but it still hesitating between a and b, with a slight preference for a, and other voters being certain. The next example proposes a belief functions accounting for that.

Example 4. *Let $|M| = 3$ and $n = 3$, we could have $m_1(\{a\}) = 0.5$, $m_1(\{a, b\}) = 0.5$, so that $\underline{P}_1(\{a\}) = 0.5$, $\overline{P}_1(\{a\}) = 1$, and $\overline{P}_1(\{b\}) = 0.5$, meaning that voting for a is definitely more probable, yet voting for b remains plausible, albeit less. Provided we have $m_2(\{b\}) = 1$ and $m_3(\{c\}) = 1$, the joint mass over scores would be $\mathcal{M}(\{(1, 1, 1)\}) = 0.5, \mathcal{M}(([0, 1], [1, 2], 1)\}) = 0.5$.*

Let us now illustrate one of our decision rules on this same example, to give the reader an idea of how this plays out.

Example 5. *Consider for instance that the preference of voter 2 are $b \succ_2 c \succ_2 a$, which is coherent with its ballot and a Hurwicz criterion with $\alpha = \frac{1}{3}$, meaning that the weight of optimism is larger. We will consider the following utilities:*

$$\forall i \in N, u_i(st(s, a_i, a_i')) - u_i(s) = \begin{cases} 1 & \text{if } st(s, a_i, a_i') \succ_i s, \\ 0 & \text{if } st(s, a_i, a_i') \sim_i s, \\ -1 & \text{else.} \end{cases}$$

We want to evaluate the deviation from b to c. Then, with the linearity of the mass functions, we compute

$$\underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0$$

and

$$\overline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$$

Indeed, note that for the set $\{(1, 1, 1)\}$ having mass 0.5, the move makes c elected instead of a, and the utility is a constant 1. For the set $([0, 1], [1, 2], 1)$ that also has mass 0.5, the move is still beneficial for the possible score $(1, 1, 1)$, but would bring negative utility for $(0, 2, 1)$, as c would be elected instead of b, a clear downside for voter 2. Thus, we get:

$$\begin{aligned} &\frac{1}{3} \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) \\ &+(1 - \frac{1}{3}) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) \\ &= \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3} \geq 0. \end{aligned}$$

The conclusion is that voter 2 will do the deviation from b to c with such a decision criterion.

Note that we restrict our uncertainty models here to belief functions, yet some natural assessments would need to opt for the richer language of generic convex sets of probabilities, such as providing partial order between the probabilities of voting for some candidates [19]. Considering such cases is however beyond the scope of the current paper.

4 Convergence with Quantitative Uncertainty on Scores

In this section we generalize existing results on convergence from Meir [14] by adding quantitative uncertainty about the score as in Example 1 and Example 2. Let us recall that a neighborhood of scores S_r with respect to a distance d is defined as in Meir [14]: $S_r(s) = \{s' \in \mathcal{S} \mid d(s, s') \leqslant r\}$, where d can be ℓ_1, ℓ_∞, earth moving distance (EMD) or another one. 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. For each voter i, we will let r_i be the support of its uncertainty, so that S_{r_i} is the biggest set of voter i's uncertainty.

4.1 Extending the Meir Framework

Let us consider any belief functions of voter i, $\mathcal{M}_i : 2^{S_{r_i}} \to [0,1]$ that sums up to one, i.e., $\sum_{E \subseteq S_{r_i}, E \neq \emptyset} m_i(E) = 1$.

We will consider two particular cases of interest. The first is the case of nested sets, i.e. $\forall i \in N, \forall k \geqslant 1, m_i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$. Indeed, we might want to describe the fact that the belief mass is depending of the distance to the true score. Second, we look at the case of *partitioned* belief function, i.e. $\forall i \in N, \forall k \geqslant 2, m_i(S_k \setminus S_{k-1}) = \beta_k^i$, 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 $(\beta_k)_{1 \leqslant k \leqslant r_i}$, meaning that our evidence decreases as we get further from the observed s.

For the rest of this section, we will consider the following utilities that make sense in terms of improved or deteriorated outputs:

$$\forall i \in N, u_i(st(s, a_i, a_i')) - u_i(s) = \begin{cases} 1 & \text{if } st(s, a_i, a_i') \succ_i s, \\ 0 & \text{if } st(s, a_i, a_i') \sim_i s, \\ -1 & \text{else.} \end{cases}$$

Theorem 1. *Voters considering uncertainty given by a nested or partitioned 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 $\forall i \in N, \forall k \geqslant 1, m_i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$. We consider a feasible move $a_i \to a_{i'}$. We remark that the hypothesis that $\underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(s, a_i, a_i')) - u_i(s)) \geqslant 0$ implies that either:

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

since, if some neighbourhoods have a negative utility, this must be compensated by a positive contribution, which happens only if all moves within a neighborhood of scores the results. 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 $\exists s \in S_{\tilde{1}}, u_i(st(s, a_i, a_i')) - u_i(s) = 0$, showing that the second case never happens. Therefore, if

$$\underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(s, a_i, a_i')) - u_i(s)) \geqslant 0$$

this means that $\forall s \in S_{r_i}, u_i(st(s, a_i, a_i')) - u_i(s) \in \{0, 1\}$, with at least one s giving the null value.

Second, note that we cannot have $\overline{\mathbb{E}}_{\mathcal{M}}(u_i(st(s, a_i, a_i')) - u_i(s)) > 0$ if $\forall s \in S_{r_i}, u_i(st(s, a_i, a_i')) - u_i(s) = 0$, and there must be a situation for which making this strategic move is a local dominance move, meaning it verifies Theorem 4 from Meir [14].

We do exactly the same reasoning for the second type of belief functions, i.e. $\forall i \in N, \forall k \geqslant 2, m_i(S_k \setminus S_{k-1}) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$ and $\forall i \in N, m_i(S_1) = \beta_1^i$.

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: $\forall i \in N, \forall k \geqslant 1, m_i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$

Theorem 2. *Voters considering a 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 assume that the weights β_k^i of our belief functions, defined as $\forall i \in N, \forall k \geqslant 1, m_i(S_k) = \beta_k^i$.We consider a feasible move $a_i \to a_{i'}$. Using the fact that the lower (an upper) expectation is positive homogeneous, i.e., $\alpha \underline{\mathbb{E}}(f) = \underline{\mathbb{E}}(\alpha f)$, we get

$$\begin{aligned} &\alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) \\ &+ (1-\alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(st(S, a_i, a_i')) - u_i(S)) \\ &= \sum_{i=1}^{r_i} m(S_i)[\alpha \inf_{s \in S}(u_i(st(s, a_i, a_i')) - u_i(s)) \\ &+ (1-\alpha) \cdot \sup_{s \in S}(u_i(st(s, a_i, a_i')) - u_i(s))] \end{aligned} \tag{6}$$

Let us denote by

$$\tilde{u_i}(S, a_i, a_i') = \alpha \inf_{s \in S}(u_i(st(s, a_i, a_i')) - u_i(s))$$

$$+(1-\alpha) \cdot \sup_{s \in S}(u_i(st(s, a_i, a_i')) - u_i(s))$$

the term associated to subset S, meaning that Equation (6) can be rewritten as $\sum_{i=1}^{r_i} m(S_i)\tilde{u_i}(S, a_i, a_i')$, and that the strategic move is decided by a weighted average of $\tilde{u_i}$ values.

Denoting $\Delta u_i(s, a_i, a_i') = u_i(st(s, a_i, a_i')) - u_i(s)$, we distinguish six possible cases:

- Case A:

$$\exists s \in S, \Delta u_i(s, a_i, a_i') = 1 \text{ and } \forall s \in S, \Delta u_i(s, a_i, a_i') \geq 0$$

- Case B:

$$\exists s \in S, \Delta u_i(s, a_i, a_i') = -1 \text{ and } \forall s \in S, \Delta u_i(s, a_i, a_i') \leqslant 0$$

- Case C:

$$\exists s \in S, \Delta u_i(s, a_i, a_i') = -1 \text{ and } \exists s \in S, \Delta u_i(s, a_i, a_i') = 1$$

- Case D:

$$\forall s \in S, \Delta u_i(s, a_i, a_i') = 0$$

- Case E:

$$\forall s \in S, \Delta u_i(s, a_i, a_i') = 1$$

- Case F:

$$\forall s \in S, \Delta u_i(s, a_i, a_i') = -1$$

However, cases E and F 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

$$\alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(st(s, a_i, a_i')) - u_i(s))$$

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

then

$$\exists r_i'\leqslant r_i \quad \text{such that}$$
$$\exists s\in S_{r_i'},\Delta u_i(s,a_i,a_i')=1$$
$$\text{and } \forall s\in S_{r_i'},\Delta u_i(s,a_i,a_i')\geq 0 \quad \text{(Case A)}$$

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, it has to exist $r_i'\leqslant r_i$ 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 neighbourhood 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 of scores that receives enough evidence. Using the result from section VIII in Meir [13] 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. □

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. Note that for the rest of this chapter the notation card is used for cardinality. If S_{r_i} is our single set with positive mass, the corresponding strategic behaviour is equivalent having the following criterion:

$$a_i'\succeq_{pig}a_i$$
$$\text{iff}\quad \mathbb{E}_{\mathcal{U}}(u_i(st(S,a_i,a_i'))-u_i(S))\geq 0$$
$$\text{iff}\quad \mathrm{card}(s\in S\mid u_i(st(S,a_i,a_i'))-u_i(S)=1)$$
$$-\mathrm{card}(s\in S\mid u_i(st(S,a_i,a_i'))-u_i(S)=-1)\geq 0$$

When this criterion is not strict, it is obvious that we have a cycle because the same voter could move back and forth between two indistinguishable states. Therefore, we consider the strict version of it, namely

$$a_i'\succ_{\mathcal{U}}a_i$$
$$\text{iff}\quad \mathrm{card}(s\in S\mid u_i(st(S,a_i,a_i'))-u_i(S)=1)$$
$$-\mathrm{card}(s\in S\mid u_i(st(S,a_i,a_i'))-u_i(S)=-1)>0$$

Proposition 3. *With a uniform pignistic criterion on a neighborhood of scores of size 1 (with respect to the ℓ_1 distance), we are not guaranteed to reach convergence.*

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

a	$\succ_1$	b	$\succ_1$	c	$\succ_1$	d
c	$\succ_2$	a	$\succ_2$	b	$\succ_2$	d
d	$\succ_3$	c	$\succ_3$	a	$\succ_3$	b
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

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

$$S_1=\{(2,2,3,3),(1,2,3,3),(2,1,3,3),$$
$$(2,2,2,3),(2,2,3,2),(3,2,3,3),$$
$$(2,3,3,3),(2,2,4,3),(2,2,3,4)\}$$

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. □

Of course, if the size of the neighborhood of scores is larger, there is not much hope for convergence either. We think this example is quite interesting, in particular when put in perspective with our result of Theorem 1 about partitioned belief functions. Indeed, this latter result indicates that it is possible to consider a probability measure and a corresponding decision rule such that convergence holds, in the case where the probability measure is defined on an algebra coarser than the one induced by single score vectors. This indicate that the counter-example is not so much about having a probabilistic model itself than about the voters being perhaps too optimistic in their movement, hinting also at the fact that decision rules such as "maximax" ones, where one relies on the upper expectation alone are unlikely to lead to convergence of voting behaviours.

5 Conclusion and Future Works

In this paper, we provide new tools to model uncertainty in strategic voting from a quantitative perspective. We think that our work substantially modifies our view on uncertainty models around broadcast polls. In particular, we have shown that they capture standard settings in a single framework, and allow to extend classical results to more complex situations and new strategic decision rules. Moreover, we have provided some convergence results in our model, generalizing results from the iterative voting literature.

In addition, we introduce a novel view on uncertainty coming from the voter's perspective in Examples 3 and 4. The uncertainty could come from the voters themselves because they may not be able to fully express their preferences. We think this opens up vast avenues of research, as to our knowledge no strategic voting frameworks considered voter induced uncertainty, perhaps because identifying a central, broadcast score is then an ill-defined problem. However, it is often the case that voters can have, at some stage, incomplete preferences or uncertain ideas about who they are going to vote for, and are perhaps expecting more information to be more decisive.

We could however reframe the question of convergence of strategic results within this framework. Yet, in contrast with previous results where voters always receive as information a precise score issued from polls, and announce an updated precise ballot, they would start here from uncertain information about the voting result. Then, they may announce in the next round an updated uncertainty model in regard of the new information. Let us be a bit more concrete and illustrate this point with the following example.

Example 6. *Let* $|M| = 4$ *and* $n = 10$, $m_1(\{a,b\}) = 1$, $m_2(\{b,c\}) = 1$, $\forall i \in \{3,4\}, m_i(\{b\}) = 1$, $\forall j \in \{5,6,7\}, m_j(\{c\}) = 1$, $\forall l \in \{8,9,10\}, m_l(\{d\}) = 1$. *The joint belief function derived from this information over* $\mathcal{S}$ *is* $\mathcal{M}(([0,1],[2,4],[3,4],3)\}) = 1$.

The question is how the belief of voters 1 and 2 are going to be affected by the broadcast of the poll, namely the joint belief $\mathcal{M}$. For voter 1, we might want to say that this belief is going to become $m_1(\{b\}) = 1$, as a has no chances of being elected: voter 1, who was hesitating, may therefore report its vote to b. The situation is much less clear for voter 2, who may not want to make further commitments which is already an issue. Of course, this situation complexifies if we now consider that voters may provide more complex belief functions, such as the one in Example 4. One possible way to address this issue could be for instance to consider the recent framework introduced by Pomeret-Coquot et al. [22], that considers game and strategic moves in which agents' information is given in the form of belief functions, and where conditioning or updating of information is done in various ways. We however leave this endeavour for future work, as it would probably require to completely rethink and revisit the framework of strategic voting. Another added difficulty with respect to the standard framework is that the current decision/act of the voter (a_i) would have to be (re)defined, as well as what it means to move from it to consider another action (a_i'). One solution could be to consider that the voter is evaluating a set of potentially optimal actions, mirroring the fact that we are using decision rules such as $\succ_{pess}$ that results in partial orders among actions.

References

[1] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI conference on artificial intelligence (AAAI 2011)*, volume 25, pages 638–643, 2011.

[2] G. Debreu et al. Representation of a preference ordering by a numerical function. *Decision processes*, 3:159–165, 1954.

[3] D. Denneberg. *Non-additive measure and integral*, volume 27. Springer Science & Business Media, 2013.

[4] T. Denoeux and P. P. Shenoy. An interval-valued utility theory for decision making with dempster-shafer belief functions. *International Journal of Approximate Reasoning*, 124:194–216, 2020.

[5] S. Destercke and D. Dubois. Special cases. *Introduction to Imprecise Probabilities*, (chapter 4):79–91, 2014.

[6] P. Dey, N. Misra, and Y. Narahari. Complexity of manipulation with partial information in voting. *Theoretical Computer Science*, 726:78–99, 2018.

[7] U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. AAAI Press/International Joint Conferences on Artificial Intelligence, 2016.

[8] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

[9] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.

[10] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.

[11] L. Hurwicz. Optimality criteria for decision making under ignorance. Technical report, Cowles Commission discussion paper, statistics, 1951.

[12] D. Kreiss and T. Augustin. Undecided voters as set-valued information–towards forecasts under epistemic imprecision. In *Scalable Uncertainty Management: 14th International Conference, SUM 2020, Bozen-Bolzano, Italy, September 23–25, 2020, Proceedings 14*, pages 242–250. Springer, 2020.

[13] R. Meir. Plurality voting under uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

[14] R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

[15] R. Meir. Strategic voting. *Synthesis lectures on artificial intelligence and machine learning*, 13(1):1–167, 2018.

[16] R. Meir. *Strategic voting*. Springer Nature, 2022.

[17] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the AAAI conference on artificial intelligence (AAAI 2010)*, volume 24, pages 823–828, 2010.

[18] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 313–330, 2014.

[19] E. Miranda and S. Destercke. Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology*, 64:44–57, 2015.

[20] V. Mousseau, H. Surugue, and A. Wilczynski. Do we care about poll manipulation in political elections? In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI 2024)*, 2024.

[21] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *American Political science review*, 87(1):102–114, 1993.

[22] P. Pomeret-Coquot, H. Fargier, and É. Martin-Dorel. Games of incomplete information: A framework based on belief functions. *International Journal of Approximate Reasoning*, 151:182–204, 2022.

[23] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceeding of the 4th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, volume 4, page 8, 2012.

[24] M. A. Satterthwaite. 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.

[25] P. Smets and R. Kennes. The concept of distinct evidence. In *IPMU 92 Proceedings. pg*, pages 789–794, 1992.

[26] M. C. Troffaes. Decision making under uncertainty using imprecise probabilities. *International journal of approximate reasoning*, 45(1): 17–29, 2007.