

Voting on Actions with Uncertain Outcomes

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam

Abstract. We introduce a model for voting under uncertainty where a group of voters have to decide on a joint action to take, but the individual voters are uncertain about the current state of the world and thus about the effect that the chosen action would have. Each voter has preferences about what state they would like to see reached once the action has been executed. That is, we need to integrate two kinds of aggregation: beliefs regarding the current state and preferences regarding the next state.

1 Introduction

Imagine a group of *agents* who have to make a collective decision about what (joint) *action* to take. This action has to be chosen from a set of available actions. Each agent has her own *preferences* over the effects of actions, i.e., over the state of the world after a given action has been executed. Unfortunately, our agents are *uncertain* about the *current state* of the world and thus about the precise *effects* (or *outcomes*) of actions. Each agent only has a (possibly different) *set of states* she considers plausible (but they all agree on what the effect of executing a given action in a given state would be). What action should they take?

This is a collective decision making problem that involves the aggregation of two kinds of information: *social* information regarding the preferences over states of the world (effects of actions) and *epistemic* information regarding the plausibility of certain states being the actual current state of the world. This combination of concerns is what we should expect to encounter in a variety of application domains, e.g., when devising mechanisms for teams of autonomous software agents to interact and agree on actions to pursue collectively. However, while social choice theory has, rightly, been argued to be relevant to multiagent systems [17, 7], the standard model of social choice only deals with the aggregation of preferences [10]. In this paper we put forward a model for voting on actions with uncertain effects that integrates this standard perspective with a simple notion of uncertainty.

We model uncertainty in the simplest possible way: a voter can only distinguish between states she considers plausible and those she does not consider plausible. This model of uncertainty is sometimes called *strict uncertainty* or *complete ignorance* [5]. It is different from much work on reasoning under uncertainty, which often assumes the availability of a probability distribution over the set of possible states [13]. Our chosen representation of voter preferences is also simpler than what is typically assumed in the literature on decision-making

under uncertainty: rather than endowing each voter with a utility function mapping possible outcomes to numerical values, we only assume that they are able rank the outcomes in terms of their relative desirability. There are two reasons for choosing such simple representations of uncertainty and preference. First, it arguably is often not realistic to expect a decision-maker to be able to provide either a fully specified probability distribution (e.g., how should she judge whether a given state has probability 15% or 20% of being the actual state of the world?) or a precise utility function (what does it mean to assign utility 20 rather than 19 to a given outcome?). Second, we want the individual components of our model to be as simple as possible, so as to be able to better focus on the analysis of their interplay.

The remainder of this paper is organised as follows. We begin, in Section 2, with a detailed discussion of the challenges associated with voting under uncertainty and present three paradoxes one might encounter in this context. After presenting our formal model in Section 3, we then focus on two specific aspects of the general problem: aggregating information regarding the current state only, without considering preferences (in Section 4); and integrating uncertainty and preference information for the special case of a single agent (in Section 5). We conclude in Section 6 by using the insights gained to make a first tentative proposal for best practices for voting on actions with uncertain outcomes.

2 Three Paradoxes

In this section we introduce three paradoxes—scenarios that show how a seemingly reasonable approach to voting on actions with uncertain effects can lead to suboptimal outcomes. Our presentation is organised in terms of the point in the aggregation process at which the uncertainty about the current state of the world is being resolved.

For each paradox we also briefly discuss similarities to related phenomena from different strands of the literature on social choice theory, namely judgment aggregation, preference aggregation and voting, and ranking sets of objects.

2.1 The Paradox of Individual Uncertainty Resolution

Suppose there are two possible states of the world, A and B . Whichever state we are in, we can execute one of two actions: to *change* to the other state or to *stay* in the current state. There are three agents, who are all uncertain about the current state. Agents 1 and 2 believe A is most plausible; agent 3 believes B is most plausible. They have to agree on one of the two actions to be executed, and they each have their own preferences about what state they would like to end up in afterwards. Agent 1 prefers A ; agents 2 and 3 both prefer B .

If each agent is left to resolve their uncertainty regarding the current state on their own, then each of them will infer their “most preferred action” based on their preferences over states and based on what they believe to be the most plausible current state. That is, agents 1 and 3 will choose to *stay*, and only

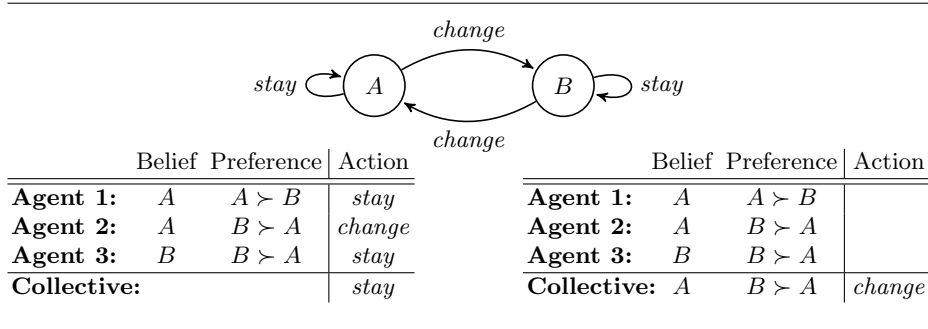


Fig. 1. The Paradox of Individual Uncertainty Resolution

agent 2 will choose to *change*. If we use the majority rule to decide between these two options, we will thus decide that the best move for the collective is to *stay* in the current state (see lefthand side of Fig. 1).

But this, arguably, is not the best course of action for this group. If instead the agents first use the majority rule to aggregate their information regarding the current state (finding *A* to be most plausible) and then to aggregate their preferences (finding *B* to be socially preferred), then the appropriate action to take would be to *change* the state (see righthand side of Fig. 1).

That is, resolving uncertainties regarding the current state individually before aggregation can lead to outcomes that, arguably, are suboptimal.

This paradox is related to the *discursive dilemma* familiar from judgment aggregation [15]. In judgment aggregation, we are asked to aggregate the views of several individuals regarding the truth or falsity of a number of formulas of propositional logic, and the discursive dilemma is a family of paradoxical situations that we may encounter in this framework. We can model our paradox as a problem of judgment aggregation as follows. Let p stand for “*A* is the most plausible current state” (and thus $\neg p$ for “*B* is the most plausible current state”); let q stand for “*A* is the most preferred next state” (and thus $\neg q$ for “*B* is the most preferred next state”). Then we should choose the action *stay* if and only if $p \leftrightarrow q$ is true (i.e., we should choose *change* if it is false). Then the following situation corresponds to the scenario described earlier:

	p	q	$p \leftrightarrow q$
Agent 1:	Yes	Yes	Yes
Agent 2:	Yes	No	No
Agent 3:	No	No	Yes
Majority:	Yes	No	?

The righthand side of Fig. 1 corresponds to what is known as the *premise-based procedure* in judgment aggregation: we use the majority rule to obtain a collective judgment on p and q and then use logical inference to decide that $p \leftrightarrow q$ must be false. The lefthand side of Fig. 1 corresponds to the *conclusion-based procedure*: to obtain a collective judgment on $p \leftrightarrow q$ we only consider the

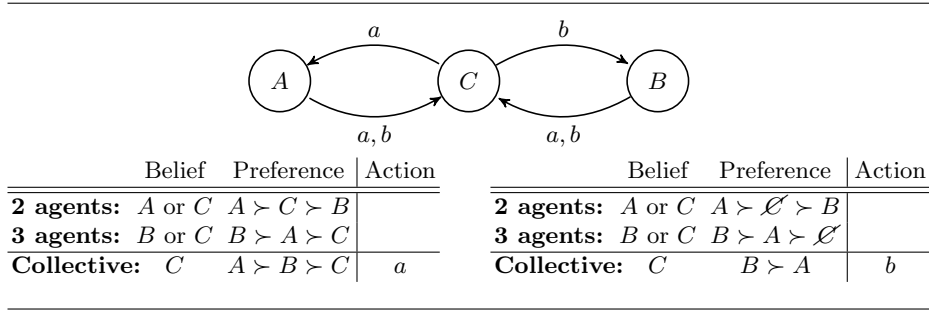


Fig. 2. The Paradox of Late Collective Uncertainty Resolution

individual judgments on the same formula and find that a majority is in favour of labelling it as being true. Thus, two seemingly reasonable forms of aggregation yield contradictory advice, i.e., we are facing a discursive dilemma.

2.2 The Paradox of Late Collective Uncertainty Resolution

Suppose there are three possible states, A , B and C , and two available actions, a and b . Executing a in state C takes us to A ; executing b in C takes us to B ; and executing either a or b in either state A or B will take us to C . We have two agents who are uncertain whether the current state is A or C , and for the next state they both prefer A over C over B ; and we have three agents who are uncertain between B and C , and they all prefer B over A over C .

Let us use the *Borda rule* to aggregate preferences: for each agent a state gets as many points as that agent ranks other states below the state in question (and we then order the states in terms of the points received). That is, A gets $2 \cdot 2 + 3 \cdot 1 = 7$ points, B gets $2 \cdot 0 + 3 \cdot 2 = 6$ points, and C gets $2 \cdot 1 + 3 \cdot 0 = 2$ points, i.e., we obtain the collective preference order $A \succ B \succ C$. To resolve the uncertainty regarding the current state, there really is only one natural choice, namely to take the state considered plausible by the highest number of agents, which is C . A , the collectively most preferred state, is reachable from C , namely by executing action a . Hence, we should execute a (see lefthand side of Fig. 2).

But now consider this: if C is the current state, which is the collectively most plausible assumption to make here, then C cannot be the *next* state, whichever action we execute (if we execute a , we end up in A ; if we execute b , we end up in B). So, arguably, our agents' preferences regarding C are not relevant. Then we have two agents who prefer A over B , and three agents who prefer B over A . Hence, the collectively most preferred state is B (under the Borda rule, as well as under any other reasonable aggregator) and we should execute action b to reach it (see righthand side of Fig. 2).

That is, by having postponed the collective uncertainty resolution until after the step of preference aggregation, we have missed the opportunity to eliminate irrelevant information from our aggregation problem which, arguably, has led to a suboptimal outcome.

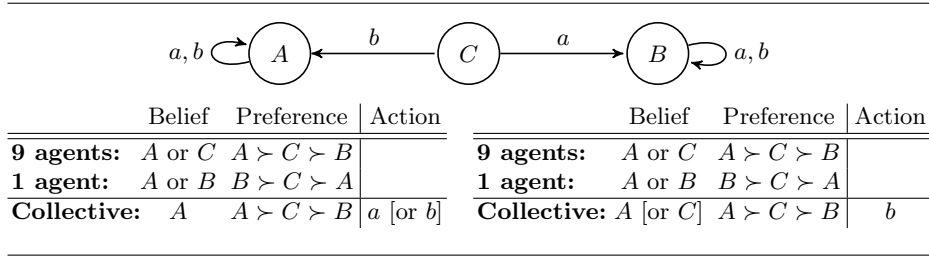


Fig. 3. The Paradox of Early Collective Uncertainty Resolution

This paradox is closely related to Arrow’s *independence of irrelevant alternatives* [2]: Arrow postulated that for any reasonable form of preference aggregation the relative collective ranking of two alternatives should only depend on their relative rankings provided by the individuals, and not on any third (“irrelevant”) alternative. Our example demonstrates that the Borda rule, when used as an aggregator for preference orders, violates this desideratum: the relative ranking of A and B adopted by the collective *does* depend on C . Our paradox also has close connections to the topic of election control by means of adding (or deleting) candidates, widely studied in computational social choice [4, 9], which in turn relies on violations of Arrow’s independence axiom: our example demonstrates how adding candidate C to a Borda election in which B was winning can result in A becoming the new winner.

2.3 The Paradox of Early Collective Uncertainty Resolution

Suppose again that there are three possible states, A , B and C , and two available actions, a and b . When in state C , action a will take us to state B , while b will take us to A . If the current state is either A or B , then neither action will change the current state. There are ten agents. Nine of them consider A and C plausible states and prefer A over C over B . The remaining agent considers A and B plausible and prefers B over C over A .

If we aggregate the information regarding the plausibility of different states, we find that A is the most plausible state (but C comes a close second). Now suppose we insist on aggressively exploiting this information to pinpoint the most plausible current state: that is, we take A to be that state. In state A it does not matter what action we execute; the next state will always be A again. Suppose that in such a case, by default, the lexicographically first action is chosen, i.e., a in this case (see the lefthand side of Fig. 3).

But this clearly is a suboptimal choice: C is almost as likely to be the current state as A . If it really is C , then executing (the chosen) action a will result in state B , while executing (the dismissed) action b will result in state A . Considering that the collective preference order is clearly $A \succ C \succ B$ (for any reasonable preference aggregation rule), executing b therefore would be a much better choice. We could obtain this outcome by delaying uncertainty resolution:

if we consider both A and C plausible current states and compute the best action for either case, we find that in the first case it makes no difference which action is chosen, while in the second case b is better (see the righthand side of Fig. 3).

That is, resolving the uncertainty regarding the current state too early, even if we take the information supplied by all agents into account, can lead to a suboptimal outcome.

There are connections here to problems analysed in the literature on *ranking sets of objects* [14, 3, 11]. The question discussed in that literature is how to extend a preference order over individual objects to a preference order over nonempty sets of such objects (with the most common interpretation of those sets being that we will eventually obtain one of the objects in the set in question, but cannot control which). One of the most basic axioms formulated in this literature postulates that, if we strictly prefer A over B , then we should (at least weakly) prefer $\{A\}$ over $\{A, B\}$ —because, whatever choice rule may get used to select from the second set, the outcome can never be worse than A . This is precisely why, in our scenario above, we should prefer action b (which can only result in A when A and C are considered the only plausible states) over action a (which might result in either A or B).

3 A Formal Model

Fix a finite set \mathcal{Q} of *states* and a finite set Σ of *actions*. A *transition function* $\delta : \mathcal{Q} \times \Sigma \rightarrow \mathcal{Q}$ determines for any given state $q \in \mathcal{Q}$ and any given action $\sigma \in \Sigma$ the state q' that will be reached after executing σ in q . In other words, δ defines the *effect* of an action for every possible state. Note that δ is a total function, i.e., every action is executable from every state—but it may well be the case that an action σ has no effect in a given state q in the sense that $\delta(q, \sigma) = q$. We use $\delta(\sigma, \mathcal{Q}) = \{\delta(q, \sigma) \mid q \in \mathcal{Q}\}$ to refer to the set of states that are *reachable* from \mathcal{Q} via σ , i.e., the set of states we might reach by executing action σ in a situation where any of the states in $\mathcal{Q} \subseteq \mathcal{Q}$ is a plausible current state.

Fix a finite set $\mathcal{N} = \{1, \dots, n\}$ of agents. We assume that each agent has complete knowledge of δ . On the other hand, an agent does not necessarily know the identity of the current state. We express an agent's uncertainty regarding the current state in terms of a (nonempty) subset of \mathcal{Q} : for each agent $i \in \mathcal{N}$, let $Q_i \subseteq \mathcal{Q}$ denote the set of states she considers possible. This is a very minimalist approach towards modelling uncertainty; other options include defining a probability distribution or a plausibility ranking on \mathcal{Q} for each agent.

Agents have preferences over states, which we model in terms of linear orders (i.e., binary relations that are irreflexive, transitive, and complete). We write \succ_i for the preference order of agent $i \in \mathcal{N}$. Let $\mathcal{L}(\mathcal{Q})$ denote the set of all linear orders on \mathcal{Q} and let $\Pi(\mathcal{Q})$ denote the set of nonempty subsets of \mathcal{Q} . That is, each agent provides us with an element of $\Pi(\mathcal{Q}) \times \mathcal{L}(\mathcal{Q})$, i.e., a pair consisting of an uncertainty set and a preference order. A *profile* is a vector of n such pairs, one for each agent, i.e., it is an element of $[\Pi(\mathcal{Q}) \times \mathcal{L}(\mathcal{Q})]^n$.

We are now ready to define our main concept: an aggregation mechanism that accepts a profile, i.e., the beliefs and preferences of each agent, and then returns a single action to be executed by the group. Our definition is relative to a given transition function δ (while \mathcal{Q} , Σ and \mathcal{N} are taken to be fixed throughout).

Definition 1. *Let δ be a transition function. Then a **social action choice function** for δ , $F_\delta : [\Pi(\mathcal{Q}) \times \mathcal{L}(\mathcal{Q})]^n \rightarrow \Sigma$, is a mapping from profiles to actions.*

Note that we insist on F_δ returning a *single* action. Alternatively, we could have defined F_δ as an *irresolute* aggregator returning a nonempty set of most preferred actions and considered the problem of tie-breaking, i.e., of choosing one best action to actually execute, as a separate problem.

A closely related problem is to derive a *weak order* (i.e., a binary relation that is reflexive, transitive, and complete) on the available actions from a given profile. This is useful, for instance, when we are not yet certain which actions will become available and we want to prepare for executing the collectively best available action. We write $\mathcal{W}(\Sigma)$ for the set of all weak orders on Σ .

Definition 2. *Let δ be a transition function. Then a **social action ranking function** for δ , $F_\delta : [\Pi(\mathcal{Q}) \times \mathcal{L}(\mathcal{Q})]^n \rightarrow \mathcal{W}(\Sigma)$, is a mapping from profiles to weak orders on actions.*

When agents are certain about the identity of the current state of the world, then social action choice functions essentially correspond to what are known as (resolute) *social choice functions* in the literature, and *social action ranking functions* correspond to *social welfare functions* [10].

4 Uncertainty Resolution in Isolation

Our model deals with problems that require both the aggregation of preferences and the aggregation of beliefs regarding the identity of the current state. The former, when viewed in isolation, is the main problem studied in classical social choice theory [10]. In this section we want to study the latter in isolation, i.e., we want to study the problem of uncertainty resolution on the basis of the reports of the individual agents. Suppose each agent reports a nonempty sets of states (a subset of \mathcal{Q}). How should we aggregate such an *uncertainty profile* into a single collective set of states that appropriately reflects the beliefs of the group?

Definition 3. *An **uncertainty resolution rule** $F : \Pi(\mathcal{Q})^n \rightarrow \Pi(\mathcal{Q})$ is a mapping from an uncertainty profile to a single nonempty set of states.*

We use $\mathbf{Q} = (Q_1, \dots, Q_n)$ to refer to a profile of sets of states, with Q_i being the set reported by agent i . Let $N_q^{\mathbf{Q}} = \{i \in \mathcal{N} \mid q \in Q_i\}$ denote the set of agents who include state q in their set under profile \mathbf{Q} .

Structurally, an uncertainty resolution rule has the same form as a voting rule based on *approval ballots* [16], of which *approval voting* is the main representative [6]: such a voting rule takes as input a set of candidates from each voter

and returns a winning candidate as output (or possibly a set of tied winners). Indeed, the most natural choices for an uncertainty resolution rule all correspond to rules discussed in the literature on voting:

- Under *approval voting* we return the states with maximal support [6]:

$$F(\mathbf{Q}) = \operatorname{argmax}_{q \in \mathcal{Q}} |N_q^{\mathbf{Q}}|$$

- Under *even-and-equal cumulative voting* each agent evenly distributes a total of weight 1 over the states they report as plausible and the state(s) with the maximal sum of weights are being returned [1]:

$$F(\mathbf{Q}) = \operatorname{argmax}_{q \in \mathcal{Q}} \sum_{i \in N_q^{\mathbf{Q}}} \frac{1}{|Q_i|}$$

- Under the *mean-based rule* we return all those states that receive at least an average amount of support [8]:

$$F(\mathbf{Q}) = \left\{ q \in \mathcal{Q} \mid |N_q^{\mathbf{Q}}| \geq \frac{|Q_1| + \dots + |Q_n|}{|\mathcal{Q}|} \right\}$$

To see that above equality correctly formalises the mean-based rule, note that $|Q_1| + \dots + |Q_n| = \sum_{q \in \mathcal{Q}} |N_q^{\mathbf{Q}}|$.

Observe that in case there is at least one state that is reported by every agent, approval voting yields the same result as taking the intersection of all individual sets of states. Let us call this the *intersection rule*.

We now want to consider uncertainty resolution rules from an axiomatic point of view and formulate desiderata to characterise appropriate rules. The simplest such desideratum is *neutrality*, a standard concept in social choice theory, which asks that the chosen rule should treat all states symmetrically. Our formulation below is closest to how neutrality has been formalised in the literature on judgment aggregation [12].

Definition 4. *An uncertainty resolution rule F is called **neutral** if, for all profiles \mathbf{Q} and all states q and q' , $N_q^{\mathbf{Q}} = N_{q'}^{\mathbf{Q}}$ implies $q \in F(\mathbf{Q}) \Leftrightarrow q' \in F(\mathbf{Q})$.*

Before we formulate further desiderata, we need to clarify the semantics of the set of states reported by an agent. Do agents report what they *know* about the identity of the current state or do they merely report their *beliefs*?

4.1 Uncertainty Resolution when Agents Report Knowledge

Assuming that *agents report knowledge* means assuming that, for each agent, the true state of the world is an element of the set reported by that agent. In particular, this means that $Q_1 \cap \dots \cap Q_n \neq \emptyset$. This is not as unrealistic an assumption as it may seem: after all, each agent has the option to report the full set \mathcal{Q} , i.e., even a completely ignorant agent can report true (albeit vacuous) information. If we *trust* in our assumption of agents reporting knowledge, then we can exclude any state not reported by all agents.

Definition 5. *Suppose agents report knowledge. An uncertainty resolution rule is called **trustful** if, for all profiles \mathbf{Q} , we have that $F(\mathbf{Q}) \subseteq Q_1 \cap \dots \cap Q_n$.*

Proposition 1. *Suppose agents report knowledge. Then an uncertainty resolution rule is both trustful and neutral if and only if it is the intersection rule.*

Proof. Clearly, the intersection rule is both trustful and neutral. For the other direction, suppose F is an uncertainty resolution rule that is trustful and neutral. By virtue of F being trustful, we have $F(\mathbf{Q}) \subseteq Q_1 \cap \dots \cap Q_n$. Now, for the sake of contradiction, assume that $F(\mathbf{Q})$ is a *proper* subset of $Q_1 \cap \dots \cap Q_n$. Then there exist a state $q \in F(\mathbf{Q})$ and another state $q' \in Q_1 \cap \dots \cap Q_n \setminus F(\mathbf{Q})$. That is, despite having $N_q^{\mathbf{Q}} = N_{q'}^{\mathbf{Q}}$ (both are equal to the full set \mathcal{N}), we do not have $q \in F(\mathbf{Q}) \Leftrightarrow q' \in F(\mathbf{Q})$, i.e., we have observed a violation of neutrality. \square

4.2 Uncertainty Resolution when Agents Report Mere Beliefs

If agents merely report beliefs regarding the identity of the current state, then the intersection rule ceases to be a viable option (as it may then return the empty set). In this case, approval voting is maybe the most natural choice. Even-and-equal cumulative voting is appropriate if we have reason to assume that agents who report small sets do so because they possess more accurate information. The mean-based rule is an attractive choice if we do not want to exclude too many alternatives before also considering preference information. Axiomatisations of all these rules are available in the literature [18, 1, 8].

5 The Single-Agent Case

In this section we discuss the special case where there is just a single agent. Even in this severely simplified scenario we still face the challenge of integrating belief and preference information. As we shall see, this special case is interesting in its own right, but it is also relevant to the more general aggregation problem discussed in this paper. The reason is that one natural approach to voting under uncertainty is to first aggregate belief and preference information independently and to then decide what action to take based on the thus obtained collective belief and collective preference order.

For the case of $n = 1$, Definitions 1 and 2 simplify to what we shall call *action choice functions* and *action ranking functions*, respectively. In the sequel, we focus on the latter, so as to be able to discuss connections to the literature on ranking sets of objects [14, 3, 11].

For any $Q \subseteq \mathcal{Q}$ and any linear order \succ on \mathcal{Q} , let \succ_Q represent the weak order returned by F_δ (i.e., \succ_Q denotes $F_\delta(Q, \succ)$). That is, we write $\sigma \succ_Q \sigma'$ to say that σ is *at least as good an action to take* as σ' . We furthermore write $\sigma \succ_Q \sigma'$ in case σ is the strictly better action to take and $\sigma \sim_Q \sigma'$ in case both actions are equally good choices, i.e., \succ_Q is the *strict part* of \succ_Q and \sim_Q is the *indifference part* of \succ_Q . Finally, \succ (without a subscript) denotes the reflexive closure of \succ (our agent's strict preference order over outcomes).

5.1 Desiderata for Action Choice Functions

We now want to formulate desiderata (or *axioms*, in the terminology of social choice theory) for an action ranking function F_δ . Recall that, given a set Q of plausible current states, every action σ gives rise to a set of plausible outcomes $\delta(\sigma, Q)$. There are two approaches to formulating desiderata:

- *Outcome-based desiderata*: First, we may choose to formulate desiderata regarding the relative ranking of actions purely in terms of the sets of plausible outcomes they correspond to.
- *Case-based desiderata*: Alternatively, we may also take into account which of the plausible current states would give rise to what outcomes and formulate our desiderata case by case.

We begin with the former approach:

Definition 6. *An action ranking function satisfies **outcome-dominance** if, for all $\sigma, \sigma' \in \Sigma$, we have $\sigma \succ_Q \sigma'$ whenever the following conditions hold:*

- (i) $q \succ q'$ for all $q \in \delta(\sigma, Q) \setminus \delta(\sigma', Q)$ and all $q' \in \delta(\sigma', Q)$
- (ii) $q \succ q'$ for all $q \in \delta(\sigma, Q)$ and all $q' \in \delta(\sigma', Q) \setminus \delta(\sigma, Q)$
- (iii) $\delta(\sigma, Q) \neq \delta(\sigma', Q)$

Outcome-dominance is known as the *Gärdenfors principle* in the literature on ranking sets of objects [3]. It is a natural requirement to impose: it classifies any change of action as an improvement if it results either in a new best possible outcome to be added or in the currently worst possible outcome to be eliminated.

An example for an action ranking function that satisfies outcome-dominance is the *max-min ordering*: Prefer σ over σ' (with Q being the plausible states) if and only if either $\max(\delta(\sigma, Q)) \succ \max(\delta(\sigma', Q))$ or both $\max(\delta(\sigma, Q)) = \max(\delta(\sigma', Q))$ and $\min(\delta(\sigma, Q)) \succ \min(\delta(\sigma', Q))$.¹ That is, under the max-min ordering we first check for which action the best possible outcome is better and in case that is not enough to differentiate the two actions, we go by the worst possible outcome for either action.

Definition 7. *An action ranking function satisfies **casewise-dominance** if $\sigma \succ_Q \sigma'$ whenever $\delta(q, \sigma) \succ \delta(q, \sigma')$ for all states $q \in Q$, and that preference is strict in at least one case.*

That is, casewise-dominance says that we should prefer an action if it does at least as well or better than a competing action for every state we consider possible. This certainly is a property we would like to see satisfied.

An example for an action ranking function that satisfies it is what we shall call a *casewise-lexicographic ordering*: Fix any strict linear order \gg on \mathcal{Q} . We want to rank σ and σ' (with Q being the plausible states). First consider the state q_1 that is maximal in Q with respect to \gg . If $\delta(q_1, \sigma) \succ \delta(q_1, \sigma')$, then prefer σ ; if $\delta(q_1, \sigma') \succ \delta(q_1, \sigma)$, then prefer σ' ; and if $\delta(q_1, \sigma) = \delta(q_1, \sigma')$, then

¹ Here max and min are defined with respect to the given order \succ on states.

postpone the decision and compare σ and σ' for state q_2 , the next best state with respect to \gg . Continue until one action is found to be preferred or until all states in Q have been exhausted (in which case you declare indifference).

Proposition 2. *Outcome-dominance does not imply casewise-dominance, nor does casewise-dominance imply outcome-dominance.*

Proof (sketch). To see this, check that casewise-lexicographic orderings violate outcome-dominance and the max-min ordering violates casewise-dominance. \square

Definition 8. *An action ranking function satisfies **outcome-relevance** if $\sigma \sim_Q \sigma'$ whenever $\delta(\sigma, Q) = \delta(\sigma', Q)$.*

That is, outcome-relevance says that we should be indifferent between any two actions with identical sets of reachable states (given our beliefs regarding the current state). Intuitively speaking, outcome-relevance is much weaker a condition than outcome-dominance, but we note that technically the latter does not imply the former. The reason is that outcome-dominance does not prescribe how to rank two actions that produce the same outcome set.

5.2 Outcome-based Desiderata and Ranking Sets of Objects

Our next result shows that we are dealing with a framework that is at least as expressive as the classical framework for *ranking sets of objects* [3].² Consider a problem of the latter kind, with m objects. To embed this into an action-ranking problem, we will show how to construct a scenario with a set \mathcal{Q} of m states (corresponding to the m objects) and a set Σ of $2^m - 1$ actions (corresponding to the $2^m - 1$ nonempty sets of objects), where for every possible nonempty set $Q \subseteq \mathcal{Q}$ of states there exists an action $\sigma \in \Sigma$ such that Q is equal to the set of plausible outcomes associated with σ .³

Proposition 3. *For any set \mathcal{Q} of states, we can construct a set Σ of actions and a state transition function δ such that for every nonempty set $Q \subseteq \mathcal{Q}$ there exists exactly one action $\sigma \in \Sigma$ with $\delta(\sigma, \mathcal{Q}) = Q$.*

Proof. Let \mathcal{Q} be an arbitrary set of states with $|\mathcal{Q}| = m$. Define $\Sigma := \Pi(\mathcal{Q})$, i.e., every action is associated with a nonempty set of states. Fix an arbitrary order on \mathcal{Q} so that $\min(Q)$ is well-defined for every set $Q \in \Pi(\mathcal{Q})$. That is, $\min(\sigma) \in \sigma$ for every action σ . Now define a transition function δ as follows:

$$\delta(q, \sigma) = \begin{cases} q & \text{if } q \in \sigma \\ \min(\sigma) & \text{otherwise} \end{cases}$$

² At this point we shall assume familiarity with the standard model used in the relevant literature [3].

³ For our specific construction, this will be the case when the set of states considered to be plausible as current states is the full set \mathcal{Q} .

Now suppose our agent considers the entire set \mathcal{Q} to be the set of plausible states. Then, given our definition of δ , after executing action σ the states she will consider possible outcomes are precisely the states in the set corresponding to σ . That is, $\delta(\sigma, \mathcal{Q}) = \sigma$. As every nonempty set $Q \subseteq \mathcal{Q}$ is in fact equal to some action σ , we are done. \square


The interest of this result stems from the fact that any of the impossibility results about satisfying certain combinations of desiderata known for ranking sets of objects [14, 3, 11] are now inherited by our model. To be sure, this is true only if we allow for arbitrary transition functions. An interesting direction for future work would be to investigate to what extent more positive results are attainable when we limit, for instance, the number of distinct follow-up states any given action may give rise to.

5.3 Impossibility of Satisfying Both Kinds of Desiderata

Our final result shows that, maybe somewhat surprisingly, casewise-dominance (which we suggested to be highly desirable) is in fact incompatible with the weakest possible outcome-related requirement:

Proposition 4. *There exists no action ranking function that satisfies both casewise-dominance and outcome-relevance.*

Proof. Consider a scenario with three states and three actions: $\mathcal{Q} = \{A, B, C\}$ and $\Sigma = \{a, b, c\}$. Suppose our agent considers all three states plausible and has the preference order $A \succ B \succ C$. Furthermore, suppose the definition of the transition function δ is given by the following figure:



Now, for the sake of contradiction, let us try to construct an action ranking function producing the weak order $\succsim_{\mathcal{Q}}$ for this scenario. We can derive the following relative rankings of actions:

- $a \succ_{\mathcal{Q}} b$. This follows from casewise-dominance: if the current state is A or C , then actions a and b have the same effect; if the current state is B , then a will result in A , while b will only result in B .
- $b \succ_{\mathcal{Q}} c$. This also follows from casewise-dominance: if the current state is A or C , then actions b and c have the same effect; if the current state is B , then b will result in B , while c will only result in C .
- $a \sim_{\mathcal{Q}} c$. This, finally, follows from outcome-relevance, given that we have $\delta(a, \mathcal{Q}) = \delta(c, \mathcal{Q}) = \{A, C\}$.

But there can be no weak order $\succsim_{\mathcal{Q}}$ that satisfies all three constraints, as the first two constraints entail that a should be strictly preferred to c , while the third requires that they should be equally preferred. \square

It is not too hard to see that this impossibility result extends to the natural adaptations of our desiderata to action *choice* functions.

6 Conclusion

We have introduced a simple model for voting under uncertainty and demonstrated its interest through three paradoxes, i.e., situations that show how certain seemingly reasonable choices in that model can lead to unexpected and undesirable outcomes. We have also seen that certain components of the model already give rise to interesting questions. Specifically, we have discussed the problem of aggregating several sets of plausible states into a single such set and we have discussed how to rank available actions when given a single set of plausible states and a single preference order. Finally, while we believe that the main contribution of this paper is of a conceptual nature, we have also established some basic technical results regarding the aforementioned components of the model.

But we have not yet answered our initial question: how should we make a decision when voting on actions with uncertain effects?

Of course, as for standard voting (without uncertainty), we cannot expect a definitive answer to this question: no method of aggregation will satisfy all desiderata that one might wish to impose. Still, our results and discussion above suggest at least some tentative guidelines for best practices. First, while full integration of uncertainty resolution and preference aggregation is desirable in principle, given our current understanding of the problem domain, the best pragmatic approach we can recommend is to first aggregate beliefs and preferences in isolation, and to then integrate the collective beliefs and the collective preferences thus obtained into a decision regarding the action to take. Second, for each of the three phases of this process, we can make some recommendations:

- *Uncertainty resolution*: To avoid the *Paradox of Early Collective Uncertainty Resolution*, uncertainty resolution should not be overly aggressive, i.e., we should not aim at excluding too many possible states. In case the agents are known to report *knowledge* (rather than mere belief) regarding the current state, the intersection rule is the only reasonable choice. Otherwise, the mean-based rule promises to offer a good compromise.
- *Preference aggregation*: It is advisable to use a social welfare function (returning a collective preference order) rather than just a voting rule (returning one or several top states) at this stage, so as to be able to report more information to the next stage. Preference aggregation is subject to the well-understood challenges of social choice, but, e.g., the Kemeny rule is often regarded as an aggregator that makes a good trade-off between desiderata [10], and it may also be a good compromise in our context.
- *Integration*: Finally, we have to integrate the collective beliefs and the collective preference order into a final decision regarding the action to choose. This corresponds to the single-agent case analysed in Section 5. Here a case-based approach appears superior to an outcome-based approach. For instance, we

could use a casewise-lexicographic rule. The “salience order” \gg used to initialise this rule could even refer back to the full profile and favour states proposed by larger numbers of individuals.

References

1. Alcalde-Unzu, J., Vorsatz, M.: Size approval voting. *Journal of Economic Theory* 144(3), 1187–1210 (2009)
2. Arrow, K.J.: *Social Choice and Individual Values*. John Wiley & Sons, 2nd edn. (1963)
3. Barberà, S., Bossert, W., Pattanaik, P.: Ranking sets of objects. In: *Handbook of Utility Theory*, vol. 2. Kluwer Academic Publishers (2004)
4. Bartholdi, III, J.J., Tovey, C.A., Trick, M.A.: How hard is it to control an election? *Mathematical and Computer Modelling* 16(8–9), 27–40 (1992)
5. Ben Larbi, R., Konieczny, S., Marquis, P.: A characterization of optimality criteria for decision making under complete ignorance. In: *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR-2010)*. pp. 172–181. AAAI Press (2010)
6. Brams, S.J., Fishburn, P.C.: *Approval Voting*. Springer, 2nd edn. (2007)
7. Brandt, F., Conitzer, V., Endriss, U.: Computational social choice. In: Weiss, G. (ed.) *Multiagent Systems*, pp. 213–283. MIT Press (2013)
8. Duddy, C., Piggins, A.: Collective approval. *Mathematical Social Sciences* 65(3), 190–194 (2013)
9. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.: Using complexity to protect elections. *Communications of the ACM* 53(11), 74–82 (2010)
10. Gaertner, W.: *A Primer in Social Choice Theory*. LSE Perspectives in Economic Analysis, Oxford University Press (2006)
11. Geist, C., Endriss, U.: Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research* 40, 143–174 (2011)
12. Grossi, D., Pigozzi, G.: Introduction to judgment aggregation. In: Bezhanishvili, N., Goranko, V. (eds.) *Lectures on Logic and Computation*, pp. 160–209. Springer-Verlag (2012)
13. Halpern, J.Y.: *Reasoning about Uncertainty*. MIT Press (2003)
14. Kannai, Y., Peleg, B.: A note on the extension of an order on a set to the power set. *Journal of Economic Theory* 32(1), 172–175 (1984)
15. List, C., Pettit, P.: Aggregating sets of judgments: An impossibility result. *Economics and Philosophy* 18(1), 89–110 (2002)
16. Merrill, III, S., Nagel, J.: The effect of approval balloting on strategic voting under alternative decision rules. *The American Political Science Review* 81(2), 509–524 (1987)
17. Shoham, Y., Leyton-Brown, K.: *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press (2009)
18. Xu, Y.: Axiomatizations of approval voting. In: Sanver, M.R., Laslier, J.F. (eds.) *Handbook on Approval Voting*, pp. 91–102. Springer-Verlag (2010)