Voting on Actions with Uncertain Outcomes

Ulle Endriss
ILLC, 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.

The Paradox of Individual Uncertainty Resolution

	Belief	Preference	Action
Agent 1	A	$A \succ B$	*stay*
Agent 2	A	$B \succ A$	*change*
Agent 3	B	$B \succ A$	*stay*
Collective			*stay*

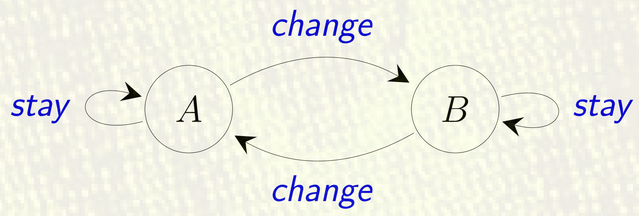

	Belief	Preference	Action
Agent 1	A	$A \succ B$	
Agent 2	A	$B \succ A$	
Agent 3	B	$B \succ A$	
Collective	A	$B \succ A$	*change*

[using *majority* for aggregation of any kind of information]

The Paradox of Late Collective Uncertainty Resolution

	Belief	Preference	Action
Agents 1–2	A or C	$A \succ C \succ B$	
Agents 3–5	B or C	$B \succ A \succ C$	
Collective	C	$A \succ B \succ C$	*left*

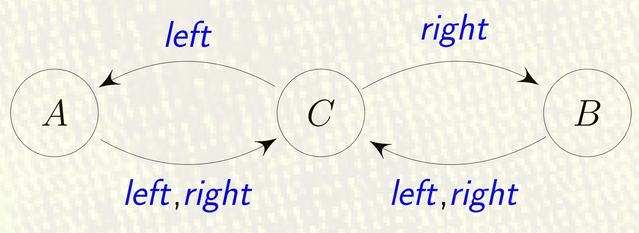

	Belief	Preference	Action
Agents 1–2	A or C	$A \succ \not{C} \succ B$	
Agents 3–5	B or C	$B \succ A \succ \not{C}$	
Collective	C	$B \succ A$	*right*

[using *approval voting* for beliefs and *Borda* for preferences]

The Paradox of Early Collective Uncertainty Resolution

	Belief	Preference	Action
Agents 1–9	A or C	$A \succ C \succ B$	
Agent 10	A or B	$B \succ C \succ A$	
Collective	A	$A \succ C \succ B$	*down*

[*approval* for beliefs, *Condorcet* for preferences, and *lexicographic tie-breaking* for actions]

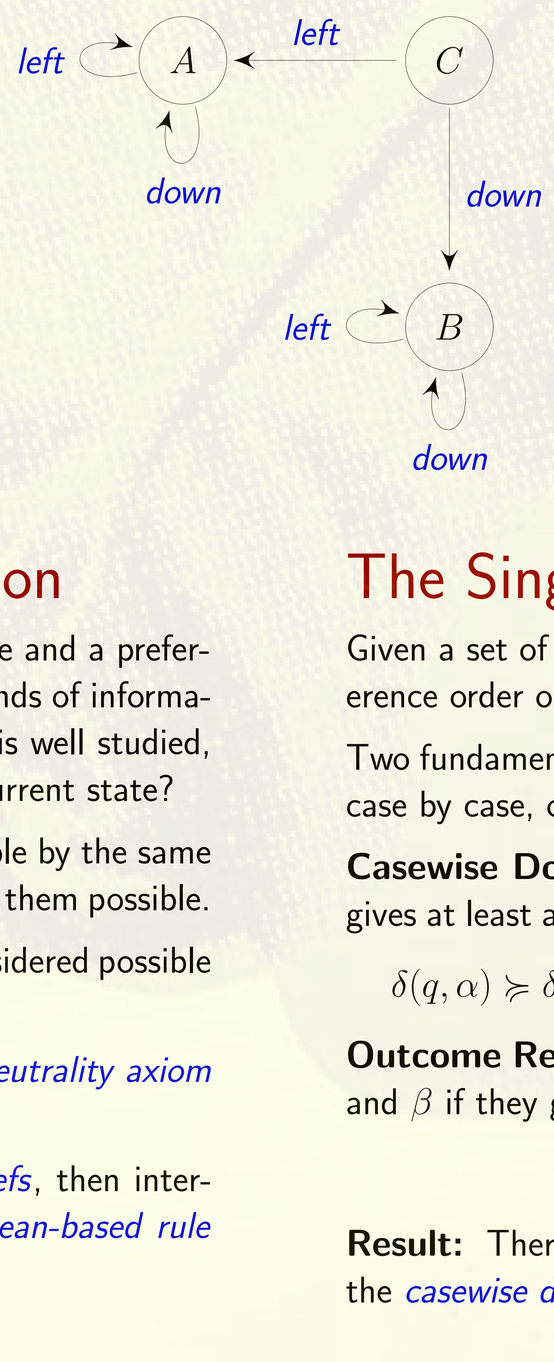

	Belief	Preference	Action
Agents 1–9	A or C	$A \succ C \succ B$	
Agent 10	A or B	$B \succ C \succ A$	
Collective	A [*or* C]	$A \succ C \succ B$	*left*

[considering also the *runner-up* for beliefs]

Uncertainty Resolution in Isolation

Each agent reports a set of states they consider possible and a preference order on outcomes. We need to *aggregate* both kinds of information and *integrate* the results. Preference aggregation is well studied, but how should we *resolve uncertainty* regarding the current state?

Neutrality Axiom: If two states are considered possible by the same agents, then we should consider either both or neither of them possible.

Knowledge Axiom: The true state of the world is considered possible by all agents (as may be many other states).

Result: The only aggregator that satisfies both the *neutrality axiom* and the *knowledge axiom* is the *intersection rule*.

When agents do not report *knowledge* but rather *beliefs*, then interesting aggregators include *approval voting* and the *mean-based rule* (accept all states with at least average support).

The Single-Agent Case

Given a set of states believed to be possible current states and a preference order on outcomes, how should we rank the available actions?

Two fundamental ways of approaching this: considering possible states case by case, or considering the set of possible outcomes as a whole.

Casewise Dominance Axiom: Prefer action α over β whenever α gives at least as good* a result as β for every state considered possible.

$$\delta(q,\alpha) \succcurlyeq \delta(q,\beta) \text{ for all } q \in Q \text{ [}^\star\text{strictly for some]} \;\Rightarrow\; \alpha \succ_Q \beta$$

Outcome Relevance Axiom: Remain indifferent between actions α and β if they give rise to the same set of possible outcomes.

$$\delta(\alpha,Q) = \delta(\beta,Q) \;\Rightarrow\; \alpha \sim_Q \beta$$

Result: There exists *no* action ranking function that satisfies both the *casewise dominance axiom* and the *outcome relevance axiom*.