

Strategic voting and the logic of knowledge

(Extended Abstract)

Hans van Ditmarsch
Logic
University of Seville
hvd@us.es

Jérôme Lang
LAMSADE
Université Paris Dauphine
lang@irit.fr

Abdallah Saffidine
LAMSADE
Université Paris Dauphine
abdallah.saffidine@gmail.com

Categories and Subject Descriptors

I.2.4 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems

General Terms

Theory

Keywords

Social choice, Dynamics, Epistemic logic

1. INTRODUCTION

A well-known fact in social choice theory is that strategic voting, also known as manipulation, becomes harder when voters know less about the preferences or votes of other voters. Standard approaches to manipulation in social choice theory [6] as well as in computational social choice [3] assume that the manipulating voter or the manipulating coalition knows perfectly how the other voters will vote. Some approaches [2] assume that voters have a probabilistic prior belief on the outcome of the vote, which encompasses the case where each voter has a probability distribution over the set of profiles. A recent paper [5] extends coalitional manipulation to incomplete knowledge, by distinguishing manipulating from non-manipulating voters and by considering that the manipulating coalition has, for each voter outside the coalition, a set of possible votes encoded in the form of a partial order over candidates. Uncertainty of voters about the uncertainties of other voters, i.e., higher-order beliefs of voters, has not been treated in full generality.

We model how uncertainty about the preferences of other voters may determine a strategic vote, and how a reduction in this uncertainty may change a strategic vote. A link between epistemic logic and voting has been given in [4]—they use knowledge graphs to indicate that a voter is uncertain about the preference of another voter. A more recent approach, within the area known as social software, is [8]. The recent [5] walks a middle way namely where equivalence classes are called information sets, as in treatments of knowledge and uncertainty in economics, but where the uncertain voter, or coalition, does not take the uncertainty of other voters into account.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2. KNOWLEDGE AND VOTING

We assume voters $\mathcal{N} = \{1, \dots, n\}$, candidates $\mathcal{C} = \{a, b, c, \dots\}$, and votes $V_i \subseteq \mathcal{C} \times \mathcal{C}$ that are linear orders. If agent i prefers candidate a to candidate b , we write $a \succ_i b$. A profile P is a collection $\{V_1, \dots, V_n\}$ of n votes, and a voting rule is a function $F : O(\mathcal{C})^n \rightarrow \mathcal{C}$ from the set of profiles to the set of candidates. We may further assume a tie-breaking mechanism. If $F(P[V_i/V_i']) \succ_i F(P)$, then V_i' is a *successful manipulation*. Given a profile P , a profile P' is an *equilibrium profile* iff no agent has a successful manipulation.

We model uncertainty about voting as incomplete knowledge about profiles. This terminology is standard in modal logic. The novelty consists in taking models with *profiles* instead of *valuations of propositional variables*.

DEFINITION 1 (KNOWLEDGE PROFILE). A profile model is a structure $\mathcal{P} = (S, \{\sim_1, \dots, \sim_n\}, \pi)$, where S is a domain of abstract objects called profile names; where for $i = 1, \dots, n$, \sim_i is an indistinguishability relation, that is, an equivalence relation; and where valuation $\pi : S \rightarrow O(\mathcal{C})^n$ assigns a profile to each profile name. A knowledge profile is pointed structure \mathcal{P}_s where \mathcal{P} is a profile model and s is a profile name in the domain of \mathcal{P} .

DEFINITION 2 (KNOWLEDGE). Given a knowledge profile \mathcal{P}_s and a proposition q , agent i knows that q if and only if q holds for all profile names in \mathcal{P} indistinguishable for i from s (i.e., for all $s' \in \mathcal{P}$ such that $s \sim_i s'$).

Propositions like ‘voter i knows the profile’ or even ‘voter i knows that P is an equilibrium profile’ have a precise formal description in this framework.

Under conditions of incomplete knowledge it may be that voter i (or coalition G) can manipulate the outcome of a profile P but does not know that, because she considers another profile (name) possible that she cannot manipulate. Such situations call for more refined notions of manipulation, that also involve knowledge. They can be borrowed from the knowledge and action literature [9, 7]. Our main interest is when voters know the manipulation.

DEFINITION 3 (KNOWLEDGE OF MANIPULATION). Given a knowledge profile \mathcal{P}_s . Voter i knows **de re** that she can strongly successfully manipulate \mathcal{P}_s if there is a vote V_i' such that for all t such that $s \sim_i t$, $F(P[V_i'/V_i]) \succ_i F(P)$, where t has profile P .


In the presence of knowledge, the definition of an equilibrium extends naturally. The trick is that for each agent, the combination of an agent i and an equivalence class $[s]_{\sim_i}$ for

that agent (for some state s in the knowledge profile) defines a *virtual agent*. Thus, agent i is multiplied in as many virtual agents as there are equivalence classes for \sim_i in the model. An equilibrium is then a combination of votes such that none of the virtual agents has an interest to deviate. An intuitively more appealing solution than virtual agents, applied in [1], is to stick to the agents we already have, but change the set of votes into a larger set of *conditional votes* — where the conditions are the equivalence classes for the agents. This we will now follow in the definition below. For *risk averse* voters (this criterion fits best our probability-free and utility-free model — it was also chosen in [5]) we can effectively determine if a conditional profile is an equilibrium without taking probability distributions into account, unlike in the more general setting of Bayesian games that it originates with.

DEFINITION 4 (CONDITIONAL EQUILIBRIUM). *Given is a knowledge profile model \mathcal{P} . For each agent i , let CV_i be the set of all conditional votes for that agent. A conditional vote is a function $CV_i : S/\sim_i \rightarrow O(C)$, i.e., a function that assigns to each equivalence class for that agent a vote. A conditional profile is a collection of n conditional votes, one for each agent. A conditional profile is an equilibrium iff no agent has a successful manipulation. A conditional profile is a strong equilibrium iff no coalition has a successful manipulation.*

3. EXAMPLE

Consider two voters a, b , four candidates $1, 2, 3, 4$, and three profile names s, t, u (for two profiles P and P') as below. The profile name s is assigned to profile P , wherein $a \succ_1 c \succ_1 b \succ_1 d$ and $d \succ_2 c \succ_2 b \succ_2 a$, etc. Profile names that are indistinguishable for a voter i are linked with an i -labelled edge. The partition for 1 on the domain is therefore $\{\{s, t\}, \{u\}\}$, and the partition for 2 on the domain is $\{\{s\}, \{t, u\}\}$.



Note that the names s and t are assigned to the same profile. However, s and t have different epistemic properties. In s , 2 knows that 1 prefers a over d , whereas in t 2 does not know that.

Consider a plurality vote with a tie-breaking rule $b \succ a \succ c \succ d$. If there had been no uncertainty, then in profile P , if 1 votes for her preference a and 2 votes for his preference d , then the tie prefers a , 2's least preferred candidate. If instead 2 votes c , a will still win. But if 2 votes b , b wins. We observe that (a, b) and (b, b) are equilibria pairs of votes, and that for 1 voting a is dominant. If there had been no uncertainty, then in profile P' pair (d, d) is the dominant equilibrium.

This situation changes when we take the uncertainty of the voters into account. There are two equilibria that we can associate with this knowledge profile model. Below, the conditional vote for 1 in the first equilibrium actually is defined as (given that $\pi(t) = P$ and $\pi(u) = P'$): $CV_1(\{t\}) = V_1$

and $CV_1(\{u\}) = V'_1$; the vote for 2 is conditional to one equivalence class — in other words, it is unconditional. The equivalent verbose formulation is more intelligible:

- (if 1 prefers a then 1 votes a and if 1 prefers d then 1 votes d , 2 votes b),
- (if 1 prefers a then 1 votes b and if 1 prefers d then 1 votes d , 2 votes b).

Unfortunately for voter 2, if the actual profile is P' so that d is his equilibrium vote, he will still not be inclined to cast that vote because he considers it possible that the profile is P , where, if 2 votes d and 1 votes a , a gets elected, voter 2's least preferred candidate. As 2 is risk averse his (known) equilibrium vote is therefore b .

If P' is the case, voter 1 has an incentive to make her true vote (i.e., her intention) known to 2, and even to declare her vote prior to 2.

4. DYNAMICS

The modal logical setting for voting and knowledge can be extended with dynamic logical operations. Three examples are: *deliberation of a coalition*, *public announcement of a proposition* (such as an agent revealing her true preference), and *declaring a vote*. These can be formalized as semantic operations $P_s \mapsto P_s|G$, $P_s \mapsto P_s|p$ (for proposition p), and $P_s \mapsto P_s|d(V_i)$, respectively. All these correspond to standard dynamic epistemic logical operations [10].

5. REFERENCES

- [1] T. Ågotnes and H. van Ditmarsch. What will they say? - Public announcement games. *Synthese*, 179(S.1):57–85, 2011.
- [2] S. Barbera, A. Bogomolnaia, and H. van der Stel. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences*, 35(2):89–103, 1998.
- [3] J. Bartholdi, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [4] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. *Proceedings of JELIA*, pages 18–30, 2004. LNCS 3229.
- [5] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. *Proceedings of AAAI*, 2011. To appear.
- [6] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.
- [7] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63:185–219, 2004.
- [8] R. Parikh, C. Tasdemir, and A. Witzel. The power of knowledge in games. *Proceedings of the Workshop on Reasoning About Other Minds*, 2011. CEUR Workshop Proceedings. Volume: 751.
- [9] J. van Benthem. Games in dynamic epistemic logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.
- [10] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.