6

An example of voting with uncertain preferences

Hans van Ditmarsch

6.1. Holandes volante encuentra sevillano

No soy un holandés que vuela. Prefiero quedarme en tierra firme y odio los vuelos. Y Ángel no es un sevillano. Me dice que viene de los alrededores de Sevilla, de Los Rosales, y que tienen que pasar varias generaciones para que alguien pueda llamarse sevillano. Por seguro, una vez Hans encontró a Ángel. Era en Salamanca, en el congreso internacional Tools for Teaching Logic en el 2006. Para mi ésta era una ocasión para conocer a los lógicos españoles desde una perspectiva de antípoda. En el 2006 todavía trabajaba en Nueva Zelanda, en la universidad de Otago. Desde el 2009, después de otro muy corto aterrizaje en Escocia, mi trabajo está en Sevilla. Y esto nunca habría sido posible sin Ángel Nepomuceno. Claramente no habría sido posible. ¿Cómo un holandés podría procurarse un proyecto de la Junta de Andalucía?

Tengo suerte en esta vida. Tengo suerte de vivir en Andalucía. Todos mis amigos de los países del norte la conocen de vacaciones y por las playas, pero nada más —esto es diferente—. Tengo suerte de conocer a Ángel, y de conocer a su familia. Aparte de no ser un holandés que vuela, soy un holandés impaciente. Ante la paciencia infinita de Ángel, todas mis batallas de un minuto o cinco minutos nada más —normalmente disfrutando un café juntos sobre el ojo vigilante y amable de Álvaro— son como olas del Atlántico que se destruyen en una roca firme de la costa de Andalucía. Ángel es la continuidad de la academia, la planificación a largo tiempo. Esto cuenta.

Ángel, pienso que nunca aprenderé a acostarme mas tarde que las diez u once de la noche. ¿Cuándo duermen los españoles? Aparte de eso, continuaré mi proceso de aprendizaje de buen sevillano. Felicidades por tu cumpleaños. Te deseo docenas más.

6.2. Introduction

A well-known result in social choice theory is that strategic voting (insincere preferences) is impossible if nothing is known about the preferences of the other voters. We consider this one example of the far more general investigation of uncertain preferences. It may be that there is no common prior in the standard sense that the preferences of all voters are common knowledge between them. A weaker form of common prior can be obtained, if each voter may be uncertain about a number of alternative profiles, and if this is common knowledge between all voters.

The case where a voter knows his own preference and is uncertain about all the preferences of all other voters, is the case above. If we also assume that all voters know their own preferences, a well-known architecture for multi-agent systems applies, namely that of an interpreted system. But in this framework many other scenarios for uncertainty can be modelled, e.g., if a voter has not made up his mind yet if he prefers candidate a over candidate b, but is already convinced that c is least preferred, this can be modelled as uncertainty between votes $a > b > c$ and $a > c > b$. Clearly, any partially ordered preference can thus be seen as uncertainty between a set of totally ordered preferences. But there are yet more complex scenarios that cannot be seen as simply another expression of partial preference: it may be that a voter cannot distinguish two situations with *identical profiles* from one another, because in the first case another voter has some uncertainty about the profile but in the second case not.

We present a detailed example of voting under uncertain preferences.

6.3. Epistemic profile

6.3.1. Voting with perfect information

Assume a finite set $\mathcal{A} = \{1, \dots, n\}$ of n *agents* or *voters*, and a finite set $\mathcal{C} = \{a, b, c, \dots\}$ of m *candidates* or *alternatives*. For each agent i a preference relation $\succ_i \subseteq \mathcal{C} \times \mathcal{C}$ is a total order on $\mathcal{C}$. If $a \succ_i b$, then voter i prefers candidate a to candidate b. Instead of $c_1 \succ_i \cdots \succ_i c_n$ we also write $i : c_1 \dots c_n$, or

i
c_1 c_2 $\dots$ c_n

Agent i's most preferred candidate is named $\max(\succ_i)$ and his least preferred candidate is named $\min(\succ_i)$. A *profile* is a set $\{\succ_1, \dots, \succ_n\}$ of $|\mathcal{A}|$ preference relations. Let $O(\mathcal{C})$ be the set of total orders of $\mathcal{C}$; $O(\mathcal{C})^n$ is the set of all profiles for $\mathcal{A}$. Par abus de langage we identify a sequence of n total orders with the set $\{\succ_1, \dots, \succ_n\}$, where the name of the agent corresponds to the rank in the sequence.

A *voting rule* is a deterministic function $F : O(\mathcal{C})^n \to \mathcal{C}$ from the set of profiles to the set of candidates. The voting rule determines which candidate wins the election. It may be the case that the voting rule is such that a given profile determines more than one candidate, which then needs a *tie breaking preference* $\succ$, that is a total order again, to determine the winner (so that F is indeed a function).

Voters are not required to vote according to their preference, they may vote *strategically* or *insincerely*. Let $p \in O(\mathcal{C})^n$ and $\succ_i \in p$. Any $\succ'_i \neq \succ_i$ is called an *insincere vote* or insincere preference, and from that perspective $\succ_i$ is the true preference. If there is a $\succ'_i$ such that $F(p[\succ'_i/\succ_i]) \succ_i F(p)$, then agent i can *manipulate* the outcome of the election, and we write $\gg_i$ instead of $\succ'_i$. In other words, if all voters vote according to their preferences candidate c will get elected, if i votes insincerely but all other voters still vote according to their preferences a candidate c' gets elected, and i prefers c' over c.

We do not require that the outcome is still better if more than one voter votes insincerely. The obvious (Nash) equilibrium notion is as follows:

> A profile is an *equilibrium profile* iff no agent can manipulate the outcome.

There may be many such equilibria for a given $O(\mathcal{C})^n$ and F.

6.3.2. Voting with imperfect information

We represent uncertainty about profiles in the form of multi-agent pointed Kripke models. A *profile model* (Kripke model) is a structure $M = (S, \{\sim_i\}_{i \in \mathcal{A}}, V)$, where S is a *domain* of abstract objects called *states* or worlds; where *valuation* $V : O(\mathcal{C})^n \to S$ maps profiles to states; and for $i = 1, \dots, n$, $\sim_i$ is an *indistinguishability relation* that is an equivalence relation. An *epistemic profile* is pointed structure (M, s) (a.k.a. pointed Kripke model) where M is a profile model and s is a state in the domain of M. If $s \sim_i s'$, $s \in V(p)$, and $s' \in V(p')$, then agent i is uncertain if the profile is p or p'; e.g. if $j : bca$ in p and $j : cba$ in p', then agent i is uncertain if agent j prefers b over c or c over b.

The universal relation $S \times S$ can be seen as the indistinguishability relation of the central authority—this is proper as this corresponds (for connected structures) to the common prior of the voting agents. When there is perfect information, the agents commonly know what the profile is. When there is imperfect information in the now defined sense, the agents commonly know what their uncertainties about the profile are. This is a nearly equally powerful assumption. Nearly but not equally: for now we bypass the issue of a probability distribution of the possible states and investigate results that hold for *any* probability distribution.

The uncertainty relation serves to *talk* about profiles, i.e., it can be used to interpret a logical language. We forego this opportunity, and merely employ the meta-level to give meaning to knowledge and ignorance. The perspective for statement about knowledge and ignorance over profiles is a given state s in a given epistemic profile M. Given that, p is true in a state s if it is a member of $V(s)$, and a proposition of the form 'Agent i knows that ϕ' is *known* (by agent i) in a state s if ϕ is true in all states t indistinguishable for i from s, i.e., in all t such that $t \sim_i s$. Finally, ignorance is absence of knowledge: agent i does not know that ϕ (is ignorant about ϕ) in a state s if ϕ is false in some state t indistinguishable for i from s.

Given uncertainty, the definition of manipulation remains the same—it now simply applies to the profile of the actual state (the point of the epistemic profile). But we have more interesting notions too. It may be that agent i can manipulate the outcome of p but does not know that, because she considers another state possible for another profile that she cannot manipulate. Two relevant additional notions are as follows.

> *Agent i knows* **de re** *that she can manipulate the outcome of the election in a given state s, if there is a preference $\succ'_i$ such that for all states t considered possible by i (for all t such that $s \sim_i t$), $F(p[\succ'_i/\succ_i]) \succ_i F(p)$, where t has profile p.*

On the other hand,

> *Agent i knows* **de dicto** *that she can manipulate the outcome of the election in a given state s, if in every state t considered possible by i (for all t such that $s \sim_i t$), there is a preference $\succ'_i$ such that $F(p[\succ'_i/\succ_i]) \succ_i F(p)$, where t has profile p.*

In plain words, if i knows de re that she can manipulate the election, she has the ability to manipulate, namely by voting according to $\succ'$. But if i knows de dicto that she can manipulate the election, she does not have that ability because the preference change required for manipulation may be different in every state she considers possible... In the continuation we give an example of such de dicto manipulability.

6.4. Example of manipulation given incomplete information

There are two voters 1, 2, and four candidates a, b, c, d. Consider a plurality vote with a tie-breaking rule $b \succ a \succ d \succ c$. I.e., the preferred candidate who gets the most votes, wins. If more than one candidate gets the most votes, the candidate preferred by the tie wins.

Perfect information Consider the profile p defined as

1	2
a	d
c	c
b	b
d	a

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 (in)sincere votes, and that for 1 voting a is dominant. This is also apparent from the voting matrix

1\2	a	b	c	d
a	a	b	a	a
b	b	b	b	b
c	a	b	c	d
d	a	b	d	d

and even more so when we express the payoffs for both voters by their ranking for the winner:

1\2	a	b	c	d
a	30	$\boxed{11}$	30	30
b	11	$\boxed{11}$	11	11
c	30	11	22	03
d	30	11	03	03

Uncertainty between two profiles We now add uncertainty to this framework. Consider another profile p', namely

1	2
d	d
c	c
b	b
a	a

The uncertainty between the two profiles can be visualized as

1	2		1	2
a	d		d	d
c	c	—2—	c	c
b	b		b	b
d	a		a	a

Voter 1 knows her own preference so she can distinguish the two profiles from one another. (Therefore, there is not a 1-link between the two states/profiles.) Voter 2 is uncertain which of p and p' is really the case. We assume that in fact, p is the real profile. It is therefore still the case that the optimal vote for 1 and 2 is b. To investigate whether 1 and 2 also *know* that, we have to take the equilibria of the other profile into account. In case of p', it is trivial that the optimal vote for 1 and 2 is d, not b—1 and 2 have identical preferences there. What should 1 and 2 vote, given this uncertainty about preferences? The outcome depends on further assumptions.

First we assume that voters are risk-averse (given a choice, avoid the option with the worst possible outcome), and that this is part of the common ground. Because 2 is risk-averse he will always vote b. Because if he votes d and the profile is p, a wins if 1 votes a. Of course, if the profile is p', a is also 1's least preferred candidate and (d,d) the equilibrium. But as 2 does not know that the profile is p, he will still vote b, just to be sure. It may not be to surprising that a profile that is uncertain for 2 may change his voting preferences. More surprising is that it also changes 1's voting behaviour. In particular, even in the case that 1 and 2 both prefer d, (b,b) is still an equilibrium. This is because 1 knows that 2's best response always is b, a result she cannot improve by voting differently. Another equilibrium for the epistemic profile with point p' is (d,d). For the epistemic profile with point p, (a,b) and (b,b) are still equilibria.

observation 1. *Some interesting theory seems behind this. The equilibria are computed by considering the ranked payoffs for both profiles. Each agent performs a worst case analysis on the* set of payoffs *for all indistinguishable states (for* that *agent). This is a perfectly general procedure. For this example, the payoff matrices are*

1\2	a	b	c	d
a	30	$\boxed{11}$	30	30
b	11	$\boxed{11}$	11	11
c	30	11	22	03
d	30	11	03	03

1\2	a	b	c	d
a	00	11	00	00
b	11	$\boxed{11}$	11	11
c	00	11	22	33
d	00	11	33	$\boxed{33}$

We observe that (d,d) is not an equilibrium given the epistemic profile with point p' because if 2 votes d, then there is a profile considered possible by 2, namely p, where 1 can improve his outcome by switching to a (and where 2 is worse off).

When probability distributions come into play If the voters are risk-seeking, the picture becomes more complex. We then need to assume a probability distribution over the states of the profile model, and it then also matters which exact payoffs we have associated to the ranking of candidates. For example, assume uniform distribution and payoffs 3, 2, 1, 0 as above.

If 2 votes b, then the outcome always is b (so also when 1 votes her preference, in either case), which has a payoff of 1, and the expected outcome for 2 is then also 1. If 2 votes d and 1 votes her preference, then the outcome for 2 is 0 if the profile was p ((a,d) / 30) and the outcome for 2 is 3 is the profile is p' ((d,d) / 33). This gives an expected payoff for 2 of 1,5. As this exceeds the expected payoff 1 of playing b, 2 *seems* to have a reason to prefer voting d. But the setting is sensitive to the exact amount some candidates are preferred over others. If the payoffs are not 3, 2, 1, 0 but, for example, 1,5, 1,25, 1, 0, then the expected payoff of 2 voting d is 0,75, less then the assured 1 when playing b.

In this setting, is 1 still advised to vote b and not a, given profile p? That is not entirely clear. Agent 1 can reason about agent 2 maximizing expected payoff. Should she therefore switch again from the formerly equilibrium b to her original preference a? Maybe, it depends on additional assumptions—see the next observation.

observation 2. *We combine the data in the two payoff matrices above in a 16×4 matrix from which we can read off the equilibria in the standard way. A pure strategy for a voter in* *this matrix is not a vote as such but a vote conditional to the epistemic state of the voter. For example, a strategy for 1 is 'If my preference is acbd I vote a and if my preference is dcba I vote d'. For convenience we abbreviate that strategy as 'ad'. Therefore, 1 has* $2^4 = 16$ *strategies. As 2 cannot distinguish the two states of the profile model from one another he still has four strategies. The resulting matrix for this example is therefore as follows. The profile where 1 votes ad as above and 2 votes b is written as* (ad, b)*, etc. We have added the payoffs of the component game matrices, the expected payoffs are of course half of that, but for the determination of equilibria this does not matter.*

$1\backslash 2$	a	b	c	d
aa	30	22	30	30
ab	41	22	41	41
ac	30	22	52	$\boxed{63}$
ad	30	22	$\boxed{63}$	$\boxed{63}$
ba	11	$\boxed{22}$	11	11
bb	22	$\boxed{22}$	22	22
bc	11	22	33	44
bd	11	22	44	44
ca	30	22	22	03
cb	41	22	33	14
cc	30	22	44	36
cd	30	22	55	36
da	30	22	03	03
db	41	22	14	14
dc	30	22	25	36
dd	30	22	36	36

From this matrix we observe that the equilibria are—as expected—(bb, b)*, and also* (ac, d)*,* (ad, c)*,* (ad, d)*, and* (ba, b)*. The* (ad, d) *equilibrium is the one compared to* (bb, b) *above. We now see, that neither 1 nor 2 has a reason to prefer* (ad, d) *over* (bb, b)*, although* (ad, d) *is preferable for both; it has a larger social welfare of 9. The strategy* (cd, c) *maximizes social welfare to 10, but it is not an equilibrium.*

This analysis models the situation as a Bayesian game for uniform distribution. This simplification is valid for binary payoffs. It seems to generalize to non-binary payoffs, but this is unclear.

Different states for the same profile We can add further uncertainty to the example by considering a third state that has the same profile p as the actual state, but that has different epistemic properties: 2 is not uncertain about the profile there, but 1 cannot distinguish this from the other state for p wherein 2 is uncertain about the profile. The profile model looks like

$$
\begin{array}{|c|c|}\hline 1 & 2 \\ \hline a & d \\ c & c \\ b & b \\ d & a \\ \hline \end{array}
\;\overset{}{\text{———}}1\text{———}\;
\begin{array}{|c|c|}\hline 1 & 2 \\ \hline a & d \\ c & c \\ b & b \\ d & a \\ \hline \end{array}
\;\text{———}2\text{———}\;
\begin{array}{|c|c|}\hline 1 & 2 \\ \hline d & d \\ c & c \\ b & b \\ a & a \\ \hline \end{array}
$$

Given that 1 knows that the profile is p, will she act differently now from the situation above where she knew that 2 was uncertain about the profile? Will the equilibria be different given that she now also considers it possible that 2 also knows the profile? We can repeat an analysis as in the observation above, but now for a 16×16 matrix.

6.5. Conclusions

This is merely one example of a fascinating novel area of the analysis of uncertain voting, taken, with kind permission, from work in progress by Hans van Ditmarsch and Jérôme Lang. Hans recently visited Paris as professeur invité, Université Paris Dauphine, as a guest of Jérôme in the LAMSADE group. He thanks the Universidad de Sevilla for their permission for this visit, and the Université Paris Dauphine for their hospitality.