This is an Accepted Manuscript of:
Ivanovska, M., Slavkovik, M. (2022). Probabilistic Judgement Aggregation by Opinion Update. In: Torra, V., Narukawa, Y. (eds) Modeling Decisions for Artificial Intelligence. MDAI 2022. Lecture Notes in Computer Science, vol 13408. P 36-37 Springer, Cham. https://doi /10.1007/978-3-031-13448-7_3

Probabilistic Judgement Aggregation by Opinion Update

Magdalena Ivanovska[1] [0000-0002-3916-3486] and Marija Slavkovik[2] [0000-0003-2548-8623]

[1] BI Norwegian Business School, Oslo, Norway
magdalena.ivanovska@bi.no
[2] University of Bergen, Bergen, Norway
marija.slavkovik@uib.no

Abstract. We consider a situation where agents are updating their probabilistic opinions on a set of issues with respect to the confidence they have in each other's judgements. We adapt the framework for reaching a consensus introduced in [2] and modified in [1] to our case of uncertain probabilistic judgements on logically related issues. We discuss possible alternative solutions for the instances where the requirements for reaching a consensus are not satisfied.

Keywords: judgement aggregation · probabilistic logic · Markov chains

1 Introduction

Judgement aggregation (JA) is concerned with aggregating categorical judgements about the truth values of logically related issues (propositions) [7,4]. An example is given in Table 1, where the rows contain judgements of agents over the issues p, q, and $p \wedge q$. As observed from the example, pooling the truth valuations on each issue does not always lead to a consistent set of collective judgements. JA designs and studies aggregators that produce a consistent outcome.

Table 1: An example of a judgement aggregation using the simple majority rule.

	p	q	$p \wedge q$
agent 1	true	true	true
agent 2	true	false	false
agent 3	false	true	false
Majority	true	true	false

Aggregation problems, however, are not always Boolean, since the judgements on whether an issue is true or false are not always certain. In order to deal with this kind of uncertainty, in [6] we define a framework that aims at aggregating judgements about the probabilities of issues, as in the example given in Table 2.

Table 2: An example of probabilistic judgement aggregation with a threshold majority rule. The numbers in each row represent the probabilities of the issues being true according to the corresponding agent.

	p	q	$p \wedge q$
agent 1	0.7	0.8	0.7
agent 2	0.6	0.7	0.5
agent 3	0.2	0.8	0.2
Majority$_{\geq 0.6}$	true	true	false

One way of aggregating judgements into a consistent collective opinion would be to modify (update) some of the individual judgements until the chosen aggregation rule produces a consistent judgement. This idea is obviously more applicable to probabilistic judgements than to categorical ones, since the modification there amounts to adjusting a probability value rather than completely changing the attitude about the truth of an issue. As can be observed in the example in Table 2, a small modification of an individual opinion (in this case agent 2's judgement on $p \wedge q$) could result in obtaining a consistent collective judgement.

In this paper we consider a setting where agents update their individual opinions to align with each other and eventually converge to a consensual opinion. We adopt the model for reaching a consensus over probability distributions described in [2] and show that it is applicable to the case of probabilistic opinions on logically related issues as well. The model presumes a confidence matrix representing the trust the agents have in each other's opinions. The opinion updating is performed based on the confidence matrix. In the cases where the repeated updates converge to a consensus, the aggregated opinion is obtained from the individual opinions through a linear function, and some desirable properties follow by definition. We discuss possible solutions to the cases where this repeated update will not lead to a consensus due to the properties of the confidence matrix and the particular opinions that are to be aggregated.

2 Framework

We use a slightly modified version of the framework defined in [6] which we include for self-sufficiency.

2.1 Probabilistic Judgement Profiles

Let $\mathcal{L}$ be a set of propositional logic formulas. An *agenda* is a finite set $\Phi \subset \mathcal{L}$,

$$\Phi = \{\varphi_1, \ldots, \varphi_m\} , \quad (1)$$

s.t. φ_i is neither a tautology nor a contradiction. We call the elements of the agenda *issues*. For example, in Table 1, we have $\Phi = \{p, q, p \wedge q\}$. We are interested in aggregating a collection of judgements on the agenda issues coming from a group of *information sources* (we will also call them *agents*) into a collective judgement representative for the group. Let $\Phi^{\cup} = \Phi \cup \{\neg\varphi \mid \varphi \in \Phi\}$. We model the information sources as sets of likelihood judgements on $\Phi^{\cup}$.

A *likelihood judgement* on the issue $\varphi \in \Phi^{\cup}$ is a simple likelihood formula of the type:

$$\ell(\varphi) \geq a \ , \tag{2}$$

where $a \in [0, 1]$. The likelihood judgement $\ell(\varphi) \geq a$ expresses that the likelihood (probability)[3] of the statement φ being true is at least a. The formula (2) is an instance of the logic of likelihood (see [3] and [5]), the language of which consists of Boolean combinations of linear likelihood formulas of the type

$$a_1\ell(\varphi_1) + \ldots + a_n\ell(\varphi_n) \geq b \ , \tag{3}$$

where a_i, b are real numbers, and φ_i are pure propositional formulas.[4] The likelihood formulas are interpreted in probability spaces $M = (W, F, \mu)$, where W is a set of possible worlds, F is a σ-algebra on W, and $\mu : F \to [0, 1]$ is a probability measure. The propositional formulas are given possible world semantics in the standard way:

$$\varphi^M = \{w \in W \mid w \models \varphi\} \ , \tag{4}$$

and the term $\ell(\varphi)$ is interpreted as $\mu(\varphi^M)$[5], i.e. as the probability of the set of worlds at which φ is true. This leads to the following interpretation of (3):

$$a_1\mu(\varphi_1^M) + \ldots + a_n\mu(\varphi_n^M) \geq b \ , \tag{5}$$

i.e. (3) is true in M if and only if (5) holds. The interpretation of Boolean combinations of formulas of type (3) is defined in the standard way.

The axiomatic system for the logic of likelihood consists of axioms for propositional reasoning, reasoning about inequalities, and reasoning about probabilities. In particular, for every propositions φ and ψ, and every likelihood formulas f and g, the following are axioms:

- (Prop) All substitution instances of tautologies in propositional logic,
- (MP) From f and $f \to g$, infer g,
- (Inq) All substitution instances of valid linear inequality formulas,
- (L1) $\ell(\varphi) \geq 0$,
- (L2) $\ell(\top) = 1$,
- (L3) $\ell(\varphi) = \ell(\varphi \wedge \psi) + \ell(\varphi \wedge \neg\psi)$,
- (L4) From $\varphi \leftrightarrow \psi$ infer $\ell(\varphi) = \ell(\psi)$.

The above set of axioms is shown to be sound and complete with respect to the above interpretation [3].

[3] In this paper we interpret likelihood as probability and we use the two terms interchangeably. Note that, however, likelihood can also be interpreted as another measure of belief, see [5].

[4] Expressions containing all the other types of inequalities or equality can be defined as abbreviations.

[5] To ensure that every φ^M is measurable, we may take $F = 2^W$.

Each of the information sources is represented as a *set of likelihood judgements* $\hat{J}$. The set $\hat{J}$ has one likelihood judgement on each of the issues in $\Phi^{\cup}$:

$$\hat{J} = \{\ell(\varphi) \geq a(\varphi) \mid \varphi \in \Phi^{\cup}\} , \tag{6}$$

where $a(\varphi) \in [0,1]$ is called a *judgement coefficient* of φ.

From a given judgement set $\hat{J}$ as defined in (6), using the above axioms, we can derive $\ell(\varphi) \leq 1 - a(\neg\varphi)$. This means that providing likelihood formulas for both φ and $\neg\varphi$ in the judgement set $\hat{J}$ is equivalent to providing intervals for the likelihood of φ. In the cases where $a(\varphi) + a(\neg\varphi) = 1$, these intervals collapse to a point, i.e. we obtain precise likelihood judgement. Judgements in $\hat{J}$ can also be Boolean, since we can represent by $\ell(\varphi) \geq 1$ that φ is true, and by $\ell(\neg\varphi) \geq 1$ that φ is false.

Given a set of n agents $N = \{1, \ldots, n\}$, a *likelihood profile*:

$$\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n) , \tag{7}$$

is a collection of sets of likelihood judgements for an agenda Φ, each representing one agent $k \in N$. We slightly abuse notation and write $\hat{J}_k \in \hat{P}$ to denote that $\hat{J}_k$ is the k-th likelihood judgement set in $\hat{P}$:

$$\hat{J}_k = \{\ell(\varphi) \geq a_k(\varphi) \mid \varphi \in \Phi^{\cup}\} , \tag{8}$$

where $a_k(\varphi) \in [0,1]$ are the judgement coefficients of the k-th agent, $k = 1, \ldots, n$. An example of a likelihood profile is given in Table 3.

Table 3: An example of a likelihood profile over the agenda $\Phi = \{p, q, p \wedge q\}$. The set of likelihood judgements of agent 1 is $J_1 = \{\ell(p) \geq 0.7, \ell(\neg p) \geq 0.2, \ell(q) \geq 0.8, \ell(\neg q) \geq 0.1, \ell(p \wedge q) \geq 0.7, \ell(\neg(p \wedge q)) \geq 0.2\}$. Similarly, for J_2 and J_3.

	p	$\neg p$	q	$\neg q$	$p \wedge q$	$\neg(p \wedge q)$
agent 1	≥ 0.7	≥ 0.2	≥ 0.8	≥ 0.1	≥ 0.7	≥ 0.2
agent 2	≥ 0.6	≥ 0.2	≥ 0.7	≥ 0.3	≥ 0.5	≥ 0.5
agent 3	≥ 0.2	≥ 0.2	≥ 0.8	≥ 0.2	≥ 0.2	≥ 0.4

The example in Table 2 is a special case of a likelihood profile for the agenda $\Phi = \{p, q, p \wedge q\}$, where each row represents a judgement set with precise likelihood judgements. For example, 0.6 in the row of agent 2 stands for $\ell(p) = 0.6$, or, equivalently, $\ell(p) \geq 0.6$ and $\ell(\neg p) \geq 0.4$.

2.2 Rationality of Probabilistic Judgement Sets

We require that the sets of likelihood judgements in the profile are *rational*. We now define what are rational likelihood judgements sets.

A probabilistic judgement set is *consistent* if it is a consistent set of formulas in the logic of likelihood. A probabilistic judgement set is not always consistent. Consider, for example, the agenda $\Phi = \{p_1 \wedge p_2, p_1 \wedge \neg p_2\}$ and a set $\hat{J}$ containing the judgements $\ell(p_1 \wedge p_2) \geq 0.4$ and $\ell(p_1 \wedge \neg p_2) \geq 0.7$. The set $\hat{J}$ is an inconsistent set of formulas, because it implies $\ell(p_1) \geq 1.1$ by axiom (L3). Furthermore, note that for a judgement set $\hat{J}$ defined as in (6) to be consistent, it has to satisfy $a(\varphi) + a(\neg\varphi) \leq 1$, for every $\varphi \in \Phi$.

A set of likelihood judgements is always *complete* in the sense that it contains a likelihood judgement for each of the issues. This assumption does not limit the freedom of not having a specific likelihood estimate for a given issue φ. To represent the absence of a specific likelihood or an "abstention" on an issue φ, we can use the tautologies $\ell(\varphi) \geq 0$ and $\ell(\neg\varphi) \geq 0$.

In classical judgement aggregation, a judgement set is rational if it is consistent and complete, which means it provides a truth value for each of the issues in the agenda, and these values are consistent. In the probabilistic case, consistency and completeness are not enough of conditions for rationality. For example, $\hat{J} = \{\ell(p_1) \geq 0.3, \ell(\neg p_1) \geq 0.5, \ell(p_1 \wedge p_2) \geq 0.4, \ell(\neg(p_1 \wedge p_2)) \geq 0.5\}$ is a consistent set. However, if we use the axioms, we can easily derive $\ell(p_1) \geq 0.4$, which is stronger than the existing $\ell(p_1) \geq 0.3$ and, as such, is a more valuable judgement. In general, we say that $\ell(\varphi) \geq a$ is a *stronger judgement* than $\ell(\varphi) \geq b$ iff $a > b$. To ensure that we always have the strongest possible judgements in the consistent judgement sets, we introduce the notion of a *final judgement*. A consistent probabilistic judgement set is *final* if it does not imply stronger judgements than the ones it contains, i.e. a judgement set $\hat{J}$ as defined by (6) is final iff $\hat{J} \vdash \ell(\varphi) \geq c$ implies $c \leq a(\varphi)$, for every $\varphi \in \Phi^{\cup}$. We say that the probabilistic judgement set $\hat{J}$ is *rational* if it is consistent and final. A profile is rational if all the judgement sets in it are rational.

3 Updating Probabilistic Judgements

One can imagine there are many different ways the agents can update their judgements, depending on the kind of information the update is based upon. Here, we handle the situation where there is no new factual information that the agents receive, but they are able to observe each other's judgements and update their own judgement sets based on this observation. Similarly as in [2], we assume that each agent has certain degrees of confidence in the other agents' opinions and in her own, and updates her probabilistic judgements upon observing the judgements of others based on these confidence degrees. More formally, let

$$t_k = (t_{k1}, \ldots, t_{kn}) \;, \tag{9}$$

where $t_{kr} \in [0, 1]$, for every r, and $\sum_{r=1}^{n} t_{kr} = 1$, be the confidence distribution of the k-th agent, $k = 1, \ldots, n$. t_{kr} is interpreted as the degree of confidence the agent k assigns to the agent r, $r = 1, \ldots, n$. We call the matrix $T = [t_{kr}]_{n \times n}$, where each row represents a confidence distribution of the corresponding agent, a *confidence matrix*.

Given a likelihood profile $\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n)$, we assume that the agent k is updating her judgements by calculating new judgement coefficients as weighted average of everyone's judgement coefficients wrt. her confidence distribution t_k:

$$\hat{J}_k^1 = \{\ell(\varphi) \geq \sum_{r=1}^{n} t_{kr} a_r(\varphi) \mid \varphi \in \Phi^{\cup}\} \ . \tag{10}$$

The updating process can iterate several times, and the result of each iteration is defined recursively:

$$\hat{J}_k^i = \{\ell(\varphi) \geq a_k^i(\varphi) \mid \varphi \in \Phi^{\cup}\} \ , \tag{11}$$

where the judgement coefficients are determined by

$$a_k^i(\varphi) = \sum_{r=1}^{n} t_{kr} a_r^{i-1}(\varphi) \ , \tag{12}$$

for $k = 1, \ldots, n$, where $a_r^0(\varphi) = a_r(\varphi)$. $\hat{P}^i = (\hat{J}_1^i, \ldots, \hat{J}_n^i)$ is the i-th update of the profile $\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n)$, for $i \in \mathbb{N}$.

Theorem 1. *Let $\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n)$ be a rational profile and $T = [t_{kr}]_{n \times n}$ be a confidence matrix. Then $\hat{P}^i = (\hat{J}_1^i, \ldots, \hat{J}_n^i)$ is a rational profile, for every $i \in \mathbb{N}$.*

The proof of the above theorem follows directly from the following proposition.

Proposition 1. *Let $\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n)$ be a rational profile and $t = (t_1, \ldots, t_n)$ be a vector of coefficients such that $t_k \in [0, 1]$, for every $k = 1, \ldots, n$, and $\sum_{k=1}^n t_k = 1$. Then the judgement set $\hat{J} = \{\ell(\varphi) \geq \sum_{k=1}^n t_k a_k(\varphi) \mid \varphi \in \Phi^{\cup}\}$ is rational.*

Proof. Consistency: Let W be a set of possible worlds and let $F = 2^W$ be the σ-algebra of all the subsets of W. Since the profile $\hat{P}$ is rational, the set $\hat{J}_k = \{\ell(\varphi) \geq a_k(\varphi) \mid \varphi \in \Phi^{\cup}\}$ is a consistent set of formulas, for every $k = 1, \ldots, n$. This means that there exist probability measures on (W, F), $\mu_k : F \to [0, 1]$, $k = 1, \ldots, n$, such that the inequalities in the sets $\{\mu_k(\varphi^M) \geq a_k(\varphi) \mid \varphi \in \Phi^{\cup}\}$ hold. Then the linear function of these measures with the components of the vector t as coefficients, $\mu = \sum_k t_k \mu_k$, is a probability measure on (W, F) for which the set of inequalities $\{\mu(\varphi^M) \geq \sum_k t_k a_k(\varphi) \mid \varphi \in \Phi^{\cup}\}$ holds. The last implies consistency of the judgement set $\hat{J} = \{\ell(\varphi) \geq \sum_k t_k a_k(\varphi) \mid \varphi \in \Phi^{\cup}\}$.

Finality: Let us denote by $a(\varphi) = \sum_k t_k a_k(\varphi)$ the likelihood coefficients of the set $\hat{J}$. Suppose that the set $\hat{J}$ is not final. This means that there exists $\varphi_i \in \Phi^{\cup}$, and $c > a(\varphi_i)$, such that $\hat{J} \vdash \ell(\varphi_i) \geq c$, i.e. that using the formulas in $\hat{J}$ and the axioms of the logic, one can derive $\ell(\varphi_i) \geq c$. Since $c > a(\varphi_i)$, this derivation needs to include axioms (L3), (L4) and some of the likelihood judgements of $\hat{J}$ referring to issues other than φ_i. In particular, φ_i must "include" some of the other issues, i.e. there must exist issues $\varphi_{i_1}, \ldots, \varphi_{i_r} \in \Phi^{\cup}$, other than φ_i, such that

$$\vdash \ell(\varphi_i) \geq \ell(\varphi_{i_1}) + \cdots + \ell(\varphi_{i_r}) \ , \tag{13}$$

and their judgement coefficients are such that:

$$a(\varphi_{i_1}) + \cdots + a(\varphi_{i_r}) \geq c \ . \tag{14}$$

Now, from (13) and the consistency of $\hat{J}_k$, for every k, we will have

$$\hat{J}_k \vdash \ell(\varphi_i) \geq a_k(\varphi_{i_1}) + \cdots + a_k(\varphi_{i_r}) \ . \tag{15}$$

From this and the finality of $\hat{J}_k$, we obtain

$$a_k(\varphi_i) \geq a_k(\varphi_{i_1}) + \cdots + a_k(\varphi_{i_r}) \ , \tag{16}$$

for every k. But then,

$$\sum_k t_k a_k(\varphi_i) \geq \sum_k t_k a_k(\varphi_{i_1}) + \cdots + \sum_k t_k a_k(\varphi_{i_r}) \ , \tag{17}$$

which with the shortened notation becomes

$$a(\varphi_i) \geq a(\varphi_{i_1}) + \cdots + a(\varphi_{i_r}) \tag{18}$$

The last, together with (14), implies $a(\varphi_i) \geq c$, which is in contradiction with the initial assumption.

4 Convergence to a Consensus

Let us denote $a_{kj} = a_k(\varphi_j)$, for $\varphi_j \in \Phi^{\cup}$, and $k = 1, \ldots, n$. Then $A = [a_{kj}]_{n \times 2m}$ is a matrix consisting of the judgement coefficients of the n agents on the set of propositions $\Phi^{\cup}$, with each row corresponding to one agent, and each column corresponding to one issue of the extended agenda $\Phi^{\cup}$. We call it a *judgement matrix* of the profile $\hat{P} = (\hat{J}_1, \ldots, \hat{J}_n)$. For example, the likelihood profile given in Table 3 is represented by the judgement matrix given in Fig.1.

$$A = \begin{bmatrix} 0.7 \ 0.2 \ 0.8 \ 0.1 \ 0.7 \ 0.2 \\ 0.6 \ 0.2 \ 0.7 \ 0.3 \ 0.5 \ 0.5 \\ 0.2 \ 0.2 \ 0.8 \ 0.2 \ 0.2 \ 0.4 \end{bmatrix}$$

Fig. 1: An example of a judgement matrix of three agents on the extended agenda $\Phi^{\cup} = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$.

If A is a judgement matrix of the profile $\hat{P}$ and T is a confidence matrix, then the matrix

$$A^1 = TA$$

will be the judgement matrix of the profile $\hat{P}^1$, we denote it by $A^1 = [a^1_{kj}]$. For example, if T is defined as:

$$T = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/4 & 3/4 & 0 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}$$

then the judgement matrix given in Fig.1 will be updated as follows:

$$A^1 = \begin{bmatrix} 0.650000 & 0.200000 & 0.750000 & 0.200000 & 0.600000 & 0.350000 \\ 0.625000 & 0.200000 & 0.725000 & 0.250000 & 0.550000 & 0.425000 \\ 0.500000 & 0.200000 & 0.766667 & 0.200000 & 0.466667 & 0.366667 \end{bmatrix}$$

In general, if we denote $A^0 = A$, from Eq.(12) we will have

$$A^i = TA^{i-1} \ , \tag{19}$$

for $i \in \mathbb{N}$, where the matrices $A^i = [a^i_{kj}]$ and $A^{i-1} = [a^{i-1}_{kj}]$ are the judgement matrices of the profiles $\hat{P}^i$ and $\hat{P}^{i-1}$, correspondingly. Applying the associativity of matrix multiplication at Eq.(19), we obtain:

$$A^i = T^i A \ , \tag{20}$$

which means that for every $i \in \mathbb{N}$, the matrix of the profile $\hat{P}^i$ can be obtained from the matrix of the initial profile $\hat{P}$ and the i-th power of the confidence matrix T.

Let us assume that the agents continue to update their judgement sets, i.e. iterations continue indefinitely, or until for some i, we obtain $A^{i+1} = A^i$, which would mean that the opinions are no longer being updated. During this updating process, we assume that a *consensus is reached* if the opinions of the agents converge to the same judgement set, i.e. there exists a judgement set $\hat{J}^*$, such that

$$\lim_{i\to\infty} \hat{J}^i_k = \hat{J}^* \ , \tag{21}$$

for every $k = 1, \dots, n$. This amounts to all the rows of the matrix A^i, we denote them by A^i_k, converging to the same row vector, i.e. *convergence to a consensus* presumes existence of a vector $a^* = (a^*_1, \dots, a^*_{2m})$, such that:

$$\lim_{i\to\infty} A^i_k = a^* \ , \tag{22}$$

for every $k = 1, \dots, n$, and equivalently, due to Eq.(20):

$$\lim_{i\to\infty} T^i_k A = a^* \ , \tag{23}$$

for every $k = 1, \dots, n$, where T^i_k is the k-th row of the matrix T^i.

Now, from Eq.(23) we can observe the following: If the rows of T^i converge to the same row vector i.e., if there exists a vector $\pi = (\pi_1, \dots, \pi_n)$ such that:

$$\lim_{i\to\infty} T^i_k = \pi \ , \tag{24}$$

for every $k = 1, \dots, n$, then the matrix product in Eq.(23) will converge to πA, hence this product will determine a vector a^* with the above requirements. Now, having the vector a^* defined as:

$$a^* = \pi A \ , \tag{25}$$

the corresponding consensual judgement set will be given by:

$$\hat{J}^* = \{\ell(\varphi_j) \geq a_j^* \mid j = 1, \ldots, 2m\} \ , \tag{26}$$

or, in terms of the input judgement sets and the notation in Section 2:

$$\hat{J}^* = \{\ell(\varphi) \geq \sum_{r=1}^{n} \pi_r a_r(\varphi) \mid \varphi \in \Phi^{\cup}\} \ . \tag{27}$$

5 A Necessary and Sufficient Condition for Reaching a Consensus

According to the discussion in the previous section, the existence of a vector π such that Eq.(24) holds is a *sufficient* condition for reaching a concensus and the consensual solution in the case this condition is satisfied is obtained by mutiplying the initial opinion matrix A by π. It is worth noticing that, if such a vector $\pi = (\pi_1, \ldots, \pi_n)$ exists, then its components are non-negative and $\sum_{r=1}^{n} \pi_r = 1$. Let us now see when such π exists.

Observe that the matrix T is a row stochastic matrix (the sum of each row is 1). This means that it can be regarded as the transition probability matrix of a time-homogeneous Markov chain with n states. With this interpretation of T, the condition in Eq.(24) means that π is the limiting distribution of T, which (since T is time-homogeneous) is also a stationary distribution, i.e. satisfies the equation $\pi T = \pi$. Hence, if a solution π to the last equation exists, then a consensus is reached, and the vector π provides the coefficients for the linear combination of individual judgements that gives the consensual solutiuon a^*, i.e. $a^* = \pi A$.

In the example of a confidence matrix of three agents given in the previous section, the corresponding stationary solution will be:

$$\pi = (1/3, 2/3, 0) \ ,$$

and the corresponding consensual vector $a^* = \pi A$, where A is the judgement matrix given in Fig.1, will be the following:

$$a^* = (0.633333, 0.200000, 0.733333, 0.233333, 0.566667, 0.400000) \ .$$

Now, the existence of a limiting distribution is equivalent to the Markov chain being irreducible and aperiodic. This means that all the agents need to form one closed communicating aperiodic class for a global consensus to be reached. In cases where T is block-diagonal, like the one in Fig.2, i.e. there are smaller groups of agents that only give positive confidence to members of their own group, a global consensus will not be reached in the way described above, that is through the stationary distribution of the confidence matrix. Here, a possible solution would be to determine the consensual opinions of each of the groups and then aggregate them, for example, by taking an average.

However, as observed in [1], it is not hard to imagine a case where a consensus obviously exists no matter of the matrix T (and the existence of a stationary distribution).

$$\begin{bmatrix} 1/3 & 2/3 & 0 & 0 \\ 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 1/4 & 3/4 \\ 0 & 0 & 1/2 & 1/2 \end{bmatrix}$$

Fig. 2: A confidence matrix of a group of four agents. Each row is a distribution of trust that an agent has in the opinions of the other agents. Agents 1 and 2 "listen" to each other and form one closed reccurent group of agents that will reach a consensus between themselves. Similarly, agents 3 and 4 communicate between each other and will reach a consensus.

For example, in the trivial case where all the agents have the same probabilistic judgement set $\hat{J}$, the consensus is the set $\hat{J}$ itself, no matter of the confidence matrix T. The authors of [1] proceed further with the above observation and derive a *necessary and sufficient condition* for a consensus to be reached that applies to any possible choice of T and A: For each reccurent class of agents, they construct a certain linear combination of the agents' probability distrubutions, and show that the consensus exists if and only if all of these linear combinations lead to the same probability distribution. For example, for the confidence matrix given in Fig.3, they calculate that the consensus is reached if and only if $\frac{3}{8}p_1 + \frac{3}{11}p_2 + \frac{4}{11}p_3 = \frac{11}{25}p_4 + \frac{14}{25}p_5 = \frac{9}{25}p_6 + \frac{16}{25}p_7$ holds for the probability distributions $p_1, \ldots, p_8$ of the agents. We refer the reader to Theorem 2 in [1] for the latter result as stating it properly here would require intruducing terminology and notation that is beyond the scope of this paper.

$$\begin{bmatrix} 1/2 & 1/4 & 1/4 & 0 & 0 & 0 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 & 0 & 0 & 0 & 0 \\ 1/4 & 1/4 & 1/2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/4 & 3/4 & 0 \\ 0 & 0 & 0 & 1/3 & 2/3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 & 0 & 0 & 0 \\ 1/3 & 0 & 0 & 1/3 & 0 & 0 & 0 & 1/3 \end{bmatrix}$$

Fig. 3: A confidence matrix of a group of eight agents.

The work in both [2] and [1] considers the case where the opinions of the agents are expressed in terms of (precise) probability distributions over a set of mutually exclusive propositions, while in our case, the opinions of the agents are effectively expressed as probability intervals over logically related issues. While the result of [2] applies directly to our case as it depends solely on the matrix T, it is not immediately clear how to apply the result of [1] to our case. One way to proceed would be to form a system of equations based on the linear combinations of distributions as defined in Theorem 2 in [1] and the initial intervals for each probability value, and try to find a solution in terms of imprecise probabilities over the issues (probabilistic judgement set). Another idea would be to look at the intersection of the sets of all possible probability distributions of each agent (if non-empty) determined by the initial probabilistic profile and try to find those probability distributions among them that satisfy the requirement for convergence given in [1]. Exploring these ideas is a future work.

There will still be many cases of a periodic or reccurent confidence matrix not satisfying the necessary and sufficient conditions for reaching a consensus as given in [1], even in the case of imprecise probabilities. A possible general solution would then be to require a modification of the confidence matrix T to an aperiodic irreducible matrix that will have a stationary solution. In practice, this means the agents to be required to redistribute their confidence degrees in a certain way that enables the resulting confidence matrix to have a stationary solution. Finding out how exactly this should be done is also a part of our future work.

6 Conclusions

In this paper we refine our framework for probabilistic judgement aggregation defined in [6] and we propose a new method for aggregating probabilistic judgement of agents based on the method for aggregating probability distributions described in [2]. In order to apply the method from [2] to our case, we prove that any linear combination of the judgements of all the agents leads to a rational judgement if the individual judgements are rational, which we consider a central result of the paper.

By defining the judgement coefficients of the collective judgement as a linear combination of the judgement coefficients of the individual judgements, we satisfy certain aggregation properties by definition: *Universal domain* will certainly hold, as J^* in Eq.(27) is well-defined for every choice of $a_r(\varphi)$, by construction. Proposition 1 proves the property of *rationality*. If all the individual judgements assign a probability estimate larger than $c \in [0, 1]$, then the linear combination of these estimates will also be larger than c, hence *unanimity* will as well be satisfied. If the matrix T has a column k that contains only 1's (while all the other elements are 0), then $\hat{J}^* = \hat{J}_k$ and the aggregation is *dictatorial*.

According to [2], the convergence to a consensus relies on the properties of the confidence matrix which can be regarded as a transition probability matrix of a time-homogeneous Markov chain with n states and hence, according to the theory of Markov chains and their properties, a consensus exists whenever this matrix has a stationary solution. However, as observed in [1], the existence of a stationary vector for the matrix T is just a sufficient, but not a necessary condition for a consensus to be reached, and we discuss how a consensual solution could be reached in cases when the properties of the matrix T do not guarantee one.

There exist other works on aggregating opinions on logically related issues by convergence to a consensus using the DeGroot framework [8]. However, the way these works express and deal with the logical relatedness of the issues is different than ours, namely, they express the opinions of agents as subjective degrees of pair-wise logical relatedness of the issues. In our case, the logical relatedness of the issues is predetermined (by an agenda setter, for example) and formalized in their representation as propositional formulas, while the opinions of the agents are probabilistic estimates of the truth of the issues.

References

1. Berger, R.: A necessary and sufficient condition for reaching a consensus using degroot's method. Journal of the American Statistical Association **76**, 415–418 (1981)
2. DeGroot, M.H.: Reaching a consensus. Journal of the American Statistical Association **69**, 118–121 (1974)
3. Fagin, R., Halpern, J.Y., Megiddo, N.: A Logic for Reasoning about Probabilities. Information and Computation **87**, 78–128 (1990). https://doi.org/10.1016/0890-5401(90)90060-U
4. Grossi, D., Pigozzi, G.: Judgment Aggregation: A Primer. Morgan and Claypool Publishers, San Rafael, CA, USA (2014). https://doi.org/10.2200/S00559ED1V01Y201312AIM027
5. Halpern, J.Y.: Reasoning about uncertainty. MIT Press (2005), https://mitpress.mit.edu/books/reasoning-about-uncertainty-second-edition
6. Ivanovska, M., Slavkovik, M.: Aggregating probabilistic judgments. In: Moss, L.S. (ed.) Proceedings Seventeenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2019. EPTCS, vol. 297, pp. 273–292 (2019). https://doi.org/10.4204/EPTCS.297.18, https://doi.org/10.4204/EPTCS.297.18
7. List, C., Puppe, C.: Judgment aggregation: A survey. In: Anand, P., Puppe, C., Pattanaik, P. (eds.) The Handbook of Rational and Social Choice. Oxford University Press, UK (2009). https://doi.org/10.1093/acprof:oso/9780199290420.003.0020
8. Parsegov, S.E., Proskurnikov, A.V., Tempo, R., Friedkin, N.E.: Novel multidimensional models of opinion dynamics in social networks. vol. 62, pp. 2270–2285 (2017). https://doi.org/10.1109/TAC.2016.2613905