

Group Decision Making via Probabilistic Belief Merging

Nico Potyka

University of Osnabrück
Germany

Erman Acar

University of Mannheim
Germany

Matthias Thimm

University of Koblenz
Germany

Heiner Stuckenschmidt

University of Mannheim
Germany

Abstract

We propose a probabilistic-logical framework for group decision-making. Its main characteristic is that we derive group preferences from agents' beliefs and utilities rather than from their individual preferences as done in social choice approaches. This can be more appropriate when the individual preferences hide too much of the individuals' opinions that determined their preferences. We introduce three preference relations and investigate the relationships between the group preferences and individual and subgroup preferences.

1 Introduction

Decision theory explains how to rationally choose between different alternatives by means of their expected utilities [Fishburn, 1969; French, 1986]. In group decision-making, we face the problem that different individuals have different opinions about the probability and the utility of alternatives [Keeney and Raiffa, 1976; Gilboa *et al.*, 2004; Chambers and Hayashi, 2006; Nehring, 2007; Gajdos *et al.*, 2008; Crès *et al.*, 2011]. There are different ways to address this problem. Some approaches focus on the dynamic aspects of the group decision making process like communication between agents [Wooldridge and Jennings, 1999; Panzarasa *et al.*, 2002]. A rather static approach is to take the individual preferences for granted and to apply voting rules to the individuals' preferences to make a group decision [Shoham and Leyton-Brown, 2008; Brandt *et al.*, 2012].

Our approach falls into the static category, but instead of aggregating the individuals' preferences, we will aggregate their beliefs that constituted these preferences. To this end, we will consider knowledge bases for individuals that contain their personal beliefs about alternatives and criteria that they consider important for their decision. We will represent the agents' beliefs by probabilistic conditionals that express subjective beliefs [Lukasiewicz, 1999; Kern-Isberner, 2001]. In general, there will be conflicts between these beliefs that we resolve by means of generalized probabilistic entailment [Potyka and Thimm, 2015] as will be explained later. In this way, we obtain an interval of expected utilities with respect

to the group beliefs. From this interval, we will derive three group preference relations: a pessimistic one, an optimistic one, and a cautious one. We assume that all agents' beliefs are available and leave strategic issues like how to manipulate the outcome by proclaiming false beliefs for future work.

We will start by explaining the probabilistic-logical framework [Lukasiewicz, 1999] and generalized probabilistic entailment [Potyka and Thimm, 2015] (Section 2). Subsequently, we introduce our decision-theoretic framework, our belief aggregation functions and preference relations (Section 3). We will then investigate how our aggregated group preferences relate to preferences of individuals and subgroups and how the parameters of generalized probabilistic entailment can be used to control the influence of large interest groups (Section 4). All proofs have been moved to an online appendix¹ to meet space restrictions.

2 Basics

We consider a relational probabilistic logic \mathcal{L} built up over a finite signature $\Sigma = (Const, Pred)$, where $Const$ is a finite set of constants and $Pred$ is a finite set of predicate symbols. A *term over \mathcal{L}* is a variable or a constant. Formulas are built up over the terms and predicate symbols in the usual way. A formula is called *ground* iff it does not contain variables. A *possible world over \mathcal{L}* is a truth assignment to the ground atoms in \mathcal{L} , similar to Herbrand interpretations. We denote the set of all possible worlds by Ω . Satisfaction of a ground formula ϕ by a possible world ω is defined in the usual propositional way and is denoted by $\omega \models \phi$, see, e. g., [Lukasiewicz, 1999] for more details. We will define the semantics of non-ground formulas later by considering their ground instances.

A probabilistic conditional over \mathcal{L} is an expression of the form $(\phi \mid \psi)[l, u]$, where $\phi, \psi \in \mathcal{L}$ and $l, u \in [0, 1], l < u$. Intuitively, $(\phi \mid \psi)[l, u]$ expresses that the conditional probability of ϕ given that ψ holds is between l and u . If $\psi \equiv \top$ we simply write $(\phi)[l, u]$ instead of $(\phi \mid \psi)[l, u]$. Semantics are given to probabilistic conditionals by means of probability distributions over possible worlds. For a probability distribution $P : \Omega \rightarrow [0, 1]$ and a ground formula ϕ , we let $P(\phi) = \sum_{\omega \models \phi} P(\omega)$. As in [Lukasiewicz, 1999], we say that P satisfies the probabilistic conditional $(\phi \mid \psi)[l, u]$ iff

$$l \cdot P(\psi) \leq P(\phi \wedge \psi) \leq u \cdot P(\psi). \quad (1)$$

¹mthimm.de/misc/gdm_ijcai2016_proofs.pdf

This corresponds to demanding that the conditional probability of ϕ given ψ is between l and u whenever $P(\psi) > 0$. If $P(\psi) = 0$, P satisfies the conditional independently of l and u . This behaviour is similar to the behaviour of implication in classical logic in that we can conclude arbitrary things from a false assumption. The definition is used frequently in probabilistic logics [Paris, 1994; Hansen and Jaumard, 2000; Lukasiewicz, 1999] because it has some technical advantages like guaranteeing that model sets are topologically closed. A *probabilistic knowledge base* \mathcal{K} over \mathcal{L} is a finite set of probabilistic conditionals over \mathcal{L} . A probability distribution P satisfies \mathcal{K} (is a *model* of \mathcal{K}) iff P satisfies all conditionals in \mathcal{K} . We denote set of all models of \mathcal{K} by $Mod(\mathcal{K})$.

Note that (1) corresponds to two linear inequalities over probability functions. With a slight abuse of notation, we will identify probability distributions P with probability vectors of size $|\Omega|$ (where the i -th component contains the probability of the i -th world with respect to an arbitrary ordering). Then linearity of (1) allows us to rewrite (1) compactly in the form $aP \leq 0$, where a is an $|\Omega|$ -dimensional row vector. If a knowledge base \mathcal{K} induces k such row vectors a_1, \dots, a_k , we let $I_{\mathcal{K}} = \{1, \dots, k\}$ denote the corresponding set of indices. Then

$$Mod(\mathcal{K}) = \{P \mid \forall i \in I_{\mathcal{K}} : a_i P \leq 0\}.$$

The *probabilistic entailment problem* [Nilsson, 1986; Hansen and Jaumard, 2000; Lukasiewicz, 1999] is defined as follows. Given a query $(\phi \mid \psi)$ consisting of two formulas $\phi, \psi \in \mathcal{L}$, compute tight upper and lower bounds on the probability of ϕ given ψ among all probability distributions P that satisfy \mathcal{K} and $P(\psi) > 0$. Formally, the bounds are defined by the optimization problems [Hansen and Jaumard, 2000]

$$\min_{P \in Mod(\mathcal{K})} / \max_{P \in Mod(\mathcal{K})} \frac{P(\phi \wedge \psi)}{P(\psi)} \quad (P(\psi) > 0). \quad (2)$$

In case that $P(\psi) = 0$, we return the empty interval $[1, 0]$. However, probabilistic entailment yields no solution if \mathcal{K} is inconsistent, that is, if $Mod(\mathcal{K}) = \emptyset$. In order to overcome this problem, we can replace the distributions that satisfy \mathcal{K} with those that minimally violate \mathcal{K} in the sense that they minimize the error in the system of inequalities $\{a_i P \leq 0 \mid i \in I_{\mathcal{K}}\}$ [Potyka, 2014; De Bona and Finger, 2015; Potyka and Thimm, 2015]. More generally, we can minimize over the set of probability distributions that satisfy a distinguished system of inequalities called *integrity constraints* that are denoted by \mathcal{IC} . If \mathcal{IC} is empty, we end up with the original formulation.

We will use *p-norms* to measure the error in the system of inequalities. For $p \geq 1$, the *p-norm* $\|\cdot\|_p : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by $\|x\|_p = \sqrt[p]{\sum_{i=1}^n |x_i|^p}$. For the limit $p \rightarrow \infty$, we obtain the *maximum norm* $\|\cdot\|_{\infty}$ that is defined by $\|x\|_{\infty} = \max\{|x_i| \mid 1 \leq i \leq n\}$. Let $f^+ : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ be defined by $f(x) = \max\{0, x\}$. The minimal violation value of \mathcal{K} with respect to $\|\cdot\|_p$ and \mathcal{IC} is defined by the optimization problem

$$\min_{P \in Mod(\mathcal{IC})} \|f^+(v(P))\|_p,$$

where $v(P)$ is the $|I_{\mathcal{K}}|$ -dimensional *violation vector* whose i -th component contains the value of $a_i P$ and $f^+(v(P))$

means component-wise application of f^+ to the vector $v(P)$. Note that the minimal violation value is 0 if and only if $a_i P \leq 0$ for all $i \in I_{\mathcal{K}}$, i. e., iff $\mathcal{K} \cup \mathcal{IC}$ is consistent. The optimal solutions of this problem are called the *generalized models* of \mathcal{K} , denoted by $GMod_{\mathcal{IC}}^{\|\cdot\|}(\mathcal{K})$ [Potyka and Thimm, 2014]. See [Potyka, 2014; Potyka and Thimm, 2015] for some intuitive explanations on the influence of the choice of p .

To improve readability, we may omit the subscript \mathcal{IC} and the superscript $\|\cdot\|$ when they are clear from the context. We can use the generalized models to generalize probabilistic entailment in a straightforward way by replacing the models with the generalized models in (2), see [Potyka and Thimm, 2014] for a detailed discussion. This problem is called the *generalized entailment problem* and can generally be solved by convex optimization techniques [Potyka and Thimm, 2015].

3 Decision-Theoretic Framework

Given a probabilistic relational language \mathcal{L} as introduced in the previous section, we consider a group decision-making problem over a set of agents $N = \{1, \dots, n\}$. A *decision base* consists of the public and individual beliefs of our agents, a set of alternatives from which the group can choose, a set of criteria on which the decision depends, and a utility function.

Definition 1 (Decision Base). A *decision base* \mathcal{D} over a set of agents $N = \{1, \dots, n\}$ is a tuple $(\mathcal{K}, \mathcal{A}, \mathcal{C}, U)$, where

- $\mathcal{K} = \langle \mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_n \rangle$ is called the *knowledge base of \mathcal{D}* . $\mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_n$ are probabilistic knowledge bases over \mathcal{L} such that $\mathcal{K}_0 \cup \mathcal{K}_i$ is consistent for all $i \in N$. \mathcal{K}_0 contains the *public (explicit) beliefs* and \mathcal{K}_i contains the *individual (explicit) beliefs of agent i* .
- $\mathcal{A} = \{a_1, \dots, a_k\}$ is a non-empty *set of alternatives*, where each alternative $a_i \in Const$ is a constant.
- $\mathcal{C} = \{C_1, \dots, C_m\}$ is a non-empty *set of criteria*, where each criterion $C \in \mathcal{C}$ is a formula that contains exactly one free variable.
- $U : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}^n$ is the *utility function* that maps each criterion to a non-negative utility vector whose i -th value represents the utility for agent i . For a criterion $C \in \mathcal{C}$, we denote agent i 's utility by $u_i(C)$, that is, $U(C) = (u_1(C), \dots, u_n(C))$.

We do not demand that agents' utilities are normalized in the sense that $\sum_{i=1}^m u_j(C_i) = \sum_{i=1}^m u_k(C_i)$ for all agents $j, k \in N$. However, there can be good reasons for making this assumption and we can do so, of course, without losing any properties in the general analysis that will follow. The reason for demanding that a criterion contains exactly one variable is that criteria serve as templates for properties that an alternative might have.

Example 1. Suppose there are two agents $N = \{1, 2\}$, which are about to choose a politician from a set of candidates $\mathcal{A} = \{peter, nicole\}$. Our agents evaluate the candidates with respect to the criteria $\mathcal{C} = \{Honest(x), Intelligent(x)\}$. \mathcal{K}_0

contains $(Prestigious(harvard))[1]$. Agent 1 believes
 $(Honest(peter))[0.2], (Honest(nicole))[0.9],$
 $(Studied(peter, harvard))[1],$
 $(Intelligent(x) \mid Studied(x, y) \wedge Prestigious(y))[0.9].$

Agent 2's knowledge base contains

$(Honest(peter))[0.9], (Honest(nicole))[0.6],$
 $(LooksIntelligent(nicole))[1],$
 $(Intelligent(x) \mid LooksIntelligent(x))[0.8].$

Furthermore, $u_1(Intelligent) = 70, u_2(Intelligent) = 80, u_1(Honest) = 40, u_2(Honest) = 95$. Note that the agents not only have different beliefs about the candidates, but also about the causes of different criteria.

Given the beliefs and preferences of the agents, which alternative should the group choose if all agents are treated equally? Even though we demand that all individual beliefs are consistent with the public beliefs, the union $\bigcup_{i=0}^n \mathcal{K}_i$ of all knowledge bases can be inconsistent. However, in order to achieve a group decision that is fair with respect to the beliefs of all agents, we must take all knowledge bases into account. We do this by applying generalized entailment.

Definition 2 (Individual Beliefs, Group Belief). Let $\mathcal{D} = (\mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U$ be a decision base. Let $\phi, \psi \in \mathcal{L}$ be formulas. For $i \in N$, agent i 's individual belief in ϕ given ψ is the interval $\mathcal{B}_i(\phi \mid \psi) = [l_i, r_i]$, where l_i and r_i denote the minimum and maximum of

$$\min_{P \in \text{Mod}(\mathcal{K}_0 \cup \mathcal{K}_i)} \frac{P(\phi \wedge \psi)}{P(\psi)} \Big/ \max_{P \in \text{Mod}(\mathcal{K}_0 \cup \mathcal{K}_i)} \frac{P(\phi \wedge \psi)}{P(\psi)} \quad (P(\psi) > 0).$$

Let $\|\cdot\|$ be a p -norm and let $\mathcal{M} = \bigsqcup_{i \in N} \mathcal{K}_i$ denote the multiset obtained from adding all agents' individual beliefs. The group belief in ϕ given ψ is the interval $\mathcal{B}_G(\phi \mid \psi) = [l, r]$, where l and r denote the minimum and maximum of

$$\min_{P \in \text{GM}_{\mathcal{K}_0}^{\|\cdot\|}(\mathcal{M})} \frac{P(\phi \wedge \psi)}{P(\psi)} \Big/ \max_{P \in \text{GM}_{\mathcal{K}_0}^{\|\cdot\|}(\mathcal{M})} \frac{P(\phi \wedge \psi)}{P(\psi)} \quad (P(\psi) > 0).$$

Remark 1. 1. Note that the public beliefs \mathcal{K}_0 serve as integrity constraints for the group beliefs and are therefore guaranteed to be maintained [Potyka and Thimm, 2015]. The agents' individual beliefs are treated equally by adding all individual knowledge bases to a multiset.

2. We will often consider unconditional beliefs. Formally this means that for $\psi \equiv \top$ we will again abbreviate $\mathcal{B}_i(\phi \mid \psi)$ by $\mathcal{B}_i(\phi)$ and $\mathcal{B}_G(\phi \mid \psi)$ by $\mathcal{B}_G(\phi)$.

We show the individual and group beliefs for Example 1 in Table 1. In the following, we will often be interested in decision bases that are conflict-free in the following sense.

Definition 3 (Conflict-free Decision Base). A decision base $\mathcal{D} = (\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U$ is called *conflict-free* iff $\mathcal{K}_0 \cup \bigcup_{i \in N} \mathcal{K}_i$ is consistent.

We want to compute the expected utility of an alternative with respect to the group beliefs. Since \mathcal{B}_G yields an interval, we define the following operations on intervals:

$$c \cdot [L, R] = [L, R] \cdot c = [c \cdot L, c \cdot R],$$

$$[L, R] + [L', R'] = [L + L', R + R'].$$

	Honest		Intelligent	
	nicole	peter	nicole	peter
\mathcal{B}_1	[0.9, 0.9]	[0.2, 0.2]	[0, 1]	[0.9, 0.9]
\mathcal{B}_2	[0.6, 0.6]	[0.9, 0.9]	[0.8, 0.8]	[0, 1]
\mathcal{B}_G	[0.6, 0.9]	[0.2, 0.9]	[0.8, 0.8]	[0.9, 0.9]

Table 1: Individual Beliefs $\mathcal{B}_1, \mathcal{B}_2$ and group belief \mathcal{B}_G for Example 1 when using the 1-norm.

	nicole	peter
EU_1	[36, 106]	[71, 71]
EU_2	[121, 121]	[85.5, 165.5]
EU_G	[100.5, 120.75]	[81, 128.25]

Table 2: Expected utilities of individuals and the group for Example 1 when using the 1-norm.

where $c, L, R', U, R' \in \mathbb{R}_{\geq 0}$. Note that we will never get negative values in our computations because both probabilities and our utilities are non-negative. We can now define our expected utilities with respect to individuals and the group.

Definition 4 (Individual, Group Expected Utility Interval). Consider a decision base $\mathcal{D} = (\mathcal{K}, \mathcal{A}, \mathcal{C}, U)$. The *individual expected utility interval of agent $i \in N$ of alternative $a \in \mathcal{A}$* is

$$EU_i(a) = \sum_{C \in \mathcal{C}} \mathcal{B}_i(C(a)) \cdot u_i(C).$$

The *group's expected utility interval of an alternative $a \in \mathcal{A}$* is

$$EU_G(a) = \frac{1}{|N|} \sum_{i \in N} \sum_{C \in \mathcal{C}} \mathcal{B}_i(C(a)) \cdot u_i(C).$$

For some expected utility interval EU , we will refer to the lower and upper bound of the expected utilities as the *pessimistic* and *optimistic* expected utilities and denote it by \underline{EU} and \overline{EU} , respectively.

Table 2 shows the expected utilities for Example 1. We can derive different preference relations from the expected utility intervals. We consider interval preference relations that are defined by means of a point (expected) utility function EU and a vagueness function σ .

Definition 5 (optimistic, pessimistic, cautious preference). Let EU, σ be functions mapping \mathcal{A} to the non-negative reals. The *interval preference relation $\succeq_{\subseteq} \mathcal{A} \times \mathcal{A}$* with respect to EU and σ is defined by $a_1 \succeq_{\subseteq} a_2$ iff $EU(a_1) \geq EU(a_2) + \sigma(a_2)$. We say that a_1 is *preferred* to a_2 iff $a_1 \succeq a_2$; a_1 is *indifferent* to a_2 denoted by $a_1 \sim a_2$ iff $a_1 \succeq a_2$ and $a_2 \succeq a_1$; a_1 is *strictly preferred* to a_2 denoted by $a_1 \succ a_2$, iff $a_1 \succeq a_2$ but $a_2 \not\succeq a_1$. If a_1, a_2 are incomparable w.r.t. \succeq , we write $a_1 \parallel a_2$. We will consider the interval preference relations shown in Table 3.

Optimistic (pessimistic) preference corresponds to comparing the upper (lower) bounds of the expected utility intervals. a is cautiously preferred over b iff the even the lower bound for a is above the upper bound for b . Note that $a \succ_X^c b$ implies both $a \succ_X^p b$ and $a \succ_X^o b$.

	\succ_X^p	\succ_X^o	\succ_X^c
EU	\underline{EU}_X	\overline{EU}_X	\underline{EU}_X
σ	0	0	$\overline{EU}_X - \underline{EU}_X$

Table 3: (*p*)essimistic, (*o*)ptimistic and (*c*)autious preference relations. $X \in N \cup \{G\}$ stands for the *individual preference of agent* $i \in N$ or the *group preference*, respectively. For instance, $a_1 \succ_X^c a_2$ iff $\underline{EU}_X(a_1) \geq \underline{EU}_X(a_2) + \overline{EU}_X(a_2) - \underline{EU}_X(a_2)$, i.e., iff $\underline{EU}_X(a_1) \geq \underline{EU}_X(a_2)$.

Example 2. From Table 2, we can see that peter \succ_1^p nicole and nicole \succ_1^o peter. All alternatives are cautiously incomparable.

4 Analysis

The purpose of this section is to investigate some interesting properties of our approach. In particular, we are interested in the influence of the interactions between agents' beliefs on the expected utilities and the corresponding preference relations. It is difficult to make an objective statement about what we should expect in the presence of conflicts between agents' beliefs, but we can think of some intuitive properties that should hold if there are no conflicts. Then we can attack the question what happens if there are minor conflicts in the decision base by continuity arguments. We will also address the questions how independencies in the agents' beliefs can be exploited and how we can control the influence of large interest groups on the group decisions. The intuitive notions used in the following overview of our results will be made precise in the rest of this section. In order to meet space restrictions, we moved the proofs to an online appendix².

- **Consensus:** If the decision base is conflict-free, the expected utility of the group will be a refinement of the individual expected utilities in the sense that it yields a subinterval of the averaged individual belief intervals.
- **Cautious Dominance:** If the decision base is conflict-free and alternative a is cautiously preferred over b by all agents, then a will be cautiously preferred over b with respect to the group preferences.
- **Cautious Condorcet-Consistency:** If the decision base is conflict-free and alternative a is cautiously preferred over all b by all agents, then a will be cautiously preferred over all b with respect to the group preferences.
- **Continuity of Expected Utilities:** Minor changes in the agents' beliefs and utilities will not result in major changes in the expected utilities. In particular, if there are only minor conflicts in the knowledge base, Consensus, Cautious Dominance and Cautious Condorcet-Consistency will still hold.
- **Decomposition of Utility:** If the agents' beliefs are independent of each other, the problem of computing expected utilities can be decomposed.

- **Modularity:** If the agents' beliefs and utilities are independent of each other, then the expected utility of the decision base will be a weighted sum over the expected utilities of independent sub decision bases.
- **Majority:** The influence of large interest groups on the aggregated group beliefs and preferences can be regulated by the choice of the norm.

We will illustrate the cases in which these properties can fail by means of examples.

Let us first investigate the connection between individual expected utilities and the group's expected utilities. Our first proposition states that if there are no conflicts in the agent's beliefs, then the expected utility of the group will be a refinement of the averaged expected utilities of the agents. The result is a consequence of the following two lemmas.

Lemma 1. Let $[l_i, r_i], [L_i, R_i]$ be non-negative intervals for $i = 1, \dots, k$ such that $[l_i, r_i] \subseteq [L_i, R_i]$. Then for all $c_i \in \mathbb{R}_{\geq 0}$, we have $\sum_{i=1}^k [l_i, r_i] \cdot c_i \subseteq \sum_{i=1}^k [L_i, R_i] \cdot c_i$.

Lemma 2. If \mathcal{D} is conflict-free, then $\mathcal{B}_G(F) \subseteq \mathcal{B}_i(F)$ for all $i \in N$. In particular, $\mathcal{B}_G(F) \subseteq \bigcap_{i \in N} \mathcal{B}_i(F)$

Proposition 1 (Consensus). If \mathcal{D} is conflict-free, then $EU_G(a) \subseteq \frac{1}{|N|} \sum_{i \in N} EU_i(a)$ for all $a \in \mathcal{A}$.

However, the following example shows that conflicts can prohibit such a relationship in general.

Example 3. Consider two agents, one alternative a and one criterion C . Let the public knowledge base be empty, let agent 1 believe $(C(a))[0]$ and let agent 2 believe $(C(a))[1]$. Let $u_1(C) = u_2(C) = 1$. If we employ the Manhattan norm, we have that $\mathcal{B}_i(C(a)) = [i - 1, i - 1]$ for $i \in N$ and that $\mathcal{B}_G(C(a)) = [0, 1]$. Hence,

$$\frac{1}{2} \sum_{i=1}^2 EU_i(a) = \frac{1}{2}([0, 0] \cdot 1 + [1, 1] \cdot 1) = \left[\frac{1}{2}, \frac{1}{2}\right]$$

$$EU_G(a) = \frac{1}{2}[0, 1] \cdot (1 + 1) = [0, 1].$$

If there are no conflicts in the decision base and all agents' cautious preferences agree for some alternatives, then this preference will be maintained in the group preference.

Proposition 2 (Cautious Dominance, Condorcet-Consistency). Suppose that \mathcal{D} is conflict-free.

1. **Cautious Dominance:** If there are $a, b \in \mathcal{A}$ such that $a \succeq_i^c b$ for all $i \in N$, then $a \succeq_G^c b$.
2. **Cautious Condorcet-Consistency:** If there is an $a \in \mathcal{A}$ such that $a \succeq_i^c b$ for all $b \in \mathcal{A}, i \in N$, then $a \succeq_G^c b$ for all $b \in \mathcal{A}$.

We conjecture that Dominance and Condorcet-Consistency also hold for optimistic and pessimistic preference, but did not find a proof so far.

Again, conflicts in the knowledge base can prohibit Dominance and Condorcet-Consistency. This can be seen from Example 3 by adding an alternative b such that the agents have the same beliefs and utilities for both a and b . Then $a \sim_1^c b$, $a \sim_2^c b$ and $a \parallel_G^c b$. Hence, in particular, both $a \succeq_1^c b$ and $a \succeq_2^c b$ hold, but $a \not\succeq_G^c b$.

²mthimm.de/misc/gdm_ijcai2016_proofs.pdf

However, by using a continuity argument, we can show that both properties remain true if the conflicts in the decision base are not too strong. Intuitively, this is the case if the knowledge base is close to a conflict-free knowledge base. In order to make this intuition more precise, we have to define how to compare knowledge bases. This question is subtle, but can be addressed appropriately by comparing knowledge bases extensionally, that is, with respect to their model sets [Paris, 1994]. The *Blaschke distance* $\|S_1, S_2\|_B$ between two convex sets of probability distributions S_1, S_2 is the smallest real number d such that for each distribution in one of the sets, there is a probability distribution in the other that has total variation distance at most d to the former. Formally, $\|S_1, S_2\|_B$ is defined by

$$\inf\{\delta \in \mathbb{R} \mid \forall P_1 \in S_1 \exists P_2 \in S_2 : \|P_1, P_2\|_1 \leq \delta \text{ and} \\ \forall P_2 \in S_2 \exists P_1 \in S_1 : \|P_2, P_1\|_1 \leq \delta\}$$

Generalized Entailment is continuous in the following sense: If the generalized models of \mathcal{K}_1 are close to the generalized models of \mathcal{K}_2 with respect to the Blaschke distance, then the entailment results will be close [Potyka and Thimm, 2015]. In order to talk about the influence of minor changes in the utilities, we define the distance between two knowledge bases over the same set of agents, alternatives, and criteria by measuring the difference in the extensional knowledge and in the utilities.

Definition 6 (Distance between Decision Bases). Let $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$, $\mathcal{D}' = ((\mathcal{K}'_0, \dots, \mathcal{K}'_n), \mathcal{A}, \mathcal{C}, U')$ be decision bases over a set of agents N . Let $\mathcal{M}, \mathcal{M}'$ be defined as in Definition 2 for \mathcal{D} and \mathcal{D}' , respectively. The distance between \mathcal{D} and \mathcal{D}' is

$$d(\mathcal{D}, \mathcal{D}') = \|GMod_{\mathcal{K}_0}(\mathcal{M}), GMod_{\mathcal{K}'_0}(\mathcal{M}')\|_B + \|U, U'\|,$$

where $\|U, U'\| = \sum_{i \in N} \sum_{C \in \mathcal{C}} |u_i(C) - u'_i(C)|$.

The next proposition intuitively states that if a decision base is close to another decision base, then the expected utilities will also be close.

Proposition 3 (Continuity of Expected Utilities). Let $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$, $\mathcal{D}_l = ((\mathcal{K}_{0,l}, \dots, \mathcal{K}_{n,l}), \mathcal{A}, \mathcal{C}, U_l)$ for $l \in \mathbb{N}$. Let EU, EU' denote some expected utility (individual or group for some arbitrary alternative). Then $EU \rightarrow EU'$ whenever $\mathcal{D}_l \rightarrow \mathcal{D}$.

Now suppose that a decision base contains only minor conflicts in the sense that it is close to a conflict-free decision base. Then Continuity implies that Consensus and Cautious Dominance will almost hold. That is, the expected utility interval for the group will not significantly exceed the averaged expected utility interval of the individual agents. Furthermore, if alternative a was significantly preferred over b (the difference $EU(a) - (EU(b) + \sigma(b))$ is not too small) then a will still be preferred over b in the group.

If our decision base decomposes into independent parts, we can reduce the problem of computing expected utilities for the group to two smaller problems. A sufficient condition is given in the following proposition. To exclude dependencies, we assume that the agents decompose into two groups that have only beliefs about distinct criteria.

Proposition 4 (Decomposition of Utility). Let $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$ be a decision base over a finite set of agents N . Suppose that $N = N_1 \uplus N_2$, $\mathcal{K}_0 = \mathcal{K}_{0,1} \uplus \mathcal{K}_{0,2}$, $\mathcal{C} = \mathcal{C}_1 \uplus \mathcal{C}_2$ and that $\mathcal{K}_{0,1} \cup \mathcal{K}_i$ contains only criteria from \mathcal{C}_1 for $i \in N_1$ and $\mathcal{K}_{0,2} \cup \mathcal{K}_i$ contains only criteria from \mathcal{C}_2 for $i \in N_2$. Let $U|_{\mathcal{C}_i}$ denote the restriction of U to \mathcal{C}_i for $i \in \{1, 2\}$ and let $\mathcal{D}_i = (\mathcal{K}_{0,i} \sqcup \langle \mathcal{K}_j \mid j \in N_i \rangle, \mathcal{A}, \mathcal{C}_i, U|_{\mathcal{C}_i})$ denote the restricted decisions bases for the agents in N_i . For $a \in \mathcal{A}$, let $EU_G(a)$ denote the expected group utility with respect to \mathcal{D} and $EU_G^i(a)$ the expected group utility of alternative a with respect to \mathcal{D}_i . Then

$$EU_G(a) = \frac{|N_1|}{|N|} EU_G^1(a) + \frac{|N_2|}{|N|} EU_G^2(a) + IU(a),$$

where $IU(a)$ denotes the agents' utility that is independent of their beliefs:

$$\frac{\sum_{i \in N_2} \sum_{C \in \mathcal{C}_1} \mathcal{B}_G^1(C(a)) \cdot u_i(C) + \sum_{i \in N_1} \sum_{C \in \mathcal{C}_2} \mathcal{B}_G^2(C(a)) \cdot u_i(C)}{|N|}.$$

Besides the computational benefits, Decomposition has also an interesting intuitive interpretation. Suppose that a subset of agents is completely indifferent with respect to a subset of criteria in the sense that the agents do not have any knowledge about these criteria and their utility for these criteria is 0. Then the expected utility of the group will be a weighted sum of the expected utilities of the subgroups as shown in the following corollary.

Corollary 1 (Modularity). In addition to the assumptions made in Proposition 4, assume that $u_i(C_j) = 0$ whenever $i \neq j$ for $i, j \in \{1, 2\}$. Then

$$EU_G(a) = \frac{|N_1|}{|N|} EU_G^1(a) + \frac{|N_2|}{|N|} EU_G^2(a), \quad (3)$$

Let us now investigate the influence of majority beliefs. That is, if a overwhelming majority of agents share the same beliefs, how does this influence the group beliefs and the corresponding expected utilities? The answer depends on the norm that we use for generalized entailment. We formalize the notion of majority beliefs by adding copies of a particular agent. Intuitively, this agent can be understood as a representative of a fraction sharing common beliefs. We are interested in what happens if the fraction grows. Given some decision base \mathcal{D} , we let \mathcal{D}_k^s denote the decision base that is obtained from \mathcal{D} by adding k copies of some agent $s \in N$. Formally, if $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$, then

$$\mathcal{D}_k^s = ((\mathcal{K}_0, \dots, \mathcal{K}_n, \underbrace{\mathcal{K}_s, \dots, \mathcal{K}_s}_{k \text{ times}}, \mathcal{A}, \mathcal{C}, U'),$$

where for $n < l \leq n + k$, $u'_l(C) = u'_s(C)$ for all $C \in \mathcal{C}$.

The following proposition states that the size of the fraction has no influence whatsoever on the outcome when using the maximum norm.

Proposition 5 (Maximum-Norm Majority Ignorance). Let $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$ be a decision base over a finite set of agents N , let $s \in N$ and let \mathcal{B}_G^k denote the group belief function w.r.t. \mathcal{D}_k^s .




Figure 1: Our group belief approach (solid lines) versus the qualitative social choice approach (dashed lines)

Suppose that \mathcal{B}_G is defined with respect to the maximum norm, and let $\phi \in \mathcal{L}$ be a formula. Then $\mathcal{B}_G^k(\phi) = \mathcal{B}_G(\phi)$ for all $k \in \mathbb{N}$. In particular, for all $a \in \mathcal{A}$, $EU_G^k(a) = EU_G(a)$.

When applying p -norms other than the maximum norm, intuition suggests that the group beliefs will converge to the majority beliefs. However, we have to rule out some exceptions. Even though we could not find a proof so far, there is some empirical evidence that the following conjecture is true.

Conjecture 1 (p -Norm Majority Sensitivity). Let $\mathcal{D} = ((\mathcal{K}_0, \dots, \mathcal{K}_n), \mathcal{A}, \mathcal{C}, U)$ be a decision base over a finite set of agents N let $s \in N$ and let \mathcal{B}_G^k denote the group belief function w.r.t. \mathcal{D}_k^s .

Suppose that \mathcal{B}_G is defined with respect to a p -norm, where $1 \leq p < \infty$. Let $\phi \in \mathcal{L}$ be a formula such that $\mathcal{B}_s(\phi) \neq [0, 1]$. Then $\mathcal{B}_G^k(\phi)$ approaches $\mathcal{B}_s(\phi)$ as $k \rightarrow \infty$.

5 Related Work

There exist various methods in the literature to make group decisions. Many frameworks and their properties have been discussed extensively in the economic literature, see, for instance, [Gilboa *et al.*, 2004; Chambers and Hayashi, 2006; Nehring, 2007; Gajdos *et al.*, 2008; Crès *et al.*, 2011]. Sometimes group decision making is regarded as a dynamic process, in which agents reason about other agents' mental states and try to increase other agent's cooperativeness towards joint decisions in their own interest, see, for instance, [Wooldridge and Jennings, 1999; Panzarasa *et al.*, 2002]. In a similar spirit, group decision making can be regarded as an argumentation problem [Amgoud and Prade, 2009; Fan *et al.*, 2014] Another popular approach is to compute agent preferences individually and to apply voting rules afterwards to make a group decision [Shoham and Leyton-Brown, 2008; Brandt *et al.*, 2012]. Our approach is closer to the latter methodology, but instead of deriving individual preferences first and then merging them into group preferences, we start by merging the individual beliefs to derive group preferences from the group beliefs. We illustrate this in Figure 1.

There are other belief merging approaches that we could apply to derive the group beliefs, see [Konieczny and Pérez, 2011] for classical approaches or [Kern-Isberner and Rödder, 2004; Adamcik, 2014; Wilmers, 2015] for some probabilistic alternatives. An approach very closely related to generalized entailment was introduced in [Daniel, 2009]. Daniel also considered probability functions that, in another sense minimally violate an inconsistent knowledge base. However, generalized entailment has some computational advantages,

see [Potyka and Thimm, 2014; 2015] for a more detailed discussion.

The properties discussed in Section 4 are inspired by the literature on social choice and probabilistic logic. In particular, Dominance and Condorcet-Consistency correspond to social-choice-theoretic properties [Shoham and Leyton-Brown, 2008], whereas Continuity and Modularity correspond to properties considered for probabilistic logics [Paris, 1994; Adamcik, 2014; Wilmers, 2015].

Our decision-theoretic framework extends the single agent framework from [Acar *et al.*, 2015] to several agents. Since in multi-agent settings very different problems arise, the properties analyzed here are very different from those considered in [Acar *et al.*, 2015]. Some properties like Continuity and Decomposition are, however, related to corresponding properties of generalized entailment [Potyka and Thimm, 2015].

6 Conclusions and Future Work

We proposed a probabilistic-logical framework for group decision making and introduced three preference relations. Our decision making approach provides several guarantees as discussed in Section 4. If there are no conflicts among the agents' beliefs, our approach satisfies the social-choice-theoretic properties Dominance and Condorcet-Consistency when using the cautious preference relation. Proving these properties for optimistic and pessimistic preference is part of future work. Continuity of our framework guarantees that Dominance and Condorcet-Consistency remain true when there are only minor conflicts. We are also planning to investigate the general case in more detail to provide more rigorous statements for the case of conflicts. The decomposition properties of our framework are independent of this question and can be used to solve larger decision problems by decomposing them into smaller, independent parts. By applying different p -norms, we can control the influence of large interest groups. Proposition 5 states that the size of the group has no influence at all when we let $p = \infty$. Understanding generalized entailment as a norm minimization problem suggests that, for other choices of p , the group belief will converge to the beliefs of growing interest groups. The reason is that the violation vector will eventually be dominated by entries from the large interest group. Currently, p -norm Majority Sensitivity for $p < \infty$ is only supported by empirical evidence and intuition, a formal proof is part of future work. In practice, $p = \infty$ might be a good choice when majority ignorance is desired, $p = 1$ or $p = 2$ when majority sensitivity is desired. A discussion of the computational implications of the choice of p can be found in [Potyka and Thimm, 2015].

We are also going to consider aggregation functions that are based on point probabilities rather than on probability intervals. To this end, we can determine the group beliefs by means of a best generalized model like the one maximizing entropy [Daniel, 2009; Adamcik, 2014; Potyka and Thimm, 2014; Wilmers, 2015]. This approach will, in particular, always yield point utilities instead of utility intervals. We are also planning to provide a direct comparison between our approach and social choice approaches.

References

- [Acar *et al.*, 2015] Erman Acar, Camilo Thorne, and Heiner Stuckenschmidt. Towards decision making via expressive probabilistic ontologies. In *Algorithmic Decision Theory - 4th International Conference, ADT 2015, Proceedings*, pages 52–68, 2015.
- [Adamcik, 2014] Martin Adamcik. *Collective Reasoning under Uncertainty and Inconsistency*. PhD thesis, University of Manchester, 2014.
- [Amgoud and Prade, 2009] Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artif. Intell.*, 173(3-4):413–436, 2009.
- [Brandt *et al.*, 2012] Felix Brandt, Vincent Conitzer, and Ulle Endriss. *Computational social choice*, pages 213–283. MIT Press, 2012.
- [Chambers and Hayashi, 2006] Christopher P Chambers and Takashi Hayashi. Preference aggregation under uncertainty: Savage vs. pareto. *Games and Economic Behavior*, 54(2):430–440, 2006.
- [Crès *et al.*, 2011] Hervé Crès, Itzhak Gilboa, and Nicolas Vieille. Aggregation of multiple prior opinions. *Journal of Economic Theory*, 146(6):2563–2582, 2011.
- [Daniel, 2009] L. Daniel. *Paraconsistent Probabilistic Reasoning*. PhD thesis, L'École Nationale Supérieure des Mines de Paris, 2009.
- [De Bona and Finger, 2015] G. De Bona and M. Finger. Measuring inconsistency in probabilistic logic: rationality postulates and dutch book interpretation. *Artificial Intelligence*, 227:140–164, 2015.
- [Fan *et al.*, 2014] Xiuyi Fan, Francesca Toni, Andrei Mocanu, and Matthew Williams. Dialogical two-agent decision making with assumption-based argumentation. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14*, pages 533–540. IFAA-MAS/ACM, 2014.
- [Fishburn, 1969] Peter C. Fishburn. *Utility Theory for Decision Making*. Robert E. Krieger Publishing Co., Huntington, New York, 1969.
- [French, 1986] Simon French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Halsted Press, New York, NY, USA, 1986.
- [Gajdos *et al.*, 2008] Thibault Gajdos, J-M Tallon, and J-C Vergnaud. Representation and aggregation of preferences under uncertainty. *Journal of Economic Theory*, 141(1):68–99, 2008.
- [Gilboa *et al.*, 2004] Itzhak Gilboa, Dov Samet, and David Schmeidler. Utilitarian aggregation of beliefs and tastes. *Journal of Political Economy*, 112:932–938, 2004.
- [Hansen and Jaumard, 2000] Pierre Hansen and Brigitte Jaumard. Probabilistic satisfiability. In Jürg Kohlas and Serafin Moral, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, volume 5 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, pages 321–367. Springer Netherlands, 2000.
- [Keeney and Raiffa, 1976] R.L. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. J. Wiley, New York, 1976.
- [Kern-Isberner and Rödder, 2004] G. Kern-Isberner and W. Rödder. Belief revision and information fusion on optimum entropy. *International Journal of Intelligent Systems, Special Issue on Uncertain Reasoning*, 19(9):837 – 857, 2004.
- [Kern-Isberner, 2001] G. Kern-Isberner. *Conditionals in nonmonotonic reasoning and belief revision*. Springer, Lecture Notes in Artificial Intelligence LNAI 2087, 2001.
- [Konieczny and Pérez, 2011] Sébastien Konieczny and Ramón Pino Pérez. Logic based merging. *J. Philosophical Logic*, 40(2):239–270, 2011.
- [Lukasiewicz, 1999] Thomas Lukasiewicz. Probabilistic deduction with conditional constraints over basic events. *Journal of Artificial Intelligence Research*, 10:380–391, 1999.
- [Nehring, 2007] Klaus Nehring. The impossibility of a paretian rational: A bayesian perspective. *Economics Letters*, 96(1):45–50, 2007.
- [Nilsson, 1986] Nils J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–88, February 1986.
- [Panzarasa *et al.*, 2002] Pietro Panzarasa, Nicholas R. Jennings, and Timothy J. Norman. Formalizing collaborative decision-making and practical reasoning in multi-agent systems. *J. Log. Comput.*, 12(1):55–117, 2002.
- [Paris, 1994] J.B. Paris. *The uncertain reasoner's companion – A mathematical perspective*. Cambridge University Press, 1994.
- [Potyka and Thimm, 2014] N. Potyka and M. Thimm. Consolidation of probabilistic knowledge bases by inconsistency minimization. In *Proceedings ECAI 2014*, pages 729–734. IOS Press, 2014.
- [Potyka and Thimm, 2015] N. Potyka and M. Thimm. Probabilistic reasoning with inconsistent beliefs using inconsistency measures. International Joint Conference on Artificial Intelligence 2015 (IJCAI'15), pages 3156–3163, 2015.
- [Potyka, 2014] N. Potyka. Linear programs for measuring inconsistency in probabilistic logics. In *Proceedings KR 2014*, pages 568–578. AAAI Press, 2014.
- [Shoham and Leyton-Brown, 2008] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, NY, USA, 2008.
- [Wilmers, 2015] George Wilmers. A foundational approach to generalising the maximum entropy inference process to the multi-agent context. *Entropy*, 17(2):594–645, 2015.
- [Wooldridge and Jennings, 1999] Michael Wooldridge and Nicholas R. Jennings. The cooperative problem-solving process. *J. Log. Comput.*, 9(4):563–592, 1999.