Group Intentions are Social Choice with Commitment*

short version

Guido Boella[1], Gabriella Pigozzi[2], Marija Slavkovik[2], and Leendert van der Torre[2]

[1] `guido@di.unito.it`
[2] `{gabriella.pigozzi, marija.slavkovik, leon.vandertorre}@uni.lu`

Abstract. An agent intends g if it has chosen to pursue goal g and it has committed itself to pursuing it. How do groups decide which is their goal? Social epistemology gives us two views on collective attitudes: according to the summative approach, a group has attitude p if all or most of the group members have the attitude p; according to the non-summative approach, for a group to have attitude p it is required that the members together agree that they have attitude p. The summative approach is used extensively in multi agent systems. We propose a formalization of non-summative group intentions, using social choice to determine the group goals. We use judgment aggregation as a decision-making mechanism and a multi-modal multi-agent logic to represent the collective attitudes, as well as the commitment and revision strategies for the groups intentions.

1 Introduction

Within the context of multi-agent systems, the concept of collective intentions is studied and formalized in (Chapter 3, [7]) and also in [12, 15, 32]. All of these theories and formalizations use the summative approach to define group beliefs and goals: a group has attitude p if all or most of the group members have the attitude p [8, 13, 23]. Alternatively, collective attitudes can be dened using the non-summative approach: a group has an attitude p if the members together agree that they have that attitude p. To the best of our knowledge, there is no formalization of non-summative collective attitudes within multi agent systems. We consider the following research question:

How can a group agree on what to believe, pursue and what to intend?

This paper summarizes our initial efforts towards formalizing non-summative group intentions using the conceptualizations proposed by Gilbert [8–11].

How can a group decide which goals to pursue? A rational agent makes decisions based on what he believes, what he knows and what he desires. Each group member can expresses whether he is for or against the group needing to achieve a given goal. He can rationalizes this view by expressing opinions on relevant reasons that justify his goal opinion. To reach non-summative group attitudes, a group can use a decision making mechanism that aggregates the members opinions to produce the group agreement of what to believe and, based on those beliefs, which goals to pursue. A group that jointly decided on a course of action is jointly committed to uphold that decision [10].

One of the roles of intentions is to “persist long enough, according to a reconsideration strategy and intentions influence beliefs upon which future practical reasoning

* An extended version of this paper, titled Group Intentions are Social Choice with Commitment, had been submitted to the Post-poceedings of COIN@MALLOW 2010

is based" (Chapter 3, [7]). Thus, a formalization of group intentions should be completed with a formalization of group intention persistence and a consideration of group intentions reconsideration strategies.

Our research question breaks down to the following sub-questions:

1. How to aggregate the individual opinions into group beliefs and goals?
2. How to represent individual opinions and non-summative group attitudes?
3. How can groups persist in their intentions?
4. How can groups reconsider their attitudes?

We need a mechanism for generating group beliefs and goals that aggregates individual opinions into collective attitudes, as studied in voting, merging and judgment aggregation [3, 16, 17]. We find the use of judgment aggregation is the most adequate.

The relation between individual goals and beliefs can be specified and analyzed in modal agent logics like BDI_{LTL} [27]. The challenge is to find an adequate representation for the individual opinions and the non-summative beliefs, goals and intentions into multi agent logic. We give an extension logic AGE_{LTL} that fuses existing modal logics to provide the adequate modalities. We use this logic to represent the group intention and reconsideration strategies.

We require that the group has a set of candidate group goals, a relevance order over this set, as well as a set of decision rules, one for each candidate goal, in the form of logic formulas, that express what is the relation between a goal and a given set of reasons. The members are required to have the ability to form and communicate "yes" or "no" opinions regarding a candidate and associated reasons. There are two modes of communicating the individual judgments: a centralized and a decentralized mode. In the decentralized mode, every individual judgment is communicated to all the agents in the group. Each agent then aggregates the individual judgments using the known mechanism and generates the group beliefs, goals and thus intentions. In the centralized mode, one of the members acts as an administrator for the group. All individual judgments are sent to the administrator who aggregates them and notifies the rest of the members what are the group beliefs and goals.

We assume that all members are aware of the judgment aggregation mechanism (and possibly tie breaking rule), the commitment and the revision strategy that are in use. We also assume that group membership does not change; neither does the aggregation mechanism, the commitment and revision strategy for each goal. The group members can communicate with each other. We further assume that all members accept the decision rules and give opinions that are logically consistent with them. Lastly, we assume that each agent is able to revise his individual judgments, on a given goal and reasons, with given information.

The generation and revision of decision rules is outside of the scope of this paper. The agents of the group may have individual goals in addition to the group ones. It is not a requirement that the group attitudes are decomposable to individual ones.

The layout of the paper is as follows. In Section 2 we discuss show how to choose group goals. We first discuss the non-summative view on collective attitudes. We then extend BDI_{LTL} with the necessary modalities for representing the non-summative group attitudes and the concepts from judgment aggregation. We introduce a judgment aggregation framework using this logic extension and show how it can be used. Sections 3 and 4 respectively study the commitment and reconsideration strategies. Related work, conclusions and outlines for future work are in Section 5.

Cohen and Levesque, in their seminal paper [5], proclaimed that intentions are choice (of a goal) with commitment. Judgment aggregation is a social choice mechanism. Following the intuition of Cohen and Levesque, (a non-summative) group intention is (a group goal determined by) social choice with commitment.

2 Non-summative group attitudes obtained by judgment aggregation

First we discuss how non-summative goals and beliefs are determined and than introduce the logic AGE_{LTL} which is used for representing these attitudes. The formal model of judgment aggregation, using this logic, is given in Section 2.3.

2.1 From individual opinions to group attitudes

Let us consider the example of a robot cleaning crew.

Example 1. Let $C = \{w_1, w_2, w_3\}$ be a crew of cleaning robots. We denote the group goal to clean the meeting room with g_1, and the reasons to adopt this goal with: there are no people in the room (p_1), the room is dirty (p_2), the garbage bin in it is full (p_3). The individual beliefs of the robots on whether g_1 should be the group goal are justified by individual beliefs on p_1, p_2, p_3 using the decision rule $(p_1 \wedge (p_2 \vee p_3)) \leftrightarrow g_1$.

Unlike the joint intention of, for example Dunin-Keplicz and Verbrugge [7], our group intention is not necessarily decomposable to individual intentions: "an adequate account of shared intention is such that it is not necessarily the case that for every shared intention, on that account, there be correlative personal intentions of the individual parties" (pg.172, [11]). Assume that robot w_1 in Example 1 is a mopper, the robot w_2 is a garbage collector and the robot w_3 sprays adequate cleaning chemicals. It can be that the individual goals of w_1 and w_2 are to clean the room. The goal of w_3 may be others, but the group agreed to pursue g_1 and he, being committed to g_1 as part of the group will spray the cleaner as an act towards accomplishing g_1.

We formalize only goals that can be achieved by the group as whole. Whether these goals can be achieved by joint actions or by a combination of individual actions is of no concern. We define a group intention to be the goal that the members agreed on, and hence are committed to, pursuing.

The robots in Example 1 can disagree on various issues when reaching a decision for a group goal. Assume that one robot believes the room is occupied and thus, according to it, the group should not pursue g_1. According to the other two robots, the group should pursue g_1. The second robot is of the opinion that the garbage bin is full and the floor is clean, while the third believes that the floor is dirty. According to the non-summative view of collective beliefs, a group believes p if the group members together agree that as a group they believe p. The question is how should the beliefs of the robots be aggregated to reach an agreement.

The problem of aggregating individual "yes" or "no" opinions, a.k.a. judgments, over a set of logically related issues is studied by judgment aggregation [17]. Judgment aggregation is modeled in general logic and it is an abstract framework that allows for various desirable social properties to be specified.

To use judgment aggregation for aggregating the opinions of the robots, one needs to represent the individual and collective judgments. A logic of belief-desire-intention is insufficient to model these doxastic attitudes. According to Gilbert, "it is not logically sufficient for a group belief that p either that most group members believe that p, or that there be common knowledge within the group that most members believe that p"(pg.189 [8]). Furthermore, "it is not necessary that any members of the group personally believe p (pg.191 [8]). Indeed, a w_1 robots judgment "yes" on $\neg p_1$ is not implied by nor it implies that robot's belief $B_{w_2}\neg p_1$.

Hakli [13] summarizes the difference between beliefs and acceptances as: (1) beliefs are involuntary and acceptances are voluntary; (2) beliefs aim at truth and acceptances depend on goals; (3) beliefs are shaped by evidence and acceptances need not be; (4) beliefs are independent of context and acceptances are context-dependant; and (5) beliefs come in degrees and acceptances are categorical. We find that an individual judgment is closer to an acceptance than to a belief and therefore represent it with an acceptance. There is a debate among social epistemologists on whether collective believes are proper believes or they are in essence acceptances [9, 20, 13]. Since we use acceptances for individual judgments, we deem most adequate to use acceptances to represent the collective judgments as well.

The collective acceptances effectively are the agreed upon group goal and group beliefs. Having group beliefs in support of group goals is in line with Castelfrachi and Paglieri who argue [2] that the goals should be considered together with their supporting "*belief structure*". In Example 1, the decision rule $(p_1 \wedge (p_2 \vee p_3)) \leftrightarrow g_1$ which is nothing else but the "*belief structure*" for g_1. We use the group beliefs to define commitment strategies in Section 3.

2.2 The logic AGE_{LTL}

AGE_{LTL} extends BDI_{LTL} with the necessary modalities for representing individual and collective judgments, group goals and also new information that prompts intention reconsideration.

To model the considered group goals we use a single K modal operator G. Thus Gg, where g is a propositional formula, is to be interpreted as "*g is a group goal*". Since we are interested in modeling the change upon new information, we also need to model these observations of new information. To this end we add the K modal operator E, reading $E\phi$ as "*it is observed that ϕ*".

To model the individual and collective judgments we use the modal operator of acceptance A_S, where S is a subset of some set of agents N. $A_S\phi$ allows us to represent both individual judgments, $S = \{i\}$, for $i \in N$ and collective judgments with $S = N$. Positive and negative introspection holds for A_S: if a group accepts p, then all the members accept that the group accepts p; also, if a group does not accept p, then all the members of the group accept that the group does not accept p. Our operator A_S is inspired by the operator of the acceptance logic of Lorini *et al* [19]. The important difference is that we do not require that the group's acceptance of p entails that all the members accept p, a property of non-summative collective belief indicated by Gilbert in [8]. The opposite, all the agents accepting p implies that the group accepts p, we ensure via the judgment aggregation mechanism.

We use the linear temporal logic to model the change of group attitudes. By using LTL we do not need to distinguish between path formulas and state formulas. Just as in BDI_{LTL}, we are able to quantify over traces, using, for example, $E\Box\neg\phi$ to denote that ϕ is observed to be impossible.

The syntax of AGE_{LTL} is presented in Definition 1. The semantics is as that for BDI_{LTL} proposed by Schield [27].

Definition 1 (Syntax). *Let Agt be a non-empty set of agents, with $S \subseteq Agt$, and L_P be a set of atomic propositions. The admissible formulae of $AGE_{LTL}$are formulae ψ_0, ψ_1 and ψ_2 of languages $\mathcal{L}_{prop}$, $\mathcal{L}_G$ and $\mathcal{L}_{AE_{LTL}}$ correspondingly:*
$\psi_0 ::= p \mid (\psi_0 \wedge \psi_0) \mid \neg\psi_0$
$\psi_1 ::= \psi_0 \mid G\psi_0$
$\psi_2 ::= \psi_0 \mid A_S\psi_1 \mid E\psi_2 \mid \boldsymbol{X}\psi_2 \mid (\psi_2\boldsymbol{U}\psi_2)$
*where p ranges over L_P and S over 2^{Agt}. Moreover, $\Diamond\phi \equiv \top\boldsymbol{U}\phi$, $\Box\phi \equiv \neg\Diamond\neg\phi$, and $\phi\boldsymbol{R}\phi' \equiv \neg(\neg\phi\boldsymbol{U}\neg\phi')$. **X**, **U** and **R** are standard operators of LTL.*

We define the intention of the group of agents S to be their acceptance of a goal, where S ranges over 2^{Agt} as

$$I_S\psi \equiv_{def} A_SG\psi.$$

AGE_{LTL} is a fusion of four modal logics and as such inherits their decidability properties [31]. The four modal logics are: two K-modal logics [4], the logic of acceptance [19] and the linear temporal logic [22].

2.3 The judgment aggregation framework

Our judgment aggregation model in AGE_{LTL} follows the judgment aggregation (JA) model in general logics of Dietrich [6]. For a general overview of JA see [17].

We presume that all the goals which the group considers to adopt are given in a set of candidate group goals $\mathcal{G} = \{Gg \mid g \in \mathcal{L}_{prop}\}$. The decision problem in judgment aggregation, in our case choosing or not a given group goal, is specified by an agenda. An agenda is a pre-defined consistent set of formulas, each representing an issue on which an agent casts his judgments. An agenda is *truth-functional* if it can be partitioned into premises and conclusions. In our case, the agendas consists of one conclusion, which is the group goal $g \in \mathcal{G}$ being considered. The relevant reasons for this group goal are premises.

Thus, an agenda $\mathcal{A} \subseteq \mathcal{L}_G$ is a consistent set such that $\mathcal{A} = \mathcal{A}^p \cup \mathcal{A}^c$, where $\mathcal{A}^p \subseteq \mathcal{L}_{prop}$, $\mathcal{A}^c \subseteq \mathcal{L}_G$ and $\mathcal{A}^p \cap \mathcal{A}^c = \emptyset$.

For a given agenda $\mathcal{A}$, each agent in the group N expresses his judgments by accepting (or not) the agenda issues. Given a set of agent names N and an agenda $\mathcal{A}$, for each issue $a \in \mathcal{A}$ the individual judgment of agent $i \in N$ is one element of the set $\{A_{\{i\}}a, A_{\{i\}}\neg a, \}$. The collective judgment of N is one element of the set $\{A_Na, A_N\neg a, \}$.

The formula $A_{\{i\}}a$ is interpreted as agent i judges a to be true, while the formula $A_{\{i\}}\neg a$ is interpreted as agent i judges a to be false. In theory, the agents can also express the judgment that they do not know how to judge a via the formula $\neg A_{\{i\}}a \wedge \neg A_{\{i\}}\neg a$. In the scenarios we consider, for simplicity, we do not allow the agents to be unopinionated, thus a judgment $\neg A_{\{i\}}a$ is taken to be the same as judgment $A_{\{i\}}\neg a$.

The goal and the reasons are logically related. These relations are represented by the decisions rules. In our model, we assume that the decision rules are a set of formulas $\mathcal{R} \subseteq \mathcal{L}_G$. For each goal $Gg \in \mathcal{G}$ there is, provided together with the agenda, a set of decision rules $\mathcal{R}_g \subseteq \mathcal{R}$. The decision rules contain all the constraints which the agent should observe when casting judgments. These constraints contain three types of information: rules describing how the goal depends on the reasons (justification rules $\mathcal{R}_g^{just}$), rules describing the constraints of the world inhabited by the agents (domain knowledge $\mathcal{R}_g^{DK}$) and rules that describe how g interacts with other candidate goals of the group (coordination rules $\mathcal{R}_g^{coord}$). Hence, the decision rules for a group goal g are $\mathcal{R}_g = \mathcal{R}_g^{just} \cup \mathcal{R}_g^{DK} \cup \mathcal{R}_g^{coord}$.

We want the reasons for a goal to rationalize, not only the choice of a goal, but also its rejection. Having collective justifications for rejecting a goal enables the agents to re-consider adopting a previously rejected group goal. To this end, we require that the justification rules have the schema $Gg \leftrightarrow \Gamma$, where $\{Gg\} = \mathcal{A}_g^c$ and $\Gamma \in \mathcal{L}_{Prop}$ is a formula such that all the non-logical symbols of Γ occur in $\mathcal{A}_g^p$ as well.

The agents express their judgments on the agenda issues, but they accept the decision rules *in toto*.

Example 2 (Example 1 revisited). Consider the cleaning crew from Example 1. $\mathcal{R}_{g_1}^{just}$ is $(p_1 \wedge (p_2 \vee p_3) \leftrightarrow Gg_1$ and $\mathcal{A}_{g_1} = \{p_1, p_2, p_3, Gg_1\}$. Suppose that the crew has the following candidate group goals as well: place the furniture in its designated location (g_2) and collect recyclables from garbage bin (g_3). The agendas are $\mathcal{A}_{g_2} = \{p_4, p_5, p_6, p_7, Gg_2\}$, $\mathcal{A}_{g_3} = \{p_3, p_8, p_9, Gg_3\}$. The justification rules are $\mathcal{R}_{g_2}^{just} \equiv (p_4 \wedge p_5 \wedge (p_6 \vee p_7)) \leftrightarrow Gg_2$ and $\mathcal{R}_{g_3}^{just} \equiv (p_8 \wedge p_9 \wedge p_3) \leftrightarrow Gg_3$. The formulas $p_4 - p_9$ are: the furniture is out of place (p_4), the designated location for the furniture is empty (p_5), the furniture has wheels (p_6) , the furniture has handles (p_7), the agents can get revenue for recyclables (p_8), there is a container for the recyclables (p_9).
An example of a domain knowledge could be $\mathcal{R}_{g_2}^{DK} \equiv \neg p_4 \rightarrow \neg p_5$, since it can not happen that the designated location for the furniture is empty while the furniture is not out of place. Group goal Gg_3 can be pursued at the same time as Gg_1, however, Gg_2 can only be pursued alone. Thus the coordination rule for all three goals is
$\mathcal{R}_{g_1}^{coord} = \mathcal{R}_{g_2}^{coord} = \mathcal{R}_{g_3}^{coord} \equiv ((Gg_2 \wedge \neg(Gg_1 \vee Gg_3)) \vee \neg Gg_2)$.

The sets of individual judgments are only admissible if they satisfy certain conditions. Let $\varphi = \{A_M\overline{a} \mid \overline{a} = a \text{ or } \overline{a} = \neg\, a, a \in \mathcal{A}\}$ be the set of all judgments from agents $M \subseteq N$ for agenda $\mathcal{A}$. We define the set of accepted decision rules $\mathcal{R}_M = \{A_M r \mid r \in \mathcal{R}\}$. The set of judgments φ is admissible if it satisfies the following conditions: for each $a \in \mathcal{A}$, either $A_M a \in \varphi$ or $A_M \neg a \in \varphi$ (completeness), and $\varphi \cup \mathcal{R}_M \not\models \bot$ (consistency).

A profile is a set of every judgment rendered for an agenda $\mathcal{A}$ by an agent in N.

Definition 2. *A profile π is a set $\pi = \{A_{\{i\}}\overline{a} \mid i \in N, \overline{a} = a \text{ or } \in \overline{a} = \neg\, a, a \in \mathcal{A}\}$.*
We define two operators over profiles:
The judgment set for agent i is $\pi \triangleright i = \{\overline{a} \mid A_{\{i\}}\overline{a} \in \pi\}$.
The set of all the agents who accepted $\overline{a}$ is $\pi \triangledown \overline{a} = \{i \mid A_{\{i\}}\overline{a} \in \pi\}$.
A profile is admissible if the judgment set $\pi \triangleright i$ is admissible for every $i \in N$.

We introduce the operators $\triangleright$ and $\triangledown$ to facilitate the explanation of the aggregation properties we present in Section 2.4.

In judgment aggregation, the collective judgment set of a group of agents is obtained by applying a judgment aggregation function to the profile. Judgment aggregation functions are defined in Definition 3.

Definition 3. *Let Π be the set of all profiles π that can be defined for $\mathcal{A}$ and N and let $\overline{\mathcal{A}} = \mathcal{A} \cup \{\neg a \mid a \in \mathcal{A}\}$. A judgment aggregation function f is a mapping $f : \Pi \mapsto 2^{\overline{\mathcal{A}}}$.*

The definition of aggregation function we propose here is identical to that commonly given in the literature [6, 17], where a judgment aggregation function is defined as $F(J_1, J_2, \ldots J_n) = J$, with J_i, $i \in N$ being all the judgment sets of the agents in N and $J \in 2^{\overline{\mathcal{A}}}$. For $J_i = \pi \triangleright i$ it holds $F(J_1, J_2, \ldots J_n) = f(\pi)$.

We define the group attitudes regarding a goal g, *i.e.*, the decision, to be the collective judgment set of the group.

Definition 4. *Given a profile π_g for a considered goal g and a judgment aggregation function f, the group N's decision regarding g is $\mathcal{D}_g = \{A_N a \mid a \in f(\pi)\}$.*

2.4 Desirable judgment aggregation properties for generating group goals

The properties of judgment aggregation (JA) are defined in terms of properties of the judgment aggregation function. Given a JA function f, we describe the most common properties found in the literature.

Anonymity. Given a profile $\pi \in \Pi$, let $\hat{\pi} = \{\pi \triangleright 1, \ldots, \pi \triangleright n\}$, be the multiset of all the individual judgment sets in π. Two profiles $\pi, \pi' \in \Pi$ are permutations of each other if and only if $\hat{\pi} = \hat{\pi}'$. f satisfies anonymity if and only if $f(\pi) = f(\pi')$ for all permutation π and π'.

Unanimity on $a \in \overline{\mathcal{A}}$. f satisfies unanimity on $a \in \overline{\mathcal{A}}$ if and only if for every profile $\pi \in \Pi$ it holds: if for all $i \in N$, $A_i a \in \pi$, then $a \in f(\pi)$.

Collective rationality. f satisfies collective rationality if and only if for all $\pi \in \Pi$, and a given $\mathcal{R}$, $f(\pi) \cup \mathcal{R} \not\models \bot$.

Constant. f is constant when there is $\varphi \in 2^{\overline{A}}$ such that for every $\pi \in \Pi$, $f(\pi) = \varphi$.

Independence. Given $\mathcal{A} = \{a_1, \ldots, a_m\}$ and $\pi \in \Pi$, let $f_1, \ldots, f_m$ be functions defined as $f_j(\pi \triangledown a_j) \in \{a_j, \neg a_j\}$. The JA function f satisfies independence if and only if there exists a set of functions $\{f_1, \ldots, f_m\}$ such that $f(\pi) = \{f_1(\pi \triangledown a_1), \ldots f_m(\pi \triangledown a_m)\}$ for each $\pi \in \Pi$.

The best known example of a judgment aggregation function that satisfies independence is the issue-wise majority function f_{maj}, defined as $f_{maj}(\pi) = \{f_j(\pi \triangledown a_j) \mid a_j \in \mathcal{A}, f_j(\pi \triangledown a_j) = a_j$ if $|\pi \triangledown a_j| \geqslant \lceil \frac{n}{2} \rceil$, otherwise $f_j(\pi \triangledown a_j\} = \neg a_j\}$. This function f_{maj} satisfies *anonymity*, *unanimity*(on all $a \in \mathcal{A}$), *completeness*, and *independence* but it does not satisfy *collective rationality*, as it can be seen on Figure 1.

All judgement aggregation functions that satisfy *anonymity*, *independence* and *collective rationality* are *constant* [17]. Two approaches have been proposed to allow non-constant judgment aggregation: a premise-based, when independence is enforced only on the premises and a conclusion-based approach, when independence is enforced only on the conclusions. In the premise based approach, the judgments on each premise are aggregated individually. The collective judgments on the conclusions are obtained by extending the collective judgments on the premises into a decision-rule consistent judgment set. Similarly, in the conclusion based approach, the judgments on each conclusion are aggregated individually and completeness is obtained by extending the so obtained collective judgments.

The example in Figure 1 illustrates the premise-based and conclusion-baded procedure. As the example here shows, the conclusion-based procedure can produce multiple collective judgment sets. Multiple sets can be obtained via the premise-based procedure as well.

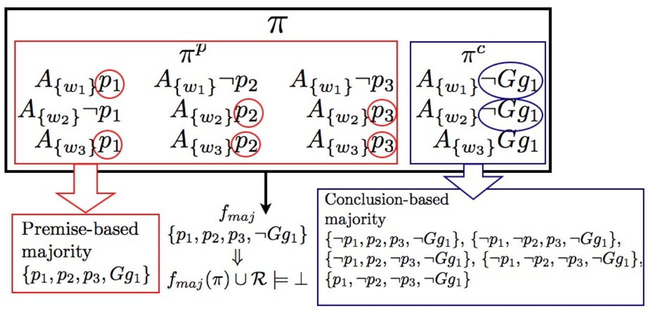

Fig. 1. A profile, issue-wise majority aggregation, premise-based and conclusion-based majority.

The JA function we can use for obtaining group goals should produce decisions that are complete and it should satisfy *collective rationality*. If the decision contains only a group goal acceptance, the group does not know why the goal was (not) adopted and consequently when to revise it. If the aggregation of an admissible profile is not consistent with the decision rules, reasons for the group goal would not be generated.

We have to choose between a conclusion-based and a premise-based procedure. In the profile in Figure 1, a premise-based procedure leads the group to adopt a conclusion that the majority (or even the unanimity) of the group does not endorse. The conclusion is the goal and a premise-based aggregation may establish a group goal which none of the agents is interested in pursuing. To avoid this we need to aggregate using a conclusion-based procedure that satisfies *unanimity on* Gg. Since we have only one goal per agenda, we use issue-based majority to aggregate he group goal judgments.

Our decision rules are of the form $g \leftrightarrow \Gamma$, hence, there exist profiles for which a conclusion-based procedure will not produce complete collective set of judgments. The conclusion-based aggregation can be supplemented with an additional procedure that completes the produced set of judgments when necessary. Such aggregation procedure is the complete conclusion-based procedure (CCBP) developed in [21]. This CCBP is *collectively rational*.

CCBP produces a unique collective judgment for the goal, but it can generate more than one set of justifications for it. This is an undesirable, especially if the agents are in a decentralized communication mode. To deal with ties, as it is the practice with voting procedures, the group determines a default set of justifications for adopting/rejecting each candidate goal, by for example, using a lexicographic order. In the centralized communication mode, the administrator can also break the ties.

The CCBP from [21] also satisfies *anonimity*. Whether this is a desirable property for a group of artificial agents depends entirely on whether the group is democratic or the opinions of some agents are more important. CCBP can be adjusted to take into consideration different weights on different agents' judgment sets.

2.5 Generation of group goals

The mental state of the group is determined by the mental states of the members and the choice of judgment aggregation function. We represent the mental state of the group by a set χ of AGE_{LTL} formulas. The set χ contains the set of all candidate goals for the group $\mathcal{G} \subseteq \mathcal{L}_G/\mathcal{L}_{prop}$ and, for each $Gg \in \mathcal{G}$, the corresponding decision rules $\mathcal{R}_g$, as well as the individual and collective acceptances made in the group regarding agenda $\mathcal{A}_g$. The set χ is common knowledge for the group members. An agent uses χ when it acts as a group member and its own beliefs and goals when it acts as an individual.

To deal with multiple, possibly mutually inconsistent goals, the group has a priority order $\gtrsim_x$ over the group goals $\mathcal{G} \subset \chi$. To avoid overburdening the language with a $\gtrsim_x$ operator, we incorporate the priority order within the decision rules $\mathcal{R}^{just}_{g_i} \equiv \Gamma_i \leftrightarrow Gg_i$. We want the decision rules to capture that if Gg_i is not consistent (according to the coordination rules) with some higher priority goals $Gg_1, \ldots, Gg_m$, then the group can accept Gg_i if and only if none of $Gg_1, \ldots, Gg_m$ is accepted. Hence, we replace the justification rule $\mathcal{R}^{just}_{g_i} \in \chi$ with $\mathcal{R}^{pjust}_{g_i} \equiv (\Gamma_i \wedge \bigwedge_j^m (A_N \neg Gg_j)) \leftrightarrow Gg_i$, where $Gg_j \in \mathcal{G}$, $Gg_j \gtrsim_x Gg_i$ and $Gg_i \wedge Gg_j \wedge \mathcal{R}^{coord}_{g_i} \models \bot$.

Example 3. Consider the goals and rules of the robot crew C from Example 2. Assume the crew has been given the priority order $Gg_1 >_\chi Gg_2 >_\chi Gg_3$. χ contains: $\mathcal{G} = \{Gg_1, Gg_2, Gg_3\}$, one background knowledge rule, one coordination rule, three justification rules, out of which two are new priority modified rules:
$\{\mathcal{G}, \neg p_4 \rightarrow \neg p_5, (Gg_2 \wedge \neg(Gg_1 \vee Gg_3)) \vee \neg Gg_2, Gg_1 \leftrightarrow (p_1 \wedge (p_2 \vee p_3)),$
$Gg_2 \leftrightarrow (p_4 \wedge p_5 \wedge (p_6 \vee p_7) \wedge A_C \neg Gg_1), Gg_3 \leftrightarrow (p_8 \wedge p_9 \wedge p_3 \wedge (A_C \neg Gg_2)\}.$

The agents give their judgments on one agenda after another starting with the agenda for the highest priority candidate goal. Once the profile π and the decision $\mathcal{D}_g$ for a goal g are obtained, they are added to χ. To avoid the situation in which the group casts judgments on an issue that has already been decided, we need to remove decided issues from $\mathcal{A}_g$ before eliciting the profile for this agenda.

The group goals are generated by executing **GenerateGoals**(χ, N).

`function` **GenerateGoals**`(`χ`, ` S`):`
`for each` $Gg_i \in \mathcal{G}$ s.t. $[\forall Gg_j \in \mathcal{G}\colon (Gg_j \gtrsim Gg_i) \Rightarrow (A_N Gg_j \in \chi$ or $A_N \neg Gg_j \in \chi)]$
  $\{\ B := (\{a \mid A_N a \in \chi\} \cup \{\neg a \mid A_N \neg a \in \chi\}) \cap \mathcal{A}_{g_i};$
   $\mathcal{A}^*_{g_i} := \mathcal{A}_{g_i}/B;$
   $\pi_{g_i} := elicit(S, \mathcal{A}^*_{g_i}, \chi);$
   $\chi := \chi \cup \pi_{g_i} \cup f^a(\pi_{g_i}); \quad \}$
`return` χ.

$GenerateGoals$ does not violate the candidate goal preference order and it terminates if $elicit$ terminates. $elicit$ requests the agents to submit complete judgment sets for $\pi_{g_i} \subset \chi$. We require that $elicit$ is such that for all returned π it holds: $\chi \cup f(\pi)) \not\models \bot$

and $\chi \cup \pi \triangleright i \not\models \perp$ for every $i \in N$. When a higher priority goal Gg_i is accepted by the group, a lower priority incompatible goal Gg_j cannot be adopted regardless of the judgments on the issues in $\mathcal{A}_{g_j}$. Nevertheless, although $elicit$ will provide individual judgments for the agenda $\mathcal{A}_{g_j}$. If the acceptance of Gg_i is reconsidered, we can obtain a new decision on Gg_j because the profile for Gg_j is available.

3 Commitment strategies

The group can choose to reconsider the group goal in presence of new information – "a joint commitment must be *terminated* jointly" (pg. 143, [10]). Whether the group chooses to reconsider depends on how committed it is to the group intention corresponding to that goal. We defined the group intention to be $I_N g \equiv A_N Gg$, *i.e.* the decision to accept g as the group goal. The level of persistence of a group in their collective decision depends on the choice of commitment strategy.

These are the three main commitment strategies (introduced by Rao and Georgeff [24]):
Blind commitment: $I_i g \to (I_i g \mathbf{U} B_i g)$
Single-minded commitment: $I_i g \to (I_i g \mathbf{U}(B_i g \vee B_i \Box \neg g))$
Open-minded commitment: $I_i g \to (I_i g \mathbf{U}(B_i g \vee \neg G_i g))$

These commitment strategies only consider the relation between the intention and the beliefs regarding g and Gg. In our model of group intentions, a commitment is to a goal acceptance. This enables intention reconsideration upon new information on either one of the agenda issues in $\mathcal{A}_g$, as well as on a higher priority goal.

The strength of our framework is exhibited in its ability to describe the groups' commitment not only to its decision to adopt a goal, but also to its decision to reject a goal. Namely, if the agents decided $I_N g_i$ and $A_N \neg Gg_j$, they are committed to both $I_N g_i$ and $A_N \neg Gg_j$. Commitment to reject g allows for g to be reconsidered and eventually adopted if the state of the world changes.

Let N be a set of agents with a set of candidate goals $\mathcal{G}$. Let $Gg_i, Gg_j \in \mathcal{G}$ have agendas $\mathcal{A}_{g_i}, \mathcal{A}_{g_j}$. We use $p \in \overline{\mathcal{A}}^p_{g_i}$ and $q_i \in \overline{\mathcal{A}}^c_{g_i}$, $q_j \in \overline{\mathcal{A}}^c_{g_j}$. The profiles and decisions are π_{g_i} and $f(\pi_{g_i})$; $Gg_j > Gg_i$, and Gg_j cannot be pursued at the same time as Gg_i.

We use the formulas $(\alpha_1)-(\alpha_5)$ to refine the blind, single-minded and open-minded commitment. Instead of the *until*, we use the temporal operator *release*: $\psi \, \mathbf{R} \, \phi \equiv \neg(\neg\psi \, \mathbf{U} \, \neg\phi)$, meaning that ϕ has to be true until and including the point where ψ first becomes true; if ψ never becomes true, ϕ must remain true forever. Unlike the *until* operator, the *release* operator does not guarantee that the right hand-side formula will ever become true, which in our case translates to the fact that an agent could be forever committed to a goal.

(α_1) $Eg_i \, \mathbf{R} \, I_N g_i$
(α_2) $\perp \mathbf{R} \, A_N \neg Gg_i$
(α_3) $(E\Box\neg g_i \vee Eg_i) \, \mathbf{R} \, A_N q_i$
(α_4) $A_N \neg q_j \, \mathbf{R} \, A_N q_i$
(α_5) $A_N p \to (E\neg p \, \mathbf{R} \, A_N q_i)$

Blind commitment: $\alpha_1 \wedge \alpha_2$.
Only the observation that the goal is achieved (Eg_i) can release the intention to achieve the goal $I_N g_i$. If the goal is never achieved, the group will always be committed to it. If a goal is not accepted, then the agents will not reconsider accepting it.

Single-minded commitment: α_3.
Only new information on the goal (either that the goal is achieved or had become impossible) can release the decision of the group to adopt /reject the goal. Hence, new information is only regarded if it concerns the conclusion, while information on the remaining agenda items is ignored.

Extended single-minded commitment: $\alpha_3 \wedge \alpha_4$.
Not only new information on Gg_i, but also the collective acceptance to adopt a more important incompatible goal Gg_j can release the intention of the group to achieve Gg_i. Similarly, if Gg_i is not accepted, the non-acceptance can be revised, not only if Gg_j is observed to be impossible or achieved, but also when the commitment to pursue Gg_j is dropped (for whatever reason).

Open-minded commitment: $\alpha_3 \wedge \alpha_5$.
A group will maintain its collective acceptances to adopt or reject a goal as long as the new information regarding all collectively accepted agenda items is consistent with $f(\pi_{g_i})$.

Extended open-minded commitment: $\alpha_3 \wedge \alpha_4 \wedge \alpha_5$.
Extending on the single-minded commitment, a change in intention to pursue a higher priority goal Gg_j can also release the acceptance of the group on Gg_i.

Once an intention is dropped, a group may need to reconsider its collective acceptances. This may cause for the dropped goal to be re-affirmed, but a reconsideration process will be invoked nevertheless.

4 Reconsideration of group attitudes

In Section 2.5 we defined the mental state of the group χ. We can now define what it means for a group to be *coherent*.

Definition 5 (Group coherence). *Given a Kripke structure $\mathcal{M}$ and situations $s \in W$, a group of N agents is coherent if the following conditions are met:*
(ρ_1): $\mathcal{M} \models \neg(A_S a \wedge A_S \neg a)$ for any $S \subseteq N$ and any $a \in \mathcal{A}_g$.
(ρ_2): If $\mathcal{M}, s \models \chi$ then $\chi \not\models \bot$.
(ρ_3): $\mathcal{M}, s \models \bigwedge \mathcal{G} \rightarrow \neg\Box\neg g$ for all $Gg \in \mathcal{G}$.
(ρ_4): Let $Gg \in \mathcal{G}$ and $\mathcal{G}' = \mathcal{G}/\{Gg\}$, then $\mathcal{M} \models (\bigwedge \mathcal{G} \wedge E\Box\neg g) \rightarrow \boldsymbol{X}(\neg Gg)$.
(ρ_5): Let $p \in \mathcal{A}_g^p$ and $q \in \{Gg, \neg Gg\}$. $Ep \wedge (Ep\, \boldsymbol{R}\, A_N q) \rightarrow \boldsymbol{X} A_N p$

The first condition ensures that no contradictory judgments are given. The second condition ensures that the mental state of the group is logically consistent in all situations. The third and fourth conditions ensure that impossible goals cannot be part of the set of candidate goals and if g is observed to be impossible in situation s, then it will be removed from $\mathcal{G}$ in the next situation. ρ_5 enforces the acceptance of the new information on the group level, when the commitment strategy so allows – after a is observed and that led the group to de-commit from g, the group necessarily accepts a.

A coherent group accepts the observed new information on a premise. This may cause the collective acceptances to be inconsistent with the justification rules. Consequently, the decisions and/or the profiles in χ need to be changed in to ensure that ρ_1 and ρ_2 are satisfied. If, however $\Box\neg g$ or g is observed, the group reconsiders χ by removing Gg from $\mathcal{G}$. In this case, the decisions and profiles are not changed.

For simplicity, at present we work with a world in which the agents' knowledge can only increase, namely the observed information is not a fluent. A few more conditions need to be added to the definition of group coherence, for our model to be able to be applicable to fluents. E.g., we need to define which observation is accepted when two subsequent contradictory observations happen.

4.1 Reconsideration strategies

For the group to be coherent at all situations, the acceptances regarding the group goals need to be reconsidered after de-commitment. Let $\mathcal{D}_g \subset \chi$ contain the group acceptances for a goal g, while $\pi_g \subset \chi$ contain the profile for g. There are two basic ways in which a collective judgment set can be reconsidered. The first way is to elicit a new profile for g and apply judgment aggregation to it to obtain the reconsidered $\mathcal{D}^*_g$. The second is to reconsider only $\mathcal{D}_g$ without re-eliciting individual judgments. The first approach requires communication among agents. The second approach can be done by each agent reconsidering χ by herself. We identify three reconsideration strategies available to the agents. The strategies are ordered from the least to the most demanding in terms of agent communication.

Decision reconsideration ($\mathcal{D}$-r). Assume that Ea, $a \in \overline{\mathcal{A}}^p_g$, $q \in \{Gg, \neg Gg\}$ and the group de-commited from $A_N q$. The reconsidered decision $\mathcal{D}^*_g$ is such that a is accepted, i.e., $A_N a \in \mathcal{D}^*_g$, and the entire decision is consistent with the justification rules, namely $\mathcal{R}^{pjust}_g \cup \mathcal{D}^*_g \not\models \bot$. If the $\mathcal{D}$-r specifies an unique $\mathcal{D}^*_g$, for any observed information and any $\mathcal{D}_g$, then χ can be reconsidered without any communication among the agents. Given the form of $\mathcal{R}^{pjust}_g$ (see Section 2.5), this will always be the case.

However, $\mathcal{D}$-r is not always an option when the de-commitment occurred due to a change in collective acceptance of a higher priority goal g'. Let $q' \in \{Gg', \neg Gg'\}$. Let the new acceptance be $A_N \neg q'$. $\mathcal{D}$-r is possible if and only if $\mathcal{D}^*_g = \mathcal{D}_g$ and $\mathcal{R}^{pjust}_g \cup \mathcal{D}_g \cup \{A_N \neg q'\} \not\models \bot$. Recall that $A_N q'$ was not in $\mathcal{A}_g$ and as such the acceptance of q' or $\neg q'$ is never in the decision for π_g.

Partial reconsideration of the profile (Partial π-r). Assume that Ea, $a \in \overline{\mathcal{A}}_g$, $Gg \in \mathcal{G}$. Not only the group, but also the individual agents need to accept a. The *Partial π-r* asks for new individual judgments be elicited. This is done to ensure the logical consistency of the individual judgment sets with the observations. New judgments are only elicited from the agents i which $A_{\{i\}} \neg a$.

Let $W \subseteq N$ be the subset of agents i s.t. $A_{\{i\}} \neg a \in \chi$. Agents i are s.t. $A_{\{i\}} a \in \chi$ when the observation is $E \neg a$. Let $\pi^W_g \subseteq \pi_g$ be the set of all acceptances made by agents in W. We construct $\chi' = \chi / \pi^W_g$. The new profile and decision are obtained by executing $GenerateGoals\,(\chi', W)$.

Example 4. Consider Example 2. Assume that $\mathcal{D}_{g_1} = \{A_C p_1, A_C \neg p_2, A_C p_3, A_C Gg_1\}$, $\mathcal{D}_{g_2} = \{A_C p_4, A_C p_5, A_C p_6, A_C p_7, A_C \neg Gg_2\}$ and $\mathcal{D}_{g_3} = \{A_C p_8, A_C p_9, A_C Gg_3\}$ are the group's decisions. Assume the group de-commits on Gg_1because of $E \neg p_2$. If the group is committed to Gg_3, the commitment on Gg_3 will not allow for $A_N p_3$ to be modified when reconsidering Gg_1. Since $A_N p_3$ exists in χ', p_3 will be excluded from the (new) agenda for g_1, although it was originally in it. $elicit$ calls only on the agents in W to complete $\pi_{g_1} \in \chi'$ with their judgment sets.

Full profile reconsideration (π-r). The full profile reconsideration is the same with the partial reconsiderations in all respects except one – now $W = N$. Namely, within the full profile revision strategy, each agent is asked to revise his judgment set by accepting the new information, regardless of whether he had already accepted it.

4.2 Combining revision and commitment strategies

Unlike the Rao and Georgeff commitment strategies [24], in our framework the commitment strategies are not axioms of the logic. We require that the commitment strategy is valid in all the models of the group and not in all the models of AGE_{LTL}. This allows the group to define different commitment strategies and different revision strategies for different goals. It might even choose to revise differently depending on which information triggered the revision. Choosing different revision strategies for each goal, or each type of new information, should not undermine the coherence of the group record χ. The conditions of group coherence of the group ensures that after every reconsideration χ must remain consistent. However, some combinations of commitment strategies can lead to incoherence of χ.

Example 5. Consider the decisions in Example 4 . Assume that initially the group chose open-minded commitment for $I_C g_1$ and blind commitment for $I_C g_3$, with goal open-minded commitment for $A_C \neg Gg_2$. If Eg_1 and thus $I_C g_1$ is dropped, then the extended open-minded commitment would allow $A_C \neg Gg_2$ to be reconsidered and eventually $I_C g_2$ established. However, since the group is blindly committed to $I_C g_3$, this change will not cause reconsideration and as a result both $I_C g_2$ and $I_C g_3$ will be in χ thus making χ incoherent.

Problems arise when $sub(\mathcal{R}^{pjust}_{g_i}) \cap sub(\mathcal{R}^{pjust}_{g_j}) \neq \emptyset$, where $sub(\mathcal{R}^{pjust}_{g})$ denotes the set of atomic sub-formulas of some goal g and $Gg_i, Gg_j \in \mathcal{G}$. Proposition 1 summarizes under which conditions these problems are avoided.

Proposition 1. *Let α' and α'' be the commitment strategies selected for g_i and g_j correspondingly. $\chi \cup \alpha' \cup \alpha'' \not\models \bot$ (in all situations):*
a) if $\phi \in sub(\mathcal{R}^{pjust}_{g_i}) \cap sub(\mathcal{R}^{pjust}_{g_j})$ and $p \in \mathcal{A}_{g_i} \cap \mathcal{A}_{g_j}$, then α_5 is either in both α' and α'' or in none;
b) if Gg_i is more important than Gg_j while Gg_j and Gg_i cannot be accepted at the same time, then $\alpha_4 \in \alpha''$.

Proof. The proof is straightforward. Namely, if the change in the group (non)acceptance of Gg_i causes the $A_N Gg_j$ to induce group incoherence, we are able to de-comit from $A_N Gg_j$. If we were not able to de-comit from $A_N Gg_j$ then group coherence is blocked. If the change in the group (non)acceptance of Gg_i is caused by an observation on a premise $p \in \mathcal{A}_{g_i} \cap \mathcal{A}_{g_j}$ then condition a) ensures that the commitment to $A_N Gg_j$ does not block group coherence. If the change on $A_N Gg_j$ is caused by a change in commitment to a higher priority goal the condition b) ensures that a commitment regarding Gg_j does not block group coherence. Condition b) allows only "goal sensitive" commitments to be selected for lower level goals.

5 Conclusions

We present a formalization of non-summative beliefs, goals and intentions in AGE_{LTL} and show how they can be generated using judgment aggregation. Our multi-agent

AGE_{LTL} logic extends BDI_{LTL}. In accordance with the non-summative view, having a group intention $I_N g$ in our framework does not imply $I_{\{i\}} g$ for each the member i. We extend the commitment strategies of Rao and Georgeff [24] to increase the reactivity of the group to new information. The commitment strategies are not axioms of the representation logic; instead they are a property of a group. We show how the group can combine different levels of commitment to different goals.

Our framework is intended for groups that engage in joint activity and it is applicable when it cannot be assumed that the agents persuade each other on a single position and goal, but it is necessary anyway that the group presents itself as a single whole from the point of view of beliefs and goals. The requirement that the group presents itself as a rational entity that has goals justified by the beliefs it holds, and is able to revise these goals under the light of new information, was held by Tuomela [28] and adopted in agent theory by Boella and van der Torre [1] and Lorini [18]. The proposal of the paper can be applied, for example, in an opensource project, where several people have to discuss online to agree on which is their position on issues and which is their goal.

We assume that the group has an order of importance for its candidate goals. Alternatively, the group can also agree on this order by expressing individual preferences. Uckelman and Endriss [29] show how individual (cardinal) preferences over goals can be aggregated. In [30] the reconsideration of individual intentions and associated plans is considered. Intentions and their role in deliberation for individual agents have been studied in a game theoretic framework by Roy [25, 26]. Icard *et al.* [14] consider the joint revision of individual attitudes, with the revision of beliefs triggering intention revision.

In our framework, the entire group observes the new information. In the future we intend to explore the case when only some members of the group observe the new information. Given that the group attitudes are established by an aggregation procedure that is, as almost all but the most trivial procedures, manipulable, we intend to explore whether an agent can have the incentive to behave strategically in rendering judgments.

References

1. G. Boella and L. van der Torre. The ontological properties of social roles in multi-agent systems: Definitional dependence, powers and roles playing roles. *Artificial Intelligence and Law Journal (AILaw)*, 15(3):201–221, 2007.
2. C. Castelfranchi and F. Paglieri. The role of beliefs in goal dynamics: Prolegomena to a constructive theory of intentions. *Synthese*, 155:237–263, 2007.
3. B. Chapman. Rational aggregation. *Politics, Philosophy and Economics*, 1(3):337–354, 2002.
4. B. F. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, Cambridge, 1980.
5. P. R. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(2-3):213–261, 1990.
6. F. Dietrich and C. List. Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33, July 2007.
7. B. Dunin-Keplicz and R. Verbrugge. *Teamwork in Multi-Agent Systems: A Formal Approach.* Wiley and Sons, July 2010.
8. M. P. Gilbert. Modeling Collective Belief. *Synthese*, 73:185–204, 1987.

9. M. P. Gilbert. Belief and acceptance as features of groups. *Protosociology: An International Journal of Interdisciplinary Research*, 16:35–69, 2002.
10. M. P. Gilbert. Acting together, joint commitment, and obligation. *Philosophische Analyse/Philosophical Analysis*, 2007.
11. M. P. Gilbert. Shared Intention and Personal Intentions. *Philosophical Studies*, 2009.
12. B. Grosz and L. Hunsberger. The dynamics of intention in collaborative activity. *Cognitive Systems Research*, 7(2-3):259–272, 2007.
13. R. Hakli. Group beliefs and the distinction between belief and acceptance. *Cognitive Systems Research*, 7(2-3):286 – 297, 2006. Cognition, Joint Action and Collective Intentionality.
14. T. Icard, E. Pacuit, and Y. Shoham. Joint revision of belief and intention. In *Proc.of the 12th international conference on principles of knowledge representation and reasoning (KR)*, 2010.
15. N. R. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artif. Intell.*, 75(2):195–240, 1995.
16. S. Konieczny and R. Pino-Pérez. Merging with integrity constraints. *Lecture Notes in Computer Science*, 1638/1999:233–244, January 1999.
17. C. List and C. Puppe. Judgment aggregation: A survey. In P. Anand, C. Puppe, and P. Pattanaik, editors, *Oxford Handbook of Rational and Social Choice*. Oxford, 2009.
18. E. Lorini and D. Longin. A logical account of institutions: From acceptances to norms via legislators. In *KR*, pages 38–48, 2008.
19. E. Lorini, D. Longin, B. Gaudou, and A. Herzig. The logic of acceptance. *Journal of Logic and Computation*, 19(6):901–940, 2009.
20. A. Meijers. Collective agents and cognitive agents. *Protosociology. Special Issue Understanding the Social:New Perspectives from Epistemology*, 16:70–86, 2002.
21. G. Pigozzi, M. Slavkovik, and L. van der Torre. A complete conclusion-based procedure for judgment aggregation. In *First International Conference on Algorithmic Decision Theory Proceedings*, pages 1–13, 2009.
22. A. Pnueli. The temporal logic of programs. In *SFCS '77: Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 46–57, Washington, DC, USA, 1977. IEEE Computer Society.
23. A. Quinton. The presidential address: Social objects. *Proceedings of the Aristotelian Society*, 76:pp. 1–27+viii, 1975.
24. A. S. Rao and M. P. Georgeff. Intentions and rational commitment. In *In Proceedings of the First Pacific Rim Conference on Artificial Intelligence (PRICAI-90)*, pages 94–99, 1993.
25. O. Roy. A dynamic-epistemic hybrid logic for intentions and information changes in strategic games. *Synthese*, 171:291–320, 2009.
26. O. Roy. Intentions and interactive transformations of decision problems. *Synthese*, 169:335–349, 2009.
27. K. Schild. On the relationship between bdi logics and standard logics of concurrency. *Autonomous Agents and Multi-Agent Systems*, 3(3):259–283, 2000.
28. R. Tuomela and K. Miller. Groups beliefs. *Synthese*, 91:285–318, 1992.
29. J. Uckelman and U. Endriss. Compactly representing utility functions using weighted goals and the max aggregator. *Artif. Intell.*, 174:1222–1246, October 2010.
30. W. van der Hoek, W. Jamroga, and M. Wooldridge. Towards a theory of intention revision. *Synthese*, 155(2):265–290, 2007.
31. F. Wolter. Fusions of modal logics revisited. In M. Kracht, M. de Rijke, and M. Zakharyaschev, editors, *Advances in Modal Logic 96*, pages 361–379. CSLI Lecture Notes, 1998.
32. M. Wooldridge and N. Jennings. The cooperative problem-solving process. *Journal of Logic and Computation*, 9(4), 1999.