

A model to support collective reasoning: Formalization, analysis and computational assessment

Jordi Ganzer^a, Natalia Criado^a, Maite Lopez-Sanchez^b, Simon Parsons^{a,d},
Juan A. Rodriguez-Aguilar^c

^a*Department of Informatics, King's College London*

^b*Faculty of Mathematics and Computer Science, University of Barcelona*

^c*Artificial Intelligence Research Institute (IIIA-CSIC)*

^d*School of Computer Science, University of Lincoln*

Abstract

Inspired by e-participation systems, in this paper we propose a new model to represent human debates and methods to obtain collective conclusions from them. This model overcomes drawbacks of existing approaches by allowing users to introduce new pieces of information into the discussion, to relate them to existing pieces, and also to express their opinion on the pieces proposed by other users. In addition, our model does not assume that users' opinions are rational in order to extract information from it, an assumption that significantly limits current approaches. Instead, we define a weaker notion of rationality that characterises coherent opinions, and we consider different scenarios based on the coherence of individual opinions and the level of consensus that users have on the debate structure. Considering these two factors, we analyse the outcomes of different opinion aggregation functions that compute a collective decision based on the individual opinions and the debate structure. In particular, we demonstrate that aggregated opinions can be coherent even if there is a lack of consensus and individual opinions are not coherent. We conclude our analysis with a computational evaluation demonstrating that collective opinions can be computed efficiently for real-sized debates.

Email addresses: jordi.ganzer_ripoll@kcl.ac.uk (Jordi Ganzer), natalia.criado_pacheco@kcl.ac.uk (Natalia Criado), maite_lopez@ub.edu (Maite Lopez-Sanchez), sparsons@lincoln.ac.uk (Simon Parsons), jar@iia.csic.es (Juan A. Rodriguez-Aguilar)

Keywords: Social Choice, Collective decision, Argumentation

1. Introduction

We live in challenging times. With the rise of artificial intelligence and the looming climate emergency, it is no exaggeration to say that humanity is facing several simultaneous existential crises. At the same time, traditional democracy is being undermined by falling participation [48, 91]. In response there have been calls [1, 39, 83] for new ways to tighten the relationship between political institutions and the citizens they represent. Social networks and the palpable impact they had in society have inspired new technologies aiming to improve this relationship — e-governance, or more specifically e-participation systems [42]. These are systems that are designed to allow citizens to propose, discuss, and even decide, government policy through online platforms.

Significant examples of deployed e-participation systems are those used by the city governments of Barcelona [30], Reykjavik [16], Madrid [29], Helsinki [25], and the one used by the French government [69]. In these systems, participants can carry out structured discussions around some issue. For example the systems deployed by Barcelona, Reykjavik, Madrid, and Helsinki are for proposals about local issues, while the French government system, *Parlement et Citoyens*, is for the discussion of potential national legislation. Turning to systems that are not tied to a specific institution, there are tools such as *consider.it* [26], *Appgree* [4] and *Baoqu* [9] whose main focus is scalability — making the systems fit for use by large numbers of participants.

Several methodologies have been already explored as a means to tackle this problem. *Social choice theory* [8, 57] is an approach for establishing the collective opinion of a group facing a choice between many alternatives. Given a set of alternatives and a set of agents who possess preference relations over the alternatives, social choice theory focuses on how to yield a collective choice that appropriately reflects the agents' individual preferences. A related approach, but one which focuses on the acceptability of a single issue, is *judgement aggregation*. This tackles the problem of whether to collectively accept a single issue once the participants have put forward their opinion on it [38, 58]. A further approach, one that focuses more on resolving conflicts in opinions, is *computational argumentation* [76]. Given a set of arguments for particular options, and a set of attack relations (conflicts) between the arguments, argumentation is concerned with identifying those arguments that

might be accepted by a rational agent (for different ideas of what makes an argument acceptable). Mixing judgement aggregation or social choice theory together with argumentation we find several proposals that structure debates using arguments and attack relationships, such as [7, 56], or attack and defence (or support) relationships [46]. These allow participants to put forward arguments, relations between arguments, and opinions about which of these arguments and relations hold. They then produce an output that is intended to reflect the collective opinion of the participants on the the debate.

This article proposes a new formal model, which we call the “relational model”, to structure the information in a debate while allowing more expressiveness than existing approaches permit, both in terms of the structure of the debate, and in terms of the opinions expressed by participants. We propose several methods for aggregating the information of the debate captured by the model to obtain an output reflecting the collective view of the participants involved. Furthermore, we analyse the performance of such methods with respect to several social choice properties adapted from those proposed in the social choice literature. Finally, we study the computational properties of the aggregation methods to establish whether it is feasible to use them in practice.

In more detail the contributions of this work are:

- *Increasing expressiveness.* Collective decision-making frameworks based on traditional argumentation models usually take as starting point a fixed argumentation structure that only model attack relationships between arguments or a combination of attack and support/defence relationships between arguments. These frameworks then allow participants to express opinions about the different arguments included in the debate [7, 46, 56]. The fixed nature of the argumentation structure, even if it is defined by the participants, represents a significant drawback for e-participation systems. Adopting a fixed argumentation structure using only attacks, or attacks plus supports/defences, reduces what participants are allowed to express. For instance, a fixed attack relationship between two arguments might be problematic for some participants who disagree on the classification of the relationship as an attack and would have defined the same relationship as a defence. To solve this problem, our model uses relationships that are not subjectively classified as attacks or defences, but only represent the connections between elements of the debate. The subjective classifica-

tion is applied individually by each participant not in terms of attack or defence, but in terms of acceptability of the connections. Thus, the structure of the debate is focused on organising relevant information, not on expressing the subjective opinion of the participants, which is added separately using other tools.

- *Going beyond abstract argumentation.* Several approaches [7, 28, 56] make use of abstract argumentation frameworks [35], or some variations of them, as in [46], to represent the elements of a debate. In such frameworks, whole arguments are the atomic elements. In work on argumentation, this limitation has led to work on “structured” [65] or “rule-based” [47] argumentation which constructs arguments out of lower level components like facts and rules. Since we believe that a debate can hinge on being able to address such lower level components, we take a similar, but more general, approach. We construct debates from two types of abstract objects: statements, which represent sentences without reasoning, and the relationships between statements, which represent the existing reasoning connecting the statements¹. We make a sharp distinction between these two types of information, statements and reasoning connecting statements, in order to allow them to be subsequently evaluated in different manners by the participants.
- *Compound and real valued opinions.* Previous work on argumentation-based approaches has only allowed participants in a debate to express opinions about either the arguments [46, 56], or about the relationships between arguments [37]². Here we allow participants to provide opinions on both statements and the relationships between them. Opinions about relationships capture participants’ acceptance, or otherwise, of the reasoning that the relationship represents; and opinions about statements reflect participants’ satisfaction with the statement itself.

Furthermore, we follow [37, 56] in allowing opinion about relationships

¹We could relate the statements with axioms, premises or conclusions, and the relationships by the rules or demonstration steps that lead from premises to conclusions, just as in structured argumentation.

²[37] is not about combining collective opinions on relationships between arguments, but it provides the groundwork for such a system by studying argumentation where the relationships between arguments have different weights.

and statements to be expressed using real values rather than the discrete values of [7, 46]. This feature allows the participants to express their opinions in a wider range of values, making the approach more flexible. (It should be noted though that most existing e-participation systems just allow users to express agreement or disagreement.)

- *A more flexible notion of coherence.* Previous work on determining collective opinions makes use of a notion of “rationality” in which an opinion is either determined to be acceptable or not acceptable (where “acceptable” has different interpretations but reflects the constraints on distributions of opinions across statements) [7, 46, 75]. We think that this is somewhat limiting. Since the opinions originate with human participants, and humans are not always consistent in their views, we feel that insisting on this rigid form can lead to lose valuable information. Hence, we propose a less restrictive notion of rationality, which we call “coherence”, to assess the degree to which an opinion is coherent, be it from an agent or from the collective aggregation. To take a simple example, an agent valuing with opposite values two related statements, one supporting the other, would be understood as not coherent, while another agent valuing the same statements with similar values would be classified as coherent. Note that we do not compare different agents’ opinions in establishing coherence, but look within the same opinion.
- *Aggregation functions exploiting dependencies.* We propose several opinion aggregation functions that use the participant’s opinion on a debate to compute a collective opinion. These proposed functions assume different forms of using the dependencies between opinions to combine them. We provide two families of function that include a function that does not use dependencies at all, and two other functions that use dependencies rather differently. These families of functions collective span the ways in which the dependencies can be taken account of, thus making it possible to choose a specific degree of use of the dependencies.
- *Formal analysis.* We assess these families of functions against a wide-ranging set of properties designed to provide a detailed characterization of their behaviour. We use several properties adapted from the social choice literature [58] to fit our model in order to assess the performance

of the functions. We carry out the same study in four scenarios that consider different assumptions for the participants’ opinions.

- *Computational analysis.* We follow the formal analysis with a computational analysis. This first computes the computational complexity of the aggregation functions, and then provides an empirical analysis of those functions when computing collective opinions for a range of scenarios that are larger (in terms of statements and number of opinions) than any online debates that we are aware of. This analysis shows that collective opinions can be computed in real time on quite modest hardware.

These contributions can be found in the following sections: section 2 provides an introduction to the relational model; and sections 3 and section 4 provide a formal definition of the model. Then, section 5 defines the problem of computing collective coherence and introduces the properties that will be used to assess aggregation functions, while section 6 defines a family of aggregation functions and section 7 uses the properties to analyse the functions in different scenarios (the proofs can be found in Appendix A). Section 8 provides the computational assessment of the model; section 9 relates the work presented in this paper to other relevant work in the literature; and section 10 summarises our conclusions.

2. Foundations of the relational model

This section introduces the main features of the relational model while the following sections provide a formalisation of the model. We start this section by describing the elements of the model and the rationale behind it, as well as the aggregation of opinions using the information enclosed within the model. We also include a discussion of how the model could be used in a participatory system.

2.1. Relational Model

The *Relational Model*, RM for short, is a model designed to represent a collective debate where participants discuss a proposal by putting forward additional information relevant for the discussion and giving their opinions about it. The RM is composed of two main parts: the structural part, organising the information (or *content*) of a debate; and the interpretative part, representing the participants’ *opinions* in a debate.

Content of a debate. The RM has two main abstract elements that capture the structure of a debate: *statements* and the *relationships* among them. *Statements* represent plain sentences that describe facts such as, for example, “Increase of house prices in the neighbourhood”. Participants will afterwards be able to express their opinion about the desirability or undesirability of the sentence. *Relationships* represent the reasoning that connects statements. Each relationship connects a set of source statements to some destination statement. We can think of relationships as logical inferences that may or may not be accepted by the participants of the debate. Such acceptance will also be included in participants’ opinions.

Commonly, debates discuss a particular subject or proposal, and may even consider a set of proposals. In the RM we consider the *target* (of a debate) as a set of statements that represents the proposal. The distinguishing feature of these statements, the target of the debate, is that none of them can be a destination statement in any relationship because they are the initiators of the debate. Thus, the target acts as root of the structure that captures the debate. This structure composed of the statements, relationships and target is called a *Directed Relational Framework*, or DRF for short.

Participants’ opinions. The opinions of the participants are encoded in the form of functions that assign values to the objects that make up the structure of the debate. The *valuation* function assigns values to statements and the *acceptance* function assigns values to the relationships. The opinion of a participant is twofold in order to obtain the two types of subjectivity involved in the debate. The first, relating to the valuation function, provides a participant with a way to express their judgement about the facts or possibilities that arise in the debate (but does not express their view about the truth of the sentence). The second, relating to the acceptance function, allows a participant to express the truth they see in each relationship that is included in the debate. Thus, the desirability or undesirability that each participant feels about each statement of the debate is represented by a positive or negative value assigned with the valuation function, and the conformity that each agent relates to the connections between the statements is represented by an acceptance value assigned by the acceptance function.

Example 2.1. Consider a debate about building a sports centre. Let’s assume a scenario where a proposal τ “Construction of a sports centre in particular location in the neighbour” is discussed. In this setting, we may

consider the reasoning that “The construction of the sport centre will imply the demolition of existing buildings that give historical relevance to the neighbourhood” is represented by the relationship r_1 that connects the target τ to a new statement s_1 “Destruction of the character of the neighbourhood”, as depicted in Figure 1.

Given these elements, participants in the debate can express their opinions by providing a value to both τ and s_1 and by assigning a degree of acceptance to the relationship r_1 . The values assigned to τ or s_1 will represent how they feel about the respective statements and the degree given to r_1 will indicate how much the participant believes in the truth of the reasoning relating both statements. Thus, someone hoping for a new sports center in the neighbour — hence valuing the target τ positively — may not want the neighbourhood to be damaged — so valuing s_1 negatively — but still believes that the construction of the sport center indeed will imply the demolition of the existing buildings that now are historically relevant (and hence would create that damage) —thus accepting the reasoning represented by r_1 .




Figure 1: Graphical representation of the relationship between proposal τ (building a sports centre) and statement s_1 (destroying the neighbourhood’s character)

When considering this example it is worth elaborating on the direction of the relationship regarding content and opinion. Observe that relationship r_1 is directed from τ towards statement s_1 so that s_1 can be interpreted as its consequence. But, conversely, the opinion about s_1 —be it either positive or negative— will also affect the opinion on τ through relationship r_1 . Thus, the opinion that a participant will have on a statement, in this case the target, will be conditioned both by the consequences of this statement and by the things from which it itself follows. In general, a direct relationship between statements has to be established and we then decide which direction to take when evaluating the framework. We may define this evaluation so as to relate one statement s (or a set of statements) to (i) its consequence statements, or,

to (ii) the statement whose opinion is affected by the opinion about s . We choose the first (i) option to reflect the direction of reasoning (from premises to conclusions) in the debate structure. However, our subsequent process aimed at opinion aggregation will traverse these relationships in the opposite (ii) direction, though we will use the same structure to represent the debate.

Direct and Indirect Opinion. Each participant involved in a debate expresses its opinion in terms of the desirability or undesirability of the different statements in the model, and the acceptability of the relationships linking these statements. As mentioned above, relationships are a key element of the RM: they indicate that the opinion about one statement affects the opinion about another statement. In particular, we will use the terms: (i) *direct opinion* to refer to the value directly given to a statement by a participant; and (ii) *indirect opinion* to refer to the values given to the related objects whose opinions may condition the direct opinion. Note that we are dealing with human opinions, and so we cannot expect that opinions are rational — contradictions or inconsistencies between the direct and indirect opinions expressed by participants may arise, just as in the example shown above. By comparing the direct opinion with the indirect opinion about each statement we can categorise “reasonable” opinions and provide a notion of *coherence* [86].

Example 2.2. Now consider three people taking part in the debate described above with the following opinions:

- **Participant 1** is a middle-aged woman with a family who lives in the neighbourhood where the sports centre is proposed to be built. She values the proposal (τ) very positively, because her family practices sports. Although she has been living in the neighbourhood for a long time, she has no childhood memories of and doesn’t care too much for the historical character of the neighbourhood, so she values statement s_1 as neutral. Finally, she assigns a small acceptance value to relationship r_1 because she acknowledges that building the sports centre will imply the demolition of some buildings, but she doesn’t believe that the character of the neighbour will be too affected by such a loss.
- **Participant 2** is a retired elderly man. He has always lived in the neighbourhood and he would like to preserve its unique features. Hence, he values statement s_1 negatively, and also target τ because he is not interested in sports and would prefer another kind of public building

instead. He fully agrees with relationship r_1 , since he considers that the buildings in the proposed location for the new sports centre are important to the neighbourhood and would be demolished if the centre is built.

- **Participant 3** is a young postgraduate student who does not practice any regular sports that can be held in the planned sports centre, so values τ quite negatively. He agrees with relationship r_1 because he acknowledges that the existing buildings could be catalogued as of special architectural interest. However, he is neutral with respect to s_1 because he does not care about preserving the character of the neighbourhood.

Opinion Aggregation. Once all users have expressed their opinions on statements and relationships, opinions must be aggregated to calculate a collective opinion. This aggregation can take into account direct opinions, indirect opinions, or a combination of both direct and indirect. In establishing suitable aggregation functions, we have to take into account that individual opinions may be incoherent. Nonetheless, we aim to design aggregation functions that can combine these “imperfect” individual opinions into a “reasonable” collective opinion.

The main contribution of the RM is to distinguish clearly between *objective* and *subjective* information in a simplified way, and to do this differently to many other argumentation models (as [7, 45, 56]), where subjective information is jointly represented with objective information in the debate structure leading to predefined relationships (such as attack or defence) that must be shared among all participants. Thus, the expressiveness of the RM is significantly higher being able to represent debates where participants do not need to agree on the relationships between different objective facts. In addition, these argumentation models usually presume that users’ opinions are rational (i.e., fully consistent). However, human debates are affected by information and/or cognitive limitations of users, which makes this rationality requirement too restrictive. We argue that our coherence notion is a weaker version of this rationality requirement that is better suited to model human debates.

2.2. The collective decision workflow

A debate can be built in many ways following the parameters of a DRF. Next we propose a simple procedure that would produce a DRF and will

serve as a support guide for the sections to come. There are four main steps to applying our model:

Step 1 - *Start debate.* A set of statements for which we intend to obtain a collective opinion are chosen as targets of the RM.

Step 2 - *Extend debate.* Participants are then allowed to put forward relationships that will represent relevant reasoning. A relationship may either be put forward in conjunction with new statements –so it connects from existing statements to a new statement– or it may connect existing statements. This step continues until no participant wishes to add a further relationship.

Step 3 - *Input opinions.* Participants express their opinions on all the relationships and statements in the RM, by providing *subjective evaluations* of them. The evaluations of statements expresses preferences over them, while the evaluation of relationships expresses agreement, or otherwise, with the reasoning represented by the relationship.

Step 4 - *Obtain collective opinion.* The opinions of the participants are merged to establish a consensus view of each statement and relationship that are in the framework. Hence we also obtain the collective opinion of the target statements.

These steps are shown graphically in Figure 2.

3. Formalising the relational model

Recall from section 2 that we conceptually distinguished between the structural part of the relational model, namely the statements made during a debate together with their causal relationships, and the opinions put forward by participants in a debate. Figure 3 depicts this distinction, together with all the basic elements presented in this section. Thus, in what follows, we formally introduce both concepts, first the structure and then the opinions, with the aid of an extension of the example shown in section 2 that we employ throughout the paper.

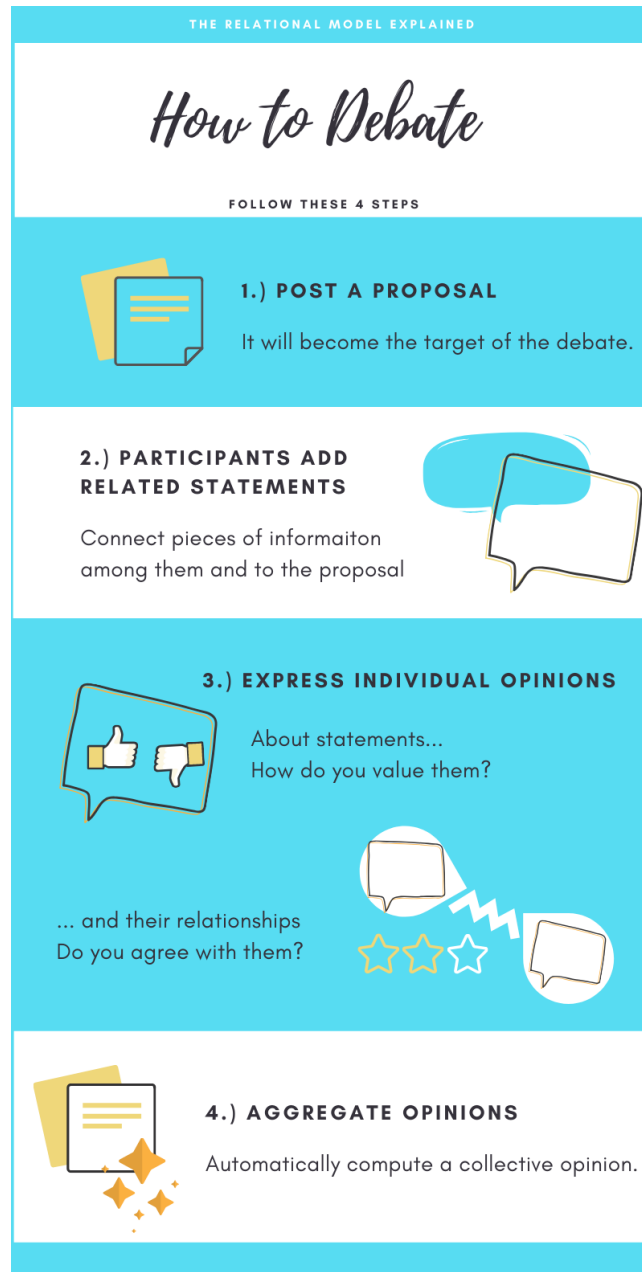


Figure 2: Participation using the relational model: debate steps graphically explained.



Figure 3: The basic elements of our relational model.

3.1. Formalising the structure

First, we introduce the formal notion of a *relational framework* to capture the relationships between statements. Our notion of relationship will consider a non-empty set of source statements and a destination statement. In general, by relating a set of (source) statements to a (destination) statement, we indicate that the source statements support inferring the destination statement, though the framework is agnostic about the form that the support and the inference mechanism takes. For instance, in our example in Figure 4, statements s_2 and s_3 support inferring s_4 . Formally:

Definition 3.1. A *relational framework* RF is a pair $\langle \mathcal{S}, \mathcal{R} \rangle$, where \mathcal{S} is a set of statements and $\mathcal{R} \subset \mathcal{P}(\mathcal{S}) \times \mathcal{S} \times \mathbb{N}$ is a relation fulfilling:

- **Acyclicity.** There are no cycles in \mathcal{R} , namely there is no subset of relationships $\{(\Sigma_0, s_1, c_1), \dots, (\Sigma_{n-1}, s_n, c_n)\} \subset \mathcal{R}$ such that $s_i \in \Sigma_i$, $i \in \{1, \dots, n-1\}$, and $s_n \in \Sigma_0$.

Since the relation \mathcal{R} is acyclic, it follows that \mathcal{R} is neither *reflexive* ($\forall s \in \mathcal{S}, (\Sigma \cup \{s\}, s, c) \notin \mathcal{R}$) nor *symmetric* ($\forall s_1, s_2 \in \mathcal{S}$, if $(\Sigma_1 \cup \{s_1\}, s_2, c_2) \in \mathcal{R}$ then $(\Sigma_2 \cup \{s_2\}, s_1, c_1) \notin \mathcal{R}$). Note that we do not impose any restriction on the transitivity of the relation \mathcal{R} .

Notice also that we include a natural number within the relation in order to differentiate relationships between the same set of statements Σ and s . From a practical perspective, this allows to signal that alternative relationships can bear on the very same statements³ (as shown in Figure 4, where target τ is related to statement s_1 through relationships r_1 and r_6).

Now, since debates are aimed at achieving a collective decision on target topics, we extend our definition above to incorporate the notion of *target* statements as follows:

Definition 3.2. A *directed relational framework* (DRF) is a tuple $\langle \mathcal{S}, \mathcal{R}, T \rangle$ such that:

- (i) $\langle \mathcal{S}, \mathcal{R} \rangle$ is a relational framework;

³The inclusion of a natural number into the specification of a relationship does not affect the formal contributions of the paper, since, as it will be shown later, relationships are grouped into (and subsequently traversed through) the sets R^+ defined in equation 1, without considering any restrictions on the statements they relate.

- (ii) $T \subset \mathcal{S}$ is a set of target statements;
- (iii) Target statements in T can only be the source of relationships, namely for any relationship $(\Sigma, s, c) \in \mathcal{R}$, $s \notin T$; and
- (iv) All non-target statements are connected to targets so that for any statement $s \in S$, $s \notin T$, there is a path $\{(\Sigma_0, s_1, c_1), \dots, (\Sigma_{n-1}, s, c_n)\} \subset \mathcal{R}$ such that $T \cap \Sigma_0 \neq \emptyset$.
- (v) Every target statement shares some common descendant with some other target statement, namely for every $\tau \in T$ there is another target $\tau' \in T$ and some statement $s \in S$ such that there is a path from τ to s and another path from τ' to s .

Note that a DRF is constrained to be a *connected acyclic* graph, albeit one that can have several targets. This reflects the idea that, since a DRF represents a single debate, every statement in that debate should have some connection to the rest of the debate.

In what follows we extend our example in section 2 to produce a graphical representation of a DRF that will help us visualise the information in a debate⁴. In this representation, statements are nodes and relationships are arcs. Recall that our example considered statements τ (Construct a new sports centre in a specific place in the neighbour) and s_1 (Destruction of the neighbour character), as well as relationship r_1 (The construction of the sport centre will imply the demolition of existing buildings that give historical relevance to the neighbourhood) connecting both. Besides that, next we consider further statements and relationships as listed in tables 1 and 2 respectively. Finally, figure 4 depicts the connections between statements through relationships. Note that r_4 is a hyperedge, connecting three statements⁵.

⁴For the sake of simplicity, we limit the example to three participants and a small number of statements and relationships.

⁵Notice that one participant could introduce an extra relationships from τ to s_5 , representing the reasoning “A new community center will give more relevance to the neighbourhood, that will increase the house price”, which is not the sum of r_2, r_3 and r_4 , but represents a whole new way to connect τ to s_5 . This shows that the transitivity allowed in the model, connecting statements from not consecutive levels of the debate via a single relationship, is not intended to represent the combined reasoning formed by the reasoning steps in between.

Statement	Description
τ	Construction of a sport centre in a particular location in the neighbour
s_1	Destruction of the neighbour character
s_2	Attraction of more affluent residents to the neighbour
s_3	Attraction of new business to the neighbour
s_4	Crime reduction in the neighbour
s_5	Property values raise in the neighbour

Table 1: Statements for the sports centre example.




Figure 4: The DRF for the sports centre example.

3.2. Formalising opinions

Now we address the formalisation of the opinions put forward by participants in a debate. We consider that opinions can be held both about statements and relationships. We therefore define two functions that capture the opinions of individuals: (i) a *valuation function* over statements; and (ii) an *acceptance function* over relationships. On the one hand, a valuation function conveys the subjective value that an individual places on each statement. On the other hand, an acceptance function expresses the subjective plausibility that an individual assigns to each relationship, representing a reasoning, connecting statements. Formally:

Definition 3.3 (Valuation function). Given a DRF $\langle \mathcal{S}, \mathcal{R}, T \rangle$, a *valuation function* $v : \mathcal{S} \rightarrow I$ maps each statement to a value in $I = [-a, a]$, $a \in \mathbb{R}^+$.

Relationship	Reasoning	Connection
r_1	The construction of the sport centre will imply the demolition of existing buildings which now give historical relevance to the neighbour	τ to s_1
r_2	The new sport centre will make the neighbour more attractive for wealthy residents because they are more interested in leisure activities	τ to s_2
r_3	A new community centre will attract more businesses to the surrounding area.	τ to s_3
r_4	Having richer residents and more businesses will increase the security measures around the neighbour and therefore reduce criminal activities.	$\{s_2, s_3\}$ to s_4
r_5	The reduction of crime will increase the price of the houses in the neighbour	s_4 to s_5
r_6	A new building will make collide the new architecture with the extended old nature of the area, changing its character	τ to s_1

Table 2: Reasoning for the sports centre example.

Given a statement $s \in \mathcal{S}$: if $v(s) = a$ we say that s counts on *full positive valuation*; if $v(s) = -a$ we say that s counts on *full negative valuation*; and if $v(s) = 0$ we say that s has *neutral valuation*.

Definition 3.4 (Acceptance function). Given a DRF $\langle \mathcal{S}, \mathcal{R}, T \rangle$, an *acceptance function* maps each relationship to a value in I^+ , $w : \mathcal{R} \rightarrow I^+ = [0, a]$.

Given a relationship $r \in \mathcal{R}$ and an acceptance function w , we will refer to the value $w(r)$ as the *acceptance degree* of r . If $w(r) = a$ we say that the acceptance function expresses *full agreement* with the relationship, whereas if $w(r) = 0$ we say that it expresses *full disagreement*. Without loss of generality, henceforth we will set a to 1, and hence $I = [-1, 1]$ and $I^+ = [0, 1]$.

Considering our running example, graphically represented in figure 4, figures 5 and 6 show the valuation functions and acceptance functions of agents 1, 2, and 3: v_1, v_2 and v_3 encode agents' valuations on statements, while w_1, w_2 and w_3 encode agents' acceptances of relationships. We consider now the description of agents' opinions in section 2.1 to exemplify how they translate into valuations and acceptances. Thus, for instance, agent one is "highly positive" on the target τ ($v_1(\tau) = 0.9$), but neutral regarding statement s_1 ($v_1(s_1) = 0$). Furthermore, agent one considers that the plausibility of relationship r_1 is "little" ($w_1(r_1) = 0.2$).

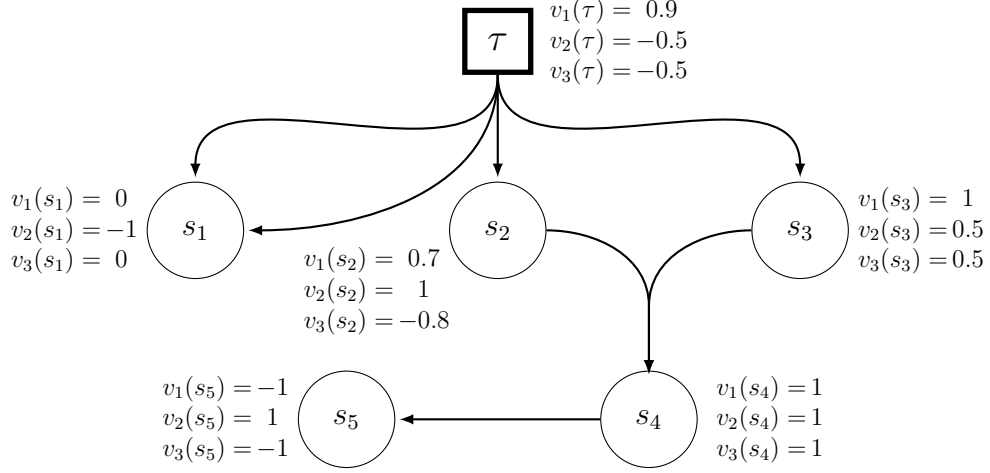


Figure 5: Agents' valuation functions.

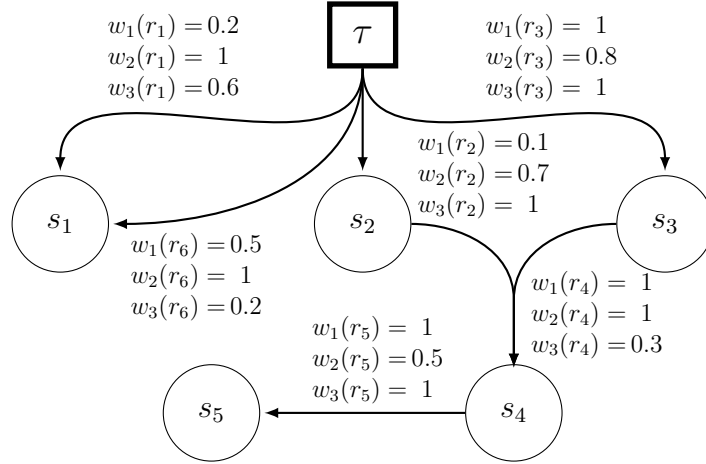


Figure 6: Agents' acceptance functions.

At this point, we are ready to formally introduce the notion of individual opinion over a DRF .

Definition 3.5 (Opinion). Given a $DRF = \langle \mathcal{S}, \mathcal{R}, T \rangle$, an *opinion* over the DRF is a pair $O = (v, w)$ such that v is a valuation function and w is an acceptance degree.

For practical and realistic purposes, we assume that for each relationship at least one agent values it different from 0. Otherwise, it would be the same, in practical terms, to have or not have the relationship in the debate.

Henceforth, we shall note the class of all opinions over a DRF as $\mathbb{O}(DRF)$.

As depicted in figures 5 and 6, each agent i involved in a debate will have its individual opinion $O_i = (v_i, w_i)$. Our goal will be to compute a collective opinion from the opinions issued by the agents.

4. Characterising coherent opinions

Previous work on the formal modelling of debates has placed restrictions on the opinions that individuals can put forward. For example, [7] interprets the opinions expressed by individuals as labels, in the sense of [11], for the arguments that they are expressing opinions about. Thus, an argument can be labelled **in**, meaning that the individual thinks that it holds, **out**, meaning that the individual thinks it does not hold, or **undec**, meaning that the individual is not sure whether it holds or not. These labellings are restricted to be *complete* labellings [11], broadly meaning that they conform to a notion of rationality where arguments are **out** if they are attacked by arguments that have been established to be **in**, and are **in** if they are only attacked by arguments that are **out**. We believe that the restrictions imposed by these notions are too restrictive for modelling human debates, as humans may express opinions that are far from rational.

Instead, we impose weaker conditions for an individual opinion to be classified as reasonable or *coherent*, along the lines of our former work in [46]. Hence, given a statement we contrast opinions expressed about that statement, the *direct opinion*, with the opinions expressed about the immediate descendants of the statement, what we call the *indirect opinion*, and look for ways in which these may be made somewhat consistent.

Informally, what we do is the following. First, we compute an estimated opinion for a statement based on the indirect opinion about it. Then we say that a set of opinions about a statement are coherent if the estimated opinion for the statement aligns with the direct opinion about the statement. This will be the case when the opinion (valuations) about the descendants is close to the overall opinion (valuation) about the statement. Considering our example in figure 5 again, consider statement τ , its descendants (s_1 , s_2 , and s_3), and the opinion of agent 2 (v_2). We observe that although the direct opinion

about τ is rather negative ($v_2(\tau) = -0.5$), the valuations for its descendants are diverse: while the valuation for s_1 is also rather negative ($v_2(s_1) = -1$), and hence in line with τ , the valuations on the other descendants are rather positive ($v_2(s_2) = 1$ and $v_2(s_3) = 0.5$), and hence not in line with τ . Thus, at first sight⁶ it would seem that the overall estimated opinion is not in line with the direct opinion.

In what follows, we first formalise our notion of *estimation* as an aggregated measure formed from the indirect opinion about a statement — i.e., the collection of values on the descendants and their relationships. This will consider valuations and acceptance degrees related to a statement and its relationships so that the more accepted a relation between a statement and its descendants, the more important the opinion about that descendant. Thereafter, we will formalise our notion of coherence by measuring how close the direct opinion about a statement is to the estimated opinion about that statement.

First, we introduce some concepts and notations that will aid us on later steps. Given a $DRF = \langle \mathcal{S}, \mathcal{R}, T \rangle$, we define the *set of relationships from* $s \in \mathcal{S}$ as the set of relationships having s in the set of initial statements. Formally,

$$R^+(s) = \{r = (\Sigma, s', c) \in \mathcal{R} \mid s \in \Sigma\}. \quad (1)$$

We will use the term *descendants of a statement s* , denoted by $D(s)$, to refer to any statement s_r connected to s by a relationship r that has s as one of the initial statements and s_r as final statement. Formally,

$$D(s) = \{s_r \in \mathcal{S} \mid r = (\Sigma, s_r, c) \in R^+(s)\}.$$

Next, we define the concept that connects direct and indirect opinion in order to characterise our notion of coherence. We will use an *estimation function* to compute an estimate of the direct opinion using the values gathered for the indirect opinion.

Definition 4.1. Given a $DRF = \langle \mathcal{S}, \mathcal{R}, T \rangle$ and $O = (v, w)$ and opinion over the DRF , the *estimation function* is a valuation function:

$$\begin{array}{ccc} e : \mathcal{S} & \longrightarrow & I \\ s & \longmapsto & e(s) = f(IO(s)) \end{array}$$

⁶Note that we are not considering acceptances at this point.

where $IO(s) = \{(v(s_r), w(r)) \mid r \in R(s), s_r \in D(s)\}$ and f is a projection:

$$f : IO(s) \longrightarrow I$$

that computes an approximate value for s from the valuations and acceptance degrees in the indirect opinion.




Figure 7: A visualization of the estimation function.

In other words, the estimation function computes an estimated value for a statement using the valuations and acceptance degrees for indirect opinion about that statement. Figure 7 is a diagrammatic representation of the estimation function. This definition is designed to be an abstract definition that allows for many estimation functions to be defined to compute different approximations for the direct opinion. It therefore specifies a broad family of estimation functions rather than any specific function.

For the rest of this paper we will use a specific estimation function to compute the estimate opinion on the results later on. We will use the weighted average of the valuations on the descendants, where the weights are the acceptance degrees on the relations leading to each descendant. In this manner, the more accepted a relation, the more valued the opinion on the descendant. Formally,

$$e(s) = \begin{cases} v(s) & \text{if } R^+(s) = \emptyset \text{ or } \sum_{r \in R^+(s)} w(r) = 0, \\ \frac{\sum_{r \in R^+(s)} w(r)v(s_r)}{\sum_{r \in R^+(s)} w(r)} & \text{otherwise.} \end{cases}$$

Notice that when a statement has no descendants we take the direct opinion as the estimated opinion, thus we cannot gather any value from an empty indirect opinion. Notice also that this specification of estimated opinion can be regarded as a general approach to compute the direct opinion using the indirect opinion. This is because we are allowing the relationships between

statements to represent any kind of reasoning, so we cannot specify a specific behaviour for the estimated value.

Informally, an opinion is characterised as coherent for a given statement when the value assigned by the participant (issuing the opinion) to the statement (i.e., its direct opinion) is aligned with the values and plausibility assigned to its descendants (i.e., its estimate opinion). Furthermore, given the continuous values allowed on the opinion, we can choose the degree of coherence by using a parameter ϵ . Formally:

Definition 4.2 (Coherence). Consider a $DRF = \langle \mathcal{S}, \mathcal{R}, T \rangle$ and an $\epsilon \in (0, 1)$ ⁷ difference. We say that opinion $O = (v, w)$ is ϵ -coherent on $s \in \mathcal{S}$ when

$$|v(s) - e(s)| < \epsilon.$$

In general, an opinion O will be ϵ -coherent if it is ϵ -coherent for every statement in \mathcal{S} . We will notate as $\mathbb{C}_\epsilon(DRF)$ the class of all the ϵ -coherent opinions. Thus, if O is an ϵ -coherent opinion then $O \in \mathbb{C}_\epsilon(DRF)$.

Example 4.1. Following the example, now we can compare the values from the expectation function and the actual value given by participant 1 to each statement, its direct opinion, see figure 8. We can see that if $\epsilon \in (0.3, 1)$ then the opinion of participant 1 for the statements s_1, s_2, s_3 and s_5 is ϵ -coherent but not for statement s_4 due to the difference between direct opinion and estimated value, which is the maximum possible. Because of this statement s_4 , the opinion of participant 1 cannot be classified as ϵ -coherent for any $\epsilon \in (0, 1)$.

⁷We choose the interval $(0,1)$ for the value of ϵ as the minimum interval that guarantees that if the direct opinion is 1 (or -1) then an opinion cannot be classified as coherent when the estimation function value is of the opposite sign, i.e., $e(s) \not\leq 0$ (or $e(s) \not\geq 0$ respectively)




Figure 8: Coherence of Agent 1

5. Formalising the collective decision making problem

As stated above, our goal is to help agents reach a collective decision on target statements. This corresponds to step 4 of the protocol in Section 2.2 and Figure 2 (as well as to the third stage in Figure 3). In Section 5.1 we cast our goal as an opinion aggregation problem. We propose to solve such problem using an aggregation function that synthesises a single opinion out of all agents' opinions. Although opinions can be aggregated in different ways, here we follow [46] in requiring that the outcome of an aggregation must be fair. In particular, Section 5.2 introduces desirable social choice properties to help analyse and compare opinion aggregation functions.

5.1. The opinion aggregation problem

First, we define our notion of an opinion profile, which brings together the opinions of the individuals involved in a debate. From hereon we use the term “agent” along with the term “individual” to refer to the participants in the debate.

Definition 5.1 (Opinion profile). Let $Ag = \{1, \dots, n\}$ be a set of n agents and a $DRF = \langle \mathcal{S}, \mathcal{R}, T \rangle$. An *opinion profile* is a collection of opinions $(O_1 = (v_1, w_1), \dots, O_n = (v_n, w_n)) \in \mathbb{O}(DRF)^n$ over the DRF such that $O_i = (v_i, w_i)$ stands for the opinion of agent i .

The problem at hand is how to aggregate the opinions in an opinion profile to produce a single opinion so that single opinion is a reasonable summary of the opinions in the opinion profile. If the opinion profile represents the views expressed by individuals in a debate, the combination should represent the collective opinion of all the individuals. The opinion aggregation function, which we formalise below, is the mechanism for establishing this collective opinion.

Definition 5.2 (Opinion aggregation function). Given a DRF and a set of n agents Ag , a function $F : \mathcal{D} \subset \mathbb{O}(DRF)^n \rightarrow \mathbb{O}(DRF)$ mapping an opinion profile to a single opinion is called an *opinion aggregation function*. Given an opinion profile P in the domain \mathcal{D} , $F(P)$ is called the *collective opinion* by F and it will be noted as $F(P) = (v_{F(P)}, w_{F(P)})$.

Thus, in terms of the components of our model, the collective opinion output by an opinion aggregation function combines the collective valuations over statements and the collective acceptances over the relationships linking them. In section 6 we define specific opinion aggregation functions that compute a collective opinion. Before that we introduce the properties that we will use to analyse the aggregation functions.

5.2. Social choice properties

Social choice theory provides formal properties to characterise aggregation methods in terms of outcome fairness [32]. In what follows, we formally adapt some of the desirable social properties of an aggregation function introduced in [7] and [46]. Besides adapting properties, we define some novel properties that characterise aggregation functions motivated by the fact that here we are considering opinions to be continuous-valued in contrast to the discrete-valued opinions considered in [7] and [46].

First, we characterise aggregation functions in terms of the opinion profiles that they can take as input. Thus, we adapt from [7] the notion of *exhaustive domain* to characterise opinion aggregation functions that are defined for any opinion profile. Thereafter, we modify this property to limit an opinion aggregation function to operate with ϵ -coherent opinion profiles.

Exhaustive Domain (ED) . An opinion aggregation function F satisfies *exhaustive domain* if its domain is $\mathcal{D} = \mathbb{O}(DRF)^n$, namely if the function can operate over all profiles.

ϵ -Coherent Domain (ϵ -CD). An opinion aggregation function F satisfies *ϵ -coherent domain* if its domain \mathcal{D} contains all ϵ -coherent opinion profiles, namely $\mathbb{C}_\epsilon(DRF)^n \subseteq \mathcal{D}$.

Note that we will sometimes refer to ϵ -Coherent Domain as “coherent domain”.

Lemma 5.1. An opinion aggregation function satisfying exhaustive domain also satisfies ϵ -coherent domain.

Proof. Straightforward, since an aggregation function taking any opinion profile will also take ϵ -coherent opinion profiles. \square

Moreover, we also define *collective ϵ -coherence* as a property characterising opinion aggregation functions that produce ϵ -coherent collective opinions. Therefore, our notion of collective ϵ -coherence here is more relaxed than the crisp notion of coherence in [46].

Collective ϵ -coherence (ϵ -CC). An opinion aggregation function F satisfies *ϵ -collective coherence* if for all $P \in \mathcal{D}$, then $F(P) \in \mathbb{C}_\epsilon(DRF)$.

In accordance with [46], here we consider ϵ -CC as the most desirable property that can be satisfied by an aggregation function, since collective coherence is the foundation of the acceptability of collective decisions [86]. Notice also that, as in [46], collective coherence can be regarded as the counterpart of the notion of *collective rationality* in [7].

Next, anonymity and non-dictatorship characterise the importance of the agents involved in a debate that yields a collective opinion. On the one hand, anonymity is a social choice property requiring that the opinions of all the agents involved in a debate are considered to be equally significant. On the other hand, non-dictatorship requires that no agent overrules the opinions of the rest of agents.

Anonymity (A) Let $P = (O_1, \dots, O_n)$ be an opinion profile in \mathcal{D} , σ a permutation over Ag , and $P' = (O_{\sigma(1)}, \dots, O_{\sigma(n)})$ the opinion profile resulting from applying σ over P . An opinion aggregation function F satisfies *anonymity* if $F(P) = F(P')$.

Non-Dictatorship (ND). An opinion aggregation function F satisfies *non-dictatorship* if no agent $i \in Ag$ satisfies that $F(P) = O_i$ for every opinion profile $P \in \mathcal{D}$.

Notice that non-dictatorship is a weaker version of anonymity since it follows directly from it — any aggregation function that satisfies anonymity will satisfy non-dictatorship.

Now we consider how an opinion aggregation function behaves when agents agree on their opinions about statements. Unanimity is the social choice property that characterises the behaviour of aggregation functions when there is agreement among agents' opinions. Here we will adapt the classic notion of unanimity in [7], here called *narrow unanimity*, and the notion of *endorsed unanimity* from [46] that helps exploit dependencies between statements. Since the notion of narrow unanimity is rather rigid, we also define two new, more flexible unanimity properties. Therefore, overall we define a family of unanimity properties, and below we study the relationships between them.

Narrow Unanimity (NU). Let $P = (O_1, \dots, O_n)$ be an opinion profile, where $P \in \mathcal{D}$. An opinion aggregation profile F satisfies *narrow unanimity* if, for any $s \in S$ such that $v_i(s) = \lambda$ for all $i \in \{1, \dots, n\}$, then $v_{F(P)}(s) = \lambda$ holds.

Narrow unanimity defines unanimity as being when all agents share the very same opinion. While this is very possible in settings where agents only have a few discrete possibilities for expressing their opinion, as in [7, 46], it is not likely to occur in the setting we are studying here, where opinions can take a wide range of values. As a result, we propose some relaxed variations which are more useful for the setting we consider. First, we say that *sided unanimity* will hold when, for each statement, either all opinions on it are positive or negative. Formally,

Sided Unanimity (SU). Let $P = (O_1, \dots, O_n)$ be an opinion profile, where $P \in \mathcal{D}$. An opinion aggregation profile F satisfies *sided-unanimity* if for every $s \in S$:

- if $v_i(s) > 0$ for all $i \in Ag$ then $v_{F(P)}(s) > 0$;
- if $v_i(s) < 0$ for all $i \in Ag$ then $v_{F(P)}(s) < 0$.

We also find a weaker version of sided unanimity to be worth distinguishing:

Weak Unanimity (WU). Let $P = (O_1, \dots, O_n)$ be an opinion profile, where $P \in \mathcal{D}$. An aggregation profile F satisfies *Weak unanimity* if, for every $s \in S$:

- if $v_i(s) = 1$ for all $i \in Ag$ then $v_{F(P)}(s) > 0$;
- if $v_i(s) = -1$ for all $i \in Ag$ then $v_{F(P)}(s) < 0$.

Although WU requires that all agents agree on fully positive (1) or fully negative (-1) valuations on statements, it does not require that the output of the opinion aggregation function takes on those same values, as required by narrow unanimity. This property has value when translating valuations expressed in a discrete model such as those in [7, 46] into our model, and so has value in allowing us to relate our model to those which came before.

From the definitions above, it follows that the three notions of unanimity are related.

Proposition 5.1 (Unanimity relationships). If an opinion aggregation function satisfies Sided Unanimity then it satisfies Weak Unanimity. If an opinion aggregation function satisfies Narrow Unanimity then it satisfies Weak Unanimity.

Proof. Clearly, if an aggregation function cannot hold the sign of the aggregation when the assumptions of Weak Unanimity are satisfied, then it is straightforward to see that will fail to satisfy Sided Unanimity.

Obviously, if an aggregation function returns a value λ when all the agents valued the statement as λ , it will return the value when it is 1 or -1 . \square

Below, we will show that we need further assumptions to prove that NU implies SU.

As a final unanimity property, we adapt the notion of endorsed unanimity from [46] to consider unanimity based on indirect opinions. In short, an opinion aggregation function will satisfy *endorsed unanimity* if, for each statement, the collective opinion on the statement is in line with the unanimous indirect opinion on it. Formally,

Endorsed Unanimity (EU). Let $P = (O_1, \dots, O_n)$ be an opinion profile such that $P \in \mathcal{D}$. An aggregation profile F satisfies *endorsed unanimity* if for every $s \in \mathcal{S}$:

- (i) if $v_i(s_d) = 1$ for any $i \in Ag$ and $s_d \in D(s)$ (called *full positive support*), then $v_{F(P)}(s) > 0$; and
- (ii) if $v_i(s_d) = -1$ for any $i \in Ag$ and $s_d \in D(s)$ (called *full negative support*), then $v_{F(P)}(s) < 0$.

Note that, this property is closely related to the notion of coherence. We will show below that restricting the domain to coherent opinion profiles will help to fulfil coherence.

Next, we introduce monotonicity properties to study how the result of an opinion aggregation function changes as opinions change. First, we adapt the notion of monotonicity from [7]: if some of the direct opinions about a statement increase (or decrease) the collective opinion should increase (or decrease) accordingly.

Monotonicity (M) Let $s \in S$ be a statement, and $P = (O_1, \dots, O_n)$ and $P' = (O'_1, \dots, O'_n)$ such that for every i $v_i(s) \leq v'_i(s)$. We say that an opinion aggregation function F is *monotonic* if $v_{F(P)}(s) \leq v_{F(P')}(s)$.

Notice that M only takes into account the direct opinion about each statement. Since we aim at handling opinion aggregation functions that merge both direct and indirect opinions, we define a variation of monotonicity taking this into account. To do this we adapt the notion of *familiar monotonicity*⁸ in [46]. In our case, familiar monotonicity requires that when the direct opinion on a statement increases, the collective opinion does not decrease provided that the opinions on the descendants of the statement do not change either. Formally:

Familiar Monotonicity (FM). Let $s \in S$ be a statement, and $P = (O_1, \dots, O_n)$ and $P' = (O'_1, \dots, O'_n)$ such that for every opinion i satisfies that $v_i(s) \leq v'_i(s)$, and, $w_i(r) = w'_i(r)$ and $v_i(s_r) = v'_i(s_r)$ for every relationship $r \in R(s)$ and its associated descendant $s_r \in D(s)$. We say that an opinion aggregation function F satisfies FM if $v_{F(P)}(s) \leq v_{F(P')}(s)$.

The following lemma establishes the relationship between monotonicity properties.

Lemma 5.2. An opinion aggregation function that satisfies monotonicity also satisfies familiar monotonicity.

Proof. Since FM assumes a restriction on the descendants' opinions that monotonicity does not, clearly fulfilling monotonicity implies the fulfilment of familiar monotonicity. \square

⁸The name derives from the fact that this form of monotonicity takes into account opinion about the descendents of a statement which make up its family.

Now we are ready to prove the relationship between narrow unanimity and sided unanimity via monotonicity.

Proposition 5.2. An opinion aggregation function that satisfies Narrow unanimity and Monotonicity also satisfies Sided unanimity.

Proof. Let F be an opinion aggregation function that fulfils Narrow unanimity and Monotonicity. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a DRF and $s \in \mathcal{S}$. Assume that for any $i \in Ag$, $v_i(s) \geq \lambda$ for a certain $\lambda > 0$. Consider the opinion profile P' such that for any i $v'_i(s) = \lambda$. Then, by Monotonicity of F , $v_{F(P)}(s) \geq v_{F(P')}(s)$ holds, and by Narrow unanimity of F , $v_{F(P')}(s) = \lambda$. Then $v_{F(P)}(s) \geq \lambda > 0$, proving that also satisfies Sided Unanimity. The proof for the negative case is analogous. \square

Finally, we introduce the property of independence, which will be used to emphasise the difference between those opinion aggregation functions that exploit direct and indirect opinions and the ones that do not. Essentially, the property states that the collective opinion on a statement will only depend on the direct opinions on it. Therefore, independence disregards indirect opinions.

Independence (I) Let be two profiles $P = (O_1, \dots, O_n)$ and $P' = (O'_1, \dots, O'_n)$, such that $P, P' \in \mathcal{D}$; and $s \in \mathcal{S}$ a statement, such that for all agents $i \in \{1, \dots, n\}$ $v_i(s) = v'_i(s)$. An aggregation function F satisfies *independence* if $v_{F(P)}(s) = v_{F(P')}(s)$.

Next result shows the relationship between monotonicity and independence.

Proposition 5.3. An opinion aggregation function that satisfies monotonicity also satisfies independence.

Proof. Let $s \in \mathcal{S}$ be a statement and $P = (O_1, \dots, O_n)$, $P' = (O'_1, \dots, O'_n)$ two opinion profiles satisfying the assumptions of independence on s , i.e., for every $i \in Ag$ $v_i(s) = v'_i(s)$, and F an aggregation function satisfying monotonicity.

For each i , the equality $v_i(s) = v'_i(s)$ is equivalent to (a): $v_i(s) \geq v'_i(s)$, and, (b): $v_i(s) \leq v'_i(s)$. Thus, assuming monotonicity from (a) we can deduce $v_{F(P)}(s) \geq v_{F(P')}(s)$, and from (b) we can deduce that $v_{F(P)}(s) \leq v_{F(P')}(s)$. Hence, we conclude that $v_{F(P)}(s) = v_{F(P')}(s)$ proving that F satisfies independence. \square

Having listed these properties, it is important to note that they are not all equal. For a multi-party discussion, we believe that the most important property is collective coherence. If an aggregation function is collectively coherent, the resulting combined opinion will be coherent regardless of the coherence of the initial opinions that are being merged. In other words, an aggregation function that satisfies collective coherence will always discover a coherent overall opinion no matter how incoherent are the opinions on which it is based. Along with collective coherence, the properties that we would like to see for an aggregation function are the two domain related properties — exhaustive domain and coherent domain — because they allow for broad applicability of the function, and naturally we would prefer exhaustive domain because of its wider reach. Finally, we regard the usual social choice properties of anonymity and non-dictatorship as essential.

Among the unanimity properties, we find sided, weak and endorsed unanimity, which allow more maneuverability when using dependencies to build the collective opinion than narrow unanimity, to be more desirable than narrow unanimity. We do not consider narrow unanimity to be desirable because of its close relationship to independence⁹. Thus, an aggregation function satisfying narrow unanimity would forbid the use of the indirect opinion the way we consider to be necessary.

Though it is natural to require some form of monotonicity, we do not consider the property M to be desirable because of its relationship to independence and the discarding of indirect opinion. Thus, in its place we prefer familiar monotonicity, which takes into consideration indirect opinion.

Finally, although we focus the design of aggregation functions towards the use of both direct and indirect opinion, we include independence, monotonicity and narrow unanimity in the set of properties we consider in order to to emphasize whether aggregation functions take account of indirect opinion or not.

6. Aggregation functions to enact collective decision making

This section defines a family of aggregation functions for the relational model. All of the aggregation functions that we define use some combination of direct and indirect opinions to generate a collective opinion. Our aim is

⁹Though it is not directly related — we need to add further minor assumptions to narrow unanimity for it to imply independence.

to explore the spectrum of aggregation functions from functions that only employ direct opinions to functions that only employ indirect opinions, so that we can later on (in section 7) compare the benefits that they yield in terms of social choice properties. This will allow us to learn how to best exploit indirect opinions to obtain collective opinions.

We start by defining an aggregation function that only aggregates direct opinions, and thus disregards indirect opinions. This function will obtain a collective opinion by averaging valuations per statement and acceptance degrees per relation from the individual opinions in an opinion profile. Formally:

Definition 6.1 (Direct aggregation). Let $\langle \mathcal{S}, \mathcal{R}, T \rangle$ be a *DRF* and $P = (O_1, \dots, O_n)$ an opinion profile over the *DRF*. The *direct aggregation* of P over the *DRF* is defined as a function $D(P) = (v_{D(P)}, w_{D(P)})$, where $v_{D(P)}(s) = \frac{1}{n} \sum_{i=1}^n v_i(s)$ and $w_{D(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r)$ for any statement $s \in \mathcal{S}$ and relationship $r \in \mathcal{R}$.

Example 6.1. Figure 9 shows the result of applying the direct function to the opinion profile of our running example, shown in figures 5 and 6.

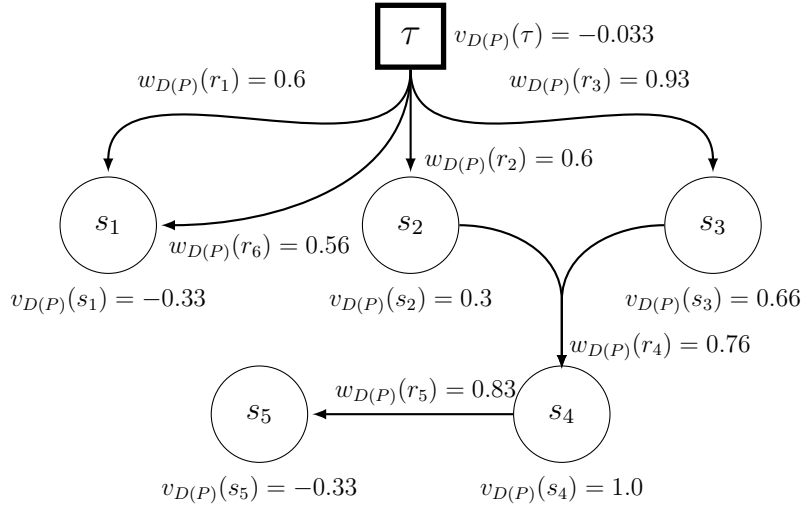


Figure 9: Direct function: aggregated valuations

Next, we define an aggregation function that only aggregates indirect opinions, and disregards direct opinions. It is thus the converse of the direct

aggregation function defined above. The aggregation of indirect opinions is based on the aggregation of the estimation functions introduced in section 4. Formally:

Definition 6.2 (Indirect aggregation). $\langle \mathcal{S}, \mathcal{R}, T \rangle$ be a *DRF* and $P = (O_1, \dots, O_n)$ an opinion profile over the *DRF*. The *indirect aggregation* of P over the *DRF* is defined as a function $I(P) = (v_{I(P)}, w_{I(P)})$, where $v_{I(P)}(s) = \frac{1}{n} \sum_{i=1}^n e_i(s)$, where e_i is an estimation function, and $w_{I(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r)$ for any statement $s \in \mathcal{S}$ and relationship $r \in \mathcal{R}$.

Notice that, while the direct aggregation function computes the average of individuals' direct opinions, the indirect aggregation function computes the average of individuals' indirect opinions. This is achieved by aggregating individuals' estimated opinions using an estimation function. We observe also that both functions calculate in the same way the aggregation of acceptance degrees. This will be the case for all the aggregation functions defined in this section, and hence the difference between them will be the way in which they aggregate valuations.

Example 6.2. Figure 10 shows the aggregated (collective) valuations obtained by the indirect aggregation function for the opinion profile of our running example shown in figure 5. For the sake of illustrating the computation, notice that $v_I(\tau) = (e_1(\tau) + e_2(\tau) + e_3(\tau))/3 = 0.079$ where $e_1(\tau) = (0.2 \cdot 0 + 0.5 \cdot 0 + 0.1 \cdot 0.7 + 1 \cdot 1)/1.8 = 0.59$, $e_2(\tau) = -0.257$, and $e_3(\tau) = -0.107$. The aggregated acceptance degrees are the same as in figure 9.

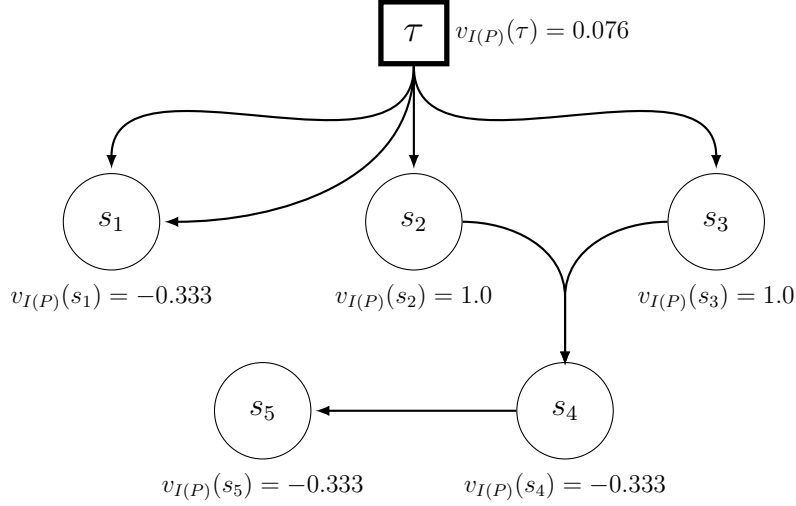


Figure 10: Indirect function: aggregated valuations

Having defined two extremes of our spectrum of functions, we introduce a family of aggregation functions based on a linear combination of the direct and indirect aggregation functions.

Definition 6.3 (α -Balanced aggregation). Let $\langle \mathcal{S}, \mathcal{R}, T \rangle$ be a *DRF* and $P = (O_1, \dots, O_n)$ an opinion profile over the *DRF*. Given the direct aggregation $D(P) = (v_{D(P)}, w_{D(P)})$, the indirect aggregation $I(P) = (v_{I(P)}, w_{I(P)})$, and $\alpha \in [0, 1]$, we define the aggregation function $B_\alpha(P) = (v_{B_\alpha(P)}, w_{B_\alpha(P)})$, where:

$$v_{B_\alpha(P)} = \alpha \cdot v_{D(P)} + (1 - \alpha) \cdot v_{I(P)}$$

$$w_{B_\alpha(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r)$$

for any statement $s \in \mathcal{S}$ and relationship $r \in \mathcal{R}$. We say that B_α is an α -balanced aggregation function.

By changing the value of α we set the importance of the direct opinion with respect to the indirect opinion. The functions resulting from definition 6.3 form a family of balanced aggregation functions: $\{B_\alpha\}_{\alpha \in [0, 1]}$. In particular, by setting α to 0 we obtain the indirect aggregation function, and by setting it to 1 we obtain the direct aggregation function.

Next, we define an aggregation function that exploits indirect opinions differently. For a given statement, the so-called *recursive aggregation* function calculates its aggregated valuation by using the collective opinion on its descendants, which in turn is recursively computed from their descendants, and so on. This recursive computing ends up when reaching statements without descendants whose indirect opinion is empty. Therefore, the recursive aggregation, unlike balanced aggregations, disregards individual valuations in the indirect opinion, and employs their collective opinions instead.

Definition 6.4 (Recursive aggregation). Let $\langle \mathcal{S}, \mathcal{R}, T \rangle$ be a *DRF* and $P = (O_1, \dots, O_n)$ an opinion profile over the *DRF*. The *recursive aggregation* of P over the *DRF* is defined as a function $R(P) = (v_{R(P)}, w_{R(P)})$, where

$$v_{R(P)}(s) = \begin{cases} \frac{1}{\sum_{r \in R^+(s)} w_{R(P)}(r)} \sum_{r \in R^+(s)} v_{R(P)}(s_r) \cdot w_{R(P)}(r) & \text{if } R^+(s) \neq \emptyset \\ v_{D(P)}(s) & \text{otherwise} \end{cases}$$

and $w_{R(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r)$ for any statement $s \in \mathcal{S}$ and relationship $r \in \mathcal{R}$.

Recall that $R^+(s)$ stands for the relationships connecting s to a descendant s_r of s through the relationship r .

Thus, the recursive function computes the average of the indirect collective opinion computed so far. In fact, we could say that, due to its recursive character, the function computes the estimated opinion for each statement in a bottom-up manner. Thus, the aggregation of opinions starts considering the direct opinions at the “leaves” of the debate, namely at the statements with no descendants, and moves up until reaching the targets.

Example 6.3. Figure 11 shows the aggregated (collective) valuations obtained by the recursive aggregation function for the opinion profile of our running example shown in figure 5. The aggregated acceptance degrees are the same as in figure 9. Again, for the sake of illustrating the computation, please notice that we start by computing $v_R(s_1) = v_D(s_1) = -0.33$ and $v_R(s_5) = v_D(s_5) = -0.33$ and, from these, we can compute $v_R(s_4) = v_R(s_5) \cdot w(r_5)/w(r_5) = -0.33$, $v_R(s_2) = v_R(s_4) \cdot w(r_4)/w(r_4) = -0.33 = v_R(s_3)$ so to finally compute $v_R(\tau) = (-0.33 \cdot 0.6 - 0.33 \cdot 0.56 - 0.33 \cdot 0.6 - 0.33 \cdot 0.93)/2.69 = -0.33$.

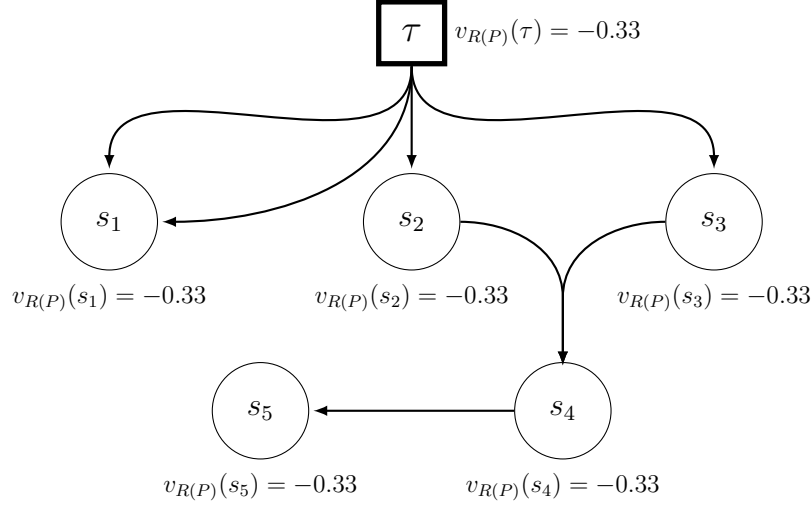


Figure 11: Recursive function: aggregated valuations

Similarly to what we did for the balanced function above, next we define a family of aggregation functions based on combining the direct and recursive aggregation functions.

Definition 6.5 (α -recursive aggregation). Let $\langle \mathcal{S}, \mathcal{R}, T \rangle$ be a *DRF* and $P = (O_1, \dots, O_n)$ an opinion profile over the *DRF*. Given the direct aggregation $D(P) = (v_{D(P)}, w_{D(P)})$, the recursive aggregation $R(P) = (v_{R(P)}, w_{R(P)})$, and $\alpha \in [0, 1]$, we define the aggregation function $R_\alpha(P) = (v_{R_\alpha(P)}, w_{R_\alpha(P)})$, where:

$$v_{R_\alpha(P)} = \alpha \cdot v_{D(P)} + (1 - \alpha) \cdot v_{R(P)}$$

$$w_{R_\alpha(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r)$$

for any statement $s \in \mathcal{S}$ and relationship $r \in \mathcal{R}$. We say that R_α is an α -recursive aggregation function.

7. Analysing opinion aggregation functions

In this section we compare the aggregation functions introduced in section 6 in terms of their satisfaction, or otherwise, of the social choice properties introduced in section 5. Our analysis will run along two dimensions: (1)

the *coherence* of an opinion profile; and (2) the *consensus* on the acceptance degrees of an opinion profile. Thus, we will consider whether agents' opinions are constrained to be coherent (the opinion profile is coherent) or not, and whether agents agree on acceptance degrees (there is consensus on acceptance degrees) or not. This results in four debate scenarios to analyse:

1. Unconstrained opinion profiles;
2. Constrained opinion profiles: assuming consensus on acceptance degrees;
3. Constrained opinion profiles: assuming coherent profiles; and
4. Constrained opinion profiles: assuming consensus on acceptance degrees and coherent profiles.

The analysis of these scenarios will help us assess the price that must be paid if the opinions stated by participating agents are not necessarily coherent. The scenarios will also help us assess that must be paid when the relationships between statements are opened for discussion by means of acceptance degrees.

Note that, in the interests of readability, we do not include the formal analysis in the body of the paper. Rather we present the results of that analysis. The formal analysis, both statements of results and the proofs of the results, can be found in Appendix A.

7.1. *Unconstrained opinion profiles*

This is the most general scenario we can consider for analysis. We assume unconstrained opinion profiles, which means that *any* opinion profile is deemed to be possible input for the aggregation functions introduced in Section 6. In other words, the domain of our aggregation functions is the class $\mathbb{O}(DRF)^n$ itself, and hence opinions need not to be coherent nor have consensus on the acceptance degrees.

Table 3 displays the social choice properties fulfilled by the functions defined in section 6 in this general case. There is one column per aggregation function and one row per social choice property. Notice that we distinguish between desirable social choice properties and other properties as discussed in section 5.2. In the table, a green square (with a tick) indicates that a property is fulfilled, while a red square (with a cross) indicates that a property is not fulfilled. As to the more general aggregation functions, α -B(alanced), and α -R(ecursive), in some cases we specify the values α for

which a given property holds. Notice that, for both families, we show the results considering $\alpha \in (0, 1)$, not considering 0 or 1, thus the cases for the extreme values represent aggregations functions already displayed in other columns.

Desirable properties	D	I	R	α -B	α -R
Collective coherence ★	✗	✗	✓	✗	$\alpha < \epsilon/2$
Exhaustive domain	✓	✓	✓	✓	✓
Coherent domain	✓	✓	✓	✓	✓
Anonymity	✓	✓	✓	✓	✓
Non-dictatorship	✓	✓	✓	✓	✓
Sided Unanimity	✓	✗	✗	✗	✗
Weak Unanimity	✓	✗	✗	$\alpha > 1/2$	$\alpha > 1/2$
Endorsed Unanimity	✗	✓	✗	$\alpha < 1/2$	✗
Familiar monotonicity	✓	✓	✗	✓	✗
Other properties					
Monotonicity	✓	✗	✗	✗	✗
Narrow Unanimity	✓	✗	✗	✗	✗
Independence	✓	✗	✗	✗	✗

Table 3: Social choice properties satisfied by aggregation functions D(irect), I(ndirect), R(ecursive), α -B(alanced), and α -R(ecursive) for: (i) a general scenario considering unconstrained opinion profiles; (ii) a scenario considering constrained opinion profiles: consensus on acceptance degrees.

Domain and anonymity. Table 3 shows that Exhaustive Domain (ED), Coherent Domain, Anonymity and Non-dictatorship are fulfilled by all the proposed opinion aggregation functions. This is because of the agnostic treatment of opinion profiles implemented by the aggregation functions. Since no constraints are imposed on opinion profiles received as input, ED is satisfied, and since no agent in an opinion profile receives a special treatment, Anonymity holds. Satisfying this family of properties is important. On the one hand, fulfilling ED ensures that any opinion profile is valid as input to a collective opinion, that is, the aggregation functions do not filter out participants' opinions prior to computing a collective opinion. On the other hand, satisfying Anonymity guarantees that all participants are equally important when calculating a collective opinion.

Collective coherence. The (D)irect and (I)ndirect functions do not satisfy collective coherence. As a result, neither does any α -Balanced aggregation function because it results from the linear combination of D and I. The result of such aggregation methods largely depends on the coherence of the opin-

ion profile at hand, which in this scenario can be as incoherent as possible. More positively, the (R)ecursive aggregation function does satisfy Collective Coherence (CC). Out of the family of recursive aggregation functions (α -R), which relies on D and R, those for which $\alpha < \epsilon/2$, where ϵ is set to assess the coherence of the output, also satisfy CC. This tells us that the closer α is to 0 (the less the use of the direct opinion), the more coherent the collective opinion obtained by an α -R function will be. The closer α is to $\epsilon/2$, the less coherent the collective opinion obtained by an α -R function will be. When α goes beyond $\epsilon/2$, the α -R function depends too much on the direct opinion (which does not satisfy CC) and CC does not hold.

Unanimity. Narrow, Sided and Weak Unanimity are not satisfied by the Indirect and Recursive aggregation functions. This is because the indirect opinion, employed by all these aggregation functions, ignores unanimity on the direct opinion of a statement and in some cases these functions can produce a result in the opposite direction. On the other hand, the Direct function, which only depends on the direct opinions of a statement, does satisfy all the unanimity properties. This benefits the Balanced and Recursive families, which satisfy Weak unanimity for some values of α . Notice that only Balanced and Recursive aggregation functions for which α is greater than $1/2$ satisfy Weak unanimity. This is to lessen the influence of the indirect opinion and sway the result towards the Direct aggregation function, which does satisfy the property. Regarding the Narrow and Sided unanimity properties, not even the influence of the Direct function is enough to guarantee that unanimity is preserved, and therefore no aggregation function in the Balanced or Recursive families fulfil them for any value of α .

On the other hand, regarding Endorsed unanimity, the situation changes for the Direct and Indirect functions. They flip sides so that the Direct function does not fulfill Endorsed unanimity, but the Indirect function does. This is because the unanimity in this case resides in indirect opinions, and hence it is in line with the Indirect function, which only depends on indirect opinions. However, this goes against the Direct function, which disregards indirect opinions, and hence unanimity on its values. Conversely to the Weak unanimity case for the Balanced family, now we need that the values of α are less than $1/2$ to sway the balanced aggregation towards the Indirect function, and hence, satisfy Endorsed unanimity. Next, although it might seem reasonable that aggregation functions in the Recursive family also fulfil Endorsed unanimity, they do not. This is caused by the recursive behaviour of these aggregation functions, which can overlook unanimity on indirect

opinions to use instead opinions deep in the debate on which there might be no unanimity. And last, due to the failure of the Direct and Recursive aggregation functions to fulfil Endorsed unanimity, so do all the aggregation functions in the Recursive family, no matter the value of α .

Monotonicity. The Familiar Monotonicity property is fulfilled by the Direct function (as a consequence of fulfilling Monotonicity), the Indirect function, and therefore by the whole family of Balanced functions that are combinations of the Direct and Indirect functions. The Recursive function, and hence the Recursive family, fails to satisfy Familiar Monotonicity because given a statement, the aggregated opinion about its descendants does not solely depend on the valuations on these descendants alone. Instead, the aggregated opinion about its descendants recursively depends on descendants down the relational framework. Thus, changes of opinion on "grandchildren" statements can cause a change of opinion independently of any change of the direct opinion.

Other properties. For the sake of completeness, next we also analyse the fulfilment of further, non-desirable, properties. As expected, Independence is not fulfilled by any of the function making use, at any degree, of indirect opinion, namely the Indirect function, the Recursive function, and the Balanced and Recursive families (for any $\alpha < 1$). Also related to the use of indirect opinions, we observe that Narrow Unanimity and Monotonicity are not fulfilled by any of these four functions. Overall, these properties can only be satisfied when employing the direct opinion alone to obtain the aggregated collective opinion. This reinforces our discussion in section 5.2 about disregarding these properties to consider alternative properties more orientated to aggregation functions that take account of dependencies.

7.2. *Constrained opinions: assuming consensus on acceptance degrees*

Assuming consensus on acceptances does not lead to any gain in the fulfilment of social choice properties with respect to those already claimed in section 7.1, so previous table 3 show the results for this scenario as well. Nonetheless, we deemed worth analysing this debate scenario because of the multiple already-in-use participation systems that do not allow participants to value differently the relationships between sentences in a debate. In our case, if we assume that participants agree on acceptance degrees, the collective opinion will only depend on valuations over sentences. This is equivalent to considering a debate where participants are allowed to value sentences, but do not express their opinions on the relationships between them.

We refer the reader to Appendix A for the proofs regarding the fulfilment of social choice properties in section 7.1. By analysing such proofs, we observe that assuming consensus on acceptances does not yield any further benefit that allows the fulfilment of any social choice property that did not hold in the analysis of section 7.1. This motivates that we have decided not to include any additional proofs for this scenario in A.2. Hence, the proofs in section 7.1 are also valid for this scenario.

7.3. *Constrained opinions: assuming coherent profiles*

In the following, we assume that the opinion profile in a debate is constrained to be coherent at some degree (according to some value $\epsilon \in (0, 1)$), so that each of the opinions in the profile is always coherent. Recall that we consider that coherence occurs when the direct and indirect opinions are in line. Therefore, assuming coherence is expected to have a positive impact on aggregation functions that exploit indirect opinions to compute a collective opinion. Here we assess the gain.

Table 4 shows the desirable properties satisfied by our aggregation functions when assuming coherence. The light green squares with check marks identify properties that are now satisfied, but were not (in table 3) when not imposing coherence. Therefore, assuming coherence yields new positive results. More precisely, table 4 shows that assuming coherence leads to the satisfaction of desirable unanimity properties for several of the functions. First, given the coherence assumption, the unanimity on the direct opinion drags the indirect opinion to become more similar to it, and therefore the Indirect function gains Weak unanimity. Now, since the Direct function also satisfies it, it follows that all α -Balanced functions now fulfil it too. Furthermore, thanks to the alignment that the coherence assumption brings between the direct and indirect opinions, the Direct function fulfils the Endorsed unanimity property. Therefore, having Endorsed unanimity fulfilled now by the Indirect and Direct functions, the aggregation functions in the Balanced family also fulfil it for any α .

Observe that unanimity and the coherence assumption work well together. Unanimity on one sentence brings together its direct and indirect opinions, making it impossible for both to be far apart, and therefore allowing the Direct and Indirect functions to fulfill more unanimity properties.

Finally, the family of Recursive function now fulfils Endorsed unanimity, though, not for any α . Depending on the degree of coherence allowed in the opinion profile, i.e. the value of ϵ , the interval of α values allowing R_α to

fulfil Endorsed unanimity will change. In this case, α has to be greater than $1/(2 - \epsilon)$, representing the need to overcome the bad result obtained by the Recursive function with respect to the Endorsed unanimity property.

Desirable properties	D	I	R	α -B	α -R
Collective coherence ★	✗	✗	✓	✗	$\alpha < \epsilon/2$
Exhaustive domain	✓	✓	✓	✓	✓
Coherent domain	✓	✓	✓	✓	✓
Anonymity	✓	✓	✓	✓	✓
Non-dictatorship	✓	✓	✓	✓	✓
Sided Unanimity	✓	✗	✗	✗	✗
Weak Unanimity	✓	✓	✗	✓	$\alpha > 1/2$
Endorsed Unanimity	✓	✓	✗	✓	$\alpha > \frac{1}{2-\epsilon}$
Familiar monotonicity	✓	✓	✗	✓	✗

Table 4: Highlighted, in light color, the fulfilment of additional desirable properties, on top of those in Table 3, when assuming coherent opinions.

7.4. Constrained opinions: assuming consensus on acceptance degrees and coherent profiles

In what follows we assume both previous constraints on the opinion profiles: coherence on the opinions and consensus on acceptance degrees. First, consensus on acceptance degrees on relationships represents a more simplified debate where participants only provide their opinions on sentences. Second, the coherence assumed on opinions aligns direct and indirect opinions. Overall, both assumptions yield major benefits in terms of the satisfaction of desired social choice properties, as we discuss next.

Table 5 shows the gain in fulfilment of desirable properties with respect to table 4. The light green squares with check marks identify properties that are now satisfied, but were not (in table 4) when not imposing consensus on acceptance degrees. Now, besides the aggregation functions in the Recursive family, which now satisfies Collective coherence for any α , the rest of aggregation functions under study also satisfy ϵ -Collective coherence when the opinion profiles are δ -coherent, $0 < \delta \leq \epsilon$.

This major improvement is because the consensus on acceptance degrees forbids the participants to value a relationship as 0, which is key to ensure collective coherence for the Direct and Indirect functions when the opinion profiles are coherent. We assume that for each relationship at least one agent

has valued it other than 0, because otherwise it would be as if the relationship did not exist, and this forces all the participants to have a positive value too.

In this manner, if we have both ϵ -coherent opinions and consensus, then all our aggregation functions can guarantee ϵ -coherent aggregated opinions, which may increase the acceptability of the results from the participants.

Notice that assuming consensus on acceptance degrees is quite reasonable. Indeed, we believe that such consensus is actually very likely to be found in many debates where the relationships are classified first and afterwards the participants are allowed to give their opinions. Our model can fit perfectly with these kinds of scenario by setting all the acceptance degrees as a constant value for every participant. Furthermore, the procedure to create the debate, and therefore the DRF, could be adapted so there is a first stage in which a collective value is established for each of the relationships, and then there is a second stage in which values are assigned to the statements.

Desirable properties	D	I	R	α -B	α -R
Collective coherence ★	✓	✓	✓	✓	✓
Exhaustive domain	✓	✓	✓	✓	✓
Coherent domain	✓	✓	✓	✓	✓
Anonymity	✓	✓	✓	✓	✓
Non-dictatorship	✓	✓	✓	✓	✓
Sided Unanimity	✓	✗	✗	✗	✗
Weak Unanimity	✓	✓	✗	✓	$\alpha > 1/2$
Endorsed Unanimity	✓	✓	✗	✓	$\alpha > \frac{1}{2-\epsilon}$
Familiar monotonicity	✓	✓	✗	✓	✗

Table 5: Highlighted, in light color, the fulfilment of additional desirable properties, on top of those in Table 4, when assuming coherent opinions and consensus on acceptance degrees.

7.5. Summary

From the analysis for each debate scenario above, we can draw the following general observations:

- The aggregation functions of the recursive family achieve collective coherence provided that they place little weight on direct opinions (or opinions are coherent and there is consensus on acceptance degrees).
- Coherence in opinion profiles favours unanimity (specifically, WU and EU), though in different ways. I and α -Balanced are fully satisfied,

while the family of recursive functions leans on the direct aggregation function to fulfil some unanimity properties with restrictions. As a result, the α -Recursive family only satisfy WU and EU under strong conditions on α , because the R function never satisfies them.

- Coherent opinion profiles are not enough for D, I, and α -Balanced functions to achieve collective coherence. They also require consensus on acceptance degrees. Recursive functions do not require such consensus (in fact, not even the non-coherent opinion profiles), and hence they are *robust* to the divergence of opinions on the relations between statements in a debate.
- While the D, I, B_α functions manage to achieve familiar monotonicity in all scenarios, the aggregation functions in the recursive family cannot even when counting on coherent opinion profiles and consensus on acceptance degrees. This is because the aggregated opinion on descendants recursively depends on descendants down the relational model. Thus, changes of opinion on “grandchildren” or further down sentences can cause a change of opinion independently of any change of the direct opinion.

Based on these general observations above, it is the task of the decision maker to decide the aggregation operator to choose considering: (1) the features of the debate scenario at hand; and (2) the desirable properties to guarantee. As a rule of thumb, since in actual-world debates we cannot assume individual rationality (coherence), we believe that recursive aggregation functions are the best choice to achieve collective rationality, though we would pay the price of losing some other valuable properties, in particular unanimity for values of α that promote a large use of the direct opinion. Otherwise, if we do not value the coherence of the collective output, or we can guarantee somehow that the opinions of participants are coherent and the participatory system at hand does not allow for divergence on acceptances degrees, the Direct function becomes the aggregation function of choice. Within such constrained settings, the Direct aggregation function fulfils almost every property considered, enve all of them in the debate scenario in section 7.4. We conclude that it seems a good trade-off to consider the Recursive family, which can behave as similar to the Direct or to the Recursive function as wanted, and set the value of α depending on the features and goals in hand.

8. Computational Analysis

The purpose of this section is twofold. First, given the opinion aggregation problem in section 5.1, we explain the complexity of the different algorithms for computing a collective decision on its target. In particular, we provide an algorithm for computing the recursive aggregation function. Thereafter, in section 8.2, we empirically analyse the use of that algorithm to solve real-world collective decision problems.

8.1. Computing Aggregation Functions

Algorithm 1 Compute recursive aggregation

```

1: function COMPUTE_RECURSIVE_AGGREGATION( $\langle \mathcal{S}, \mathcal{R}, T \rangle, (O_1, \dots, O_n)$ )
2:   for each relationship  $r \in \mathcal{R}$  do ▷ Compute averaged acceptances
3:     aggregated_acceptance[r] ← average_acceptances( $w_1(r), \dots, w_n(r)$ )
4:      $\mathcal{H}(\langle \mathcal{S}, \mathcal{R}, T \rangle) \leftarrow \text{DRF\_to\_B\_hypergraph}(\langle \mathcal{S}, \mathcal{R}, T \rangle)$  ▷ Generate B-hypergraph
       representation of DRF
5:     sorted_sentences ← reverse(topological_sorting( $\mathcal{H}(\langle \mathcal{S}, \mathcal{R}, T \rangle)$ )) ▷ Compute
       topological sorting of the DRF B-hypergraph
6:     for  $s$  in sorted_sentences do ▷ Compute aggregated valuations
7:       valuation[s] ← 0 ▷ To accumulate aggregated valuations over descendants
8:       normaliser[s] ← 0 ▷ To normalise aggregated valuations over descendants
9:       compute_relationships  $R(s)$  to descendants
10:      if  $R(s) \neq \emptyset$  then ▷ if  $s$  has descendants
11:        for each relationship  $r \in R(s)$  do
12:           $s_r \leftarrow$  descendant from relationship  $r$ 
13:          valuation[s] ← valuation[s] + aggregated_valuation[ $s_r$ ] · aggregated_acceptance[r]
14:          normaliser[s] ← normaliser[s] + aggregated_acceptance[r]
15:          valuation[s] ← valuation[s] / normaliser[s]
16:        else ▷  $s$  has no descendants
17:          valuation[s] ← average_valuations( $v_1(s), \dots, v_n(s)$ )
18:          aggregated_valuation[s] ← valuation[s]
19:  return aggregated_valuation, aggregated_acceptance

```

All of the aggregation functions proposed in Section 6 can be calculated by tractable algorithms. For example, the direct function calculates the average for all statements and relationships in a *DRF* $\langle \mathcal{S}, \mathcal{R}, T \rangle$ considering the direct opinions in an opinion profile $P = (O_1, \dots, O_n)$. Hence, its complexity is given by $\mathbf{O}((|\mathcal{R}| + |\mathcal{S}|) \times |P|)$, where $|\mathcal{R}|, |\mathcal{S}|$ are the number of relationships

and statements, respectively; and $|P|$ is the number of opinions in an opinion profile. Computing the indirect and balanced functions can be done by calculating the aggregated acceptance of each relationship as an average and by calculating the aggregated valuation of each statement as the average of the estimation function, which in turn is an average of the indirect opinions for that statement. Hence, their complexity is given by $\mathbf{O}(|\mathcal{R}| \times |\mathcal{S}| \times |P|)$. The calculation of the recursive function can be done by calculating the aggregated acceptance of each relationship as an average and calculating the aggregated valuation of each statement by starting with statements with no descendants and using these results to calculate the aggregated valuation of the statements directly connected to them. For example, algorithm 1 contains the pseudocode for the recursive function. In particular, the algorithm starts by computing aggregated acceptances ($w_{R(P)}$) as a weighted average (lines 2-3), which has a complexity of $\mathbf{O}(|\mathcal{R}| \times |P|)$. Then, the algorithm computes aggregated valuations ($v_{R(P)}$) starting from the statements with no descendants. In order to do that, we first perform a topological sorting of the DRF. This can be achieved by: (1) transforming the graph associated to the DRF into an acyclic B-hypergraph¹⁰, denoted by $\mathcal{H}(\langle S, \mathcal{R}, T \rangle)$ (line 4); and (2) then performing the topological sorting over the B-hypergraph (line 5). Starting from the sentences without descendants, the algorithm computes aggregated valuations until reaching the statements in T (lines 5-18). The calculation of the topological sorting for B-hypergraphs has been studied in [44]¹¹ and has a complexity of $\mathbf{O}(|\mathcal{R}| \times |\mathcal{S}|)$. The calculation of the aggregated valuations has a complexity of $\mathbf{O}(|\mathcal{S}| \times \max(|\mathcal{R}|, |P|))$. In

¹⁰B-hypergraphs are a particular type of hypergraph with efficient algorithms for path finding[44]. Obtaining a B-hypergraph from a DRF is straightforward. In fact, the graph associated with a DRF is a B-hypergraph with the exception of the relationships that connect the very same statements. For instance, consider relationships r_1 and r_6 in figure 4 linking τ to s_1 . Since in a hypergraph there cannot be two or more hyperedges over the very same nodes, we will only consider one single hyperedge. In our example, it suffices to consider either r_1 or r_6 . We do not lose anything by doing this simplification because we want to obtain the topological sorting of a DRF, and hence considering one of the relationships connecting the very same statements is enough.

¹¹In particular, [44] provides an algorithm to calculate the inverse topological sorting in a F-hypergraph. Any given B-hypergraph can be transformed into a symmetric F-hypergraph by changing the direction of the hyperedges. Note the inverse topological sorting of the symmetric F-hypergraph coincides with the topological sorting in the original B-hypergraph.

real debates the number of opinions is usually significantly higher than the number of relationships, hence, the complexity of the calculation of aggregated valuations is given by $\mathbf{O}(|\mathcal{S}| \times |P|)$. Thus, the total complexity of the recursive functions is $\mathbf{O}(|\mathcal{R}| \times |P|)$.

In <https://bitbucket.org/jariiii/workspace/projects/DRF> we provide a publicly-available implementation of algorithm 1 together with all the aggregation functions defined in this paper. Furthermore, we also provide guidelines to reproduce the experiments reported in section 8.2 below.

8.2. Empirical Analysis

In what follows we empirically analyse the time required by our implementation of algorithm 1 to compute collective decisions. Our purpose is to investigate whether our approach can handle collective decision making in practice.

Based on the analysis above, we generated debates, which required to the synthetic generation of DRFs and opinion profiles. On the one hand, we artificially generated DRFs whose sentences are the nodes of a directed acyclic B-hypergraph and whose hyperedges represent the relationships between sentences. We chose the number of sentences in our synthetic DRFs within $\{100, 150, 200\}$ to represent small, medium and large scenarios. Regarding the relationships between sentences, we considered two parameters:

- *Density of relationships.* Given a relationship $r = (\Sigma, s)$, we say that the number of sentences in Σ is the density of r . Since each relationship is represented as a hyperedge in a B-hypergraph, the density of relationships amounts to the size of the tails of hyperedges in the hypergraph. The (average) density of relationships in our artificial DRFs took values within $\{1, 2, 3\}$. We set the density value to 1 to generate DRFs for which there is a one-to-one connection between sentences, and so each DRF is in fact a DAG. As to the other two density values (2 and 3), they allow us to generate DRFs where each relationship has two sentences connected to one sentence, and three sentences connected to one sentence respectively.
- *Density of number of relationships,* namely the average number of sentences to which each sentence is connected to through relationships. This corresponds to the average out degree of each sentence in the DRF. We chose values for this parameter within $\{1, 2.5, 5\}$ to generate DRFs with low, medium, and high density of relationships.

On the other hand, to finish generating a debate, we must generate opinion profiles. We generated profiles with number of opinions within $\{10^6, 3 \cdot 10^6, 5 \cdot 10^6\}$, to represent the largest known actual-world scenarios¹². The values for valuations and acceptances were randomly generated within $[-1, 1]$ and $[0, 1]$ respectively.

All the computations of collective decisions for our artificially generated debates were performed on an Ubuntu 16.04 box with an Intel(R) Core(TM) i7-8700K CPU @ 3.70GHz, with 31GiB system memory, and 8th Gen Core Processor Host Bridge/DRAM R. Furthermore, our experiments only considered the recursive aggregation function (specified in algorithm 1) because it is the more expensive out of those introduced in section 5.

We performed three types of analysis:

- *Sensitivity to number of participants.* Figure 12 shows that the time to compute collective decisions increases as the number of participants increases. The figure shows the results for a medium density of number of relationships and a low density of relationships (which is the most expensive case as we discuss below).
- *Sensitivity to density of number of relationships.* Figure 13 shows that the time to compute collective decisions increases as the density of number of relationships grows. Notice though most actual-world scenarios would lie between the low and medium cases, and hence it would take less than one second to solve even the largest debate.
- *Sensitivity to density of relationships.* Figure 14 shows that the time to compute collective decisions decreases as the density of relationships increases. The figure shows the results for $3 \cdot 10^6$ opinions and a medium density of number of relationships. This tells you that in fact, our

¹²To the best of our knowledge, the Brexit discussion on UK [71] constitutes the largest such discussion: news outlets reported when the number of supporters passed 2 million [13] and the numbers kept growing during the 6 month period that the discussion was open. By the time it closed, there were 6,103,056 participants [71]. Contrasting numbers of participants can be found for other popular initiatives such as an environmental proposal in Parlement et Citoyens which had 51,493 votes [68] and in the participatory budgeting process in Helsinki [25] with 54,246 registered people, which represents 10% of the city voters. Note that the Parlement et Citoyens and Helsinki debates are probably more representative of real online debates than the Brexit example, where participants were, in effect, just voting on a specific proposal.

algorithm needs more time when relationships between sentences are one-to-one.

Overall, notice that computing collective decisions in all the artificially generated debates took less than 1.6 seconds. Therefore, we can conclude that our opinion aggregation functions can be employed to cope with large-scale debates in real time.

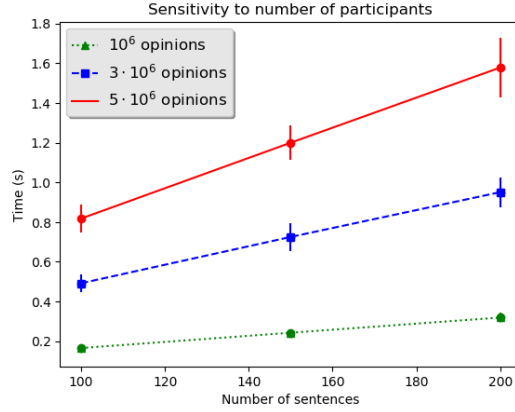


Figure 12: Sensitivity to number of participants in a DRF. The results show the computational time required by the recursive aggregation function as the number of sentences grow.




Figure 13: Sensitivity to density of relationships in a DRF. Density is considered in terms of the average out degree of sentences: low (2.0), medium (5.0), high (10.0). The results show the computational time required by the recursive aggregation function for $3 \cdot 10^6$ opinions as the number of sentences grow.

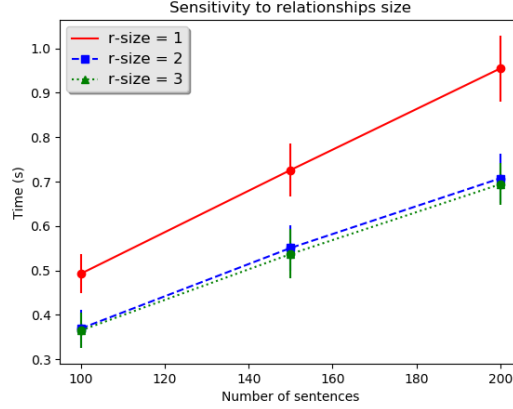


Figure 14: Sensitivity to the number of sentences in the source of a relationship (size of the tail of hypergraph edges). The results show the computational time required by the recursive aggregation function as the number of sentences grow and as the size of hyperedges (tails) grows.

9. Related Work

In this section we review the work from the literature that is closest to the work in this paper. That includes work which has focused on support for online discussions, on approaches for computing the outcome of a set of arguments, and on analysing the behaviour of discussions from the standpoint of social choice theory. We discuss work in these three areas separately below, acknowledging that some work could validly be considered under more than one heading.

9.1. Tools for online discussion

As mentioned above, we see this work as being inspired by work on online discussion forums such as Decidim Barcelona [30], Better Reykjavík [16], Decide Madrid [29], and Parlement et Citoyens [69], where participants can carry out a structured discussion around some topic, typically a policy proposal. These particular sites allow participants to offer arguments for and against a proposal, and vote(/support) for it in the context of a specific institution. However, some of the citizen participation tools behind some of these sites, such as Decidim [31] for Barcelona or Consul [27] for Madrid, have been used in other cities and organisations. Thus, for example, City of Helsinki [25] applies an instance of Decidim and the New York City Participatory Budgeting [67] is based on Consul. Other participatory tools have

also proliferated outside the context of public institutions. There, we find: `consider.it` [26], Appgree [4], and Baoqu [9] whose main focus is scalability – making the systems fit for use by large numbers of participants –; Loomio [50, 59], where participants can both comment on proposals, albeit in an unstructured way, and also vote on them; or Kialo [51], which organises debates in an structured way. What distinguishes our work from all of these approaches is that, as discussed in Section 2, we provide a much more expressive framework than any existing participation system. First, the DRF makes it possible to express opinions about both statements and relationships between statements, and, similarly to [37, 56, 12], it allows opinions to be expressed using real values, unlike most e-participation systems, which limit to allowing users to express agreement or disagreement with arguments. Second, our framework has been conceived to support relationships involving multiple statements, hence going beyond the limiting pairwise relationships supported by current e-participation systems.

There is also a long-standing line of work which develops tools to map the structure of arguments on some topic. This line of work draws from a range of sources, nicely summarized in [82], and is exemplified by [21, 79, 84, 87]. The focus of this work is on drawing the relationships between arguments as a means of helping people understand the scenarios, and as [14, 49, 80] point out, the resulting maps can be used to support group decision-making. However, all these approaches deal only with visual *representation* of the arguments — there is no attempt to compute a summary of the discussion. In contrast, our focus is on using the results of debate as input to a computational process, rather than providing support for the debate itself. In that sense our work could be viewed as a post-processing stage that could be applied in conjunction with any of the tools listed above to support structured discussion. Indeed, the generality of the relational model means that it is well-fitted to such a task.

There are other approaches that allow for structured argument-based discussions and aim to compute the outcome of the discussion. One notable body of research here is Klein’s work on the Deliberatorium [52, 53] which allows for the presentation of arguments and their interactions, and aggregates the opinions. Unlike our work, there is no analysis of the properties of the aggregation with respect to social choice principles.

In contrast to the work discussed so far, where participants in the debate have the task of structuring their arguments into the correct format, [19] considers extracting arguments from natural language texts and constructing

a formal argumentation representation from them. Such a representation can then be summarised as discussed in [78]. [12] discusses how this kind of approach can be combined with the argumentation-visualization methods described above.

9.2. Computational argumentation

Computational argumentation [76] has a lengthy history within artificial intelligence, going back at least as far as [40, 63].

At the time of writing, work in this area is split into two broad groups. First, historically, is work which is concerned with the internal structure of arguments — what arguments are constructed from, and how this construction takes place. Early examples of this work include those cited above, along with [41] and [54], [60], [70], and [74]. This line of work has reached its current endpoint with *structured argumentation* systems like logic-based argumentation [15], assumption-based argumentation [36] and structured argumentation systems such as ASPIC+ [65], and DeLP [47]. Second is the line of work on *abstract argumentation*, begun by [35], which focuses much less on the internal structure of arguments, and instead is mainly concerned with the relationships between arguments. This has led to a large body of work expanding on [35], for example [10, 64, 89]. In [35], the focus is solely on “attack” relations, where arguments are in conflict, and subsequent work has expanded the scope to consider “support” relations as well [3, 22, 23].

There is another way to, broadly, classify work on argumentation into two groups. One line of work, again exemplified by [35], focuses on argumentation as a mechanism for extracting consistent points of view from an inconsistent knowledge base. The other line of work deals with how arguments combine, or accrue, in favour of, or against some conclusion. This distinction cuts across the structured/abstract distinction with, for example [10] being concerned with consistency in abstract argumentation, and [65] dealing with consistency in structured argumentation. On the other hand, [15, 73, 88] discuss accrual in structured argumentation, while [22] looks at accrual in abstract argumentation.

The work mentioned above uses argumentation as a mechanism for a single entity to come to a conclusion. However, as [85, 90] and others have pointed out, argumentation is also a natural mechanism for multiple entities to use to reach consensus on some topic. As a result, argumentation has been used [2, 62] in multiagent systems as a mechanism for *rational interaction* [61] for a particular meaning of “rational”. That is “rational” in the sense

that each stage in the interaction is supported by well-founded reasons. Here we build upon this prior work in rational interaction. Our approach allows participants in a discussion to put forward opinions about some topic under discussion, opinions that may be either in favour or against the topic.

Our work connects to several of these themes in argumentation. First, since we are interested in aggregating the opinion from a number of participants, our work is clearly related to the use of argumentation in multiagent interaction. Second, the fact that individual steps in the participants’ reasoning process are represented in our approach means that our work is connected to work on structured argumentation. (We would argue that it is more abstract, since the relationships that connect statements are not restricted to be rules as in [36, 47, 65]). Third, our work connects with the idea of argumentation as a means of extracting a coherent view from a number of conflicting opinions. The fact that this coherent view takes into account the votes of participants also gives our work a fourth connection with the argumentation literature, in its relation to work on accrual. This voting aspect also places our work in close relation to that of social argumentation [56, 75], and previous work on collective argumentation¹³ [7, 20, 46] which we will discuss more below. Finally, one might view our work as being about the combination of different sets of arguments, one for each person who votes on the arguments or the relationships between them. From that perspective, our work also connects with that of [28], which takes as input different sets of arguments and relationships between them, and outputs consistent sets of arguments, thus “merging” the input sets. See [18] for a survey of work on this topic, and [24] for an excellent short overview of developments to date.

9.3. *Social choice theory*

Given a set of alternatives and a set of agents who possess preference relations over the alternatives, social choice theory focuses on how to yield a collective choice that appropriately reflects the agents’ individual preferences [8]. With this aim, social choice theory has extensively explored many ways of aggregating agents’ individual preferences [43]. Since there is a consensus in the literature on the desirable properties that a “fair” way of aggregating preferences should satisfy (e.g. no single agent can impose their view on the

¹³Note that our work, and [7, 20, 46] has little commonality with the “collective argumentation” of [17], which is concerned with argumentation in which relationships exist between sets of arguments.

aggregate; if all agents agree, the aggregate must reflect the agreement; and so on [43]), aggregation functions can be characterised and compared in terms of the desirable properties they satisfy. Notice though that social choice theory counts on multiple *negative* results, namely impossibility results showing the incompatibility of certain sets of desirable properties (e.g. Arrow’s famous impossibility theorem [5]).

Much of the work in social choice theory has placed little emphasis on the structure of the objects over which agents are expressing their preferences. However, there is a growing body of research that takes the subject of the preferences to be arguments in some form or other. Here the foundational work was that of [77], later developed in [7], which considered — from one perspective — the very same problem that we tackle here: given a topic of discussion and a set of agents expressing their individual opinions about the statements made in the discussion, how can the agents reach a collectively rational decision?. The way that this is tackled in [7], as [18] explains, is as a version of the “merging” problem mentioned above. That is, [7] consider each participant to have a set of arguments, and the relationships between them, and an opinion about which arguments are labelled as being acceptable and which are not. The problem they then solve is how to compute a set of labels for the arguments that reflect the opinions of all the participants such that the aggregation of opinions satisfies desirable social choice properties. The same problem was, subsequent to [77] but before the extended treatment in [7], considered by [20], and more recently by us in [46]. As [6] points out, [7] and [20] take different approaches, with [7] considering the opinions as votes, resolved by taking the plurality for individual arguments, while [20] offers a range of operators that yield a labelling which confirms to the constraints of argumentation semantics¹⁴ while also not disagreeing with the opinions of any participant. [6] compares the plurality approach with one of the operators [20] using human participants.

The recent work in [24] can be viewed as an extension of the line of work in [7, 20]. Like [7, 20], [24] proposes methods for aggregating a collection of individual argumentation frameworks, corresponding to participants in a debate, into a single argumentation framework that appropriately reflects the views of the group as a whole. [24] investigates the properties of the aggregation rules introduced in the paper, and, like ours, the work employs

¹⁴The resulting labellings are either admissible or complete [35].

techniques from social choice theory in this analysis. However, the aim is different. In [24] the aim is to analyse aggregation rules in terms of their preservation of *semantic properties* of argumentation framework, while our focus in this paper is on *social choice properties* of the aggregation operators.

From the perspective of argumentation, the major difference between [7, 20, 24] and our work in [46] and this paper is that we do not start from a set of opinions that are well-formed in an argumentation sense. [46], which does directly use the notion of arguments and labels (though the work is based on a variation on bipolar argumentation [3]) does not, unlike [7, 20], start from a “legal” [11] labelling. Similarly here we do not require any consistency of opinions. In both cases this is because we want to represent human opinions that may not be rational in an argumentation-theoretic sense.

Finally, a rather different line of work (and one that seems to have no knowledge of [7, 20, 24]) is that of [75], which extends the QuAD framework first introduced in [12] to deal with multiple agent debates. Similarly to our work in [45, 46], the QuAD-V framework in [75] allows pro and con arguments (attackers and defenders in our terminology) and agents’ votes over arguments (labels). However, the main focus of the work is not on computing a collective opinion, but on the agents contribute with individually rational opinions, a weaker version of our notion of coherent labelling defined in [46]. This contrasts with our focus on the design of aggregation functions that satisfy desirable social choice properties, particularly collective rationality (strict rationality in Rago and Toni’s terms), without requiring agents to be individually rational.

Beyond that, the major difference between the work in this paper and what has been done before is the richness of the representation. Here there are three main extensions. First, [7, 20, 24, 46] and [75] all deal with abstract arguments. Here we deal with structured objects, and, as already mentioned, these are objects that are more general than the usual object studied in structured argumentation since we place no real constraints on the kind of reasoning captured by the relationships that hold between statements. Second, unlike previous work, we allow opinions to be expressed both about individual statements and the relations between them. This combines what has been studied in [7, 20, 24, 46, 75], where opinions are expressed about individual arguments, but not the relations between them (these are assumed to be fixed), and [56], where opinions are expressed about the relationships between arguments but not about the arguments themselves. Third, we allow opinions to be real-valued. In this we move away from the labellings

studied in [7, 20, 46], and their grounding in argumentation semantics, and towards the kind of representation allowed in [56, 75].

Finally, following the assumptions already established in [46], we neither assume independence between arguments as a fundamental postulate as is the case in [7, 20, 24], nor do we require the resulting aggregation to agree with an argumentation semantics. We will deal with these differences in turn.

Dropping the independence of arguments should come as no surprise since [7] questions the necessity of assuming independence because of the dependencies between arguments that come already encoded in the form of relationships such as attack. Despite the importance of independence as a fundamental property in the judgement aggregation literature because of its theoretical value in proving strategy-proofness and strategic manipulation¹⁵, we are not alone in regarding independence as too strong a property. This is because together with mild further conditions, it implies dictatorship [55] and because it is also considered as not very plausible [66]. This explains why relaxing independence has been subject of much research (e.g. [34, 55, 66, 72]). This paper goes beyond relaxing independence. Rather, in this paper we introduce several opinion aggregation functions that use the participants’ opinions to compute a collective opinion while considering *dependencies* between statements. This is in line with our former work in [46], but here we allow to express dependencies between multiple statements per relationship, while [46] (and all the work which starts with abstract argumentation frameworks such as [7, 20, 24, 75]) constrain relationships, and thus dependencies, to exist between pairs of arguments, hence limiting expressiveness.

Turning to the fact that the result of our aggregations do not match an standard argumentation semantics, we substitute the notion of “coherence of opinions” for the form of rationality embodied by those semantics. We do this for reasons that we have already touched on above with respect to the input opinions — we feel that insisting on an output that is rational in an argumentation theoretic sense is not necessarily realistic given that we start from opinions that are put forward by human participants who may not be not consistent in their views. Instead of forcing the output of aggregation to be rational in an argumentation theoretic sense, we instead compute a measure which assesses how much concordance there is between

¹⁵If the independence criterion is not satisfied, then the function aggregating judgements is not immune to strategic manipulation [33].

related opinions, and assess our novel aggregation functions by the degree to which they can assure that their output ensures collective coherence.

Finally, from a pure social choice (not a combined argumentation and social choice) perspective notice, that it is common in the literature on judgement aggregation and preference aggregation to impose properties on the objects under aggregation in order that aggregation operators can guarantee desirable properties. For instance, in the case of distance-based aggregators, the *Kemeny* rule [38] only considers consistent judgement sets, and hence disregards those which are not, whereas premise-based aggregators [38] typically make assumptions on the agenda to guarantee consistency and completeness. Our work does not rely on this structuring of the target objects of the aggregation operators. Instead, we have introduced aggregation operators capable of guaranteeing collective coherence when opinions are unconstrained. This is motivated by the need for disregarding rationality when humans are involved in debates, since their opinions may eventually contain contradictions and inconsistencies.

10. Conclusions and future work

There is an emerging interest in the use of information and communication systems to allow citizens to participate in the governance process. These e-participation systems aim to strengthen the collaboration between governments and citizens, and to empower citizens in participating in policy decision-making. The majority of such systems allow citizens to express their views on particular issues, but they either fail to aggregate the different views in a meaningful way, or they aggregate in a way that limits users to just voting in favour of or against an issue. Current research in the areas of social choice and argumentation could help to formally represent richer human debates in e-participation systems and to define computational processes for aggregating users views. However, approaches belonging to these areas have two main drawbacks that significantly decrease their potential as a solution for real e-participation systems. Firstly, they assume that participants agree on the structure of a debate, that is the pieces of information that are relevant to particular issues and the relationships among them. Secondly, they assume that users opinions are rational, which is too restrictive for human participants.

To address these limitations, we have proposed a model to represent and reason about human debates. Our model allows users not only to introduce

new pieces of information to the discussion and relate them to existing pieces but also to express disagreement with the relationships put forward by other users. In addition, our model does not assume that users opinions are rational and we define a weaker notion of rationality to characterise *coherent* user opinions. Considering the degree of coherence of individual users opinions and the level of consensus that users have about the debate structure among users, we analyse the outcomes of different opinion aggregation functions in terms of social choice properties. Our analysis demonstrates that the recursive aggregation method could compute a coherent collective opinion even if individual opinions are incoherent and there is a lack of consensus on the debate structure. As we impose more restrictions on the coherence of individual opinions and consensus among users on the debate structure, more aggregation methods also compute coherent collective opinions. We conclude our analysis with a computational evaluation in which we study the computational cost of aggregating collective opinions and experimentally demonstrate that collective opinions can be computed efficiently for real-sized debates.

There are several lines of future work that we have in mind. On the theoretical side, we are investigating methods to analyse the quality of a debate represented in the relational model. For example, we are studying the use of systematic incoherence in participant opinions as a way to identify structural problems in a debate. On the practical side, we are planning to evaluate the practical impact of using the relational model in real online debates. We have already evaluated [81] an earlier debate model [46] using data from Decidim Barcelona, and intend to do the same for the relational model.

Acknowledgements

This work was supported by the UK Engineering & Physical Sciences Research Council (EPSRC) under grant #EP/P010105/1 and by projects LOGISTAR (H2020-769142), AI4EU (H2020-825619), Crowd4SDG (H2020-872944), CI-SUSTAIN (PID2019-104156GB-I00), PGC2018-096212-B-C33, 2017 SGR 172, and 2017 SGR 00341. Jordi Ganzer was supported by a PhD studentship from King's College London.

Appendix A. Formal proofs and results

In the following, we prove all the formal results presented in section 7 regarding the satisfaction of social choice properties by the opinion aggregation functions introduced in section 6. The section is divided into four parts, one per debate scenario as analysed in section 7.

1. Unconstrained opinion profiles;
2. Constrained opinion profiles: assuming consensus on acceptance degrees;
3. Constrained opinion profiles: assuming coherent profiles; and
4. Constrained opinion profile: assuming consensus on acceptance degrees and coherent profiles.

Furthermore, for each scenario, our results will be grouped by aggregation function according to the following order: Direct aggregation, Indirect aggregation, Recursive aggregation, Balanced family aggregation and Recursive family aggregation

A.1. Unconstrained opinion profiles

In this section we analyse the social choice properties fulfilled by the aggregation functions introduced in Section 6: assuming unconstrained opinions profiles (*any* opinion profile is deemed to be possible input for the aggregation functions). The results of this section are summarised in table 3 in section 7.1.

Proposition A.1. The aggregation function D satisfies the following properties:

- (i) Exhaustive Domain and Coherent Domain;
- (ii) Anonymity and Non-Dictatorship;
- (iii) Monotonicity and Familiar Monotonicity;
- (iv) Narrow Unanimity, Sided Unanimity and Weak Unanimity; and
- (v) Independence.

And does not satisfy:

- (vi) Collective coherence; and
- (vii) Endorsed Unanimity.

Proof. (of proposition A.1)

- (i) Exhaustive Domain is straightforward and Collective Domain follows directly.
- (ii) Anonymity and Non-Dictatorship. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a DRF and σ a permutation over a set of agents $Ag = \{1, \dots, n\}$. We must show that D maintains the same collective opinion over the permuted opinion profile $P' = (O_{\sigma(1)}, \dots, O_{\sigma(n)})$, i.e. that $D(P) = D(P')$. This is the case because the next two equalities hold:

$$v_{D(P)}(s) = \frac{1}{n} \sum_{i=1}^n v_i(s) = \frac{1}{n} \sum_{i=1}^n v_{\sigma(i)}(s) = v_{D(P')}(s);$$

$$w_{D(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r) = \frac{1}{n} \sum_{i=1}^n w_{\sigma(i)}(r) = w_{D(P')}(r).$$

Therefore, Anonymity holds and Non-Dictatorship follows from it as we discussed in section 5.2.

- (iii) Monotonicity and Familiar Monotonicity. Let s be a statement and P and P' two opinion profiles satisfying the Monotonicity assumptions in the definition of the property in section 9.3, i.e. $P = (O_1, \dots, O_n)$ and $P' = (O'_1, \dots, O'_n)$ are such that $v_i(s) \leq v'_i(s)$ for every agent $i \in \{1, \dots, n\}$. Then, from the definition of D , we obtain the aggregated valuation on s is:

$$v_{D(P)}(s) = \frac{1}{n} \sum_{i=1}^n v_i(s) \leq \frac{1}{n} \sum_{i=1}^n v'_i(s) = v_{D(P')}(s)$$

Therefore, D satisfies Monotonicity. Hence, from this and lemma 5.2, Familiar Monotonicity also holds.

- (iv) Narrow Unanimity, Sided Unanimity and Weak Unanimity. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a DRF and a statement $s \in \mathcal{S}$ such that $v_i(s) = \lambda$ for every agent in $Ag = \{1, \dots, n\}$. The aggregated opinion on s is:

$$v_{D(P)}(s) = \frac{1}{n} \sum_{i=1}^n v_i(s) = \frac{1}{n} \sum_{i=1}^n \lambda = \lambda$$

Hence, Narrow unanimity is fulfilled by D . As discussed in section 9.3, Weak Unanimity follows from Narrow Unanimity. Furthermore, according to proposition 5.2 Sided Unanimity follows from Narrow Unanimity and Monotonicity.

- (v) Independence follows directly from the fact that D satisfies Monotonicity and from proposition 5.3.
- (vi) Collective Coherence. To prove that it does not hold, it suffices to find a DRF and an opinion profile for which there is no collective coherence. Thus, consider the example depicted below in figure A.15.




Figure A.15

If we check coherence for statement s , we obtain that:

$$|v_{D(P)}(s) - e_{D(P)}(s)| = v(s) - v(a) = 2 > \epsilon.$$

for any $\epsilon \in (0, 1)$, and hence collective coherence does not hold for this profile.

- (vii) Endorsed Unanimity. Using the opinion profile depicted in figure A.15, we observe that even with full negative support on s (i.e. $v(a) = -1$), the result of the aggregation is the opposite ($v_{D(P)}(s) = 1$). Therefore, this opinion profile also serves as a counterexample to prove that D does not satisfy Endorsed Unanimity.

□

Proposition A.2. The aggregation function I satisfies the following properties:

- (i) Exhaustive Domain and Coherent Domain;
- (ii) Anonymity and Non-Dictatorship;
- (iii) Endorsed Unanimity; and
- (iv) Familiar Monotonicity.

And does not satisfy:

- (v) Collective coherence;
- (vi) Narrow Unanimity, Sided Unanimity and Weak Unanimity;
- (vii) Monotonicity; and
- (viii) Independence.

Proof. (of proposition A.2)

- (i) Exhaustive and Coherent domain are straightforward.
- (ii) Anonymity and Non-Dictatorship. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a DRF and σ a permutation over the agent in $Ag = \{1, \dots, n\}$. We must show that I maintains the same collective opinion over the permuted opinion profile $P' = (O_{\sigma(1)}, \dots, O_{\sigma(n)})$, i.e. that $I(P) = I(P')$.

For any $i \in \{1, \dots, n\}$ there is only one $j \in \{1, \dots, n\}$ such that $\sigma(j) = i$, and hence in terms of expectation functions we know that $e_i = e_{\sigma(j)}$. Using that, we can show that $I(P) = I(P')$ as follows:

$$v_{I(P)}(s) = \frac{1}{n} \sum_{i=1}^n e_i(s) = \frac{1}{n} \sum_{i=1}^n e_{\sigma(i)}(s) = v_{I(P')}(a);$$

$$w_{I(P)}(r) = \frac{1}{n} \sum_{i=1}^n w_i(r) = \frac{1}{n} \sum_{i=1}^n w_{\sigma(i)}(r) = w_{I(P')}(r)$$

- (iii) Endorsed Unanimity. Let s be a sentence and $P = (O_1, \dots, O_n)$ an opinion profile satisfying that $v_i(s_r) = 1$ for any agent i and descendant $s_r \in D(s)$ of sentence s . Since the expectation over s is:

$$e_i(s) = \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} w_i(r) v_i(s_r) = \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} w_i(r) = 1,$$

then the aggregated value for s is:

$$v_{I(P)}(s) = \frac{1}{n} \sum_i e_i(s) = 1.$$

Analogously, if we assume that $v_i(s_r) = -1$ for any agent i and descendant $s_r \in D(s)$ of sentence s , we would obtain that $v_{I(P)}(s) = -1$. Since $v_{I(P)}(s) > 0$ when there is full positive support (and $v_{I(P)}(s) < 0$ for negative support), Endorsed Unanimity holds.

- (iv) Familiar Monotonicity. It is straightforward to see that the output of the Indirect aggregation function, which uses an expectation function, depends only on the values on descendants and their relationships. So, it is clear that a different opinion profile maintaining the same values for descendants and their relationships will not change the output of the function.
- (v) Collective Coherence. To prove that it does not hold, it suffices to find a *DRF* and an opinion profile for which there is no collective coherence. Thus, consider the example depicted below in figure A.16. Here $v_{I(P)}(s) = 1$ and $v_{I(P)}(a) = -1 = v_{I(P)}(b)$. Now, if we check coherence for statement s , we obtain that $|v_{I(P)}(s) - e_{I(P)}(s)| = 2 > \epsilon$ for any $\epsilon \in (0, 1)$, and hence collective coherence does not hold for this profile.




Figure A.16

- (vi) Narrow Unanimity, Sided Unanimity and Weak Unanimity. Next we build a *DRF* and an opinion profile for which Weak Unanimity does

not hold despite satisfying the assumptions. Consider the example in figure A.17 with opinion profile $P = (O = (v, w))$. Although $v(s) = 1$, $v_{I(P)}(s) = -1$ instead of greater than 0, and hence I does not satisfy Weak unanimity. As discussed in section 5.2, an aggregation function satisfying either Narrow Unanimity or Sided Unanimity also satisfies Weak Unanimity. Thus, since Weak Unanimity does not hold, neither do Narrow Unanimity and Sided Unanimity.




Figure A.17

- (vii) Monotonicity. Next we build a *DRF* and an opinion profile for which Monotonicity does not hold despite satisfying the assumptions. Consider the opinion profile in figures A.18 and A.19 for the same *DRF*. The two profiles $P = (O = (v, w))$ and $P' = (O' = (v', w'))$ only differ on the valuation of a : $v(a) = 1$ and $v'(a) = -1$.

Clearly, $x = v(s) \leq v'(s) = x$, thus satisfying the assumptions of monotonicity. However, since the aggregated valuations on s are: $v_{I(P)}(s) = 1$ and $v_{I(P')}(s) = -1$, it does not satisfy that $v_{I(P)} \leq v_{I(P')}$, and hence Monotonicity does not hold.




Figure A.18




Figure A.19

- (viii) Independence. Next we build a *DRF* and an opinion profile for which Independence does not hold despite satisfying the assumptions. Consider the opinion profiles P and P' in figures A.20 and A.21. Although

$v(s) = v'(s)$ for those profiles, the aggregated valuations on s do not match: $1 = v_{I(P)}(s) \neq v_{I(P')}(s) = 0$. Therefore, Independence does not hold.




Figure A.20




Figure A.21

□

Proposition A.3. The aggregation function R satisfies the following properties:

- (i) Collective Coherence;
- (ii) Exhaustive Domain and Coherent Domain;
- (iii) Anonymity and Non-Dictatorship;

And does not satisfy:

- (iv) Narrow Unanimity, Sided Unanimity and Weak Unanimity;
- (v) Endorsed Unanimity;
- (vi) Familiar Monotonicity, so neither Monotonicity;
- (vii) Independence.

Proof. (i) Collective Coherence. Since $v_{R(P)} = e_{R(P)}$, the collective opinion for R is exactly the result of applying the estimation function, and hence collective coherence follows because $|v_{R(P)}(s) - e_{R(P)}(s)| = 0 < \epsilon$ for any $\epsilon \in (0, 1)$ and any sentence $s \in \mathcal{S}$.

- (ii) Exhaustive Domain and Coherent Domain. Straightforward.
- (iii) Anonymity and Non-Dictatorship. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a DRF and σ a permutation over the agents in $Ag = \{1, \dots, n\}$. We must show that R maintains the same collective opinion over the permuted opinion profile $P' = (O_{\sigma(1)}, \dots, O_{\sigma(n)})$, i.e. that $R(P) = R(P')$.

We consider first the sentences $s \in \mathcal{S}$ with no descendants such that $R^+(s) = \emptyset$. Since these have no descendants, R computes the collective opinion on them using D . As shown by proposition A.1, since D satisfies anonymity, it will also hold for R when considering sentences with no descendants. Thus, since these sentences, which are used at the beginning of the recursive process run by R , will not change through permutations, the collective opinion over any sentence will be the same after permutations. Therefore, anonymity holds for R , and from this Non-Dictatorship.

- (iv) Weak, Narrow and Sided Unanimity. The example of DRF depicted in figure A.22 with opinion profile $P = (O = (v, w))$ will be enough to show that R does not satisfy Weak unanimity. Although $v(s) = 1$, and hence the assumptions for Weak Unanimity hold, $v_{R(P)}(s) = -1$ influenced by the valuation on b . Since $v_{R(P)}(s)$ is not positive, Weak unanimity does not hold for R , and consequently neither Side unanimity nor Narrow unanimity.




Figure A.22

- (v) Endorsed Unanimity. Consider again the opinion profile depicted in figure A.22. Clearly, since $v(a) = 1$, s has full positive support, but $v_{R(P)}(s) = -1$. Since $v_{R(P)}(s)$ is not positive, Endorsed Unanimity does not hold.
- (vi) Familiar Monotonicity and Monotonicity. We build a DRF and two opinion profiles for which Familiar Monotonicity does not hold despite

satisfying the assumptions. Consider the two opinion profiles $P = (O = (v, w))$ and $P' = (O' = (v', w'))$ depicted in figures A.23 and A.24 respectively. Considering s , these two profiles satisfy the assumptions of Familiar Monotonicity: $v(s) \leq v'(s)$ and the values on the indirect opinion are the same. However, P and P' differ on the value on b : $v(b) = 1$ and $v'(b) = -1$. This leads to a change of value on the aggregated value on s . Thus, $v_{R(P)}(s) \not\leq v_{R(P')}(s)$, and R fails at satisfying Familiar Monotonicity. By lemma 5.2, Monotonicity does not hold either.




Figure A.23




Figure A.24

- (vii) Independence. Straightforward from the example employed in proposition A.2 to prove lack of independence.

□

Next, we provide the proofs for the analysis of the families of α -balanced aggregation functions $\{B_\alpha\}_{\alpha \in (0,1)}$ and α -recursive aggregation functions $\{R_\alpha\}_{\alpha \in (0,1)}$. Before that, we first introduce some general lemmas that will be useful to build the proofs of the propositions for both families. To ease notation, these general lemmas that follow consider two generic aggregation functions F and G , as well as a generic aggregation function $H = \alpha F + (1 - \alpha)G$ instead of $v_H(P) = \alpha v_F(P) + (1 - \alpha)v_G(P)$. Hereafter, the following lemmas establish the social properties fulfilled by H .

Lemma A.1. Let F and G be two opinion aggregation functions satisfying Exhaustive domain. For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Exhaustive domain.

Proof. Straightforward from the fact that both F and G satisfy Exhaustive domain. \square

Lemma A.2. Let F and G two opinion aggregation functions satisfying Anonymity over domain \mathcal{D} . For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Anonymity over domain \mathcal{D} .

Proof. For any given opinion profile P and its permuted profile P' , if $F(P) = F(P')$ and $G(P) = G(P')$, then it follows that $H(P) = H(P')$. \square

Lemma A.3. Let F and G two opinion aggregation functions satisfying Familiar Monotonicity over domain \mathcal{D} . For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Familiar Monotonicity on domain \mathcal{D} .

Proof. Let $P = (O_1 = (v_1, w_1), \dots, O_n = (v_n, w_n))$ and $P' = (O'_1 = (v'_1, w'_1), \dots, O'_n = (v'_n, w'_n))$ be a profile satisfying the assumptions of familiar monotonicity for a statement s , i.e. $v_i(s) \leq v'_i(s)$ for any i and for any $r \in R^+(s)$ then $w_i(r) = w'_i(r)$ and $v_i(s_r) = v'_i(s_r)$. Since F and G satisfy familiar monotonicity, then $v_{F(P)}(s) \leq v_{F(P')}(s)$ and $v_{G(P)}(s) \leq v_{G(P')}(s)$. Thus, since $H = \alpha F + (1 - \alpha)G$, it follows directly that $v_{H(P)}(s) \leq v_{H(P')}(s)$, so familiar monotonicity holds for H . \square

Lemma A.4. Let F and G two opinion aggregation functions satisfying Sided unanimity on domain \mathcal{D} . For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Sided unanimity on \mathcal{D} .

Proof. Since Sided unanimity holds for F and G , we know that given any opinion profile P of agents $\{1, \dots, n\}$, i.e. if for any $i \in \{1, \dots, n\}$ $v_i(s) > 0$ then $v_F(s) > 0$ and $v_G(s) > 0$, and since $v_H = \alpha v_F + (1 - \alpha)v_G$, it also follows that $v_H(s) > 0$. Likewise for the negative case, so Sided Unanimity holds for H . \square

Lemma A.5. Let F and G two opinion aggregation functions satisfying Weak unanimity on the domain \mathcal{D} . For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Weak unanimity over domain \mathcal{D} .

Proof. Since Sided unanimity holds for F and G , we know that given any opinion profile P of agents $\{1, \dots, n\}$, for any $i \in \{1, \dots, n\}$, if $v_i(s) = 1$, then $v_F(s) > 0$ and $v_G(s) > 0$. Since $v_H = \alpha v_F + (1 - \alpha)v_G$, it also follows that $v_H(s) > 0$, and hence Weak Unanimity holds for H . Analogously for the negative case. \square

Lemma A.6. Let F and G two opinion aggregation functions satisfying Endorsed unanimity on domain \mathcal{D} . For any $\alpha \in (0, 1)$, aggregation function $H = \alpha F + (1 - \alpha)G$ also satisfies Endorsed unanimity on \mathcal{D} .

Proof. Since Endorsed unanimity holds for F and G , we know that given any opinion profile P of agents $\{1, \dots, n\}$, for any $i \in \{1, \dots, n\}$ and descendant $s_r \in D(s)$ of sentence s , if $v_i(s_r) > 1$, then $v_F(s) > 0$ and $v_G(s) > 0$. Since $v_H = \alpha v_F + (1 - \alpha)v_G$, it also follows that $v_H(s) > 0$, and hence Endorsed Unanimity holds for H . Analogously for the full negative support case. \square

We are now ready to prove the results for α -balanced aggregation functions in $\{B_\alpha\}_{\alpha \in (0,1)}$.

Proposition A.4. The family of α -balanced aggregation functions $\{B_\alpha\}_{\alpha \in (0,1)}$ satisfies the following properties:

- (i) Exhaustive Domain and Coherent Domain;
- (ii) Anonymity and Non-Dictatorship;
- (iii) Weak Unanimity for $\alpha > \frac{1}{2}$;
- (iv) Endorsed Unanimity for $\alpha < \frac{1}{2}$;
- (v) Familiar Monotonicity;

and does not satisfy:

- (vi) Collective coherence;
- (vii) Narrow Unanimity, nor Sided Unanimity;
- (viii) Monotonicity;
- (ix) Independence.

Proof. (i) Exhaustive Domain and Coherent Domain follow from propositions A.1 and A.2, and from lemma A.1.

- (ii) Anonymity and Non-Dictatorship follow from propositions A.1 and A.2, and from lemma A.2.

- (iii) Weak Unanimity. Let $P = (O_1, \dots, O_n)$ be an opinion profile over a *DRF* for the agents in $Ag = \{1, \dots, n\}$, and $s \in \mathcal{S}$ a sentence such that $v_i(s) = 1$ for any i . By proposition A.1, we know that Weak unanimity holds for the Direct aggregation function, and hence $v_{D(P)}(s) = \frac{1}{n} \sum_{i \in Ag} v_i(s) = 1$. Now we turn our attention to I , the indirect function. The worst scenario occurs when $v_{I(P)}(s) = -1$ because aggregating this value to $v_{D(P)}(s)$ might prevent that $v_{B_\alpha(P)}(s) > 0$, and thus that Weak unanimity holds. The DRF and an opinion profile depicted in figure A.25 exemplifies this case.




Figure A.25

Since $v_{D(P)}(s) = 1$ and $v_{I(P)}(s) = -1$, $v_{B_\alpha(P)}(s) = \alpha - (1 - \alpha) = 2\alpha - 1$. Thus, if we set α so that $\alpha > \frac{1}{2}$, then we ensure that $v_{B_\alpha(P)}(s) > 0$, and Weak unanimity holds. The proof is analogous for the negative case of Weak unanimity.

- (iv) Endorsed Unanimity. Let s be a sentence and $P = (O_1, \dots, O_n)$ an opinion profile satisfying that $v_i(s_r) = -1$ for any agent i and descendant $s_r \in D(s)$ of sentence s . In other words, s has full negative support. It follows that $v_{I(P)}(s) = -1$. Likewise for our proof for Weak unanimity above, we consider the worst case, which would occur when $v_{D(P)}(s) = 1$. Figure A.25 depicts a DRF and single-opinion profile illustrating this case. Since $v_{D(P)}(s) = 1$ and $v_{I(P)}(s) = -1$, $v_{B_\alpha(P)}(s) = \alpha - (1 - \alpha) = 2\alpha - 1$. Thus, if we set α so that $\alpha < \frac{1}{2}$, then we ensure that $v_{B_\alpha(P)}(s) < 0$, and Endorsed unanimity holds. The proof is analogous for the positive case (full positive support) of Endorsed unanimity.
- (v) Familiar Monotonicity follows from propositions A.1 and A.2, and from lemma A.3.
- (vi) Collective coherence. To prove that it does not hold, it suffices to find a *DRF* and an opinion profile for which there is no collective coherence. Thus, consider the DRF with one-opinion profile depicted below in

figure A.26. Computing the aggregations for the Direct and Indirect functions, we have that $v_{D(P)}(s) = 1$, $v_{D(P)}(a) = 0$, and, $v_{I(P)}(s) = 0$ and $v_{I(P)}(a) = -1$. Therefore, $v_{B_\alpha(P)}(s) = \alpha$ and $v_{B_\alpha(P)}(a) = (-1)(1 - \alpha)$. And hence, the coherence at sentence s is: $|v_{B_\alpha(P)}(s) - e_{B_\alpha(P)}(s)| = |v_{B_\alpha(P)}(s) - v_{B_\alpha(P)}(a)| = 1$. Thus, we conclude that, for any $\epsilon \in (0, 1)$, ϵ -coherence cannot be satisfied regardless of the value of α . Therefore, B_α does not satisfy ϵ -coherence.




Figure A.26

- (vii) Sided Unanimity, Narrow Unanimity. To prove that neither of these properties holds, it suffices to find a *DRF* and an opinion profile for which there is no Sided unanimity. In particular, we will show that for any $\alpha \in (0, 1)$ we can find a *DRF* and an opinion profile for which Sided unanimity and Narrow Unanimity do not hold. Consider then the *DRF* with single-opinion profile in figure A.27, where $x \in (0, 1)$ is such that $0 < x < \frac{1-\alpha}{\alpha}$. Since $v(s) = x > 0$, the assumptions for Sided unanimity hold at sentence s . Now, since $v_{D(P)}(s) = x$ and $v_{I(P)}(s) = -1$, it follows that $v_{B_\alpha(P)}(s) = \alpha x + (1 - \alpha)(-1) = \alpha x + \alpha - 1 < \alpha \frac{1-\alpha}{\alpha} + \alpha - 1 = 0$. Since $v_{B_\alpha(P)} \not\geq 0$, Sided unanimity fails at s , and as it is single-opinion profile Narrow unanimity fails too. The proof goes analogously for the negative case of Sided unanimity.




Figure A.27

- (viii) Monotonicity. It suffices to find a *DRF* and an opinion profile for which there is no Monotonicity. Consider the two single-opinion profiles P and P' over the very same *DRF* in figures A.28 and A.29. We will check Monotonicity at sentence s , where the conditions for unanimity hold because $v(s) \leq v'(s)$. Computing B_α we obtain that $v_{B_\alpha(P)}(s) = 1$

and $v_{B_\alpha(P')}(s) = 2\alpha - 1$. To fulfil Monotonicity both expressions must satisfy that $v_{B_\alpha(P)}(s) \leq v_{B_\alpha(P')}(s)$, namely that $1 \leq 2\alpha - 1$. This is only possible for $\alpha \geq 1$. Therefore, for any $\alpha \in (0, 1)$ Monotonicity does not hold.




Figure A.28




Figure A.29

$$1 \leq 2\alpha - 1$$

- (ix) Independence. For any $\alpha \neq 1$, B_α does not fulfil Independence due to its dependence on I .

□

Proposition A.5. The family of α -recursive aggregation functions $\{R_\alpha\}_{\alpha \in (0,1)}$ satisfies the following properties:

- (i) Collective Coherence for $\alpha < \frac{\epsilon}{2}$;
- (ii) Exhaustive Domain and Coherent Domain;
- (iii) Anonymity and Non-Dictatorship;
- (iv) Weak Unanimity for $\alpha > \frac{1}{2}$;

and does not satisfy:

- (v) Sided unanimity, so neither Narrow Unanimity;
- (vi) Endorsed Unanimity;
- (vii) Familiar Monotonicity, so neither Monotonicity;

(viii) Independence.

Proof. (i) Collective Coherence. Given $\epsilon > 0$ and a DRF, we must prove that $|v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s)| < \epsilon$ for any sentence $s \in \mathcal{S}$. First, we develop the difference between valuation and estimation for the collective opinion:

$$\begin{aligned}
v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s) &= v_{R_\alpha(P)}(s) - \frac{\sum_{r \in R^+(s)} w_{R_\alpha(P)}(r) v_{R_\alpha(P)}(s_r)}{\sum_{r \in R^+(s)} w_{R_\alpha(P)}(r)} \\
&= [\alpha v_{D(P)}(s) + (1 - \alpha) v_{R(P)}(s)] - \frac{\sum_{r \in R^+(s)} w_{D(P)}(r) [\alpha v_{D(P)}(s_r) + (1 - \alpha) v_{R(P)}(s_r)]}{\sum_{r \in R^+(s)} w_{R_\alpha(P)}(r)} \\
&= \alpha \left[v_{D(P)}(s) - \frac{\sum_{r \in R^+(s)} w_{R_\alpha(P)}(r) v_{D(P)}(s_r)}{\sum_{r \in R^+(s)} w_{D(P)}(r)} \right] + \\
&\quad + (1 - \alpha) \left[v_{R(P)}(s) - \frac{\sum_{r \in R^+(s)} w_{R(P)}(r) v_{R(P)}(s_r)}{\sum_{r \in R^+(s)} w_{R(P)}(r)} \right] \\
&= \alpha [v_{D(P)}(s) - e_{D(P)}(s)] + (1 - \alpha) [v_{R(P)}(s) - e_{R(P)}(s)] \\
&= \alpha (v_{D(P)}(s) - e_{D(P)}(s)).
\end{aligned}$$

Notice that we get rid of $v_{R(P)}(s) - e_{R(P)}(s)$ because it is zero. Now the coherence of R_α directly depends on the coherence of the direct aggregation function D and α . Figure A.30 depicts a DRF with a single-opinion profile representing a worst-case scenario for D because $v_{D(P)}(s) - e_{D(P)}(s) = 2$. Considering the coherence of R_α , we have that $|v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s)| = \alpha |v_{D(P)}(s) - e_{D(P)}(s)| \leq 2\alpha$ for any profile P . Therefore, we must ensure that $\alpha < \frac{\epsilon}{2}$ so that $|v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s)| < \epsilon$ holds for any profile of the domain, and hence Collective coherence holds for R_α .




Figure A.30

(ii) Exhaustive Domain and Coherent Domain directly follow from propositions A.1 and A.3 and lemma A.1.

- (iii) Anonymity and Non-Dictatorship follow directly from propositions A.1 and A.3 and lemma A.2.
- (iv) Weak unanimity. To prove Weak unanimity we can resort to the proof built to prove Weak unanimity for B_α in proposition A.4. We simply have to substitute B_α for R_α .
- (v) Sided Unanimity and Narrow Unanimity. To prove that neither of these properties holds, it suffices to find a *DRF* and an opinion profile for which there is no Sided unanimity (nor Narrow unanimity). Consider the DRF and the single-opinion profile depicted in figure A.31, where $x \in (0, 1)$ is such that $0 < x < \frac{1-\alpha}{\alpha}$. Since $x > 0$, the assumption for Sided Unanimity holds at s . However, $v_{R_\alpha(P)}(s)$ is not positive, since $v_{R_\alpha(P)}(s) = x\alpha - 1 + \alpha < \frac{1-\alpha}{\alpha}\alpha - 1 + \alpha = 0$, and hence Sided (and Narrow) unanimity does not hold.




Figure A.31

- (vi) Endorsed Unanimity. Next we build a DRF and an opinion profile for which Endorsed Unanimity does not hold. Consider the DRF and the opinion profile P depicted in figure A.32. The assumptions for Endorsed unanimity hold at s because s has full negative support. However, $v_{R_\alpha(P)}(s)$ is not negative: since $v_{D(P)}(s) = 1$ and $v_{R(P)}(s) = 1$, we obtain that $v_{R_\alpha(P)}(s) = 1$ for any $\alpha \in (0, 1)$. Therefore, Endorsed Unanimity does not hold.




Figure A.32

- (vii) Familiar Monotonicity and Monotonicity. We build a *DRF* and an opinion profile for which Familiar Monotonicity does not hold despite satisfying the assumptions. Consider the DRF and single-opinion profile

P in figure A.32 together with another single-opinion profile P' in figure A.33. Clearly, the assumptions of Familiar Monotonicity are fulfilled at s because $v_i(s) \leq v'_i(s)$ and the descendant of s has the same value. However, we will show that $v_{R_\alpha(P)}(s) \leq v_{R_\alpha(P')}(s)$ is not true. For both profiles we have that $v_{D(P)}(s) = v_{D(P')}(s) = 1$, and, $v_{R(P)}(s) = 1$ and $v_{R(P')}(s) = 1 - x$. Thus, for any $\alpha \in (0, 1)$: $v_{R_\alpha(P)}(s) = \alpha + (1 - \alpha) = 1$ and $v_{R_\alpha(P')}(s) = \alpha + (1 - \alpha)(1 - x) = 1 - x(1 - \alpha) < 1$ for any $x \in (0, 1)$. Therefore, $v_{R_\alpha(P)}(s) > v_{R_\alpha(P')}(s)$ and Familiar Monotonicity is not satisfied. By lemma 5.2, Monotonicity does not hold either.




Figure A.33

- (viii) Independence. Clearly, R_α will not fulfil Independence for any $\alpha \neq 1$ because of its dependence on R .

□

A.2. Constrained opinion profiles: assuming consensus on acceptance degrees

This section relates to section 7.2, where we assume that opinion profiles share consensus on their acceptance degrees on relationships, i.e. for each relationship $r \in \mathcal{R}$ of a DRF all the agents agree on their acceptance degrees: $w_i(r) = w_j(r) \forall i, j \in Ag$.

In previous section A.1, each proof and counterexample used to demonstrate that an aggregation function does or does not satisfy a property uses opinion profiles composed by one single agent. Thus, those proofs serve as well in this section when assuming consensus on acceptance degrees. For this reason, adding this assumption does not change any of the properties fulfilled by the aggregation functions in the general case (table 3), and therefore there are no further desirable properties gained in this scenario with respect to the more general scenario thoroughly analysed in section A.1.

A.3. Constrained opinion profiles: assuming coherent profiles

This section corresponds to the results displayed in table 4 in section 7.3. We prove the results regarding the social choice properties satisfied by the

aggregation functions introduced in Section 6 when assuming the domain of the aggregation functions to be ϵ -coherent for some $\epsilon \in (0, 1)$. This means that we consider that our aggregation functions take in coherent opinion profiles.

Since in the previous section many properties have been proven for the general case, we will not need to prove them again for this more restrictive scenario. For each opinion aggregation function, we will prove only those results regarding social choice properties that change by the addition of the coherence assumption and disprove again, this time for coherent domains, those properties which are yet not satisfied.

Proposition A.6. For any $\epsilon \in (0, 1)$, D over an ϵ -coherent domain satisfies Endorsed Unanimity.

Proof. Let s a statement in a DRF and $R^+(s)$ the set of relationships r from s to its descendants s_r . Let P be an ϵ -coherent profile for $\epsilon \in (0, 1)$ with full positive support on s , i.e. $v_i(s_r) = 1$ for any i and descendant $s_r \in D(s)$. Then:

$$e_i(s) = \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} v_i(s_r) w_i(r) = \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} w_i(r) = 1$$

By the ϵ -coherence of P we have that:

$$|v_i(s) - e_i(s)| < \epsilon \implies v_i(s) > e_i(s) - \epsilon = 1 - \epsilon.$$

Therefore, for any $\epsilon \in (0, 1)$ we can ensure that $v_i(s) > 0$ for any i and the conditions for Sided Unanimity hold. Now, since D satisfies Sided Unanimity (by proposition A.1), we obtain that $v_D(s) > 0$, and hence D fulfils Endorsed Unanimity. \square

Proposition A.7. D over a δ -coherent domain, where $\delta \in (0, 1)$, still does not satisfy ϵ -Collective coherence for any $\epsilon \in (0, 1)$.

Proof. Consider the DRF and δ -coherent opinion profile P depicted in figure A.34 and any $\delta \in (0, 1)$. We will show that the collective opinion yield by the direct function for this example is never ϵ -coherent for any $\epsilon \in (0, 1)$.

Clearly, this profile is δ -coherent for any $\delta > 0$. Computing the direct function at s we obtain that: $v_{D(P)}(s) = -1$, $v_{D(P)}(a) = 0$ and $w_{D(P)}(r) = \frac{1}{2}$. Now, if we check collective coherence at s , we see that: $|v_{D(P)}(s) - e_{D(P)}(s)| =$

$|-1 - 0| = 1 > \epsilon$. Thus, since 1 is larger than any ϵ value that we take in $(0, 1)$, D does not satisfy ϵ -Collective coherence.




Figure A.34

□

Proposition A.8. For $\epsilon \in (0, 1)$, I over an ϵ -coherent domain satisfies Weak Unanimity.

Proof. Consider a DRF with a statement $s \in \mathcal{S}$ and $P = (O_1 = (v_1, w_1), \dots, O_n = (v_n, w_n))$ an opinion profile such that $v_i(s) = 1$ for every i . Hence, the conditions for Weak unanimity hold. If the profile P is ϵ -coherent, where $\epsilon \in (0, 1)$, then we can conclude that for any i : $1 - \epsilon < e_i(s) < 1 + \epsilon$, being $1 - \epsilon > 0$ for any $\epsilon \in (0, 1)$. Now, computing v_I at s we get:

$$v_{I(P)}(s) = \frac{1}{n} \sum_i e_i(s) > \frac{1}{n} \sum_i 1 - \epsilon > 0$$

Since $v_{I(P)}(s) > 0$, Weak unanimity holds. The proof for the negative case of Weak unanimity is analogous.

□

Proposition A.9. For any $\delta \in (0, 1)$, I over an δ -coherent domain still does not satisfy ϵ -Collective coherence for any $\epsilon \in (0, 1)$.

Proof. To prove that this property does not hold, it suffices to find a DRF and an opinion profile for which there is no ϵ -Collective coherence. Consider the DRF and opinion profile P in figure A.35. Clearly, opinions O_1 and O_2 of P are δ -coherent for any $\delta > 0$. Now, we compute the indirect function for all the statement: $v_{I(P)}(s) = \frac{-1}{2}$, $v_{I(P)}(a) = \frac{1}{2}$, $v_{I(P)}(b) = 0$, and, $w_{I(P)}(r_1) = \frac{1}{2} = w_{I(P)}(r_2)$. If we check coherence at s we see that:

$$|v_{I(P)}(s) - e_{I(P)}(s)| = |v_{I(P)}(s) - v_{I(P)}(a)| = \left| \frac{-1}{2} - \frac{1}{2} \right| = 1 > \epsilon.$$

Thus, since 1 is larger than any ϵ value that we take in $(0, 1)$, I does not satisfy ϵ -Collective coherence.




Figure A.35

□

Proposition A.10. For any $\epsilon \in (0, 1)$, I over an ϵ -coherent domain still does not satisfy:

- (i) Sided Unanimity, and therefore Narrow Unanimity;
- (ii) Monotonicity.

Proof. (i) Sided unanimity and narrow unanimity. To prove that these properties not hold, it suffices to find a DRF and an opinion profile for which there is no Sided unanimity. Consider the DRF and one-opinion profile depicted in figure A.36 such that $\epsilon \in (0, 1)$ and x, y such that $0 < x < y < \epsilon$. The assumptions of Sided unanimity are fulfilled at s . We check that the opinion in the profile is ϵ -coherent because $|v(s) - e(s)| = |x - y + \epsilon| < \epsilon$. However, $v_{I(P)}(s) = y - \epsilon < 0$, instead of positive, and hence Sided Unanimity is not satisfied. As in previous proofs, as the counterexample is a single-opinion, Narrow unanimity does not hold either.




Figure A.36

- (ii) Monotonicity. To prove that it does not hold, it suffices to find a DRF with opinion profiles satisfying the Monotonicity assumptions are satisfied, and yet Monotonicity is not. In fact, we will create a generic counterexample for any ϵ -coherent domain. Consider the $DRF = \langle \mathcal{S}, \mathcal{R}, \tau \rangle$

depicted in figure A.37 with $\mathcal{S} = \{s, a\}$ and $\mathcal{R} = \{r\}$. Also in the figure, let $P = ((v, w))$ be an opinion profile of one single agent such that $v(s) = x$, $w(r) = 1$, and $v(a) = y >$, where $-1 < y < x < 1$ and $0 < x - y < \epsilon$.




Figure A.37

Since $v(s) - e_O(s) = v(s) - v(a)$, then the profile P is clearly ϵ -coherent, and hence $P \in \mathbb{C}_\epsilon(DRF)$. We compute the collective opinion using I at s as: $v_{I(P)}(s) = e_O(s) = v(a) = y$.

Now, consider another profile $P' = (O = (v', w))$ over the same DRF, shown in figure A.38, such that $v'(s) = x + \frac{\omega}{3}$ and $v'(a) = y - \frac{\omega}{3}$, where $\omega > 0$, such that $x - y + \omega \leq \epsilon$, $x + \frac{\omega}{3} \leq 1$ and $y - \frac{\omega}{3} \geq -1$. Clearly $v'(s) > v(s)$ and P is also ϵ -coherent, i.e.:

$$v(s) - e_{O'}(s) = (x + \frac{\omega}{3}) - (y - \frac{\omega}{3}) < \epsilon.$$

Nonetheless, $v_{I(P')}(s) = e_{O'}(s) = y - \frac{\omega}{3} < y$, which means that $v_{I(P)}(s) \not\leq v_{I(P')}(s)$, and hence this example cannot satisfy Monotonicity for any ϵ .




Figure A.38

□

Proposition A.11. For any $\epsilon \in (0, 1)$ and considering the domain to be ϵ -coherent, R does not fulfil the following properties:

- (i) Weak Unanimity, neither Sided nor Narrow Unanimity;
- (ii) Endorsed Unanimity;
- (iii) Familiar Monotonicity, and therefore Monotonicity.

Proof. (i) Weak Unanimity, Sided unanimity, Narrow unanimity. To prove that neither of these properties hold, it suffices to build a DRF and opinion profile to show that Weak Unanimity does not hold. This is sufficient because Weak unanimity is a weaker case of Sided and Narrow unanimity, proposition 5.1 tells us that Sided and Narrow unanimity will not hold if Weak Unanimity does not. Consider the DRF and opinion profile $P = ((v, w))$ in picture A.39 such that $w(r) = 1$ for any relationship $r \in \mathcal{R}$, $\epsilon \in (0, 1)$ and $\delta \in (0, \epsilon)$ and $m \in \mathbb{N}$ so that $m\delta \geq 1 > (m - 1)\delta$.




Figure A.39

Clearly, the outcome of the recursive function at each sentence is obtained from the value of the recursive function at the previous sentence, i.e.:

$$v_{R(P)}(a_m) = v_{R(P)}(a_{m-1}) = \dots = v_{R(P)}(a_1) = v_{R(P)}(s),$$

which actually is the value $v(a_m) = 1 - m\delta \leq 0$. So, this is an opinion profile ϵ -coherent fulfilling the assumptions of Weak unanimity at sentence s because $v(s) = 1$. However, the value of the recursive function at s is negative. Therefore, R does not fulfil Weak unanimity.

- (ii) Endorsed Unanimity. We build a DRF and opinion profile to show that Endorsed unanimity does not hold from the example in the previous proof. Figure A.40 shows our example, which extends the one in figure A.39 with an additional sentence a . Since $v(a_i) - v(a_{i-1}) = \delta$, likewise in the proof above, we have an ϵ -coherent opinion profile. Since $v(s) = 1$ the assumption for Endorsed unanimity at a is satisfied, but since $v(a) = 1 - m\delta \leq 0$, Endorsed unanimity does not hold.




Figure A.40

(iii) Familiar Monotonicity and Monotonicity. Consider the opinion profiles P and P' over the same DRF depicted in figures A.41 and A.42 respectively. Since $v(s) = v(a) = v(b) = 1$, P is ϵ -coherent. By setting $0 < x < \epsilon$, we also obtain that P' is ϵ -coherent. Therefore, both P and P' are ϵ -coherent and the assumptions for familiar monotonicity hold at s . However, since $1 = v_{R(P)}(s) > v_{R(P')}(s) = 1 - x$, Familiar monotonicity cannot hold. By lemma 5.2 Monotonicity does not hold either.




Figure A.41




Figure A.42

□

Proposition A.12. For any $\epsilon \in (0, 1)$ and considering the domain to be ϵ -coherent, then the family $\{B_\alpha\}_{\alpha \in (0, 1)}$ satisfies:

- (i) Weak Unanimity; and
- (ii) Endorsed Unanimity.

Proof. (i) Weak Unanimity follows from propositions A.1, A.8 and A.5.

(ii) Endorsed Unanimity follows from propositions A.2, A.6 and A.6.

□

Proposition A.13. For any $\delta \in (0, 1)$, B_α over a δ -coherent domain still does not satisfy ϵ -Collective coherence for any $\epsilon \in (0, 1)$.

Proof. First, we show that the ϵ -coherence condition for B_α depends on the functions employed in its definition, namely on D and I :

$$\begin{aligned}
|v_{B_\alpha(P)}(s) - e_{B_\alpha(P)}(s)| &= \left| v_{B_\alpha(P)}(s) - \frac{\sum_{r \in R^+(s)} (\alpha v_{D(P)}(rs) + (1 - \alpha) v_{I(P)}(s_r)) w_{D(P)}(r)}{\sum_{r \in R^+(s)} w_{D(P)}(r)} \right| \\
&= \left| (\alpha(v_{D(P)}(s) + (1 - \alpha)(v_{I(P)}(s))) - (\alpha e_{D(P)}(s) + (1 - \alpha)e_{I(P)}(s))) \right| \\
&= \left| \alpha(v_{D(P)}(s) - e_{D(P)}(s)) + (1 - \alpha)(v_{I(P)}(s) - e_{I(P)}(s)) \right|
\end{aligned}$$

Thus, since D and I do not satisfy ϵ -collective coherence for any δ -coherent profile (by propositions A.7 and A.9 respectively), neither will B_α satisfy the property for any $\alpha \in (0, 1)$. Indeed, consider for instance the DRF and δ -coherent opinion profile P , with any $\delta \in (0, 1)$, in figure A.35 as employed in proposition A.9. If we compute ϵ -collective coherence for B_α at sentence s we obtain that:

$$|v_{B_\alpha(P)}(s) - e_{B_\alpha(P)}(s)| = \left| \alpha(-1 - 0) + (1 - \alpha)\left(-\frac{1}{2} - \frac{1}{2}\right) \right| = |-1| = 1 > \epsilon$$

for any $\epsilon \in (0, 1)$. So, B_α does not fulfill ϵ -coherence for any $\alpha \in (0, 1)$.

□

Proposition A.14. For any $\epsilon \in (0, 1)$ and considering the domain to be ϵ -coherent, B_α does not fulfil the following properties for any $\alpha \in (0, 1)$:

- (i) Sided Unanimity and Narrow Unanimity;
- (ii) Monotonicity; and
- (iii) Independence.

Proof. (i) Sided Unanimity. It suffices to build a DRF and an opinion profile for which Sided unanimity does not hold for any values of α and ϵ , where $\alpha, \epsilon \in (0, 1)$. The next counterexample serves too see that Narrow unanimity fails too.

Consider the set $A = \{(x, y) \in (0, 1) \mid 0 < y < \epsilon \text{ and } 0 < x < y - \alpha y\}$. We check first, that this set is actually not empty. For $\alpha \in (0, 1)$, $y - \alpha y > 0$, thus $y > \alpha y > 0$. So for $y \in (0, \epsilon)$, there are $x \in (0, 1)$ satisfying $x < y - \alpha y$.




Figure A.43

Now, we consider the DRF and opinion profile depicted in figure A.43 where x and y are values from A , namely $(x, y) \in A$. Since $|v(s) - e(s)| = |x - (x - y)| = |y| < \epsilon$, the opinion profile in the figure is ϵ -coherent, and satisfies the assumptions for Sided Unanimity at s because $v(s) = x > 0$. However,

$$\begin{aligned} v_{B_\alpha(P)}(s) &= \alpha v_{D(P)}(s) + (1 - \alpha) v_{I(P)}(s) \\ &= \alpha x + (1 - \alpha)(x - y) \\ &= x - y + \alpha y < 0 \end{aligned}$$

since $(x, y) \in A$. So, clearly this example shows that Sided Unanimity does not hold for the family B_α in an ϵ -coherent profile.

- (iv) Independence. For any $\alpha \in (0, 1)$, B_α does not fulfil Independence due to its dependence on I .
- (iii) Monotonicity. Straightforward from the fact that B_α does not fulfil Independence for any $\alpha \in (0, 1)$ and from proposition 5.3.

□

Proposition A.15. For any $\delta \in (0, 1)$, R_α over a δ -coherent domain satisfies ϵ -Collective coherence for $\alpha \leq \frac{\epsilon}{2}$.

Proof. As seen before in proposition A.3, the collective coherence of R_α entirely depends on the collective coherence of D , i.e.:

$$|v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s)| = \alpha |v_{D(P)}(s) - e_{D(P)}(s)|$$

Thus, finding the worst case scenario for D will give us the condition on α that ensures that R_α satisfies ϵ -collective coherence for any ϵ . Next, we consider an example showing that $|v_{D(P)}(s) - e_{D(P)}(s)|$ can be as close to 2 as wanted depending on the number of agents.




Figure A.44

Let P be the δ -coherent opinion profile over a DRF depicted in figure A.44, for any $\delta \in (0, 1)$. For any $i > 1$: $v_i(s) = 1$, $v_i(a) = -1$ and $w_i(r_1) = 0$; whereas $v_1(s) = 1$, $v_1(a) = 1$ and $w_1(r_1) = 1$. We check the condition for collective coherence at s to find that:

$$|v_{R_\alpha(P)}(s) - e_{R_\alpha(P)}(s)| = \alpha \left(1 + \frac{m-2}{m}\right) < \epsilon$$

if $\alpha < \frac{\epsilon}{1 + \frac{m-2}{m}}$. Thus, by taking $\alpha < \frac{\epsilon}{2} < \frac{\epsilon}{1 + \frac{m-2}{m}}$ we ensure that R_α satisfies ϵ -coherence for the worst case. Therefore, for any δ -coherent opinion profile, $\delta \in (0, 1)$, choosing $\alpha < \frac{\epsilon}{2}$ will ensure that R_α satisfies ϵ -collective coherence for any $\epsilon \in (0, 1)$. □

Proposition A.16. Let $\epsilon \in (0, 1)$ such that the domain of R_α is an ϵ -coherent domain, then the family $\{R_\alpha\}_{\alpha \in (0, 1)}$ satisfies:

- (i) Weak Unanimity for $\alpha > \frac{1}{2}$; and
- (ii) Endorsed unanimity for $\alpha > \frac{1}{2-\epsilon}$.

Proof. (i) Weak Unanimity. Consider a DRF with sentences \mathcal{S} , P an opinion profile over the DRF and $s \in \mathcal{S}$ a sentence such that $v_i(s) = 1$ for any agent i . We know that

$$v_{D(P)}(s) = \frac{1}{n} \sum_{i \in Ag} v_i(s) = 1.$$

Now we turn our attention to R , the Recursive function. We consider the worst scenario for R_α , which happens when $v(s) = 1$ and $v_{R(P)}(s) = -1$. The DRF and profile depicted in figure A.22 above shows that, in fact, this scenario exists with $v_{D(P)}(s) = 1$ and $v_{R(P)}(s) = -1$, and hence $v_{R_\alpha(P)}(s) = \alpha + (1 - \alpha)(-1) = 2\alpha - 1$. To fulfil Weak unanimity, we need that $v_{R_\alpha(P)}(s) > 0$ holds, but we also know that $v_{R_\alpha(P)}(s) \geq 2\alpha - 1$. Therefore, we can guarantee Weak Unanimity by choosing $\alpha > \frac{1}{2}$. The proof for the negative case of Weak Unanimity goes analogously.

(ii) Endorsed Unanimity. To prove this property we will build a customised DRF and opinion profile to demonstrate the worst case that we can find when fulfilling the assumptions of Endorsed unanimity.

Consider a DRF and let $P = (O_1 = (v_1, w_1), \dots, O_n = (v_n, w_n))$ be an ϵ -coherent profile with full positive support on statement $s \in \mathcal{S}$.

First, we consider the worst case where $v_{R(P)}(s) = -1$ can be achieved when s has full positive support. Figure A.45 depicts a *DRF* and an opinion profile illustrating this situation.




Figure A.45

By choosing $0 < x < \epsilon$ and $m \in \mathbb{N}$ such that $mx > 2 \geq (m-1)x$, this example shows an ϵ -coherent profile where $v(a_1) = 1$ (full positive support) and $v_R(s) = -1$. Next, we move to the general setting considered by the proof, an opinion profile with n agents, knowing that the worst case for this property is possible. Since $v_i(s_r) = 1$ for any descendant $s_r \in D(s)$ and any agent i , the estimation function on s will

be $e_i(s) = 1$ for any agent. Therefore, from the coherence condition at s we conclude that

$$1 - \epsilon < v_i(s) < 1 + \epsilon.$$

Consider that for every i , $v_i(s) = 1 - \delta_i$ such that $0 \leq \delta_i < \epsilon$. This clearly satisfies the previous inequality. Now we take $\delta = \max_i \{\delta_1, \dots, \delta_n\}$ to create a new opinion profile $P' = (O'_1 = (v'_1, w_1), \dots, O'_n = (v'_n, w_n))$ such that $v_i(a) = v'_i(a)$ for $a \in \mathcal{S} \setminus \{s\}$ and $v'_i(s) = 1 - \delta$ for any i . Then, since D fulfils Monotonicity and Narrow unanimity, we know that $v_{D(P')}(s) \leq v_{D(P)}(s)$ and $v_{D(P')}(s) = 1 - \delta$ respectively. And, from the example in figure A.45 we know that for any ϵ -coherent opinion profile $v_{R(P)}(s) \geq -1$. Therefore,

$$\begin{aligned} v_{R_\alpha(P)}(s) &= \alpha v_{D(P)}(s) + (1 - \alpha) v_{R(P)}(s) \\ &\geq \alpha v_{D(P')}(s) + (1 - \alpha)(-1) \\ &= (1 - \delta)\alpha - (1 - \alpha) = (2 - \delta)\alpha - 1 \end{aligned}$$

So, if we set $\alpha \in (0, 1)$ so that $(2 - \delta)\alpha - 1 > 0$, R_α will satisfy Endorsed Unanimity. Since $\delta < \epsilon$, as close as possible, imposing $\alpha \geq \frac{1}{2 - \epsilon} > \frac{1}{2 - \delta}$ the property is satisfied.

□

Proposition A.17. For any $\epsilon \in (0, 1)$, and considering the domain to be ϵ -coherent, R_α does not fulfil the following properties for any $\alpha \in (0, 1)$:

- (i) Sided unanimity, and therefore Narrow unanimity;
- (ii) Familiar Monotonicity, and therefore Monotonicity; and
- (iii) Independence.

Proof. (i) Sided Unanimity and Narrow Unanimity. It suffices to build a DRF and an opinion profile for which Sided Unanimity does not hold for any $\alpha, \epsilon \in (0, 1)$.

Consider the DRF and opinion profile P depicted in figure A.46, where: $x \in (0, 1)$ is such that $0 < x < \frac{1 - \alpha}{\alpha}$, $0 < \delta < \epsilon$; $m \in \mathbb{N}$ satisfies $(m - 1)\delta \leq 1 + < m\delta$; and for any $r \in \mathcal{R}$, $w(r) = 1$.




Figure A.46

Clearly, P is an ϵ -coherent because $|v(a_m) - e(a_m)| = 0$, and for any $i < m$, $|v(a_i) - e(a_i)| = v(a_i) - v(a_{i+1}) = \delta < \epsilon$, and $|v(s) - e(s)| = x - x + \delta < \epsilon$. Furthermore, P satisfies the assumptions of Sided unanimity at s since $v(s) = x > 0$. It is straightforward to see that $v_{D(P)}(s) = x$ and $v_{R(P)}(s) = v_{R(P)}(a_m) = -1$. Hence, $v_{R_\alpha(P)}(s) = x\alpha + (1 - \alpha)(-1) = x\alpha + \alpha - 1$. But since $x < \frac{1-\alpha}{\alpha}$, we conclude that $v_{R_\alpha(P)}(s) < \frac{1-\alpha}{\alpha}\alpha + \alpha - 1 = 0$. This proves that Sided unanimity is not fulfilled, and as it is a single-opinion profile Narrow unanimity fails too. We can proceed analogously for the negative case of Sided Unanimity.

- (ii) Familiar monotonicity. The counterexample employed in proposition A.11 to show that Familiar monotonicity does not hold for R serves here as well to prove that R_α does not satisfy Familiar monotonicity for any $\alpha \in (0, 1)$. From opinion profiles P and P' depicted in figures A.41 and A.42 respectively, we extract that $v_{D(P)}(s) = v_{D(P')}(s) = 1$ and $1 = v_{R(P)}(s) > v_{R(P')}(s) = 1 - x$. Therefore, it follows that $v_{R_\alpha(P)}(s) > v_{R_\alpha(P')}(s)$, hence proving that Familiar monotonicity does not hold.
- (iii) Independence. For any $\alpha \in (0, 1)$, function R_α does not fulfil Independence due to its dependence on R .

□

A.4. Constrained opinion profiles: assuming Consensus on acceptance degrees and coherent profiles

Next, we show the results regarding our fourth, and last, scenario. We now assume that opinion profiles are both ϵ -coherent, for some $\epsilon \in (0, 1)$, and agree on their acceptance degrees over relationships. The results that follow are summarised in table 5 in section 7.4.

Likewise in previous sections, next we only prove per aggregation function those properties that either were partially satisfied or not satisfied at all in previous scenarios, but do hold in this new scenario. We do not prove those

properties for which the proofs in the previous sections serves as well for this scenario.

Proposition A.18. Let be a *DRF* and an opinion profile $P = (O_1, \dots, O_n)$. For any $s \in \mathcal{S}$, assume that for each $r \in R^+(s)$ $w_i(r) = \lambda_r \in (0, 1]$ for any i , then:

- (i) For any $\epsilon \in (0, 1)$, if $0 < \delta \leq \epsilon$ and the domain \mathcal{D} is δ -coherent then $D(P)$ is ϵ -coherent, so satisfies ϵ -Collective coherence.
- (ii) For any $\epsilon \in (0, 1)$, if $0 < \delta \leq \epsilon$ and the domain \mathcal{D} is δ -coherent then $I(P)$ is ϵ -coherent, so satisfies ϵ -Collective coherence.
- (iii) For any $\epsilon \in (0, 1)$, if $0 < \delta \leq \epsilon$ and the domain \mathcal{D} is δ -coherent then $B_\alpha(P)$ is ϵ -coherent for any $\alpha \in (0, 1)$, so satisfies ϵ -Collective coherence.
- (iv) For any $\epsilon \in (0, 1)$, if $0 < \delta \leq \epsilon$ and the domain \mathcal{D} is δ -coherent then $R_\alpha(P)$ is ϵ -coherent for any $\alpha \in (0, 1)$, so satisfies ϵ -Collective coherence.

Proof. (i) Collective Coherence of D . Let $s \in \mathcal{S}$. We assume that for any i , $|v_i(s) - e_i(s)| < \delta \leq \epsilon$. Next we calculate the coherence condition for D at sentence s :

$$\begin{aligned}
|v_{D(P)}(s) - e_{D(P)}(s)| &= \left| \frac{1}{n} \sum_i v_i(s) - \frac{1}{\sum_{r \in R^+(s)} w_{D(P)}(r)} \sum_{r \in R^+(s)} w_{D(P)}(r) v_{D(P)}(s_r) \right| \\
&= \left| \frac{1}{n} \sum_i v_i(s) - \frac{1}{\sum_{r \in R^+(s)} \lambda_r} \sum_{r \in R^+(s)} \lambda_r \left(\frac{1}{n} \sum_i v_i(s_r) \right) \right| \\
&= \left| \frac{1}{n} \sum_i \left(v_i(s) - \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} w_i(r) v_i(s_r) \right) \right| \\
&= \left| \frac{1}{n} \sum_i \left(v_i(s) - e_i(s) \right) \right| \leq \frac{1}{n} \sum_i |v_i(s) - e_i(s)|
\end{aligned}$$

Thus, by δ -coherence of the domain we obtain that:

$$|v_{D(P)}(s) - e_{D(P)}(s)| \leq \frac{1}{n} \sum_i |v_i(s) - e_i(s)| < \frac{1}{n} \sum_i \delta \leq \epsilon$$

This proves that the collective opinion by D is ϵ -coherent.

- (ii) Collective Coherence of I . We prove collective coherence for I similarly to the proof above for D . Let $s \in \mathcal{S}$. We assume that for any i , $|v_i(s) - e_i(s)| < \delta \leq \epsilon$. We compute the condition for the collective coherence of I at sentence s as follows:

$$\begin{aligned} |v_{I(P)}(s) - e_{I(P)}(s)| &= \left| \frac{1}{n} \sum_i e_i(s) - \frac{1}{\sum_{r \in R^+(s)} w_{I(P)}(r)} \sum_{r \in R^+(s)} w_{I(P)}(r) v_{I(P)}(s_r) \right| \\ &= \left| \frac{1}{n} \sum_i e_i(s) - \frac{1}{\sum_{r \in R^+(s)} \lambda_r} \sum_{r \in R^+(s)} \lambda_r \left(\frac{1}{n} \sum_i e_i(s_r) \right) \right| \\ &= \left| \frac{1}{n} \sum_i \left(e_i(s) - \frac{1}{\sum_{r \in R^+(s)} w_i(r)} \sum_{r \in R^+(s)} w_i(r) e_i(s_r) \right) \right| \\ &= \left| \frac{1}{n} \sum_i \frac{\sum_{r \in R^+(s)} w_i(r) v_i(s_r) - \sum_{r \in R^+(s)} w_i(r) e_i(s_r)}{\sum_{r \in R^+(s)} w_i(r)} \right| \\ &= \left| \frac{1}{n} \sum_i \sum_{r \in R^+(s)} \frac{w_i(r) (v_i(s_r) - e_i(s_r))}{\sum_{r \in R^+(s)} w_i(r)} \right| \end{aligned}$$

So, by δ -coherence of the domain, we obtain that:

$$\begin{aligned} |v_{D(P)}(s) - e_{D(P)}(s)| &\leq \frac{1}{n} \sum_i \sum_{r \in R^+(s)} \frac{w_i(r) |v_i(s_r) - e_i(s_r)|}{\sum_{r \in R^+(s)} w_i(r)} \\ &< \frac{1}{n} \sum_i \frac{\sum_{r \in R^+(s)} w_i(r) \delta}{\sum_{r \in R^+(s)} w_i(r)} \\ &= \frac{1}{n} \sum_i \delta = \delta \leq \epsilon \end{aligned}$$

This proves that the collective opinion by I is ϵ -coherent.

- (iii) Collective Coherence of B_α . We have just proven that D and I satisfy ϵ -collective coherence assuming consensus on acceptance degrees and a δ -coherent domain with $\delta < \epsilon$. It directly follows that for any $\alpha \in (0, 1)$, then B_α on a δ -coherent domain also satisfies ϵ -collective coherence.
- (iv) Collective Coherence of R_α . We have just proven that D satisfies ϵ -collective coherence assuming consensus on acceptance degrees and a δ -coherent domain. ϵ -collective coherence also holds for R under the same assumptions following proposition A.3 (see collective coherence for R). Hence, it follows that for any $\alpha \in (0, 1)$, R_α on a δ -coherent domain also satisfies ϵ -collective coherence.

□

References

- [1] N. Allen and S. Birch. Process preferences and British public opinion: Citizens' judgements about government in an era of anti-politics. *Political Studies*, 63(2):390–411, 2015.
- [2] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In *MultiAgent Systems, 2000. Proceedings. Fourth International Conference on*, pages 31–38. IEEE, 2000.
- [3] L. Amgoud, C. Cayrol, M.-C. Lagasquie-Schiex, and P. Livet. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems*, 23(10):1062–1093, 2008.
- [4] Appgree. <http://www.appgree.com/>, 04 2020. last visited.
- [5] K. J. Arrow and E. S. Maskin. *Social choice and individual values: Third edition*. 2012. ISBN 9780300179316. URL <http://www.scopus.com/inward/record.url?eid=2-s2.0-84903007957&partnerID=tZ0tx3y1>.
- [6] E. Awad, J.-F. Bonnefon, M. Caminada, T. W. Malone, and I. Rahwan. Experimental assessment of aggregation principles in argumentation-enabled collective intelligence. *ACM Transactions on Internet Technology (TOIT)*, 17(3):1–21, 2017.
- [7] E. Awad, R. Booth, F. Tohmé, and I. Rahwan. Judgement aggregation in multi-agent argumentation. *Journal of Logic and Computation*, 27(1):227–259, 2017.
- [8] H. Aziz, F. Brandt, E. Elkind, and P. Skowron. Computational social choice: The first ten years and beyond. *Computer Science Today*, 10000, 2017.
- [9] Baoqu. <https://baoqu.ctrlz.es/>, 04 2020. last visited.
- [10] P. Baroni and M. Giacomin. Semantics of abstract argument systems. In *Argumentation in artificial intelligence*, pages 25–44. Springer, 2009.
- [11] P. Baroni, M. Caminada, and M. Giacomin. An introduction to argumentation semantics. *The knowledge engineering review*, 26(4):365–410, 2011.

- [12] P. Baroni, M. Romano, F. Toni, M. Aurisicchio, and G. Bertanza. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation*, 6(1):24–49, 2015.
- [13] BBC news. “Cancel Brexit” petition passes 2m signatures on Parliament site. <https://www.bbc.com/news/uk-politics-47652071>, 03 2019. last visited 04/2020.
- [14] N. Benn and A. Macintosh. Argument visualization for eparticipation: towards a research agenda and prototype tool. In *International Conference on Electronic Participation*, pages 60–73. Springer, 2011.
- [15] P. Besnard and A. Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128(1-2):203–235, 2001.
- [16] Better Reykjavík. <http://reykjavik.is/en/better-reykjavik-0>, 04 2020. last visited.
- [17] A. Bochman. Collective argumentation and disjunctive logic programming. *Journal of logic and computation*, 13(3):405–428, 2003.
- [18] G. Bodanza, F. Tohmé, and M. Auday. Collective argumentation: A survey of aggregation issues around argumentation frameworks. *Argument & Computation*, 8(1):1–34, 2017.
- [19] E. Cabrio and S. Villata. A natural language bipolar argumentation approach to support users in online debate interactions. *Argument & Computation*, 4(3):209–230, 2013.
- [20] M. Caminada and G. Pigozzi. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems*, 22(1): 64–102, 2011.
- [21] C. S. Carr. Using computer supported argument visualization to teach legal argumentation. In *Visualizing argumentation: Software tools for collaborative and educational sense-making*, pages 75–96. Springer, 2003.
- [22] C. Cayrol and M. Lagasquie-Schiex. Gradual valuation for bipolar argumentation frameworks. In *Symbolic and quantitative approaches to reasoning with uncertainty*, pages 366–377. Springer, 2005.

- [23] C. Cayrol and M.-C. Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 378–389. Springer, 2005.
- [24] W. Chen and U. Endriss. Preservation of semantic properties in collective argumentation: The case of aggregating abstract argumentation frameworks. *Artificial Intelligence*, 269:27–48, 2019.
- [25] City of Helsinki. <https://www.hel.fi/helsinki/en/administration/participate/channels/participation-model/>, 04 2020. last visited.
- [26] Consider.it. <https://consider.it/>, 04 2020. last visited.
- [27] Consul. <http://consulproject.org/en/>, 05 2020. last visited.
- [28] S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasquie-Schiex, and P. Marquis. On the merging of dung’s argumentation systems. *Artificial Intelligence*, 171(10):730–753, 2007.
- [29] Decide Madrid. <https://decide.madrid.es/>, 04 2020. last visited.
- [30] Decidim Barcelona. <https://www.decidim.barcelona/>, 04 2020. last visited.
- [31] Decidim.org. <https://decidim.org/>, 05 2020. last visited.
- [32] F. Dietrich. A generalised model of judgment aggregation. *Social Choice and Welfare*, 28(4):529–565, 2007.
- [33] F. Dietrich and C. List. Strategy-proof judgment aggregation. *Economics & Philosophy*, 23(3):269–300, 2007.
- [34] F. Dietrich and P. Mongin. The premiss-based approach to judgment aggregation. *Journal of Economic Theory*, 145(2):562–582, 2010.
- [35] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.

- [36] P. M. Dung, R. A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170(2):114–159, 2006.
- [37] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2):457–486, 2011. ISSN 00043702. doi: 10.1016/j.artint.2010.09.005.
- [38] U. Endriss and H. Moulin. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 399–426. Cambridge University Press, 2016.
- [39] J. Font, M. Wojcieszak, and C. J. Navarro. Participation, representation and expertise: Citizen preferences for political decision-making processes. *Political Studies*, 63:153–172, 2015.
- [40] J. Fox, D. Barber, and K. Bardhan. Alternatives to Bayes? a quantitative comparison with rule-based diagnosis. *Methods of Information in Medicine*, 19(04):210–215, 1980.
- [41] J. Fox, P. Krause, and M. Elvang-Gøransson. Argumentation as a general framework for uncertain reasoning. In *Uncertainty in Artificial Intelligence, 1993*, pages 428–434. Elsevier, 1993.
- [42] A. Fung and E. O. Wright. Deepening democracy: Innovations in empowered participatory governance. *Politics & Society*, 29(1):5–41, 2001.
- [43] W. Gaertner. *A primer in social choice theory: Revised edition*. Oxford University Press, 2009.
- [44] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42(2-3):177–201, 1993.
- [45] J. Ganzer-Ripoll, M. López-Sánchez, and J. A. Rodríguez-Aguilar. A multi-agent argumentation framework to support collective reasoning. In *International Workshop on Conflict Resolution in Decision Making*, pages 100–117. Springer, 2016.

- [46] J. Ganzer-Ripoll, N. Criado, M. Lopez-Sanchez, S. Parsons, and J. A. Rodriguez-Aguilar. Combining social choice theory and argumentation: Enabling collective decision making. *Group Decision and Negotiation*, 28(1):127–173, 2019.
- [47] A. J. García and G. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(1): 95–138, 2004.
- [48] M. Hooghe and A. Kern. The tipping point between stability and decline: Trends in voter turnout, 1950–1980–2012. *European Political Science*, 16:535–552, 2017.
- [49] L. Iandoli, I. Quinto, A. De Liddo, and S. B. Shum. Socially augmented argumentation tools: Rationale, design and evaluation of a debate dashboard. *International Journal of Human-Computer Studies*, 72(3):298–319, 2014.
- [50] S. K. Jackson and K. M. Kuehn. Open source, social activism and “necessary trade-offs” in the digital enclosure: A case study of platform co-operative, loomio. org. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society*, 14(2):413–427, 2016.
- [51] Kialo. <https://www.kialo.com/>, 06 2020. last visited.
- [52] M. Klein. Enabling large-scale deliberation using attention-mediation metrics. *Computer Supported Cooperative Work (CSCW)*, 21(4-5):449–473, 2012.
- [53] M. Klein and G. Convertino. A roadmap for open innovation systems. *Journal of Social Media for Organizations*, 2(1):1, 2015.
- [54] P. Krause, S. Ambler, M. Elvang-Goransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11(1):113–131, 1995.
- [55] J. Lang, M. Slavkovik, and S. Vesic. Agenda separability in judgment aggregation. In *AAAI*, pages 1016–1022, 2016.

- [56] J. Leite and J. Martins. Social abstract argumentation. In *IJCAI*, volume 11, pages 2287–2292. AAAI Press, 2011. ISBN 978-1-57735-515-1. doi: 10.5591/978-1-57735-516-8/IJCAI11-381. URL <http://dl.acm.org/citation.cfm?id=2283779>.
- [57] C. List. Democratic deliberation and social choice. *The Oxford Handbook of Deliberative Democracy*, page 463, 2018.
- [58] C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(01):89–110, 2002.
- [59] Loomio. <https://www.loomio.org/>, 06 2020. last visited.
- [60] R. P. Loui. Defeat among arguments: a system of defeasible inference. *Computational Intelligence*, 3(3):100–106, 1987.
- [61] P. McBurney. *Rational Interaction*. PhD thesis, Department of Computer Science, University of Liverpool, 2002.
- [62] P. McBurney and S. Parsons. Dialogue games for agent argumentation. In *Argumentation in artificial intelligence*, pages 261–280. Springer, 2009.
- [63] R. McGuire, L. Birnbaum, and M. Flowers. Opportunistic processing in arguments. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 58–60, 1981.
- [64] S. Modgil and M. Caminada. Proof theories and algorithms for abstract argumentation frameworks. In *Argumentation in Artificial Intelligence*, pages 105–129. Springer, 2009.
- [65] S. Modgil and H. Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013.
- [66] P. Mongin. Factoring out the impossibility of logical aggregation. *Journal of Economic Theory*, 141(1):100–113, 2008.
- [67] New York City Participatory Budgeting. <https://pbnyc.participatorybudgeting.org/budgets>, 05 2020. last visited.

- [68] Parlement & Citoyens. Draft law for the recovery of biodiversity, nature and landscapes. <https://parlement-et-citoyens.fr/consultation/projet-de-loi-pour-la-reconquete-de-la-biodiversite-de-la-nature-et-des-paysa> presentation/presentation-et-suivi-6, 12 2015. last visited 04/2020.
- [69] Parlement & Citoyens. <https://parlement-et-citoyens.fr/>, 04 2020. last visited.
- [70] S. Parsons. Normative argumentation and qualitative probability. In *Proceedings of the International Joint Conference on Qualitative and Quantitative Practical Reasoning*, Bonn, 1997.
- [71] Petitions: Brexit. “Revoke Article 50 and remain in the EU” at Petitions UK Government and Parliament. <https://petition.parliament.uk/archived/petitions/241584>, 08 2019. last visited 04/2020.
- [72] G. Pigozzi, M. Slavkovik, and L. van der Torre. Independence in judgment aggregation. In *Proceedings of the Ninth International Meeting of the Society for Social Choice and Welfare, SCW’08*, 2008.
- [73] H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the 10th International Conference on Artificial Intelligence and Law*, pages 85–94, Bologna, Italy, 2005.
- [74] H. Prakken and G. Sartor. Argument-based logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 1997.
- [75] A. Rago and F. Toni. Quantitative argumentation debates with votes for opinion polling. In *International Conference on Principles and Practice of Multi-Agent Systems*, pages 369–385. Springer, 2017.
- [76] I. Rahwan and G. R. Simari, editors. *Argumentation in artificial intelligence*, volume 47. Springer, 2009.
- [77] I. Rahwan and F. Tohmé. Collective argument evaluation as judgement aggregation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, Toronto, 2010.

- [78] P. Rajendran, D. Bollegala, and S. Parsons. Assessing weight of opinion by aggregating coalitions of arguments. In *Proceedings of the 6th International Conference on Computational Models of Argument*, pages 431–438, 2016.
- [79] C. Reed and G. Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools*, 13(04):961–979, 2004.
- [80] C. Rinner. Argumentation mapping in collaborative spatial decision making. In *Collaborative geographic information systems*, pages 85–102. IGI Global, 2006.
- [81] M. Serramia, J. Ganzer, M. Lopez-Sanchez, J. A. Rodriguez-Aguilar, N. Criado, S. Parsons, P. Escobar, and M. Fernandez. Citizen support aggregation methods for participatory platforms. In *Frontiers in Artificial Intelligence and Applications: Artificial Intelligence Research and Development.*, volume 319, pages 9–18. IOS Press, 2019.
- [82] S. B. Shum. The roots of computer supported argument visualization. In *Visualizing argumentation*, pages 3–24. Springer, 2003.
- [83] G. Smith. *Democratic innovations: Designing institutions for citizen participation*. Cambridge University Press, 2009.
- [84] D. Suthers, A. Weiner, J. Connelly, and M. Paolucci. Belvedere: Engaging students in critical discussion of science and public policy issues. In *Proceedings of the 7th World Conference on Artificial Intelligence in Education*, pages 266–273. Washington, DC, 1995.
- [85] K. P. Sycara. Persuasive argumentation in negotiation. *Theory and decision*, 28(3):203–242, 1990.
- [86] P. Thagard. *Coherence in thought and action*. MIT Press, 2002.
- [87] T. Van Gelder. Enhancing deliberation through computer supported argument visualization. In *Visualizing argumentation*, pages 97–115. Springer, 2003.
- [88] B. Verheij. Accrual of arguments in defeasible argumentation. In *Proceedings of the Second Dutch/German Workshop on Nonmonotonic Reasoning*, pages 217–224, Utrecht, 1995.

- [89] G. A. Vreeswijk. Abstract argumentation systems. *Artificial intelligence*, 90(1-2):225–279, 1997.
- [90] D. Walton and E. C. Krabbe. *Commitment in dialogue: Basic concepts of interpersonal reasoning*. SUNY press, 1995.
- [91] M. E. Warren. What can democratic participation mean today? *Political theory*, 30(5):677–701, 2002.