

Empirical Evaluation of Abstract Argumentation: Supporting the Need for Bipolar and Probabilistic Approaches[☆]

Sylwia Polberg^a, Anthony Hunter^a

^a*University College London, Department of Computer Science, 66-72 Gower Street, London WC1E, United Kingdom*

Abstract

In dialogical argumentation, it is often assumed that the involved parties will always correctly identify the intended statements posited by each other and realize all of the associated relations, conform to the three acceptability states (accepted, rejected, undecided), adjust their views whenever new and correct information comes in, and that a framework handling only attack relations is sufficient to represent their opinions. Although it is natural to make these assumptions as a starting point for further research, dropping some of them has become quite challenging.

Probabilistic argumentation is one of the approaches that can be harnessed for more accurate user modelling. The epistemic approach allows us to represent how much a given argument is believed or disbelieved by a given person, offering us the possibility to express more than just three agreement states. It comes equipped with a wide range of postulates, including those that do not make any restrictions concerning how initial arguments should be viewed. Thus, this approach is potentially more suitable for handling beliefs of the people that have not fully disclosed their opinions or counterarguments with respect to standard Dung's semantics. The constellation approach can be used to represent the views of different people concerning the structure of the framework we are dealing with, including situations in which not all relations are acknowledged or when they are seen differently than intended. Finally, bipolar argumentation frameworks can be used to express both positive and negative relations between arguments.

In this paper we will describe the results of an experiment in which participants were asked to judge dialogues in terms of agreement and structure. We will compare our findings with the aforementioned assumptions as well as with the constellation and epistemic approaches to probabilistic argumentation and bipolar argumentation.

Keywords: Dialogical argumentation, probabilistic argumentation, abstract argumentation

[☆]This research is funded by EPSRC Project EP/N008294/1 "Framework for Computational Persuasion". We thank the reviewers for their valuable comments that helped us to improve this paper.

Contents

1	Introduction	4
2	Dung’s Argumentation Framework	6
3	Bipolar Argumentation	8
4	Probabilistic Argumentation	10
4.1	Epistemic Probabilistic Argumentation	11
4.1.1	Preferential Postulates	11
4.1.2	Explanatory Postulates	12
4.1.3	Value Postulates	12
4.1.4	Multi-type Postulates	13
4.1.5	Properties of Epistemic Postulates	13
4.1.6	Relation to Labeling-based Semantics	14
4.2	Constellation Probabilistic Argumentation	15
5	Experiment Description	16
5.1	Dialogues & Tasks	16
5.2	Recruitment	18
6	Results	18
6.1	Argument Graphs	21
6.1.1	Dialogue 1	23
6.1.1.1	Intended Graph Analysis	23
6.1.1.2	Participant-Sourced Graphs Analysis	24
6.1.2	Dialogue 2	31
6.1.2.1	Intended Graph Analysis	31
6.1.2.2	Participant-Sourced Graphs Analysis	32
6.2	Postulate Satisfaction	39
6.2.1	Dialogue 1	39
6.2.2	Dialogue 2	44
6.3	Statements vs Relations	48
6.3.1	Dialogue 1	49
6.3.1.1	Effect of a Relation on its Source	49
6.3.1.2	Effect of an Argument on the Relations it Carries Out	52
6.3.2	Dialogue 2	53
6.3.2.1	Effect of a Relation on its Source	53
6.3.2.2	Effect of an Argument on the Relations it Carries Out	57
6.4	Changes in Beliefs	60
7	Related Work	62
8	Conclusions	64
9	Statistics Appendix	
9.1	Dialogue 1	
9.1.1	Postulate Satisfaction	
9.1.2	Statements vs. Relations: Effect of a Relation on its Source	
9.1.3	Statements vs. Relations: Effect of an Argument on the Relations it Carries Out	
9.2	Dialogue 2	
9.2.1	Postulate Satisfaction	

9.2.2 Statements vs. Relations: Effect of a Relation on its Source

9.2.3 Statements vs. Relations: Effect of an Argument on the Relations it Carries Out . . .

9.3 Dialogues 1& 2: Changes in Beliefs

1. Introduction

At the heart of abstract argumentation lies Dung’s framework developed in [1], which treats arguments as abstract atomic entities that can be connected through an attack relation. Since its introduction, this framework has been endowed with numerous new semantics as well as generalized in various ways. These generalizations include structures that can represent new types of relations, such as support, as well as those that handle properties such as preferences or probabilities of arguments or relations [2]. Despite the available approaches, dialogical argumentation appears to rely heavily on frameworks handling only attacks between arguments, such as Dung’s framework. It also makes certain assumptions that are supported by the argumentation theory, but not always by empirical results. Argumentation, in many ways, simplifies human reasoning and appears to consider people to be in principle rational, good reasoners, that may be simply uninformed. Unfortunately, this approach might not always be adequate.

One of the core concepts of defeasible reasoning and therefore abstract argumentation is the fallibility of human perception. Thus, we need to be able to reason even with incomplete information and be prepared to retract our conclusions in the face of new data. From a certain perspective, most of the abstract argumentation approaches can be quite conservative in their handling of these issues. Although the defeasibility of arguments and notions such as undercutting attack are widely acknowledged, at the same time there is an assumption that a universal attack relation exists. In other words, it is believed that every person participating in a given dialogue will interpret arguments and the relations between them in exactly the same way, and every time new information is presented, it will be understood and linked to the existing arguments in the “intended and correct” way. However, this does not seem realistic in various scenarios. For example, the transcript of a TV debate “as it is” can be perceived differently from its processed version, i.e. one in which arguments are identified, organized and clearly presented w.r.t. the chosen argument structure. Not all relevant pieces of information have to be explicitly stated and it can happen that people taking part in the debate may be purposefully ambiguous. Consequently, we may have to deal with arguments that have implicit premises or conclusions and thus run the risk of participants interpreting them differently. Therefore, we need to acknowledge that people can view the nature of the framework associated with a given dialogue differently and that we should be able to represent such differences and uncertainties. Furthermore, these issues can themselves become a part of the discussion, which due to the aforementioned assumptions is rarely considered.

The fact that dialogical argumentation relies so heavily on attack-based frameworks for modelling purposes has also affected the way we perceive dialogues on a higher level. This approach became rather heavily conflict-centered, by which we understand that people participating in a dialogue view each other as opponents whose arguments should be defeated and knowledge “rectified”. Although this view may very well work in a court room, it has negative effects in the context of, for example, physical or mental health, as seen in [3]. For these applications, it may be more fruitful to consider dialogue parties as partners rather than opponents and use a more support-oriented approach in the discussion. Thus, rather than giving arguments against the current views of our dialogue partners, we try to motivate them or give them arguments for changing their behaviour and opinions. A given piece of information can also be framed in a positive or negative way, which affects the way people react to it [4, 5]. Hence, not only negative, but also positive relations between arguments can play a role in dialogical argumentation, and the use of various types of support should be considered.

Finally, the applicability of argumentation semantics in dialogical argumentation is not sufficiently verified [6]. The use of two or three values, as promoted by Dung’s semantics, oversimplifies the varying degrees to which we may agree or disagree with a given argument. For example, while a mother of an 8 year old girl might disagree with her staying the night at her best friend’s house, this is probably nowhere near to her disagreement with the child going alone to a heavy metal concert. Thus, there is a need for expressing varying degrees of agreement with a given argument. Another issue lies in the nature of the semantics themselves. For example, successful persuasion is often seen as synonymous with winning a dialogue game w.r.t. a given semantics, such as grounded or admissible. However, it is easy to see that in the cases such as a doctor persuading a patient to stop smoking or to go on a diet, holding a dialectically winning position and actually convincing a person to do or not to do something can be two completely different things. A

person taking part in a dialogue might also exhibit a number of perception biases or reasoning paradoxes and, due to inability or lack of cooperation, withhold information from the doctor. Hence, he or she does not have to reason in a way that adheres to the current argumentation semantics [7, 8].

To summarize, there is a need for empirical evaluation of argumentation approaches, which so far have received limited attention [9, 6, 10]. Although this task can be quite challenging, as what we are dealing with is human perception and judgment which are not always perfect, it is necessary. One has to realize that due to the nature of abstract argumentation, the validity of its approaches is typically inherited from other methods or obtained through novelty and technical correctness. For example, in order to defend the introduction of a new framework, we can explain how it can be instantiated with a given formalism, and show that the answers the framework produces exhibit certain desirable properties connected to the formalism we instantiated it with. The validity of the new framework thus depends on the validity of the formalisms it is linked to. Another possible way to argue in favour of a new approach is through the abstract argumentation itself, where the new semantics or framework is shown to be considerably different or more expressive than the existing ones. This method is often paired with presenting a realistic motivating example that argues in favor of the new framework or semantics. Although both of these strategies promote new ideas and innovation, there comes a time when they need to be verified in real life. Without empirical evidence, we can accidentally increase a gap between applying argumentation and *successfully* applying argumentation in real life situations.

In this paper we describe the results of an empirical study in which participants are presented with dialogues separated into five stages. At every stage, they are asked to declare why and how much they agree or disagree with the presented statements and how they view the relations between them. The purpose of this study was to investigate certain aspects of abstract argumentation as used by laypeople, rather than to make a number of initial assumptions that we prove or disprove based purely on the behaviour of experts. We have been able to draw a number of important observations and gained some evidence concerning various formalisms available in abstract argumentation. However, we would like to note that due to the exploratory nature of our experiments, our results should be treated as indicative and as a basis for further studies, rather than as an indisputable proof for or against a given argumentation approach.

Observation 1 The data supports the use of the **constellation approach** to probabilistic argumentation – people may interpret statements and relations between them differently, and not necessarily in the intended manner. The constellation approach can represent our uncertainty about the argument graphs describing our opponents views.

Observation 2 People may explicitly declare that two statements are connected, however, they might **not be sure of the exact nature of the relation** between them. We therefore also need to express the uncertainty that a person has about his or her own views, which can potentially be addressed with the constellation approach or with the introduction of a suitable framework.

Observation 3 The data supports the use of **epistemic approach** to probabilistic argumentation:

- people may assign levels of agreement to statements going beyond the 3-valued Dung’s approach,
- the epistemic postulates, in contrast to the standard semantics, can be highly adhered to and due to their nature, allow us to analyze why classical semantics may fail to explain the participants behaviour,
- the extended epistemic postulates allow us to model situations where the perceived “strength” of a relation might not necessarily be tightly connected with the level of agreement assigned to its source.

Observation 4 The data supports the use of **bipolar argumentation frameworks** – the notion of defence does not necessarily account for all of the positive relations between the statements viewed by the participants.

Observation 5 The data supports the use of **bipolar argumentation in combination with the prudent/careful approaches** – many additional attacks perceived by the participants can be explained by the existing notions of indirect conflicts in these settings.

Observation 6 The data shows that **people use their own personal knowledge** in order to make judgments and might not necessarily disclose it.

Observation 7 The data shows that presenting a new and correct piece of information that a given person was not aware of **does not necessarily lead to changing that person’s beliefs**.

This paper is organized as follows. In Sections 2 to 4 we review the necessary background on Dung’s argumentation frameworks, bipolar argumentation frameworks and probabilistic argumentation. In Section 5 we explain the set up of our experiment and analyze its results in Section 6. This includes the analysis of 1) the argument graphs created for the experiment as well as those provided by the participants, 2) the satisfaction rates of various epistemic postulates on the provided frameworks, 3) the connection between the level of agreement assigned to a given statement and the way the relations it carries out are perceived, and 4) changes in participants’ opinions throughout the experiment. We close this paper with the discussion on the works related to our study and pointers for future work.

2. Dung’s Argumentation Framework

We start with a review of abstract argumentation as proposed by Dung in [1]. He introduced the following, straightforward framework, that can be easily depicted using a graph where nodes play the role of arguments and edges represent conflicts:

Definition 2.1. A **Dung’s abstract argumentation framework** (AF for short) is a pair $F = (\mathcal{A}, \mathcal{R})$, where \mathcal{A} is a set of **arguments** and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ represents an **attack** relation.

An argument $A \in \mathcal{A}$ is an **attacker** of $B \in \mathcal{A}$ iff $(A, B) \in \mathcal{R}$. By abuse of notation, we say that a set of elements attacks another element if it contains an appropriate attacker. The way we decide which arguments can be accepted or rejected (or neither) is called a semantics. Depending on whether a set of arguments or a labeling is returned, we deal with the extension–based and labeling–based semantics [11, 12].

An **extension** is a set of arguments $\mathcal{E} \subseteq \mathcal{A}$ that satisfies the requirements imposed by a given semantics. The classical semantics [1] are built on the notion of defence:

Definition 2.2. Let $F = (\mathcal{A}, \mathcal{R})$ be a Dung’s framework. An argument $A \in \mathcal{A}$ is **defended** by a set $\mathcal{E} \subseteq \mathcal{A}$ in F ¹ if for every $B \in \mathcal{A}$ s.t. B attacks A , there exists $C \in \mathcal{E}$ that attacks B . A set $\mathcal{E} \subseteq \mathcal{A}$ is:

- (cf) **conflict–free** in F iff for no $A, B \in \mathcal{E}$, A is an attacker of B .
- (ad) **admissible** in F iff it is conflict–free in F and defends all of its members.
- (pr) **preferred** in F iff it is maximal w.r.t. set inclusion admissible in F .
- (co) **complete** in F iff it is admissible in F and all arguments defended by \mathcal{E} are contained in \mathcal{E} .
- (st) **stable** in F iff it is conflict–free in F and for every $A \in \mathcal{A} \setminus \mathcal{E}$ there exists an argument $B \in \mathcal{E}$ that attacks A .
- (gr) **grounded** in F iff it is the least complete extension of F .

¹Defence is often also referred to as acceptability: we say that A is acceptable w.r.t. \mathcal{E} iff \mathcal{E} defends A .




Figure 1: Sample argument graph

Example 1. Consider the Dung’s framework $F = (\mathcal{A}, \mathcal{R})$ with $\mathcal{A} = \{A, B, C, D, E\}$ and the attack relation $\mathcal{R} = \{(A, B), (C, B), (C, D), (D, C), (D, E), (E, E)\}$, as depicted in Figure 1. It has eight conflict-free extensions in total: $\{A, C\}, \{A, D\}, \{B, D\}, \{A\}, \{B\}, \{C\}, \{D\}$ and \emptyset . As B is attacked by an unattacked argument, it cannot be defended against it. Consequently, $\{A, C\}, \{A, D\}, \{A\}, \{C\}, \{D\}$ and \emptyset are our admissible sets. From this $\{A, C\}, \{A, D\}$ and $\{A\}$ are complete. We end up with two preferred extensions, $\{A, C\}$ and $\{A, D\}$. However, only $\{A, D\}$ is stable. Finally, $\{A\}$ is the grounded extension.

More types of extension-based semantics have been proposed in the recent years [12] and even though we will not consider all of them here, we would like to recall certain notions that will be useful in the next sections. As observed in the **prudent and careful semantics** [13, 14], based on the interplay of attack and defence, there might be additional positive and negative indirect interactions between the arguments that go beyond direct attack and defence:

Definition 2.3. Let $F = (\mathcal{A}, \mathcal{R})$ be a Dung’s framework and let $A, B \in \mathcal{A}$ be two arguments:

- A **indirectly attacks** B iff there exists an odd-length path from A to B in F .
- A **indirectly defends** B iff there exists an even-length path from A to B in F . The length of this path is not zero.
- A is **controversial** w.r.t. B iff A indirectly attacks and indirectly defends B .

Additionally, for arguments $A, B, C \in \mathcal{A}$, the pair (A, B) is **super-controversial** w.r.t. C iff A indirectly attacks C and B indirectly defends C .

Originally, these notions were used to impose stronger restrictions on the conflict-free sets. Although we will not recall the semantics themselves, the aforementioned notions will be useful to us in the context of this work. In particular, we will test whether certain elements of the argument graphs sourced from the participants of our experiments could potentially be the manifestations of the above definitions.

Example 1 (Continued). In the considered framework, we can observe that each of C, D and E function as indirect attackers and defenders of E . In addition to that, we can consider C and D as indirectly attacking each other. They also serve as their own defenders. We can observe that C and D also respectively indirectly attack and defend B .

Let us now focus on the labeling-based semantics, in particular those corresponding to the classical extension-based ones. An **argument labeling** is a total function $\mathcal{L} : \mathcal{A} \rightarrow \{\text{in}, \text{out}, \text{und}\}$ [11, 12]. By $\text{in}(\mathcal{L}), \text{out}(\mathcal{L})$ and $\text{und}(\mathcal{L})$ we denote the arguments mapped respectively to in, out and und(ecided) by \mathcal{L} . We will often write a labeling as a triple (I, O, U) , where I, O and U are sets of arguments mapped to in, out and und. We can now introduce the notion of legality, on which many semantics are based.

Definition 2.4. Let $F = (\mathcal{A}, \mathcal{R})$ be a Dung’s framework and let $\mathcal{L} : \mathcal{A} \rightarrow \{\text{in}, \text{out}, \text{und}\}$ be a labeling:

- $X \in \text{in}(\mathcal{L})$ is **legally in** iff all its attackers are in $\text{out}(\mathcal{L})$.
- $X \in \text{out}(\mathcal{L})$ is **legally out** iff it has an attacker in $\text{in}(\mathcal{L})$.
- $X \in \text{und}(\mathcal{L})$ is **legally und** iff not all of its attackers are in $\text{out}(\mathcal{L})$ and it does not have an attacker in $\text{in}(\mathcal{L})$.

Definition 2.5. Let $\mathcal{L} : \mathcal{A} \rightarrow \{\text{in}, \text{out}, \text{und}\}$ be a labeling:

(lab-cf) \mathcal{L} is **conflict-free** iff every $A \in \text{out}(\mathcal{L})$ is legally out and there are no arguments $A, B \in \text{in}(\mathcal{L})$ s.t. A is an attacker of B .

(lab-ad) \mathcal{L} is **admissible** iff every $A \in \text{in}(\mathcal{L})$ is legally in and every $A \in \text{out}(\mathcal{L})$ is legally out.

(lab-co) \mathcal{L} is **complete** iff it is admissible and every $A \in \text{und}(\mathcal{L})$ is legally und.

The preferred, stable and grounded labelings are obtained from the complete ones by using the constraints from Table 1.

Restriction on a complete labeling \mathcal{L}	Semantics
No argument $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{und}$	stable (ST)
Maximal no. of $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{in}$	preferred (PR)
Maximal no. of $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{out}$	preferred (PR)
Maximal no. of $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{und}$	grounded (GR)
Minimal no. of $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{in}$	grounded (GR)
Minimal no. of $A \in \mathcal{A}$ s.t. $\mathcal{L}(A) = \text{out}$	grounded (GR)

Table 1: Relation between different labelings

The properties of the labeling-based semantics and their correspondence to the classical extension-based family have already been studied in [11, 12]:

Theorem 2.6 ([12]). *Let $F = (A, R)$ be a Dung's framework and $\mathcal{E} \subseteq A$ be a σ -extension of F , where $\sigma \in \{\text{cf}, \text{ad}, \text{co}, \text{gr}, \text{pr}, \text{st}\}$. Let $\mathcal{E}^+ = \{B \mid \text{there exists } A \in \mathcal{E} \text{ s.t. } A \text{ is an attacker of } B\}$. Then $(\mathcal{E}, \mathcal{E}^+, A \setminus (\mathcal{E} \cup \mathcal{E}^+))$ is a σ -labeling of F . Let \mathcal{L} be a σ -labeling of F , where $\sigma \in \{\text{cf}, \text{ad}, \text{co}, \text{gr}, \text{pr}, \text{st}\}$. Then $\text{in}(\mathcal{L})$ is a σ -extension of F .*

Example 1 (Continued). Let us come back to the previously analyzed framework. Its admissible labelings are visible in Table 2. We can observe that one admissible extension can be associated with more than one labeling. However, out of the possible interpretations, only \mathcal{L}_3 , \mathcal{L}_9 and \mathcal{L}_{13} are complete. They are now also in one-to-one relation with the complete extensions. \mathcal{L}_3 is the grounded labeling, while \mathcal{L}_9 and \mathcal{L}_{13} are preferred. Only \mathcal{L}_{13} is stable; we can observe that this is the only labeling in which no argument is assigned und.

Extension	\emptyset	$\{A\}$	$\{A\}$	$\{C\}$	$\{C\}$	$\{D\}$	$\{D\}$	$\{A, C\}$	$\{A, C\}$	$\{A, D\}$	$\{A, D\}$	$\{A, D\}$	$\{A, D\}$	
	#	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_3	\mathcal{L}_4	\mathcal{L}_5	\mathcal{L}_6	\mathcal{L}_7	\mathcal{L}_8	\mathcal{L}_9	\mathcal{L}_{10}	\mathcal{L}_{11}	\mathcal{L}_{12}	\mathcal{L}_{13}
Labeling	A	und	in	in	und	und	und	und	in	in	in	in	in	in
	B	und	und	out	out	und	und	und	und	out	und	und	out	out
	C	und	und	und	in	in	out	out	in	in	out	out	out	out
	D	und	und	und	out	out	in	in	out	out	in	in	in	in
	E	und	und	und	und	und	und	out	und	und	und	out	und	out

Table 2: Admissible labelings of the framework $(\{A, B, C, D, E\}, \{(A, B), (C, B), (C, D), (D, C), (D, E), (E, E)\})$

3. Bipolar Argumentation

In Dung's framework, from direct attacks we can derive defence, which can be seen as a type of a positive indirect relation between arguments. However, defence does not account for all the possible forms of support between arguments, and a structure going beyond attack was required. Consequently, the notion of abstract support and the bipolar argumentation framework [15, 16, 17] were introduced, followed by the deductive [18], necessary [19, 20] and evidential supports [21, 22, 23], with the latter two developed in their own dedicated frameworks. Although there are significant differences between these supports and the way that their dedicated frameworks model them, particularly in the context of support cycles, they have also

been partially recreated in bipolar argumentation frameworks in order to perform a comparative study [17]. Further results on this topic can be found in [23, 24]. The bipolar argumentation framework itself is defined as follows:

Definition 3.1. The **bipolar argumentation framework** (BAF for short) is a tuple $(\mathcal{A}, \mathcal{R}, \mathcal{S})$, where \mathcal{A} is a set of **arguments**, $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ represents the **attack** relation and $\mathcal{S} \subseteq \mathcal{A} \times \mathcal{A}$ the **support**.

We say that there is a sequence of supports between arguments $A, B \in \mathcal{A}$ if there is a sequence of arguments (A, C_1, \dots, C_n, B) s.t. $ASC_1, C_1SC_2, \dots, C_nSB$.

One of the ways bipolar argumentation frameworks model support on the semantics level is by transforming it into attack [17, 25]. The base framework is extended with indirect conflicts, stemming from the interplay of existing attacks and supports. The resulting structure can then be evaluated like a Dung’s framework, particularly when semantics that are at least complete are concerned. The type of support we try to model affects what sort of additional indirect attacks are created. Although this approach does not account for all possible support semantics [24], it is sufficient for the purpose of this work. Typically, the following indirect conflicts are distinguished:

Definition 3.2. Let $BF = (\mathcal{A}, \mathcal{R}, \mathcal{S})$ be a BAF. The indirect attacks of BF are as follows:

- there is a **supported attack** from A to B iff there exists an argument C s.t. there is a sequence of supports from A to C and $(C, B) \in \mathcal{R}$.
- there is a **secondary attack**² from A to B iff there exists an argument C s.t. there is a sequence of supports from C to B and $(A, C) \in \mathcal{R}$.
- there is an **extended attack**³ from A to B iff there exists an argument C s.t. there is a sequence of supports from C to A and $(C, B) \in \mathcal{R}$.
- there is a **mediated attack** from A to B iff there exists an argument C s.t. there is a sequence of supports from B to C and $(A, C) \in \mathcal{R}$.
- there is a **super–mediated attack** from A to B iff there exists an argument C s.t. there is a sequence of supports from B to C and a direct or supported attack from A to C .
- there is a **super–extended attack**⁴ from A to B iff there exists an argument C s.t. there is a sequence of supports from C to A and a direct or secondary attack from C to B .

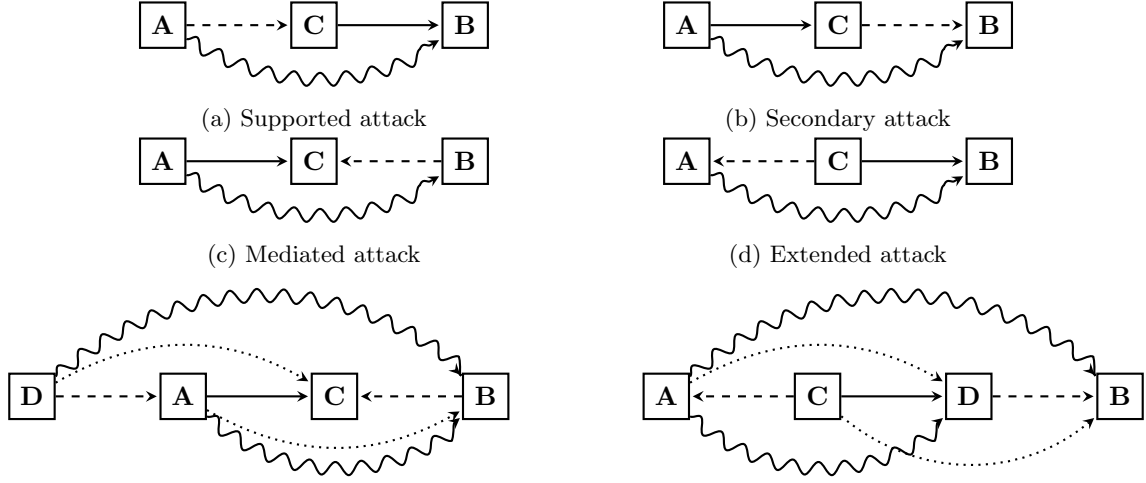
Additionally, the interactions between these auxiliary conflicts can also lead to the creation of new attacks. In particular, the super–mediated attack is a mixture of supported and mediated attacks (more details can be found in [24]). Nevertheless, the above notions are sufficient for the remainder of this report. We can see them exemplified in Figure 2.

We would like to stress that even though there are many types of conflicts available, it does not mean that all of them need to be used – the choice depends on what we intend to use a given BAF for. Usually, only some of the conflicts are studied at a time, particularly if we consider specialized forms of support rather than just the abstract type. What needs to be stated explicitly is that BAFs were meant as research frameworks for analyzing the different types of support and the consequences of their interplay with attack. Therefore, there is no “absolute” way to choose what sort of indirect attacks need to be taken into account and different interpretations of support might call for different attacks. In our experiment, we will focus on finding the existing notions that would allow us to reproduce the relations identified by the participants, rather than on fixing the interpretation of the used supports and stating that given indirect conflicts should have been used.

²This attack was also referred to as diverted in [16].

³We recall only one form of the extended attack, as other ones are already subsumed by the direct and secondary attacks.

⁴In [25] this attack is referred as an n+-attack.



(e) Super-mediated attacks (squiggly edges) and existing supported and mediated attacks (dotted edges) (f) Super-extended attacks (squiggly edges) and existing secondary and extended attacks (dotted edges)

Figure 2: Indirect attacks in BAFs. Solid edges represent direct attacks, dashed edges represent support, and squiggly or dotted edges represent indirect attacks.

Example 2. Consider the bipolar argumentation framework $BF = (\mathcal{A}, \mathcal{R}, \mathcal{S})$ with $\mathcal{A} = \{A, B, C, D, E\}$, $\mathcal{R} = \{(C, B), (C, D), (D, C), (E, E)\}$ and $\mathcal{S} = \{(A, B), (D, E)\}$, depicted in Figure 3. We can create the following indirect conflicts for this framework. Since D supports E , which is a self-attacker, we can create a supported attack (D, E) and a (super) mediated one (E, D) . These conflicts also lead to the super-mediated one (D, D) . Due to the same support, we can create a secondary attack (C, E) and a (super) extended one (E, C) . As C attacks B , which is supported by A , we can create a (super) mediated attack (C, A) .

If we decide to use all of the aforementioned attacks, then the Dung's framework associated with BF is $F = (\mathcal{A}, \mathcal{R}')$, where $\mathcal{R}' = \mathcal{R} \cup \{(C, A), (C, E), (D, D), (D, E), (E, C), (E, D)\}$. The sets \emptyset and $\{C\}$ are its admissible and complete extensions. $\{C\}$ is the stable and preferred extension, while \emptyset is grounded. Following the approach from [17], we can treat these extensions as the extensions of BF .

If we consider only the secondary and (super) extended attacks, we obtain a Dung's graph with the set of attacks $\mathcal{R}'' = \mathcal{R} \cup \{(C, E), (E, C)\}$. The admissible extensions of this framework are $\emptyset, \{A\}, \{C\}, \{D\}, \{A, C\}, \{A, D\}, \{B, D\}$ and $\{A, B, D\}$. Out of this, $\{A\}, \{A, C\}$ and $\{A, B, D\}$ are complete, with the first set being grounded and the other two preferred. Only $\{A, C\}$ is stable.




Figure 3: Sample bipolar argument graph. Solid edges represent attack, dashed edges represent support and squiggly edges represent indirect attacks.

4. Probabilistic Argumentation

The proposals for probabilistic argumentation extend Dung's framework to address various aspects of uncertainty arising in argumentation. The two main approaches to probabilistic argumentation are the

constellations and the epistemic approaches [26]. In the **constellations approach**, the uncertainty is in the topology of the graph. This approach is useful when one agent is not sure what arguments and attacks another agent is aware of. This can be captured by a probability distribution over the space of possible argument graphs, where each graph has a chance of being the real model of the agent. In the **epistemic approach**, the topology of the argument graph is fixed, but there is uncertainty as to the degree to which each argument is believed. It thus provides us with a more fine-grained approach towards modelling the acceptability of arguments. In this section we provide a brief revision of these two probabilistic formalisms and for further reading concerning the differences between them we refer the readers to [26, 27].

4.1. Epistemic Probabilistic Argumentation

We now turn to the epistemic approach to probabilistic argumentation [28, 26, 29, 30, 31]. In this section we review and extend the results from [32].

Definition 4.1. A **mass distribution** over arguments \mathcal{A} is a function $P : 2^{\mathcal{A}} \rightarrow [0, 1]$ s.t. $\sum_{\mathcal{E} \subseteq \mathcal{A}} P(\mathcal{E}) = 1$. The **probability of an argument** $A \in \mathcal{A}$ is $P(A) = \sum_{\mathcal{E} \subseteq \mathcal{A} \text{ s.t. } A \in \mathcal{E}} P(\mathcal{E})$.

The probability of a single argument is understood as the belief that an agent has in it, i.e. we say that an agent believes an argument A to some degree when $P(A) > 0.5$, disbelieves an argument to some degree when $P(A) < 0.5$, and neither believes nor disbelieves an argument when $P(A) = 0.5$. This belief can be interpreted in various ways [26], for example, if we assume that an argument has a standard premise–conclusion structure, then the degree of belief in an argument can be seen as the degree to which the agent believes the premises and the conclusion drawn from those premises.

Just like argument graphs come equipped with different types of semantics that are meant to capture various intuitions as to what is a “good” collection of arguments (and/or attacks), the probabilistic frameworks are accompanied by a number of postulates that capture the properties of a probability distribution. While classical semantics tend to represent a number of properties at the same time, a single postulate focuses on a single aspect at a time. Thus, while there are many epistemic postulates, they tend to be quite simple. They also allow a more detailed view on the participant behaviour and can allow us to analyze the cases in which classic semantics may fail to explain it. Consequently, they can provide more feedback to argumentation systems, such as for computational persuasion, than normal semantics do. For example, the low performance of the argumentation semantics such as complete does not really inform the system what aspect of the participant reasoning does not meet the semantics requirements. In contrast, through the use of the postulates we could analyze whether the issue lies within conflict–freeness, defense or not accepting/rejecting arguments that should be accepted/rejected.

We will recall some of the existing postulates for the epistemic semantics and introduce certain argument–centered alternatives of the properties from [27], namely the preferential postulate PRE, strict STC, protective PRO, restrained RES, discharging DIS, guarded GRD, trusting TRU, anticipating ANT, demanding DEM, binary BIN and n–valued VALⁿ. The available postulates can be grouped in various ways. In what follows, we separate them into the following, not necessarily disjoint, types: preferential (rational), explanatory and value families.

4.1.1. Preferential Postulates

The first family of postulates focuses on resolving a conflict between two arguments and deciding how believed or disbelieved the attacker and attackee should be. They are primarily of the form “*for every* $A, B \in \mathcal{A}$, *if* $(A, B) \in \mathcal{R}$, *then* X *holds*”, where X specifies the conditions on the beliefs assigned to the attacker and the attackee. The intuitions behind the introduced postulates are as follows. Let us assume that we believe an argument attacking another argument. The question now is what should be our belief in the attackee. The least restrictive, PRE postulate, allows us to believe the attackee as long as it is believed more than the attacker. If we follow RAT, then we do not believe the attackee, and if STC, then we disbelieve it. In a dual manner, the PRO postulate ensures that if we believe the attackee, we disbelieve its attackers. The RES postulate strengthens PRO and STC and is equivalent to saying that if two arguments are in conflict, at least one of them should be disbelieved. Hence, we no longer have the option to be undecided about an argument.

Definition 4.2. A probability mass distribution P for F is:

- (PRE) preferential** if for every $A, B \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, if $P(A) > 0.5$ and $P(B) > 0.5$, then $P(A) < P(B)$.
- (RAT) rational** if for every $A, B \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, $P(A) > 0.5$ implies $P(B) \leq 0.5$.
- (STC) strict** if for every $A, B \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, $P(A) > 0.5$ implies $P(B) < 0.5$.
- (PRO) protective** if for every $A, B \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, $P(B) > 0.5$ implies $P(A) < 0.5$.
- (RES) restrained** if for every $A, B \in \mathcal{A}$, if $(A, B) \in \mathcal{R}$ then $P(A) \geq 0.5$ implies $P(B) < 0.5$.

4.1.2. Explanatory Postulates

We now come to the explanatory family. The purpose of the postulates of this type is to demand that the degree of belief assigned to an argument is justified by the degrees of belief associated with the arguments related to it. They are roughly of the form “for every argument $B \in \mathcal{A}$, if $P(B)$ meets a condition X then it has an attacker A s.t. $P(A)$ meets condition Y ” or “for every argument $B \in \mathcal{A}$, if for all of its attackers A , $P(A)$ meets condition Y , then $P(B)$ meets a condition X ”. For instance, DIS demands that if an argument is disbelieved, then it possesses a believed attacker. GRD postulate relaxes this requirement by allowing the use of undecided attackers as well. On the other hand, the TRU property requires us to believe an argument when we have no reason against it, i.e. when all of its attackers are disbelieved. ANT modifies TRU by saying that lack of believed attackers is a good reason to believe the attackee. Finally, the DEM property requires that a completely disbelieved argument has to be paired with a completely believed attacker, and a completely believed argument can only be attacked by completely disbelieved arguments.

Definition 4.3. A probability mass distribution P for F is:

- (DIS) discharging** if for every $B \in \mathcal{A}$, if $P(B) < 0.5$ then there exists an argument $A \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$ and $P(A) > 0.5$.
- (GRD) guarded** if for every $B \in \mathcal{A}$, if $P(B) < 0.5$ then there exists an argument $A \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$ and $P(A) \geq 0.5$.
- (TRU) trusting** if for every $B \in \mathcal{A}$, if $P(A) < 0.5$ for all $A \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, then $P(B) > 0.5$.
- (ANT) anticipating** if for every $B \in \mathcal{A}$, if $P(A) \leq 0.5$ for all $A \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, then $P(B) > 0.5$.
- (DEM) demanding** if for every $A \in \mathcal{A}$, if $P(A) = 1$, then for every $B \in \mathcal{A}$ s.t. $(B, A) \in \mathcal{R}$, $P(B) = 0$, and if $P(A) = 0$, then $\exists B \in \mathcal{A}$ s.t. $(B, A) \in \mathcal{R}$ and $P(B) = 1$.

4.1.3. Value Postulates

All of the previously listed postulates are relatively general and tell us whether to believe or disbelieve an argument, but not to what degree. This is where the value family of postulates comes in. The FOU and SFOU postulates tell us how much we should believe initial arguments. OPT and SOPT provide us with lower bounds for the degrees of belief we have in an element based on the beliefs we have in its attackers. The BIN postulate prohibits any indecisiveness we may have about the arguments. MIN, NEU, MAX permit the use of only 0, 0.5 and 1 degrees of belief respectively, while TER allows for all three. Finally, with VAL^{*n*} postulate we can distinguish those distributions that assign to their arguments no more than n distinct values in total.

Definition 4.4. A probability mass distribution P for F is:

- (SFOU) semi-founded** if $P(A) \geq 0.5$ for every initial $A \in \mathcal{A}$.
- (FOU) founded** if $P(A) = 1$ for every initial $A \in \mathcal{A}$.

(SOPT) semi-optimistic if $P(A) \geq 1 - \sum_{B \in \{A\}^-} P(B)$ for every $A \in \mathcal{A}$ that is not initial.

(OPT) optimistic if $P(A) \geq 1 - \sum_{B|(B,A) \in \mathcal{R}} P(B)$ for every $A \in \mathcal{A}$.

(BIN) binary if for no $A \in \mathcal{A}$, $P(A) = 0.5$.

(TER) ternary if $P(A) \in \{0, 0.5, 1\}$ for every $A \in \mathcal{A}$.

(NEU) neutral if $P(A) = 0.5$ for every $A \in \mathcal{A}$.

(MAX) maximal if $P(A) = 1$ for every $A \in \mathcal{A}$.

(MIN) minimal if $P(A) = 0$ for every $A \in \mathcal{A}$.

(VALⁿ) n-valued if $|\{x \mid \exists A \in \mathcal{A} \text{ s.t. } P(A) = x\}| \leq n$.

4.1.4. Multi-type Postulates

Finally, we have the postulates that can be seen as shared between the families. In particular, on the intersection of the rational (preferential) and value types are the COH, INV and JUS postulates. COH gives us the upper bound of the belief we might have in an argument based on its (strongest) attacker. By combining this upper bound with the lower bound from the OPT postulate, we receive the JUS property. The INV postulate requires that the belief in the attackee is dual to the belief in the attacker. Please note that the DEM, SFOU and FOU postulates can be seen as shared between the value and explanatory families, rather than belonging to just one of them. We leave it to the reader to classify them as he or she sees fit.

Definition 4.5. A probability mass distribution P for F is:

(COH) coherent if for every $A, B \in \mathcal{A}$ s.t. $(A, B) \in \mathcal{R}$, $P(A) \leq 1 - P(B)$.

(INV) involutory if for every $A, B \in \mathcal{A}$, if $(A, B) \in \mathcal{R}$ then $P(A) = 1 - P(B)$.

(JUS) justifiable if it is coherent and optimistic.

4.1.5. Properties of Epistemic Postulates

In Proposition 4.6 and Figure 4 we show some of the relations between the presented postulates, extending the results available in [33]. These properties make it more explicit which postulates are more restrictive than the others and show additional connections between the different families of postulates we have presented.

Proposition 4.6. Let F be an argument graph and let \mathcal{P}_X be the collection of all distributions on F satisfying postulate X . The following holds:

$$\begin{aligned} \mathcal{P}_{OPT}(F) &= \mathcal{P}_{SOPT}(F) \cap \mathcal{P}_{FOU}(F) & \mathcal{P}_{TER}(F) \cap \mathcal{P}_{TRU}(F) &\subseteq \mathcal{P}_{FOU}(F) \\ (\mathcal{P}_{COH}(F) \cap \mathcal{P}_{TER}(F)) &= (\mathcal{P}_{PRO}(F) \cap \mathcal{P}_{STC}(F) \cap \mathcal{P}_{TER}(F)) \\ \mathcal{P}_{INV}(F) \cap \mathcal{P}_{FOU} &\subseteq \mathcal{P}_{DEM}(F) \end{aligned}$$

Example 1 (Continued). Let us come back to the argument graph $F = (\{A, B, C, D, E\}, \{(A, B), (C, B), (C, D), (D, C), (D, E), (E, E)\})$ from Figure 1. In Table 3 we have listed a number of probability distributions and marked which postulates are or are not satisfied by them under F . For the sake of simplicity, we have only focused on the resulting argument probabilities, not on the full description of the distributions.

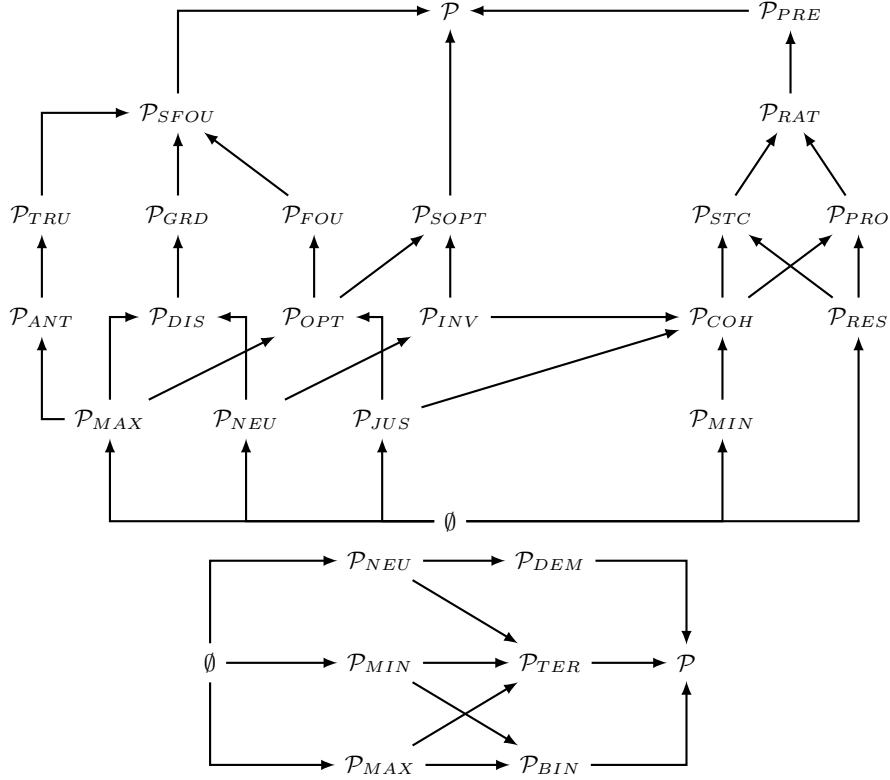


Figure 4: Classes of probability functions where $\mathcal{P}_{\mu_1} \rightarrow \mathcal{P}_{\mu_2}$ denotes $\mathcal{P}_{\mu_1} \subseteq \mathcal{P}_{\mu_2}$

#	A	B	C	D	E	PRE	RAT	STC	PRO	RES	COH	JUS	INV	DIS	GRD	TRU	ANT	DEM	SFOU	FOU	SOPT	OPT
P_1	0.6	0.5	0.2	0.4	0.8	×	×	×	×	×	×	×	×	×	×	×	×	✓	✓	×	×	×
P_2	0.3	0.5	0.6	0.2	0.5	✓	✓	×	✓	×	×	×	×	×	×	×	×	✓	×	×	×	×
P_3	0.7	0.4	0.6	0.4	0.5	✓	✓	✓	✓	×	×	×	×	✓	✓	✓	×	✓	✓	×	✓	×
P_4	1	0.5	1	0	0.4	✓	✓	×	✓	✓	×	×	×	×	×	×	×	✓	✓	✓	×	×
P_5	1	0	1	0	0.5	✓	✓	✓	✓	×	✓	✓	×	✓	✓	✓	×	✓	✓	✓	✓	✓
P_6	1	0	0.5	0.4	0.5	✓	✓	✓	✓	×	✓	×	×	×	✓	×	×	✓	✓	✓	×	×
P_7	0.5	0	1	0	0.5	✓	✓	✓	✓	×	✓	×	×	✓	✓	×	×	✓	✓	×	✓	×
P_8	0.5	1	0.5	0.5	0.5	✓	✓	✓	×	×	×	×	×	✓	✓	×	×	×	✓	×	✓	×
P_9	0.6	1	0.4	0.7	0.5	✓	×	×	×	×	×	×	×	✓	✓	✓	✓	×	✓	×	✓	×
P_{10}	0.5	0.5	0.5	0.5	0.5	✓	✓	✓	✓	×	✓	×	✓	✓	✓	×	×	✓	✓	×	✓	×

Table 3: Satisfaction of epistemic postulates on Dung’s framework from Figure 1.

4.1.6. Relation to Labeling-based Semantics

In the previous sections we could have observed that the epistemic approach offers a wide variety of postulates describing certain properties that a given probability function may or may not possess. In this section, we would like to show how some of the combinations of these postulates can capture the intuitions behind the standard labeling based semantics. This information will become useful once we analyze the postulate adherence in our experiment.

Definition 4.7. Let $F = (\mathcal{A}, \mathcal{R})$ be an argument graph and let P be a probability distribution. The

epistemic labeling associated with P is $\mathcal{L}_P = (I, O, U)$, where $I = \{A \in \mathcal{A} \mid P(A) > 0.5\}$, $O = \{A \in \mathcal{A} \mid P(A) < 0.5\}$, and $U = \{A \in \mathcal{A} \mid P(A) = 0.5\}$. A labeling \mathcal{L} and a probability function P are *congruent*, denoted $\mathcal{L} \sim P$, if for all $A \in \mathcal{A}$ we have that $\mathcal{L}(A) = \text{in} \Leftrightarrow P(A) = 1$, $\mathcal{L}(A) = \text{out} \Leftrightarrow P(A) = 0$, and $\mathcal{L}(A) = \text{und} \Leftrightarrow P(A) = 0.5$.

Please note that if $\mathcal{L} \sim P$, then $\mathcal{L} = \mathcal{L}_P$, i.e. if a labeling and a probability function are **congruent** then this labeling is also an epistemic labeling of this function.

Proposition 4.8. *Let P be an epistemic distribution and let $F = (\mathcal{A}, \mathcal{R})$ be a Dung’s framework.*

- $P \in \mathcal{P}_{RAT}(F) \cap \mathcal{P}_{DIS}(F)$ iff $\mathcal{L}_P \in \text{cf}(F)$.
- $P \in \mathcal{P}_{PRO}(F) \cap \mathcal{P}_{DIS}(F)$ iff $\mathcal{L}_P \in \text{ad}(F)$.
- $P \in \mathcal{P}_{PRO}(F) \cap \mathcal{P}_{STC}(F) \cap \mathcal{P}_{DIS}(F) \cap \mathcal{P}_{TRU}(F)$. iff $\mathcal{L}_P \in \text{co}(F)$.

If we were to consider retrieving the extension–based semantics through epistemic postulates, the above results should be rephrased in the following manner: if a probability distribution satisfies the given postulates, then the set of *in* arguments of the associated labeling is an extension of a given type, and if a set of arguments is an extension of a given type, then there exists a probability distribution satisfying the given postulates s.t. the set of believed arguments coincides with the extension. The only relaxation we are aware of is in the case of conflict–free semantics, where a set of arguments is conflict–free iff the associated probability distribution is rational.

Example 1 (Continued). For every labeling listed in Table 2, we can create a congruent probability function and show that it satisfies the required postulates. For example, we can consider the admissible labeling \mathcal{L}_6 and its congruent distribution s.t. $P(A) = P(B) = P(E) = 0.5$, $P(C) = 0$ and $P(D) = 1$. The protective and discharging postulates are easily satisfied, however, the trusting and strict ones are not – in particular, A should have been believed and B, E disbelieved. Addressing these issues would produce a distribution that is associated with the complete labeling \mathcal{L}_{13} .

Let us now look at the distributions in Table 3. All P_3, P_5, P_7, P_8 and P_{10} are both rational and discharging. It is easy to verify that the associated labelings $\mathcal{L}_{P_3} = \mathcal{L}_{P_5} = \{A : \text{in}, B : \text{out}, C : \text{in}, D : \text{out}, E : \text{und}\}$, $\mathcal{L}_{P_7} = \{A : \text{und}, B : \text{out}, C : \text{in}, D : \text{out}, E : \text{und}\}$, $\mathcal{L}_{P_8} = \{A : \text{und}, B : \text{in}, C : \text{und}, D : \text{und}, E : \text{und}\}$ and $\mathcal{L}_{P_{10}} = \{A : \text{und}, B : \text{und}, C : \text{und}, D : \text{und}, E : \text{und}\}$ are conflict–free. All of them, with the exception of \mathcal{L}_{P_8} , are also admissible – at the same time, we can observe that P_8 is not protective. Finally, only \mathcal{L}_{P_3} and \mathcal{L}_{P_5} are complete, which is reflected by P_3 and P_5 being the only of the listed functions that are additionally strict and trusting.

Therefore, as we can observe, epistemic probability, understood as the degree of belief, can also express various levels of agreement in the Dung’s sense, i.e. accepting, rejecting and being undecided about an argument.

4.2. Constellation Probabilistic Argumentation

Here we review the constellation approach to probabilistic argumentation from [34], which extends the methods from [35] and [36]. The constellation approach allows us to represent the uncertainty over the topology of the graph. Each subgraph of the original graph is assigned a probability which is understood as the chances of it being the actual argument graph of the agent. It can be used to model what arguments and attacks an agent is aware of. If our uncertainty is about which arguments appear in the graph, then only the full (induced) subgraphs of the argument graph have a non–zero probability. If we are only uncertain about which attacks appear, then it is the spanning subgraphs of the argument graph that can have a non–zero probability.

Definition 4.9. Let $F = (\mathcal{A}, \mathcal{R})$ and $F' = (\mathcal{A}', \mathcal{R}')$ be two argument graphs. F' is a subgraph of F , denoted $F' \sqsubseteq F$, iff $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{R}' \subseteq (\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}$. The set of subgraphs of F is $\text{Sub}(F) = \{F' \mid F' \sqsubseteq F\}$. A subgraph $(\mathcal{A}', \mathcal{R}')$ is **full** iff $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{R}' = (\mathcal{A}' \times \mathcal{A}') \cap \mathcal{R}$. A subgraph $(\mathcal{A}', \mathcal{R}')$ is **spanning** iff $\mathcal{A}' = \mathcal{A}$ and $\mathcal{R}' \subseteq \mathcal{R}$.

Definition 4.10. A **subgraph distribution** is a function $P^c : \text{Sub}(F) \rightarrow [0, 1]$ with $\sum_{F' \in \text{Sub}(F)} P^c(F') = 1$. A subgraph distribution P^c is a **full subgraph distribution** iff if $(\mathcal{A}', \mathcal{R}')$ is not a full subgraph, then $P^c((\mathcal{A}', \mathcal{R}')) = 0$. A subgraph distribution P^c is a **spanning subgraph distribution** iff if $(\mathcal{A}', \mathcal{R}')$ is not a spanning subgraph, $P^c((\mathcal{A}', \mathcal{R}')) = 0$.

The constellation semantics can be seen as a two-level construct. Each graph in the subgraph distribution is evaluated with a given base semantics, the results of which are later paired with the probabilities of the frameworks producing them, thus leading to constellation semantics. Although we will not be using this particular method for evaluating argument graphs in this paper, we briefly recall and exemplify it so that it is more clear that it is different from the epistemic approach. Further details can be found in [26, 27].

Determining the probability that a set of arguments or a labeling follows the semantics of a particular type (e. g. grounded, preferred, etc.) is done by collecting the probabilities of the subgraphs producing the desired extensions or labelings. In a similar fashion, we can derive the probability of an argument being accepted in a labeling of a given type.

Definition 4.11. For $\mathcal{E} \subseteq A$ and $\sigma \in \{\text{cf}, \text{ad}, \text{co}, \text{pr}, \text{gr}, \text{st}\}$, the probability that $\mathcal{L} : \mathcal{E} \rightarrow \{\text{in}, \text{out}, \text{und}\}$ is a σ -labeling is:

$$P_\sigma(\mathcal{L}) = \sum_{F' \in \text{Sub}(F) \text{ s.t. } \mathcal{L} \in \sigma(F')} P^c(F')$$

Definition 4.12. Given a semantics $\sigma \in \{\text{ad}, \text{co}, \text{pr}, \text{gr}, \text{st}\}$, the probability that $A \in \mathcal{A}$ is assigned an in status in a σ -labeling is

$$P_\sigma(A) = \sum_{F' \in \text{Sub}(F) \text{ s.t. } \mathcal{L} \in \sigma(F') \text{ and } A \in \text{in}(\mathcal{L})} P(F')$$

Example 3. Consider the graph $F = (\{A, B\}, \{(A, B)\})$. Its subgraphs are $F_1 = (\{A, B\}, \{(A, B)\})$, $F_2 = (\{A, B\}, \emptyset)$, $F_3 = (\{A\}, \emptyset)$, $F_4 = (\{B\}, \emptyset)$ and $F_5 = (\emptyset, \emptyset)$. Out of them, F_1, F_3, F_4 and F_5 are full, and F_1 and F_2 are spanning. Consider the following subgraph distribution P^c : $P^c(F_1) = 0.09$, $P^c(F_2) = 0.81$, $P^c(F_3) = 0.01$ and $P^c(F_4) = 0.09$ and $P^c(F_5) = 0$. The probability of a given set being a grounded extension is as follows: $P_{\text{gr}}(\{A, B\}) = P^c(F_2) = 0.81$; $P_{\text{gr}}(\{A\}) = P^c(F_1) + P^c(F_3) = 0.1$; $P_{\text{gr}}(\{B\}) = P^c(F_4) = 0.09$; and $P_{\text{gr}}(\{\}) = P^c(F_5) = 0$. Therefore $P_{\text{gr}}(A) = 0.91$ and $P_{\text{gr}}(B) = 0.9$.

For further reading concerning the constellation approach to probabilistic argumentation we refer the readers to [34, 35, 36, 37]. Computational results can be found in [38, 39, 40], and in [41, 42] we can find approaches combining the constellation probabilities with generalizations of Dung’s framework. Finally, in [43] we can find a characterization of one of the versions of the constellation approach in terms of probabilistic logic.

5. Experiment Description

In this section we explain how the experiment was carried out. The purpose of our study was to gather the opinions on dialogues concerning the topic of flu vaccines. To this end, we prepared two separate dialogues and asked the participants we recruited online to judge at least one of them. In the following sections we will show the dialogues we have used, explain the tasks given to the participants and describe the recruitment process.

5.1. Dialogues & Tasks

The dialogues presented to the participants can be seen in Tables 4 and 5. The first dialogue discusses the use of the flu vaccine by hospital staff members, while the second concerns the safety of the children flu vaccine. Every dialogue is split into five steps in which two fictional discussants, person 1 (**P1**) and person 2 (**P2**), take turns in presenting their opinions. A given statement, once uttered, is visible throughout the rest of the dialogue. Although the general topic is the same, i.e. concerns flu vaccinations, we can observe that the statements used in the dialogues do not overlap. The dialogues have been created based on the

Table 4: Dialogue 1 between people P1 and P2. This dialogue starts with P1 claiming that hospital staff do not need to receive flu shots, to which P2 objects. The two counterarguments of P1 are then defeated by P2. The table presents at which steps a given statement was visible, who uttered it and what was its content.

Steps	Person	Statement	Content
1 to 5	P1	A	Hospital staff members do not need to receive flu shots.
1 to 5	P2	B	Hospital staff members are exposed to the flu virus a lot. Therefore, it would be good for them to receive flu shots in order to stay healthy.
2 to 5	P1	C	The virus is only airborne and it is sufficient to wear a mask in order to protect yourself. Therefore, a vaccination is not necessary.
3 to 5	P2	D	The flu virus is not just airborne, it can be transmitted through touch as well. Hence, a mask is insufficient to protect yourself against the virus.
4 to 5	P1	E	The flu vaccine causes flu in order to gain immunity. Making people sick, who otherwise might have stayed healthy, is unreasonable.
5	P2	F	The flu vaccine does not cause flu. It only has some side effects, such as headaches, that can be mistaken for flu symptoms.

Table 5: Dialogue 2 between people P1 and P2. This dialogue starts with P1 claiming that the vaccine is not safe, to which P2 objects, and the discussion proceeds to revolve around whether it contains mercury-based compounds or not. The table presents at which steps a given statement was visible, who uttered it and what was its content.

Steps	Person	Statement	Content
1 to 5	P1	A	The flu vaccine is not safe to use by children.
1 to 5	P2	B	The flu vaccine does not contain poisonous components and is safe to use.
2 to 5	P1	C D	The vaccine contains some mercury compounds. The mercury compounds are poisonous and therefore the vaccine is not safe to use.
3 to 5	P2	E F	The child vaccine does not contain any mercury compounds. The virus is only accompanied by stabilizers and possibly trace amounts of antibiotics used in its production.
4 to 5	P1	G	The vaccine contains a preservative called thimerosal which is a mercury-based compound.
5	P2	H	Children receive the nasal spray vaccine and thimerosal has been removed from it over 15 years ago.

information found on the Centers for Disease Control and Prevention (CDC USA) and National Health Services (NHS UK) websites.

During the experiment, at every stage of the dialogue the participants were presented with three tasks - **Agreement**, **Explanation** and **Relation** - with an additional **Awareness** question at the end of the dialogue:

- **Agreement:** the participants were asked to state how much they agree or disagree with a given statement. They were allowed to choose one of the seven options (*Strongly Agree*, *Agree*, *Somewhat Agree*, *Neither Agree nor Disagree*, *Somewhat Disagree*, *Disagree*, *Strongly Disagree*) or select the answer *Don't Know*.
- **Explanation:** the participants were asked to explain the chosen level of agreement for every statement. In particular, we requested them to provide us with any reason they may have for disagreement that was not mentioned in the dialogue and to explain reasons for changing their opinions compared to the previous step in the discussion (if applicable).

- **Relation:** the participants were asked to state how they viewed the relation between the statements. For every listed pair, they could say whether one statement was *A good reason against*, *A somewhat good reason against*, *Somewhat related, but can't say how*, *A somewhat good reason for*, *A good reason for* the other statement or select the answer *N/A* (i.e. that the statements were unrelated). The questions were kept with the flow of the dialogue, i.e. the source of a given relation is preceded by the target in the presented discussion. For example, we would ask how statement **F** is related to **A**, but not the other way around. Thus, the temporal aspects of the dialogues are taken into account by the graph, similarly as is done in the analysis of online discussions [44, 45]. This also reduced the number of questions we had to ask and therefore simplified the task for the participants.
- **Awareness:** the participants were asked which of the presented statements they were familiar with prior to the experiment. By this we understand that, for example, someone they know may have expressed the opinion contained in a given statement, they might have read it in some source, or generally heard about it before. To put it simply, we wanted to know what statements the participants were aware of, independently of whether they agreed or disagreed with them.

5.2. Recruitment

The recruitment of the participants was done using the Amazon Mechanical Turk (AMT for short) and the survey itself was carried out using the SurveyMonkey, which are both common platforms for experiments of this type. In addition to the tasks explained in the previous section, the participants were subjected to an additional language exercise, two attention checks and a comprehension test in order to ensure that they had sufficient skills to complete the test and that they worked honestly. The language exercise was on an intermediate level in terms of difficulty. The attention checks presented the participants with two sentences requesting them to choose particular answers. They were meant to disqualify the participants who were too distracted during the experiment or simply resorted to random clicking in order to complete the survey. Finally, in the comprehension exercise, they were asked to select the answers that described what the statements presented in the dialogue concerned. Prior to the survey, the participants were warned of such tests. We also provided them with explanations and examples of the tasks they would be asked to complete, including what should be understood as a reason for or against or what we meant by being aware of a given statement. We also requested the participants not to use Google or Wikipedia in order to verify the statements in the dialogues. At the very end of these instructions the participants obtained the code used to unlock the actual questionnaire without which they could not proceed. Moreover, it was also necessary for them to accept the terms and conditions of our experiment, which included familiarizing themselves with these explanations. Hence, we took all reasonable steps in informing the participants what to expect.

We ran the survey until we had collected 80 survey responses (40 per dialogue) in which the participants had sufficiently high scores in the language, attention and comprehension tests. As we have noted previously, the participants - if they wished to do so - were allowed to judge both of the dialogues. We have found 11 people that have completed both tests, which means we have recruited a total of 69 unique participants. The total number of entries (not participants) was equal to 156, this brings us to an acceptability rate of around 51%. The analysis of the demographics data on our participants can be found in the auxiliary appendix at [46]. This research project has been approved by the designated ethics officer in the Computer Science Department at University College London. The participants could withdraw from the experiment at any point in time and could refuse to provide any piece of information they deemed too private.

6. Results

In this section we will review the results of our study. We will first describe the argument graphs obtained from our participants as well as the ones we had in mind when creating the dialogues. Given these structures, we will discuss if and how the declared levels of agreement satisfy the epistemic postulates we have recalled in Section 4.1. Finally, we will discuss the changes in participants' beliefs throughout the dialogues.

Before we start, we would like to introduce some necessary notions. As we have explained previously, in addition to declaring that one statement is a good reason for or against another statement, a possible answer in the relation task was *Somewhat related, but can't say how*. In order to be able to represent this dependency that a participant could not classify, we propose the notion of a tripolar argument graph:

Definition 6.1. The **tripolar argument graph** is a tuple $TF = (\mathcal{A}, \mathcal{R}, \mathcal{S}, \mathcal{D})$ where \mathcal{A} is a set of arguments, $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ is the set of attacks, $\mathcal{S} \subseteq \mathcal{A} \times \mathcal{A}$ is the set of supports and $\mathcal{D} \subseteq \mathcal{A} \times \mathcal{A}$ is the set of dependencies.

Please note that we introduce the tripolar graph only for representation purposes and we do not define any semantics for it. When required, we can simply extract a Dung's or bipolar graph from it. Its primary purpose is simply to mark relations which, potentially due to confusion or the difficulty of the experiment, the participants had problems classifying. Consequently, unlike in [47, 48, 49], it does not necessarily mean that the relation is really neither an attack nor a support, or that the relation might not exist.

In the following sections the ability to compare two graphs will be important for us. To this end, we introduce various types of subgraphs. To put it simply, if all arguments and links contained in one framework are present in another framework, we say that the first framework is a subgraph of another. If the links are perceived in the same manner (e.g. if it is an attack in one graph, it is also attack in the other), we are dealing with a correct subgraph, which is the most commonly considered type. We can then relax this notion in a number of ways. For example, if we allow links such as attacking or supporting to become classified as dependent, we can use the notion of a confusion subgraph. In a dual fashion, by allowing the dependent links to be classified as attacking or supporting, we create a precision subgraph. Finally, if we only concern ourselves with verifying whether the arguments connected in one framework are connected in another, independently of what the nature of that connection might be, we can use the notion of a lenient subgraph:

Definition 6.2. Let $TF = (\mathcal{A}, \mathcal{R}, \mathcal{S}, \mathcal{D})$ and $TF' = (\mathcal{A}', \mathcal{R}', \mathcal{S}', \mathcal{D}')$ be two tripolar argument graphs. We say that:

- TF' is a **correct subgraph** of TF iff $\mathcal{A}' \subseteq \mathcal{A}$, $\mathcal{R}' \subseteq \mathcal{R}$, $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathcal{D}' \subseteq \mathcal{D}$.
- TF' is a **confusion subgraph** of TF iff $\mathcal{A}' \subseteq \mathcal{A}$, $\mathcal{R}' \subseteq \mathcal{R} \cup \mathcal{D}$ and $\mathcal{S}' \subseteq \mathcal{S} \cup \mathcal{D}$ and $\mathcal{D}' \subseteq \mathcal{D}$.
- TF' is a **precision subgraph** of TF iff $\mathcal{A}' \subseteq \mathcal{A}$, $\mathcal{R}' \subseteq \mathcal{R}$ and $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathcal{D}' \subseteq \mathcal{D} \cup \mathcal{R} \cup \mathcal{S}$.
- TF' is a **lenient subgraph** of TF iff $\mathcal{A}' \subseteq \mathcal{A}$ and $(\mathcal{R}' \cup \mathcal{S}' \cup \mathcal{D}') \subseteq (\mathcal{R} \cup \mathcal{S} \cup \mathcal{D})$.

We say that a tripolar framework TF is **clarified** if $\mathcal{D} = \emptyset$.

Given the fact that the new types of subgraphs were introduced as a way to relax the classical correct subgraph notion, there certain relations between them. We can observe that every correct subgraph is also a confusion and a precision subgraph, and every confusion subgraph or precision subgraph is a lenient subgraph. However, not every confusion subgraph meets precision requirements and not every precision subgraph meets confusion requirements.

Example 4. Let us consider the graphs $TF_1 = (\{A, B, C, D\}, \{(A, B), (B, C)\}, \{(A, C)\}, \emptyset)$, $TF_2 = (\{A, B, C, D\}, \{(A, C)\}, \{(A, B), (B, C)\}, \emptyset)$, $TF_3 = (\{A, B, C, D\}, \{(A, B), (B, C)\}, \{(A, C), (A, D)\}, \emptyset)$ and $TF_4 = (\{A, B, C, D\}, \{(A, B)\}, \{(A, C), (A, D)\}, \{(B, C)\})$. We can observe that any framework is its own subgraph of any type. In addition to that, TF_1 is a correct subgraph of TF_3 . It is a confusion subgraph of both TF_3 and TF_4 , a precision subgraph of TF_3 and a lenient subgraph of all of the listed frameworks. On the other hand, TF_2 is not a correct, confusion or precision subgraph of any other framework apart from itself. It is however a lenient subgraph of all of the listed structures. TF_3 is not a correct subgraph of any other framework apart from itself and is a confusion and lenient subgraph of TF_4 . TF_4 is only its own correct and confusion subgraph, and a precision and lenient subgraph of TF_3 and TF_4 .

In addition to being able to represent the answers provided by the participants, we would also like to be able to measure the disagreement between the various frameworks we obtain at a given stage of the

dialogue. For this purpose, we create a notion distance between frameworks conceptually similar to the ones from [49, 50, 51]. The distance is measured in terms of links by which two given frameworks differ. In particular, if the given arguments are seen as unrelated in one structure, but as attacking/supporting or dependent in the other, we set the difference to 2 and 1 respectively. If a link has opposing polarities (i.e. attacking vs supporting) in the frameworks, the distance is set again to 2. If it is attacking or supporting in one and dependent in another, the distance is set to 1. Finally, if the nature of a given relation is seen in the same way, then the difference is naturally 0. The distance between the frameworks is then simply a sum of all such differences:

Definition 6.3. Let $TF = (\mathcal{A}, \mathcal{R}, \mathcal{S}, \mathcal{D})$ and $TF' = (\mathcal{A}, \mathcal{R}', \mathcal{S}', \mathcal{D}')$ be two tripolar frameworks defined over the same set of arguments s.t. $\mathcal{R} \cap \mathcal{S} \cap \mathcal{D} = \emptyset$ and $\mathcal{R}' \cap \mathcal{S}' \cap \mathcal{D}' = \emptyset$. Let $Rel^{TF} = \mathcal{R} \cup \mathcal{S} \cup \mathcal{D}$ and $Rel^{TF'} = \mathcal{R}' \cup \mathcal{S}' \cup \mathcal{D}'$ be the sets of all relations in both structures. The difference between TF and TF' on edge α is defined in the following manner:

- If $\alpha \notin Rel^{TF} \cap Rel^{TF'}$:

$$\text{diff}(TF, TF', \alpha) = \begin{cases} 0 & \text{if } \alpha \notin Rel^{TF} \cup Rel^{TF'} \\ 1 & \text{if } \alpha \in \mathcal{D} \cup \mathcal{D}' \\ 2 & \text{otherwise} \end{cases}$$

- If $\alpha \in Rel^{TF} \cap Rel^{TF'}$:

$$\text{diff}(TF, TF', \alpha) = \begin{cases} 1 & \text{if } \alpha \in \mathcal{D} \wedge \alpha \notin \mathcal{D}' \\ 1 & \text{if } \alpha \notin \mathcal{D} \wedge \alpha \in \mathcal{D}' \\ 2 & \text{if } \alpha \in \mathcal{S} \wedge \alpha \in \mathcal{R}' \\ 2 & \text{if } \alpha \in \mathcal{R} \wedge \alpha \in \mathcal{S}' \\ 0 & \text{otherwise} \end{cases}$$

The distance between TF and TF' is then defined as:

$$d(TF, TF') = \sum_{\alpha \in Rel^{TF} \cup Rel^{TF'}} \text{diff}(TF, TF', \alpha)$$




Figure 5: Difference measure between attacking, supporting, dependent and nonexistent relations

The introduced notion of a distance between the frameworks meets the requirements of a metric, i.e. the following holds:

Lemma 6.4. Let $TF_1 = (\mathcal{A}, \mathcal{R}_1, \mathcal{S}_1, \mathcal{D}_1)$, $TF_2 = (\mathcal{A}, \mathcal{R}_2, \mathcal{S}_2, \mathcal{D}_2)$ and $TF_3 = (\mathcal{A}, \mathcal{R}_3, \mathcal{S}_3, \mathcal{D}_3)$ be tripolar frameworks defined over the same sets of arguments and s.t. for $i \in \{1, 2, 3\}$, $\mathcal{R}_i \cap \mathcal{S}_i \cap \mathcal{D}_i = \emptyset$. The following holds:

1. $d(TF_1, TF_2) \geq 0$.

2. $d(TF_1, TF_2) = 0 \iff TF_1 = TF_2$ ⁵.
3. $d(TF_1, TF_2) = d(TF_2, TF_1)$.
4. $d(TF_1, TF_3) \leq d(TF_1, TF_2) + d(TF_2, TF_3)$.

Finally, in order to be able to analyze which postulates may or may not be satisfied by the participants, we need to map the declared levels of agreement to actual numerical values. By this we understand that if an argument X is assigned an answer such as *Strongly Disagree* by a participant, then there is a numerical value y associated with *Strongly Disagree* s.t. for the purpose of this analysis we can say that $P(X) = y$. From such assignments, a full probability distribution over the powerset of arguments can be reproduced [31]. However, given the way the probabilities are used by the epistemic postulates that we have recalled, knowing that $P(X) = y$ is sufficient for our purposes. We propose the following assignments:

- *Strongly Agree* – 6/6.
- *Agree* – 5/6.
- *Somewhat Agree* – 4/6.
- *Neither Agree nor Disagree* – 3/6.
- *Somewhat Disagree* – 2/6.
- *Disagree* – 1/6.
- *Strongly Disagree* – 0/6.

Additionally, the *Don't know* answer is available in the questionnaire. Unless stated otherwise, for calculation purposes we assign it the same value as *Neither Agree nor Disagree* (i.e. 3/6). It is a possible way of shifting from a four-valued setting (in, out, undecided and unknown) to a three-valued one (in, out, undecided), where inability to make a decision is treated similarly as indecisiveness.

Finally, we create the graphs based on the answers in relation task in the following manner. The *Reason for* answer we will mark as support, *Reason against* as attack, and the *Somewhat related, but can't say how* as dependent. The declared strength of the attack or support relation will only become relevant once we analyze its connection to the agreement level assigned to its source in Section 6.3.

6.1. Argument Graphs

In this section we will analyze the argument graphs we have received from the participants based on their responses to the relation tasks. In what follows we will distinguish the following types of graphs:

- the **intended graph** is meant to depict the minimal set of relations we considered reasonable for a given set of arguments.
- the **augmented graph** is obtained by adding to the intended graph the indirect relations from the prudent/careful or bipolar argumentation approaches (Definitions 2.3 and 3.2).
- the **participant-sourced graphs**:
 - the **declared graph** is constructed from the answers given to us by the participants in the Relation tasks, where the *Reason for* answer we will mark as support, *Reason against* as attack, and the *Somewhat related, but can't say how* as dependent. The arguments correspond to the statements appearing at a given step of the dialogue.

⁵In this context we say that $TF_1 = TF_2$ iff $\mathcal{R}_1 = \mathcal{R}_2$, $\mathcal{S}_1 = \mathcal{S}_2$ and $\mathcal{D}_1 = \mathcal{D}_2$.

- the **expanded graph** is constructed from the declared graph, extended with the statements extracted from the answers that the participants have provided in the Explanation task.
- the **common graph** is the declared graph created by the highest number of participants at a given step in the dialogue.

The intended graph simply depicts certain minimal constraints we had in mind for this experiment. For example, we had intended the statement “*Hospital staff members are exposed to the flu virus a lot. Therefore, it would be good for them to receive flu shots in order to stay healthy.*” to attack “*Hospital staff members do not need to receive flu shots.*”. Please note that this graph does not necessarily contain all of the possible interactions between the statements. For example, we decided against putting any supporting links in. The purpose of this graph was merely to encompass the minimal constraints we had intended the participants to recognize. Consequently, this does not necessarily mean that it contains all the possible relations that could be extracted from the dialogue.

The augmented graph is built from the intended graph through the addition of additional relations, similarly as indirect edges can be added to bipolar graphs for evaluation (see Section 3 and [24] for further details). The purpose of this graph is to check whether the relations declared by the participants in addition to the ones stated in the intended graph can in fact be reproduced through the use of notions from Definitions 2.3 and 3.2. In particular, we will augment the intended graph with attacks using the definitions of indirect conflicts from the aforementioned definitions and with supports using the definition of indirect defense. We will then observe that these notions do not necessarily give us all the possible relations defined by the participants and analyze how the ones that are not accounted for could be explained.

At this point we would like to stress that in this report, we will follow the abstract interpretation of supports, i.e. one in which it is understood simply as a positive relation between arguments without any further requirements [15]. As stated in Section 3, we will not impose any restrictions as to which indirect conflicts should be used. The reason for this is two-fold. First of all, there is no consensus as to which indirect conflicts should be associated with a given specialized form of support, such as deductive or necessary [24, 25, 52]. Second of all, we are not aware of any empirical study verifying that people indeed associate with a given type of support the indirect attacks ascribed to it in theory. There also appears to be no study testing whether a given supporting edge is assigned the same interpretation both by theory and by actual people. All of these are extremely interesting questions, however, they are beyond the scope of this particular study.

We would also like to explain our approach to reproducing supports obtained from the participants through defense in the augmented graph. First of all, both defenses and supports are forms of positive relations between arguments. Hence, interactions between them are unavoidable, particularly given the results in [24, 25]. Taking into account that our participants are not argumentation experts and their task was to mark positive relations in the dialogue, what they have identified could fall into either support or defense category. This will always be the case in any experiment that takes into account people from all walks of life, not just specialists in our area, and we cannot discard parts of their answers and extract as supports only those relations that in our opinion are “real”. What we therefore show in Sections 6.1.1.2 and 6.1.2.2 is that even though many of the positive links identified by the participants can be reproduced through defense, there are support edges that cannot be explained in this way.

The expanded graph contains an auxiliary statement P that embodies the additional reasons for and against the statements in the dialogues that the participants may have given in their explanations and that have not appeared in the dialogue. For example, P could contain a statement such as “*Masks are insufficient to prevent the flu, gloves are also needed.*” or “*The vaccine does cause flu, everyone I know gets very sick after the shot.*”.

In this section we will perform two types of analyses. We first consider the intended graphs created for both of the dialogues (Sections 6.1.1.1 and 6.1.2.1) and see how they are reflected by the graphs obtained from the participants. Based on this information we create the notions of the total and core samples, used throughout the rest of this paper. We then focus on looking at the declared and common graphs in Sections 6.1.1.2 and 6.1.2.2. In particular, we explain up to what extent the common graphs can be reproduced from

the intended graphs by the use of indirect defenses and attacks from the prudent/careful and bipolar setting (see Definitions 2.3 and 3.2).

6.1.1. Dialogue 1

6.1.1.1. Intended Graph Analysis

The minimal graphs we had wanted the participants to acknowledge at every stage of Dialogue 1 are presented in Figure 6. The first thing we would like to analyze is how the intended graphs are related to the ones declared by the participants.

In the first chart in Figure 7 we can see for how many participants the intended graph was a subgraph of a given type of their declared graph. The portion of the participants who have correctly recognized the intended relations at a given step of the dialogue does not exceed 50%, with the worst performance of 30% at the final step of the dialogue. The results look somewhat more interesting once we look at the confusion subgraph, which - as the name suggests - assumes the participants may become confused and specify the link as simply dependent rather than attacking or supporting. The biggest difference in the number of participants between the correct and confusion subgraphs can be seen in the last step of the dialogue, which might point to participants' exhaustion or to the complexity of the pieces of information presented in the dialogue. Finally, we can consider looking at the lenient subgraphs, in which it only matters that the statements we have seen as related are also seen as such by the participants, independently of the nature of this relation. The high number of participants whose declared graphs contained the intended graph in this lenient approach tells us that the participants had no major problems in recognizing the related statements, however, marking the type of this connection has caused some issues. All of the declared graphs that have not met the leniency requirements missed at least one of the (C, B) , (E, B) or (F, E) attacks. Nevertheless, in all of the approaches we can see that the lowest number of participants satisfying a given subgraph relation can be seen in the last two steps of the dialogue in which statements E and F were presented.

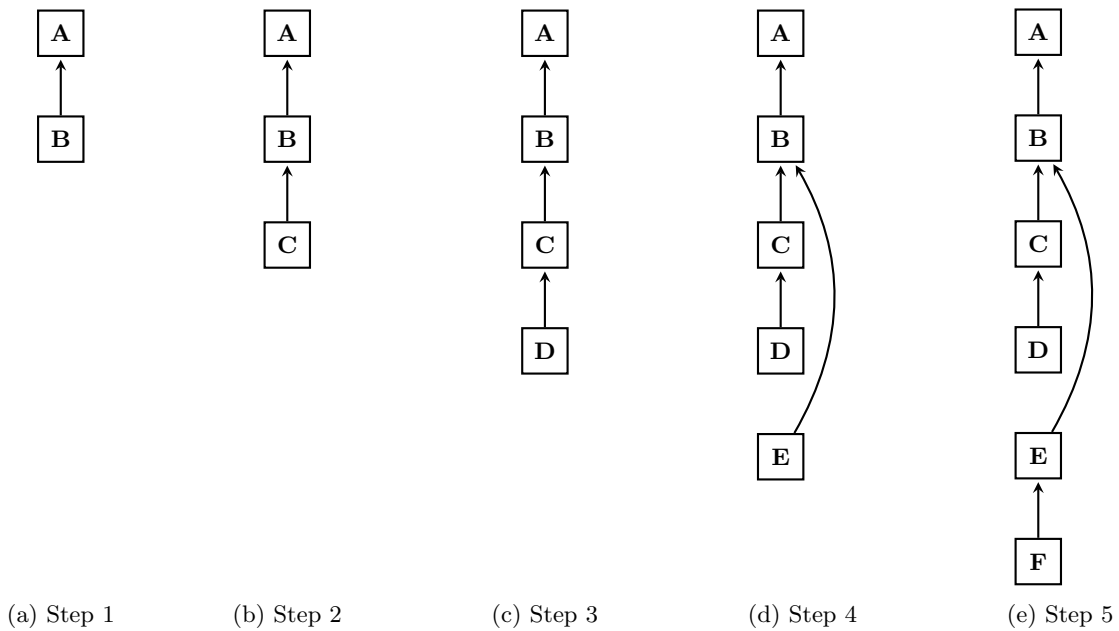


Figure 6: The intended argument graphs for Dialogue 1. Solid edges represent the attack relation.

Although the first chart in Figure 7 presents how many participants recognized the intended graphs at a given step in the dialogue, it does not show how many participants recognized the intended graphs at a given number of steps. For example, even though 18 (resp. 20) participants recognized the intended graph correctly at the first step (resp. second step) of the dialogue, it does not necessarily mean that 18 participants recognized the intended graph correctly at both of these steps. Consequently, in the second chart in Figure

7 we present at how many steps the participants have recognized the intended graph in a particular way (correct, confusion, lenient). The lenient requirements are quite easily satisfied, just like in the previous case. However, we can observe that a relatively small number of participants (7 w.r.t. the correct and 12 w.r.t. the confusion subgraph) recognized the intended graph in all dialogue stages. We believe this indicates that the relation task might have been more confusing than we had anticipated. Consequently, we have decided to report the results from all of the participants – we will refer to them as the **total sample** from now on – as well as from those that have recognized the intended framework in at least four out of five dialogue stages w.r.t. the confusion subgraph approach – we will refer to them as the **core sample**. The samples contain respectively 40 and 16 participants.




Figure 7: Analysis of intended graph containment in Dialogue 1

6.1.1.2. Participant-Sourced Graphs Analysis

In this section we would like to take a closer look at the graphs declared by the participants throughout the first dialogue. Although we have already analyzed whether the participants have recognized the intended graph, we have not yet seen what relations - in addition to the ones in Figure 6 – they have also observed. We can observe that at every step of the dialogue, the number of unique declared graphs we have obtained is smaller than the number of participants (auxiliary data appendix is available at [46]). This is a natural effect of the fact that different participants could answer the relation tasks in the same manner at some of the dialogue stages. Nevertheless, the further we are in the dialogue, the more common the graphs that have only been declared by one participant become. As seen in Figure 8, also the bigger the chances that the framework declared by a participant will not be clarified (i.e. it will contain dependent links, see Definition 6.2), which was to be expected given the increasing number of edges at every stage. Although at every step we can find a graph that has been declared by the largest number of participants, the obtained structure rarely accounts for more than 50% of the (total) sample. Thus, not every participant perceives the dialogue in the same manner, in part because they do not appear to be always aware of the same facts and arguments (we will discuss this issue further in Section 6.4). This disparity in knowledge and perception lends support to the use of the constellation approach in opponent modelling. In particular, if we treat the portion of participants declaring a given graph as a probability, what we obtain at every stage of the dialogue is actually a subgraph distribution.

As we have stated before, at every stage of the dialogue we can find a framework that has been declared by the largest number of participants. We present these structures in Figure 9. Only in the first step of the dialogue we obtain 2 equally common graphs, one in which (B, A) edge is seen as attacking and one in




Figure 8: Portion of unclarified graphs (i.e. with nonempty set of dependencies) obtained in Dialogue 1

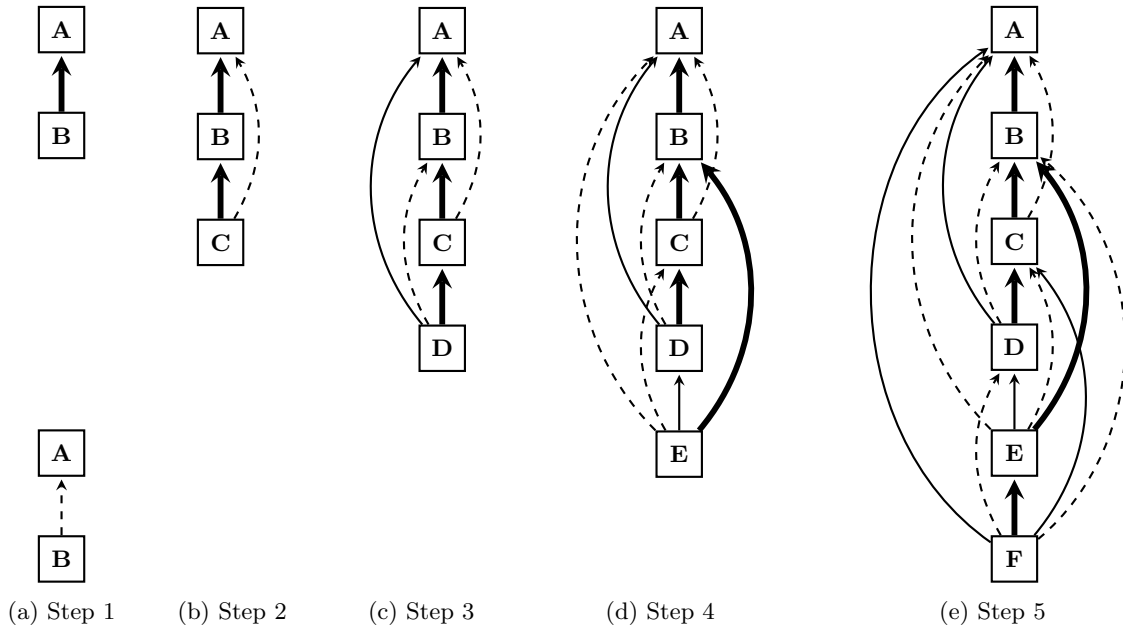


Figure 9: The common argument graphs for Dialogue 1 based on the total sample. With the exception of the bottom framework in Step 1, they are also common in the core sample. The thicker edges represent the relations appearing in the intended graph. Solid edges stand for attack and dashed for support links.

which it is supporting (both are declared by 18 participants). These proportions change notably once we look at the core sample, where the (B, A) link is primarily attacking. We believe this behaviour can be put down to the initial confusion caused by the experiment, particularly that afterwards the edge is seen mostly as attacking by the participants from both samples.

We can observe that w.r.t. the intended graphs, the participants have declared significantly more relations in the common graphs. However, as we will see, many of these edges can in fact be explained by the existing notions of indirect relations. We will first look at the common graphs from the perspective of the prudent/careful semantics (Definition 2.3). In other words, we perform a conflict-centered analysis, where the positive edges between the statements are seen only as a result of the interplay between the attacks. By extending the intended graphs with additional conflicts stemming from the indirect attacks and additional supports corresponding to the indirect defenses, we obtain the augmented graphs visible in Figure 10. The augmented graphs obtained in the first three steps are identical with the common frameworks declared

by the participants. The differences between them start appearing once the statement E is introduced in the fourth step of the dialogue. In particular, in the augmented framework, the (E, D) attack and (E, C) support are unaccounted for. Once F is presented in the fifth step of the dialogue, also the (F, D) support and (F, C) attack are missing.

This situation could be interpreted as a sign that the intended framework is too conservative and that the missing edges should have been included in it in the first place. However, there are also other possible interpretations that do not resort to modifying the intended graph. According to Definition 2.3, E is controversial w.r.t. D as it is an attacker of B , which is defended by D . Similarly, F is controversial w.r.t. C , as F attacks A which is defended by C . Consequently, it might be the case that the (E, D) and (F, C) attacks are manifestations of these issues⁶. The missing (E, C) and (F, D) supports could then be explained in two ways. If we were to use the aforementioned controversy as a basis for additional conflicts and add them to the (intended or augmented) graph, the absent supports could potentially be created by allowing indirect defense to also take into account these new conflicts in their definition. This could perhaps give rise to “controversial defense”, the existence of which we are not aware of in the literature. Another possible, and perhaps more natural approach, is to consider the statements that indirectly attack (defend) the same statement as positively related. In this respect, we can see E and C as working towards the same goal, which is defending A and attacking B . Nevertheless, we are not aware of this aspect being explored in the context of attack-based graphs.

Let us now look again at the common graphs, but from the perspective of bipolar argumentation. Like in the previous case, we extend the intended graph with support relations corresponding to the indirect defenses from Definition 2.3. However, instead of extending the graph with the indirect conflicts in the prudent/careful sense, we will use the ones offered by bipolar argumentation. The obtained augmented graphs are visible in Figure 11. By adding the secondary or supported attacks (in this particular case, they overlap) to the intended graphs, we recreate the augmented structures visible in Figure 10. By considering the (super) mediated attacks that adhere to the flow of the dialogue (i.e. the source of the attack has to appear later than the target), we can create the (E, D) and (F, C) attacks, which were missing in the previous approach. Consequently, we are now only missing the (E, C) and (F, D) supports. This issue can be addressed in two ways. One method is to repeat the process of adding defense-based support to the graph. This turns E into a supporter of C (it provides defense against D) and F into a supporter of D (it provides defense against E). Another approach is to consider (E, C) as a supporting link that should have been included in the intended graph in the first place. In such a situation, the (F, C) conflict can be reproduced by a secondary attack, while the (F, D) support would still require a repetition of the step in which we add the defense-based support.

We can observe that both approaches offer a way of reconstructing the common graphs (with the exception of the $(\{A\}, \emptyset, \{(B, A)\}, \emptyset)$ at step 1) from the intended one. The purely conflict-based approach would require the introduction of certain auxiliary notions. The bipolar approach uses methods already available in the literature, but in turn requires some of the procedures to be repeated.

In addition to the presented graphs, in Table 6 we can see how often a given relation was declared as attacking, supporting, dependent or nonexistent (in [46] we can find the same data w.r.t the dialogue steps). There are various differences between how a given relation was perceived by the participants belonging to the core and total samples. However, we can observe that for a given link, the most commonly associated type by one sample is the same as in the other. The only exceptions are the (F, A) and (F, C) links, which were considered primarily attacking in the core sample and dependent in the total sample. There appear to be more differences in what is the second and the third most common type. For example, a relation primarily seen as attacking and secondly as supporting in the total sample is more likely to be seen as primarily attacking and secondly dependent in the core one. However, we believe these differences are an effect of the method we have used to select the core sample and are most likely not indicative of any particular behaviour patterns. Finally, we can note that if we extracted graphs based on the data found in Table 6 for the core sample, we would recreate the common graphs for that sample. In the total case, we would not obtain the

⁶Please note that controversy could also lead the D attacking E and C attacking F , but such relations could not have been defined by the participants due to the experiment set up.

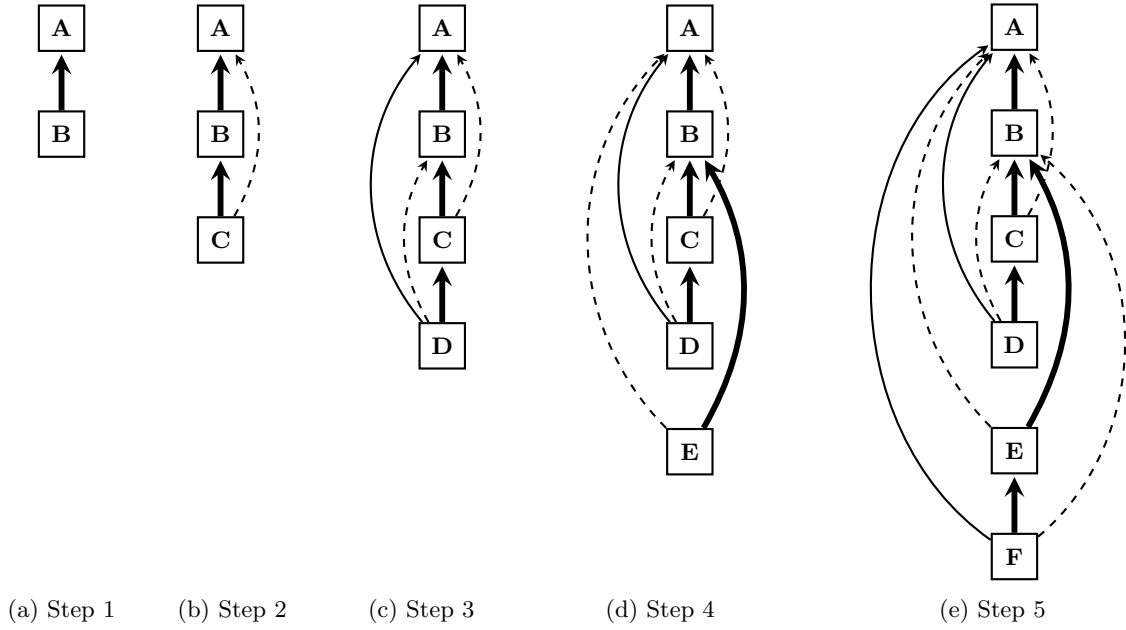


Figure 10: The intended argument graphs for Dialogue 1 extended with indirect conflicts and defenses from Definition 2.3. Solid edges represent the attack relation and dashed edges represent the defense-generated support relation. Thick edges represent the relations appearing in the intended graphs.

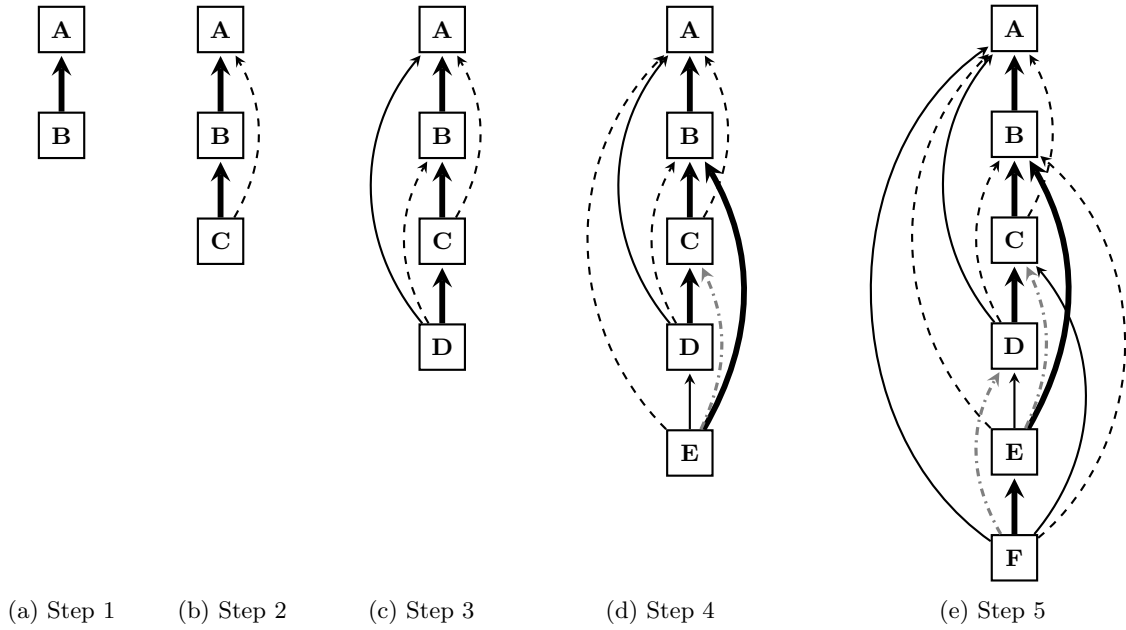


Figure 11: The intended argument graphs for Dialogue 1 extended with support coming from indirect defenses from Definition 2.3, secondary/supported attacks and super-mediated attacks following the dialogue flow. Dash dotted gray edges represent the additional supports obtainable if we again add defense-based support. Solid edges represent the attack relation and dashed edges represent the support relation. Thick edges represent the relations appearing in the intended graphs.

additional support graph from step 1, and due to the issue with the (F, A) and (F, C) links, the common graphs would be precision subgraphs of the structures extracted from the table.

In [46] we have also calculated the average distance from one framework to other declared frameworks

Relation	Core Sample				Total Sample			
	in \mathcal{R}	in \mathcal{S}	in \mathcal{D}	N/A	in \mathcal{R}	in \mathcal{S}	in \mathcal{D}	N/A
(B, A)	85	1.25	13.75	0	60.50	31.50	8	0
(C, A)	12.50	84.38	3.13	0	29.38	63.75	5.63	1.25
(C, B)	89.06	3.13	7.81	0	68.75	20	7.50	3.75
(D, A)	81.25	6.25	12.50	0	53.33	30.83	15	0.83
(D, B)	2.08	97.92	0	0	10	84.17	5.83	0
(D, C)	91.67	0	8.33	0	66.67	25	8.33	0
(E, A)	21.88	62.50	15.63	0	30	45	16.25	8.75
(E, B)	81.25	3.13	15.63	0	53.75	16.25	21.25	8.75
(E, C)	12.50	71.88	15.63	0	28.75	43.75	20	7.5
(E, D)	68.75	3.13	28.13	0	53.75	10	27.50	8.75
(F, A)	50	6.25	43.75	0	27.50	30	40	2.5
(F, B)	12.50	68.75	18.75	0	10	62.50	25	2.5
(F, C)	62.50	6.25	31.25	0	30	27.50	37.50	5
(F, D)	0	68.75	31.25	0	0	55	37.50	7.5
(F, E)	81.25	0	18.75	0	52.50	22.50	20	5

Table 6: Occurrences of the declared relations in Dialogue 1 (values are expressed as %)

obtained from the participants, understood as the average of all distances from a given structure to the remaining 39 ones in the total sample and 15 in the core one, i.e.:

Definition 6.5. Let n be the number of participants and let $\{G_1, \dots, G_n\}$ be the frameworks they have declared. The average distance from a framework G_i , where $1 \leq i \leq n$, to other frameworks is defined as:

$$avg_dist(G_i) = \frac{\sum_{j=1}^n d(G_j, G_i)}{n-1} \quad (1)$$

The results for the total sample can be seen in Figure 12, while the overall summary of our findings is presented in Table 7. In this table we present the average distance from the common framework at a given step to all of the other frameworks as well as the general analysis of all of the obtained average distances. We provide the minimum, maximum and median of the obtained values, and the overall average distance (i.e. the average of the averages).

We can observe that the further we are in the dialogue, the higher these values become, which is natural given the increasing size of the frameworks. The average distance calculated for the common framework tends to be smaller than the overall average and less or equal to the obtained median. In the core sample, all of the averages for the common frameworks are also the minimum ones from all of the obtained values. We can therefore observe that the fact that this framework is defined by the largest number of participants and is close to or identical with the framework induced by Table 6 is reflected in the obtained distances between the structures declared by the participants.

	Total Sample					Core Sample				
	Common Framework	Minimum	Maximum	Average	Median	Common Framework	Minimum	Maximum	Average	Median
Step 1	1.03	0.92	1.03	1.02	1.03	0.33	0.33	1.80	0.52	0.33
Step 2	1.67	1.67	4.54	2.40	1.92	0.40	0.40	3.87	0.72	0.40
Step 3	3.44	3.44	7.33	4.75	4.23	1.33	1.33	7.60	2.22	1.33
Step 4	6.74	6.74	13.36	8.78	8.39	2.07	2.07	11.93	3.44	2.07
Step 5	11.33	10.62	18.77	13.40	12.62	5.13	5.13	13.13	7.17	6.87

Table 7: Analysis of the average distance values for the declared frameworks for Dialogue 1, where average distance is formulated in Definition 6.5. We include the average distance from the common framework to other frameworks, minimum and maximum average distances amongst all the averages, median, and overall average.

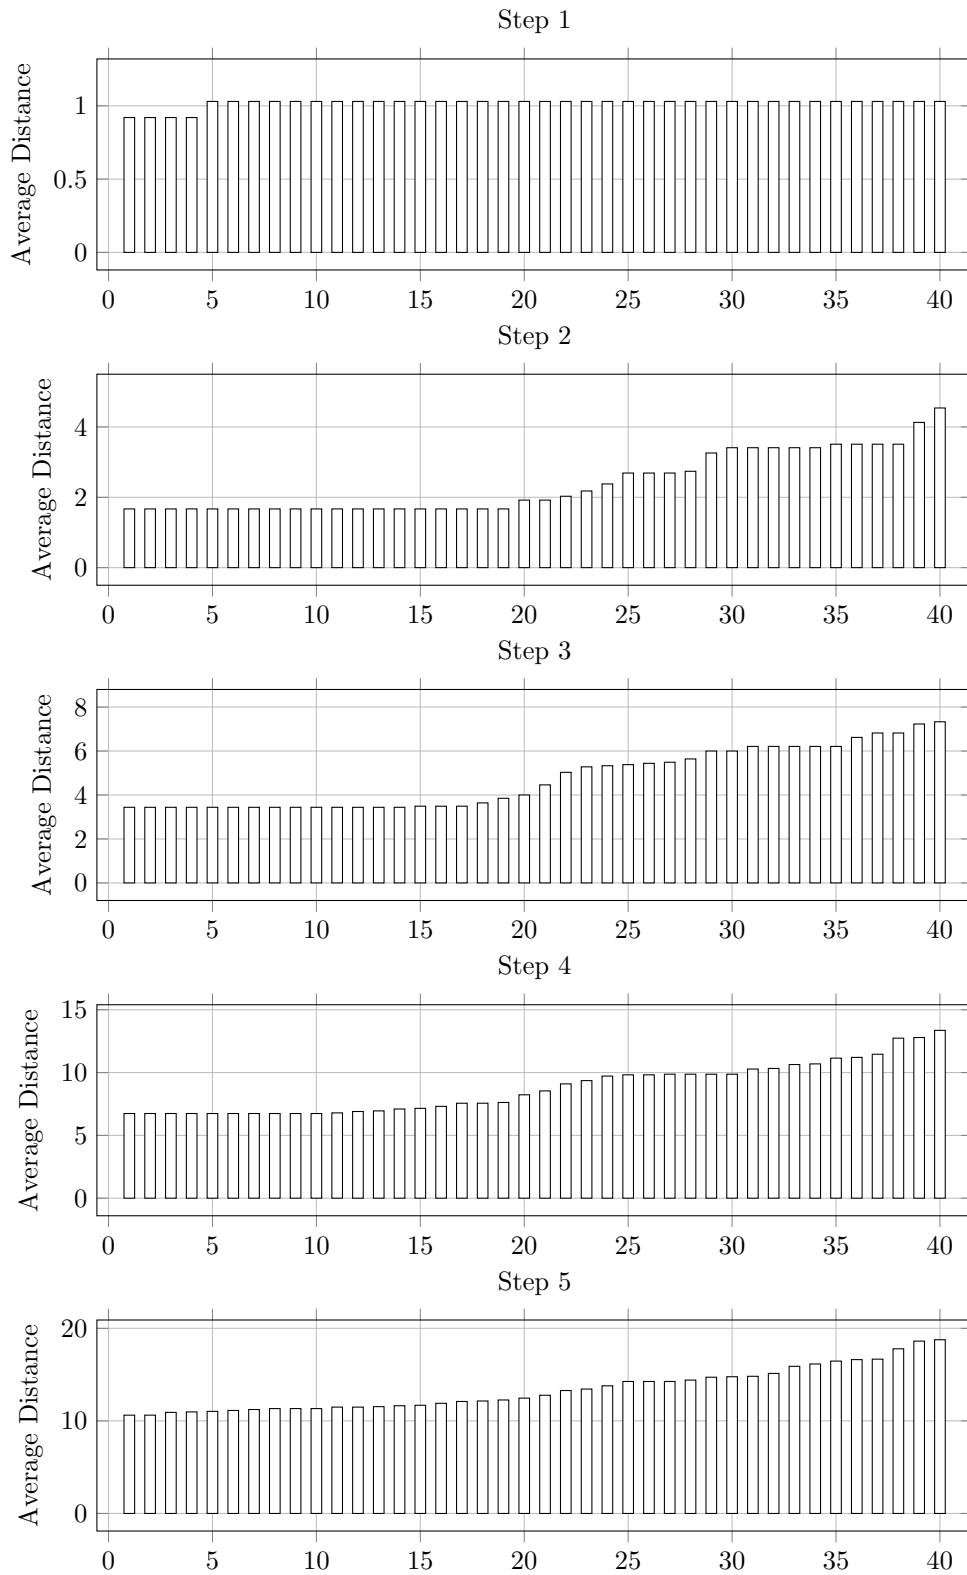


Figure 12: The list of average distances from the graph declared by a given participant to the graphs of the remaining participants at a given step in Dialogue 1.

6.1.2. Dialogue 2

6.1.2.1. Intended Graph Analysis

The graphs depicting the minimal set of relations we had intended the participants to recognize in the second dialogue are presented in Figure 13. We will now discuss the reason why the (G, F) edge is grayed out. Argument G was meant to be a counterargument for E and F . In the first case the contradiction is more obvious. In the second instance, the fact that the virus is accompanied only by stabilizers and antibiotics means it is not accompanied by thimerosal, which is only a preservative. Thus, the contradiction depends on the distinction between stabilizers and preservatives. Based on the explanations provided by some of the participants we can however observe that these two notions were occasionally confused, which additional research has shown to be a common situation in reality. Consequently, thimerosal was seen as an example of a stabilizer and as a result, G could have been in fact understood as supporting F rather than attacking it. Hence, declaring this relation differently was a conscious and somewhat justified decision, not a result of misunderstanding the exercise or an unintentional choice (a “misclick”). Nevertheless, removing this link from the intended graph does not significantly affect the core sample and would allow one more person in w.r.t. the original approach. This indicates that the majority of people marking (G, F) differently than intended have also had other issues that prevented them from entering the core sample. Thus, this issue deserves attention on its own, and for now we will evaluate the answers without modifying the initially set methodology.

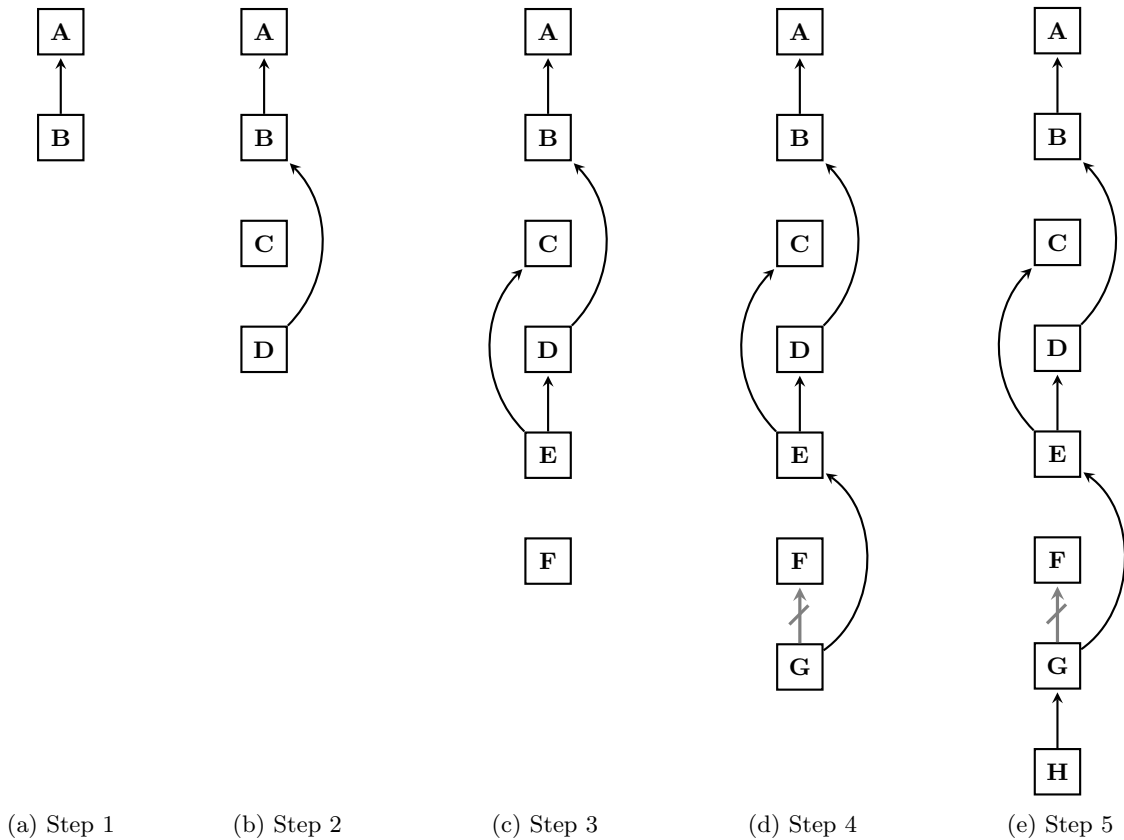


Figure 13: The intended argument graphs for Dialogue 2. Solid edges represent the attack relation. The crossed gray edges represent attacks that some participants have interpreted as supports in a way that can be considered justified.

Similarly as in the case of Dialogue 1, in Figure 14 we present how many participants recognized the intended graph at a given step of the dialogue, and how many participants recognized the intended graph in a given number of steps. We create the core sample for Dialogue 2 by gathering participants who have

acknowledged the intended graph w.r.t. the confusion subgraph relation in at least 4 stages. Hence, from now on we will be working with the total sample of size 40 and the core sample of size 15.

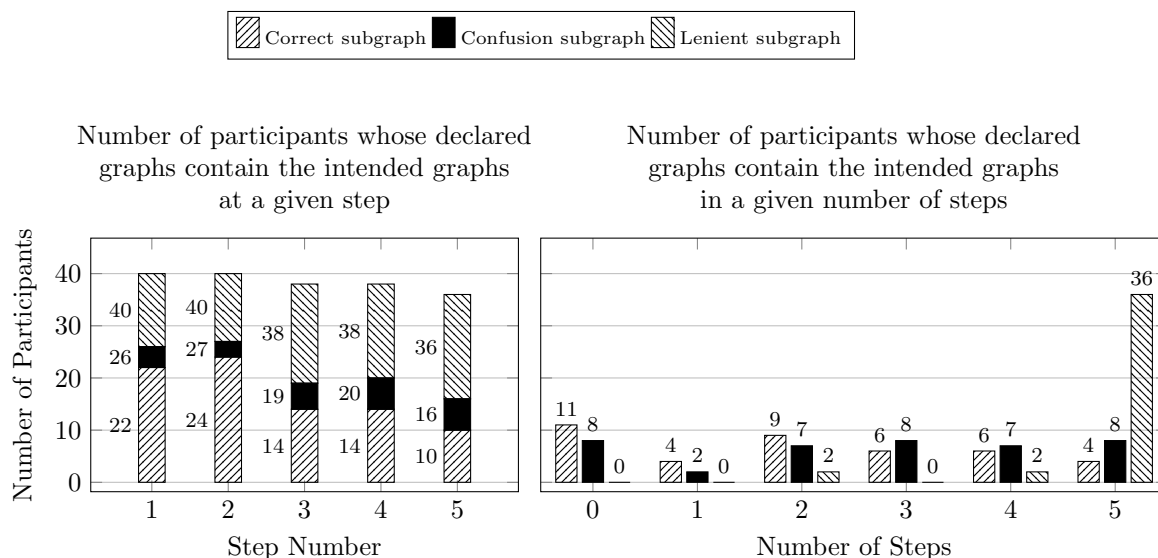


Figure 14: Analysis of intended graph containment in Dialogue 2

6.1.2.2. Participant-Sourced Graphs Analysis

Let us now look at the graphs given to us by the participants of Dialogue 2 (full data can be found in the data appendix in [46]). Just like in the previous dialogue, the further we are in the discussion, the more single-participant graphs we obtain. However, in this case, the maximal dispersion occurs during the fourth rather than fifth step, in which the participants tend to agree a bit more again. As the dialogue progresses, the number of clarified graphs decreases, as seen in Figure 15. Although there are some differences between the decrease in this dialogue and in Dialogue 1, particularly concerning the second stage, the end states are not that different. It therefore appears that participants in both samples had similar issues in deciding on the nature of some of the edges.

In Figure 16 we have depicted the common graphs at every stage of the dialogue, where thick edges represent the relations contained in the intended graph. Just like in Dialogue 1, at the very first step the participants are asked only about the (B, A) relation. Although more people have seen the (B, A) relation as attacking rather than supporting (22 vs 14 in the total sample), both groups are important enough to be reported. Many people who have marked (B, A) as supporting did not make it to the core sample, in which the attacking interpretation is dominant.

Similarly as we had done in the case of the first dialogue, we will now analyze whether the relations contained in the common graphs, but not in the intended ones, could possibly be explained with additional notions such as indirect attacks and defenses. Let us first consider creating the augmented graphs with the use of the indirect attacks and defenses from the prudent/careful setting (see Definition 2.3). The results are visible in Figure 17. Starting from step 2, there are differences between the common and the augmented frameworks. They are primarily caused by C and F being initially disconnected from other statements. In total, we are missing eight edges – four attacks (C, B) , (F, A) , (F, C) and (F, D) , and four supports – (C, A) , (D, C) , (F, B) and (F, E) . We will now try to explain them following the conflict-centered perspective.

Let us start with the (C, B) conflict and the (D, C) and (C, A) supports. Statement C simply presents a “fact” that vaccine contains a particular substance, which in this case is a mercury compound. As such, it does not contradict any previous information. However, it becomes conflicting once it is paired with the information that mercury compounds are poisonous, contained in statement D . This supporting (D, C) link possibly reflects this pairing and makes the (C, B) attack perfectly reasonable. One can therefore argue that




Figure 15: Portion of unclarified graphs (i.e. with nonempty set of dependencies) obtained in Dialogue 2

either (C, B) or both (C, B) and (D, C) should have been included in the intended graph. If only (C, B) was added to the graph, we could potentially reproduce the (D, C) support by treating statements working towards the same goal as positively related, similarly as proposed in the case of (E, C) and (F, D) supports in Dialogue 1. Adding only (C, B) or both (C, B) and (D, C) would also lead to the introduction of the (C, A) support through indirect defense.

A similar analysis could be carried out for the relations associated with F . Statement F serves as a further backing for E , which was recognized as support by the participants. Consequently, including the (F, E) support and (F, C) and (F, D) attacks in the intended graph (or, based on the aforementioned explanation, just the (F, C) and (F, D) attacks) would lead to the reproduction of the remaining missing relations through the use of indirect attack or defense.

Let us now look at the intended graphs augmented with supports corresponding to the indirect defenses from the prudent/careful setting and with indirect attacks associated with the bipolar setting. By using secondary/supported attacks and defense-induced support, we obtain the augmented frameworks identical to the ones present in Figure 17. Thus, again we are missing the same set of relations as in the conflict-centered analysis: four attacks – (C, B) , (F, A) , (F, C) and (F, D) – and four supports – (C, A) , (D, C) , (F, B) and (F, E) . In this case, the (super) mediated attacks do not offer any additional insight. By considering (super) extended attacks, we can retrieve the (C, B) attack, though only starting from the fourth step of the dialogue. In the same fashion we can also obtain the (F, A) , (F, C) and (F, D) conflicts at the last stage of the discussion. Nevertheless, these changes occur later than the arguments are introduced and our main issue is not resolved. In the conflict-centered analysis, we had proposed adding certain attacks (or both attacks and supports) to the intended graph in order to bridge the gap between the common graph and the augmented graph. In this, more support-centered, analysis, we can consider adding the (D, C) and (F, E) supports to the intended graph. Adding the (D, C) support would allow us to recreate the (C, B) conflict as an extended attack. The inclusion of the (F, E) support would lead to the recreation of the missing (F, C) and (F, D) attacks as supported attacks in the augmented graph.

We are thus left with the recreation of the (F, A) attack and the (C, A) and (F, B) supports. In order to do so, we can repeat the procedure of adding indirect defenses and conflicts, i.e. we extend the augmented graph in the way we extended the intended one. Then F becomes a supported and secondary attacker of A through the secondary/supported attack (E, A) or the supported attack (F, D) . Additionally, F is now an indirect defender of B through the supported attack (F, D) . Finally, (C, A) can be recreated through indirect defense generated by the the extended conflict (C, B) .

We can observe that although certain supports can be recreated through indirect defense, this is not the only possible reason why the participants have declared these relations. For example, we can also observe the similarity between statement A and the conclusion of D and the fact that G is a more detailed reiteration of C . If we view (F, E) support as a relation that should have been in the intended graph, then given the fact that E is a defense-mimicking supporter of B , the (F, B) support could be seen as a particular mixture

stemming from the support transitivity. Consequently, it is possible to be recreate some of the missing supports using a different methodology.

To summarize, we can observe that recreation of the common graph from the intended one in Dialogue 2 requires the addition of at least 2 relations carried out by C and F . However, just like in Dialogue 1, we would still need to introduce auxiliary notions when recreating the common graphs via the prudent/careful approach only and perform a two-step recreation via the bipolar approach.

In addition to the graph reproduction analysis, in Table 8 we list how often a given relation was perceived as attacking, supporting, dependent or nonexistent during the dialogue (more detailed tables can be found in [46]). Similarly as in the case of Dialogue 1, we can observe that the dominating assignment for a given edge is the same in both core and total samples, though the second and third most common choices often differ. By focusing on the most common edges, we can reconstruct the graphs from Figure 16.

Relation	Core Sample				Total Sample			
	in \mathcal{R}	in \mathcal{S}	in \mathcal{D}	N/A	in \mathcal{R}	in \mathcal{S}	in \mathcal{D}	N/A
(B, A)	85.33	2.67	12	0	66.50	26.50	7	0
(C, A)	13.33	73.33	13.33	0	27.50	58.75	13.13	0.63
(C, B)	75	3.33	21.67	0	68.13	12.50	16.25	3.13
(D, A)	11.67	81.67	6.67	0	27.50	65.63	6.88	0
(D, B)	85	3.33	11.67	0	76.25	12.50	9.38	1.88
(D, C)	11.67	60	25	3.33	27.50	51.88	16.25	4.38
(E, A)	73.33	8.89	17.78	0	58.33	26.67	14.17	0.83
(E, B)	8.89	86.67	4.44	0	18.33	73.33	8.33	0
(E, C)	91.11	0	8.89	0	65	17.50	12.50	5
(E, D)	84.44	0	15.56	0	61.67	21.67	13.33	3.33
(F, A)	73.33	11.11	15.56	0	58.33	29.17	11.67	0.83
(F, B)	8.89	84.44	6.67	0	13.33	75.83	10	0.83
(F, C)	82.22	0	17.78	0	53.33	30.83	12.50	3.33
(F, D)	82.22	0	17.78	0	55.83	29.17	12.50	2.50
(F, E)	2.22	88.89	8.89	0	11.67	75.83	9.17	3.33
(G, A)	16.67	63.33	20	0	23.75	58.75	16.25	1.25
(G, B)	60	10	30	0	61.25	16.25	21.25	1.25
(G, C)	10	83.33	6.67	0	25	63.75	8.75	2.50
(G, D)	16.67	70	13.33	0	21.25	65	11.25	2.50
(G, E)	86.67	3.33	10	0	70	16.25	10	3.75
(G, F)	80	6.67	13.33	0	63.75	18.75	15	2.50
(H, A)	73.33	0	26.67	0	45	25	27.50	2.50
(H, B)	13.33	73.33	13.33	0	12.50	67.50	17.50	2.50
(H, C)	86.67	0	13.33	0	60	20	12.50	7.50
(H, D)	73.33	6.67	20	0	52.50	22.50	22.50	2.50
(H, E)	0	93.33	6.67	0	5	77.50	12.50	5
(H, F)	0	73.33	26.67	0	7.50	72.50	15	5
(H, G)	80	6.67	13.33	0	50	32.50	10	7.50

Table 8: Occurrences of the declared relations in Dialogue 2 (expressed as %)

In [46] we have also calculated the average distance from one framework to other declared frameworks obtained from the participants, understood as the average of all distances from a given structure to the remaining 39 ones in the total sample and 14 in the core one (see also Definition 6.5). The results in the case of the total sample are visible in Figure 18, with the overall summary of our findings visible in Table 9. In this table we present the average distance from the common framework at a given step to all of the other frameworks as well as the general analysis of all of the obtained average distances. We provide the minimum, maximum and median of the obtained values, and the overall average distance (i.e. the average

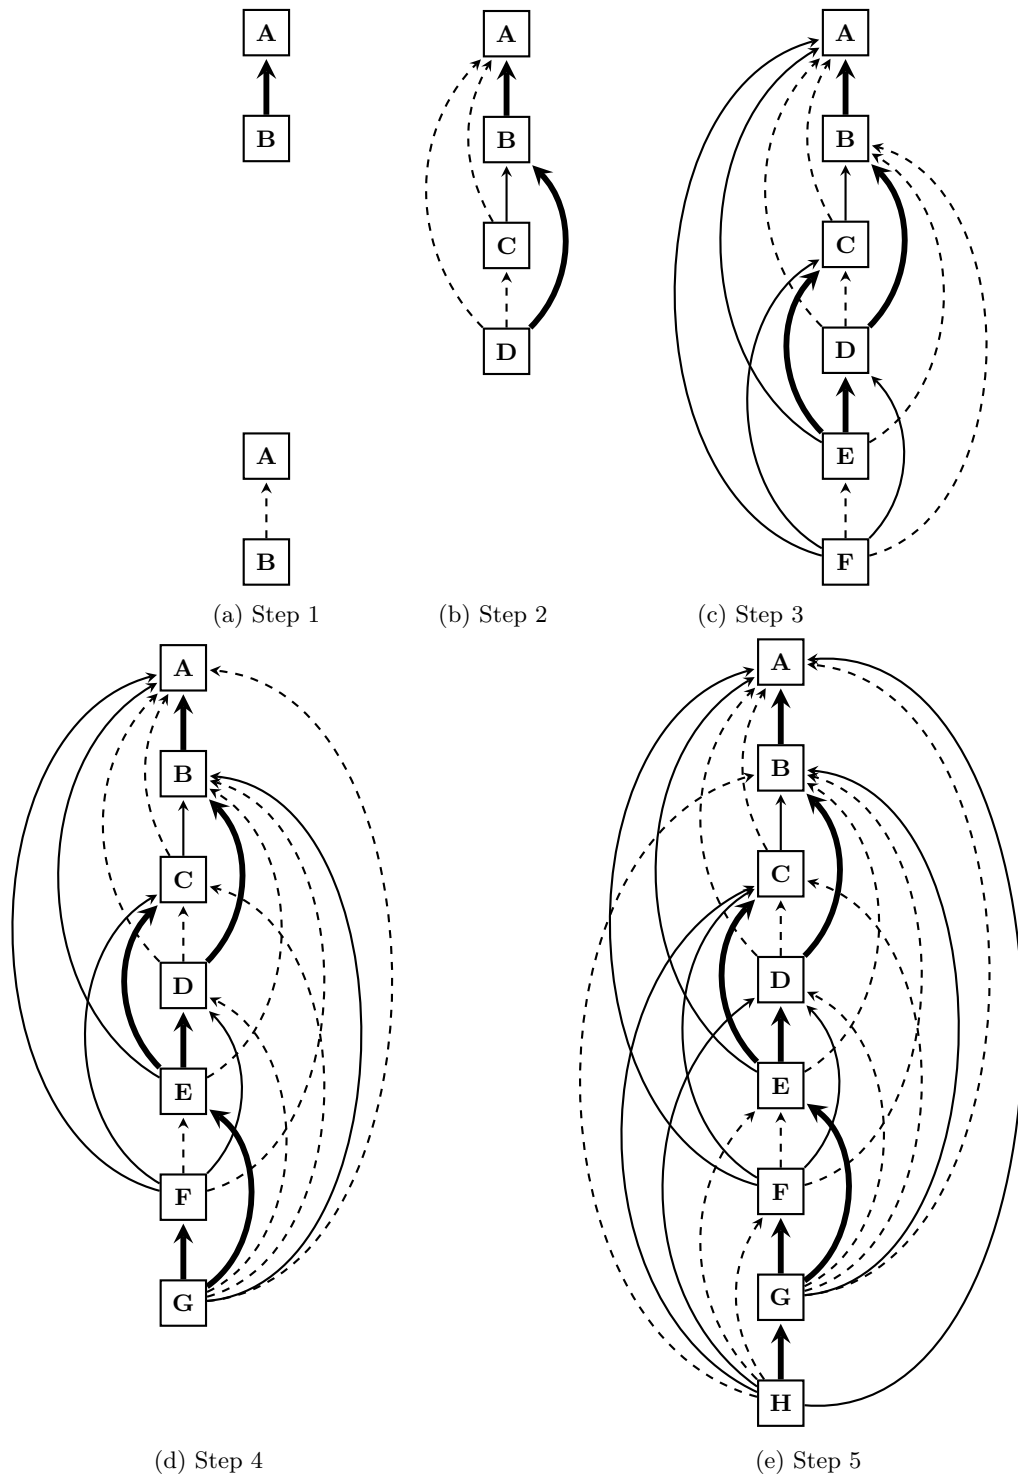


Figure 16: The common argument graphs for Dialogue 2 based on the total sample. With the exception of the bottom framework in Step 1, they are also common in the core sample. The thicker edges represent the relations appearing in the intended graph. Solid edges stand for attack and dashed for support links.

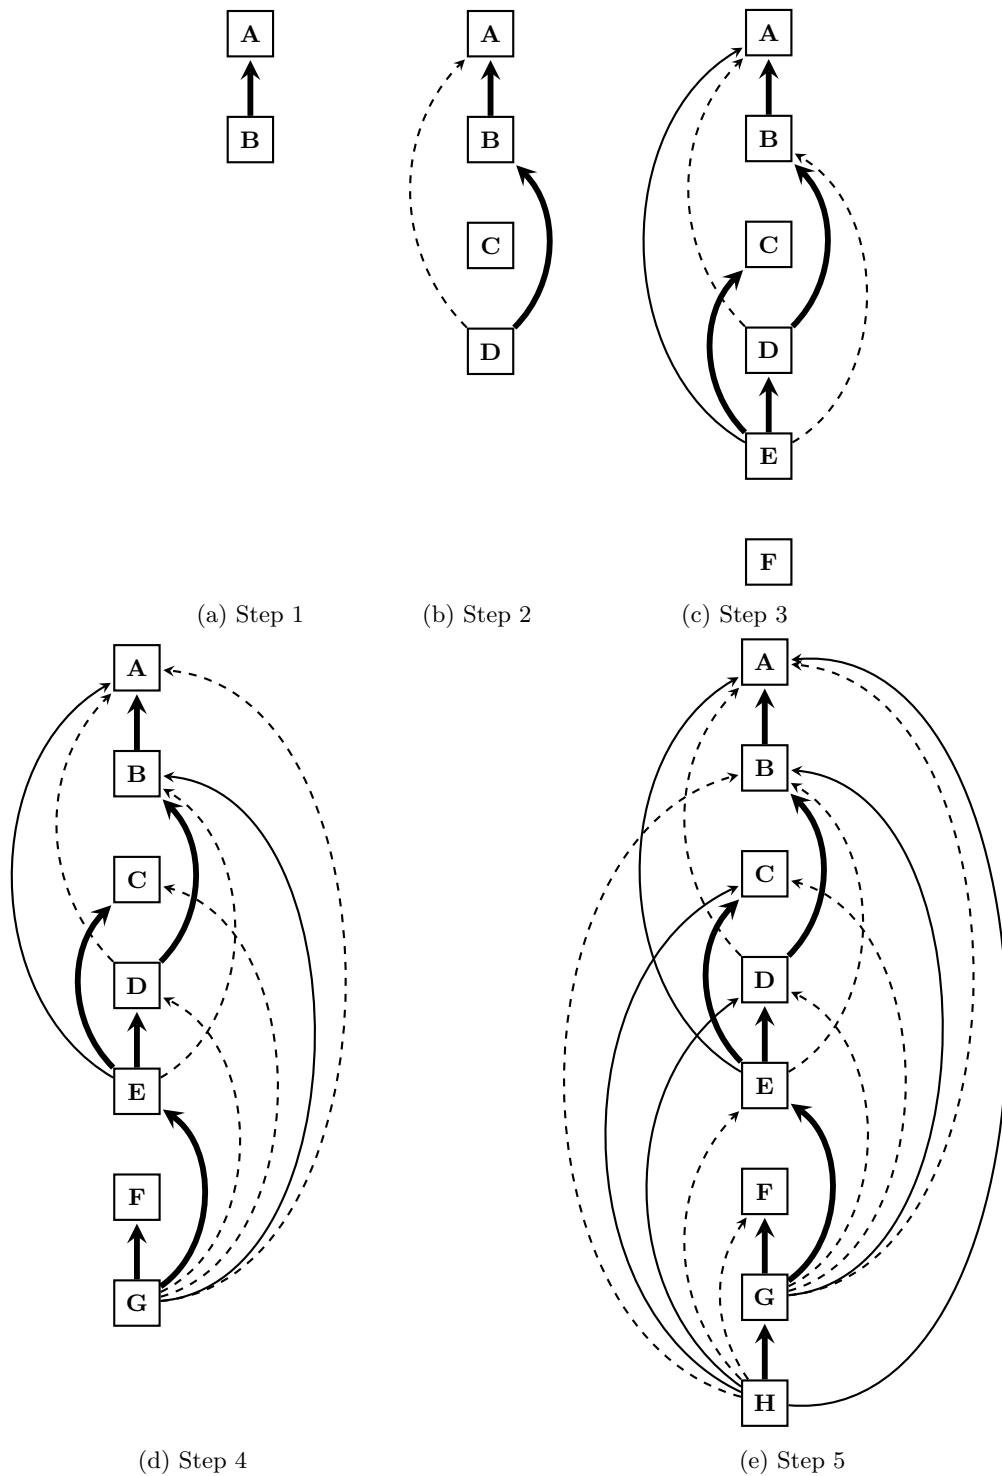


Figure 17: The intended argument graphs for Dialogue 2 extended with indirect conflicts and defenses from Definition 2.3. The thicker edges represent the relations appearing in the intended graph. Solid edges stand for attack and dashed ones for defense-generated support.

of the averages).

The same observations as in the case of Dialogue 1 can be made for Dialogue 2. The average distance calculated for the common framework tends to be smaller than the overall average and less or equal to the obtained median. In the case of the core sample, all of the averages for the common frameworks are also the minimum ones from all of the obtained values. This is also true in the case of the total sample, with the exception of the third step of the dialogue. The fact that the common framework is defined by the largest number of participants and is identical with the framework induced by Table 8 is reflected in the obtained distances between the structures declared by the participants.

	Total Sample					Core Sample				
	Common Framework	Minimum	Maximum	Average	Median	Common Framework	Minimum	Maximum	Average	Median
Step 1	0.82	0.82	1.23	0.97	0.82	0.29	0.29	1.86	0.48	0.29
Step 2	3.38	3.38	7.74	4.74	4.21	1.93	1.93	8.14	2.95	2.00
Step 3	9.82	9.77	18.08	13.07	13.41	5.00	5.00	15.57	7.70	5.79
Step 4	13.28	13.28	25.49	17.91	16.87	6.64	6.64	21.64	10.15	8.29
Step 5	16.26	16.26	37.13	22.59	22.29	7.21	7.21	25.64	11.43	8.36

Table 9: Analysis of the average distance values for the declared frameworks for Dialogue 2, where average distance is formulated in Definition 6.5. We include the average distance from the common framework to other frameworks, minimum and maximum average distances amongst all the averages, median, and overall average.

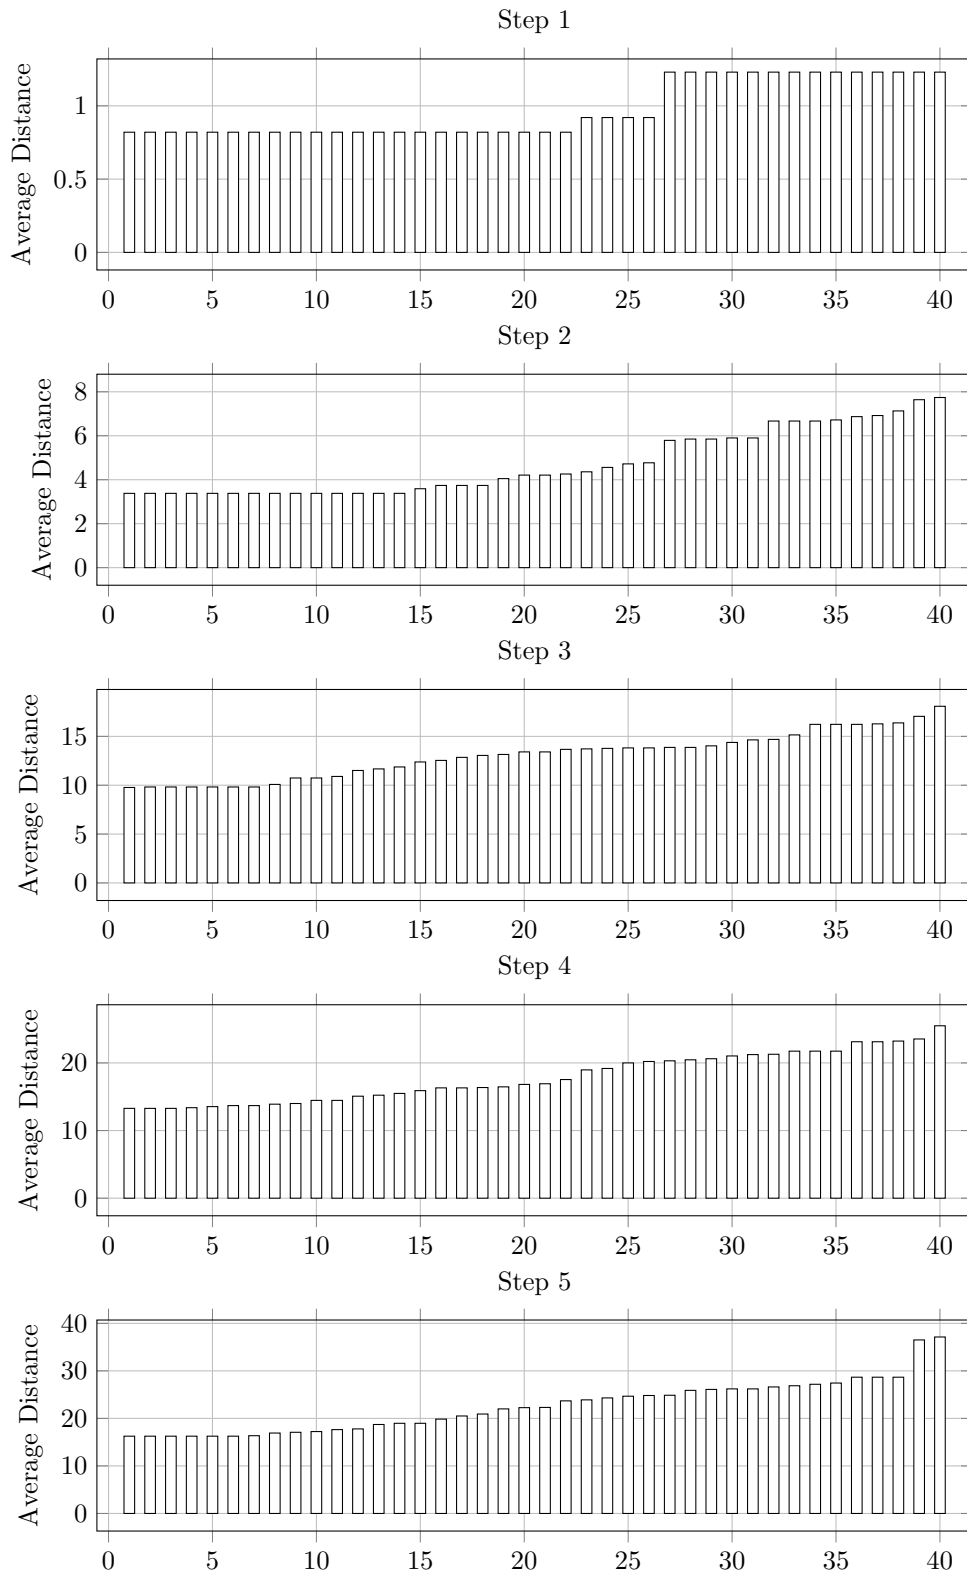


Figure 18: The list of average distances from the graph declared by a given participant to the graphs of the remaining participants at a given step in Dialogue 2.

6.2. Postulate Satisfaction

In this section we analyze if and how the postulates we recalled in Section 4.1 are reflected in the behaviours of the participants. In order to do so, we will evaluate the probability distributions extracted from the agreement tasks against the intended, expanded, declared and common graphs (see the end of the introduction of Section 6 on page 21 concerning how these distributions were obtained). Given the fact that the postulates are defined only for Dung’s frameworks, we need to consider the attack subgraphs of the aforementioned graphs in order to perform the analysis. Although this means that the frameworks contain information not yet harnessed by the postulates, in Sections 6.1.1.2 and 6.1.2.2 we could have observed that many of the supporting links could be explained with indirect defence or lead to additional indirect conflicts. Thus, we believe that the effect of the positive relations in the obtained graphs to be reasonably, even if not completely, reflected by the negative relations. Therefore, the analysis performed only on the attack graph is still valuable.

The results will be presented in the following manner. For every participant, we calculate the proportion of steps on which a given postulate was satisfied according to a given graph. For example, if we have five dialogue steps and the obtained statement probabilities conformed to the rational postulate on four of them, we obtain the adherence rate of 80% (or 0.8). The values we report are the averages of these scores on the total and core samples. We then analyze the differences between the results obtained between the graphs and samples. The additional statistical analysis can be found in Section 9.

6.2.1. Dialogue 1

The results concerning the average adherence rates of the postulates we have recalled in Section 4.1 on the intended, declared, expanded and common graphs are presented in Figures 20 (total sample) and 21 (core sample). Let us first look at the total sample. In the majority of cases, the results obtained on the intended and common graphs do not differ significantly, which can be seen as supporting our graph augmentation approach from Section 6.1.1.2 (see also Table 10 in Section 9.1.1). All of the results appear to be closely related on postulates that, in general, have a high satisfaction rate or belong to the preferential family. Thus, the adherence rates between the graphs are closely connected for postulates such as preferential, strict, protective and restrained, coherent, involutory and semi-optimistic. Since the ternary and binary postulates do not depend on the structure of the graph in question, they do not contribute to this particular observation. Nevertheless, once the satisfaction rate goes down and we look at the valued or explanatory postulates, the results begin to diverge and it is mostly the intended and common graphs that still behave similarly (see Table 10 in Section 9.1.1).

We can thus observe good adherence to the postulates and a similarity between the results on all of the graphs in the case the preferential branch of our epistemic postulates. The obtained satisfaction rates slowly decrease as the more specialized versions are considered, reaching the minimum on the justifiable postulate. Low adherence to this property is most likely caused by the optimistic property, which is also rarely satisfied. In principle, postulates considered shared between preferential and value based families do not perform as well as the ones classified as only preferential. We would also like to observe that the preferential postulate is highly satisfied not just due to its “leniency”, but also because it reflects the reasoning of various participants. For example, some of the participants satisfying the preferential, but not the rational postulate in the first step of the dialogue, justified their agreement with both of the statements by saying that e.g. “A: I agree that Hospital staff members don’t need to receive flu shots unless they want to. They should not be required to.” and “B: It is also true that hospital workers are exposed to the flu a lot so if they choose to it would be beneficial for them to receive the shot.”, or “It’s been shown, I think, that mysteriously, nurses and doctors are less prone to catching the diseases they treat in general.” and “Extra precautions couldn’t hurt during a particularly nasty flu season.”. To summarize, these results mean that participants are in general capable of reasoning well when it comes to considering the conflicts between two statements.

Lower satisfaction rate of the postulates and/or bigger dissimilarities between the results on different graphs (see Table 10 in Section 9.1.1) can be observed particularly in the case of postulates that are highly affected by the personal knowledge of the participants and by accuracy of modelling the perceived attacks. For example, these are the cases in which the results for the expanded graph – which is precisely the declared

graph augmented with statements extracted from explanations - differ from the results for the declared graph. What was considered an initial argument may not have been seen in the same way by the participants, thus reducing the adherence to the optimistic postulate. This, along with the reluctance of the participants to choose the extreme values (see the ternary postulate), may explain the lower performance of the founded property. In particular, by relaxing the founded restriction of complete belief to the semi-founded one, which simply requests non-disbelief, we obtain much better results. All of the guarded, discharging, trusting and anticipating postulates are intended to grasp the idea of believing an argument when there is no reason against it and disbelieving an argument only if there is a reason for it. Consequently, they are highly dependent on the accuracy of the graphs we create for the participants in terms of their knowledge, and are as such more difficult to measure and estimate than the postulates of the preferential family. The fact that the demanding postulate appears to be the least affected by it might be connected to the ternary postulate – again, if the extreme values are not chosen, this property is satisfied more easily.

We can observe that the semi-optimistic postulate is commonly satisfied, even twice as often as the coherent one. This has certain important implications. The purpose of these properties is to set respectively lower and upper bounds on the belief we have in a statement in the face of its attackers. For example, given the performance of the protective and strict postulates, it appears that while the participants are capable of choosing what to believe and what to disbelieve, they can be more “generous” with their belief than anticipated by some of the postulates.

We can now consider the results obtained for the core sample. We can observe that the average satisfaction rate of the postulates that the participants appeared to adhere to in the total sample increased in the core sample. In other words, the preferential, rational, strict, protective and restrained postulates are, on average, satisfied more often. The results obtained between the graphs appear to be much more closely connected than in the total sample (see Table 10 in Section 9.1.1). Significant differences between the adherence rates can be observed primarily in the case of declared and expanded graph. They are also bigger than in the case of the total sample, i.e. the average of the differences between adherence rates of participants between the declared and expanded graphs is bigger in the total sample only in the case of the involuntary and demanding postulates. This may be caused by the participants being more forthcoming with their personal knowledge and thus the expanded graphs being more precise. It would also explain the somewhat better results on the explanatory and optimistic, founded and justifiable postulates, which are more dependent on such factors. Finally, we can consider the binary and ternary postulates. In both samples, the binary postulate has a high adherence rate, which means that the participants were capable of deciding whether to believe or disbelieve the presented statements using their own knowledge and the information found in the dialogue. Thus, the text itself could not have been overly complicated for the participants. The low adherence to the ternary postulate in both of the samples, combined with the low indecisiveness of the participants, shows that people often exhibit various degrees of belief and disbelief and are reluctant to completely accept or completely reject a given statement.

In addition to the postulates available in Figures 20 and 21, we have also considered a slight modification of the VAL^n property. This postulate tells us how many distinct probabilities have been used in a given probability distribution. We can recall that from every user we obtain five distributions, one per every step of the dialogue. However, in some dialogue steps the number of used arguments is not sufficient to make any important observations w.r.t. the VAL^n postulate. Consequently, rather than counting the distinct values in a single distribution, we count the distinct values used in all five distributions defined by a participant⁷. The results are presented in Figure 19 (see also the auxiliary data appendix at [46] for a more detailed table). We can observe that the vast majority of the participants did not use 3 or less values during the dialogue, independently of whether we consider the total or the core sample. The preferred options appear to be 4 and 6 (total sample) or 4 and 5 (core sample) values. These results are significantly different from what we would expect if the participants were providing us with random answers and show that what we obtained were in fact conscious choices⁸.

⁷In this analysis, we treat *Don't know* the same as *Neither Agree nor Disagree* due to the fact that they both evaluate to $3/6$.

⁸Under the assumption that every value is equally likely to be selected, if the participants were answering randomly through-

The above observations indicate that the two (or three) valued perception of acceptance exhibited by the classical Dung’s semantics oversimplifies human modelling and might be insufficient to model the beliefs of the majority of participants. This also confirms the findings of [9]. We may also recall that rational, strict, protective, discharging and trusting postulates were related to classical conflict-free, admissible and complete semantics. While the performance of the preferential branch indicates that in general the statements accepted by the participants are not conflicting, the discharging and trusting results show that admissibility and completeness may be more problematic to satisfy. Thus, breaking down the classical semantics into smaller, separate properties, might give us a more detailed insight into the reasoning of the participants and provide more feedback in situations in which classical semantics are not satisfied. Consequently, more fine-grained approaches than Dung’s semantics should be considered in order to provide a more accurate human modelling.




Figure 19: Usage of different levels of agreement in Dialogue 1

out Dialogue 1 then the probability of them using exactly 7 values throughout the dialogue is 0.704, probability of them using exactly 6 values is 0.272, of using 5 values is 0.024, of using 4 is 0.00048, and of 3 or less is 0.0000015. Exact multinomial test shows there are significant differences between the obtained and the expected probabilities (obtained p-value was small enough that it was rounded up to 0 by the used software).




Figure 20: Postulate satisfaction on the total sample in Dialogue 1




Figure 21: Postulate satisfaction on core sample Dialogue 1

6.2.2. Dialogue 2

In Figures 22 and 23 we can observe the average adherence rates to the epistemic postulates in the total and core samples in Dialogue 2. Let us focus on the total sample first. Some of the observations we have made in the case of Dialogue 1 carry over to Dialogue 2. Similarly as in the first dialogue, the results obtained for the declared, expanded and common graphs appear to be similar in the case of preferential family (see Table 19). However, unlike in the previous case, the behaviour of the intended graph on this family appears to be more distinct from the rest. Fortunately, with the exception of the preferential postulates, the similarities between the results on the common and intended graphs, present in first dialogue, hold here as well. This behaviour may indicate that the intended graphs were too conservative in terms of the considered attacks w.r.t. the common graphs (see also Section 6.1.2.2). We can also note that similarly as in the first dialogue, the adherence to the rational family of postulates decreases the more specialized the property is. Again, the justifiable postulate is barely adhered to.

The differences between the results on the declared graph and other graphs in the case of explanatory and value-based families are more visible than in the first dialogue (see also Table 10 in Section 9.1.1 and Table 19 in Section 9.2.1). In principle, the largest number of differences can be observed between the intended and declared graphs, and exceeds our findings in the first dialogue. Additionally, many of the postulates we have considered do not perform as well as in the first dialogue in terms of the declared graphs. Nevertheless, the rational, semi-optimistic, semi-founded, guarded, discharging and particularly demanding properties appear to achieve better results than in the previous dialogue. We believe some of this behaviour (including also the lower performance of the strict and restrained postulates) can be explained with the behaviour of the binary property. The low satisfaction rate of this particular postulate means that the participants had problems in deciding whether to believe or disbelieve the presented statements using their own knowledge and the information found in the dialogue. This dialogue presented more specialized information than the first dialogue and without access to, for example, Wikipedia, the participants had problems verifying the statements. The presence of the undecided assignments appears to have benefited some of the postulates that are easily satisfied with neutral distributions. We can observe that all of the properties that have experienced an increase in adherence belong to this group. At the same time the undecided assignments harmed those postulates that try to force more decisiveness, such as strict or protective properties, or force belief or disbelief in the face of undecided attackers (e.g. restrained or anticipating postulates). Consequently, in addition to the postulates vulnerable to participants who do not share all of their knowledge (see the analysis of Dialogue 1), we can also identify those that are vulnerable or benefiting from the indecisive participants.

We can now consider the results obtained for the core sample. The average adherence rates of the postulates in the core sample are quite close to the ones from the total sample. The results obtained for different graphs also appear to be more closely connected (see Table 19 in Section 9.2.1). The most visible changes can be observed in the case of the explanatory family of postulates (with the exception of the demanding property), which have all increased their values. The results w.r.t. the expanded graphs on the guarded and discharging postulates become even more distinct from the declared graphs in terms of the average difference of participant responses between these graphs. It therefore appears that the participants were more forthcoming with their personal knowledge, in particular with the additional reasons against believing certain statements. We therefore again identify the properties more vulnerable to such information.

In Figure 24 we can observe the number of participants who have chosen a given number of distinct probabilities throughout the dialogue ⁹ (see also the auxiliary data appendix at [46] for a more detailed table). Similarly to Dialogue 1, the majority of participants choose 4 or more values, though in this case 5 values appear more often in the total sample than 4 values, and they are picked equally often in the core sample. This increase may be related to a wider use of the *Neither Agree nor Disagree* and *Don't know* values, which is visible in the difference of the adherence to the binary postulate between the dialogues. Also in this case it is easy to see that the selection of the values by the participants is clearly not random¹⁰ and thus the tendency of the participants to choose more than 3 values is their conscious choice.

⁹The presented results treated *Don't know* as the same as *Neither Agree nor Disagree*.

¹⁰Under the assumption that every value is equally likely to be selected, if the participants were answering randomly through-




Figure 22: Postulate satisfaction on the total sample in Dialogue 2

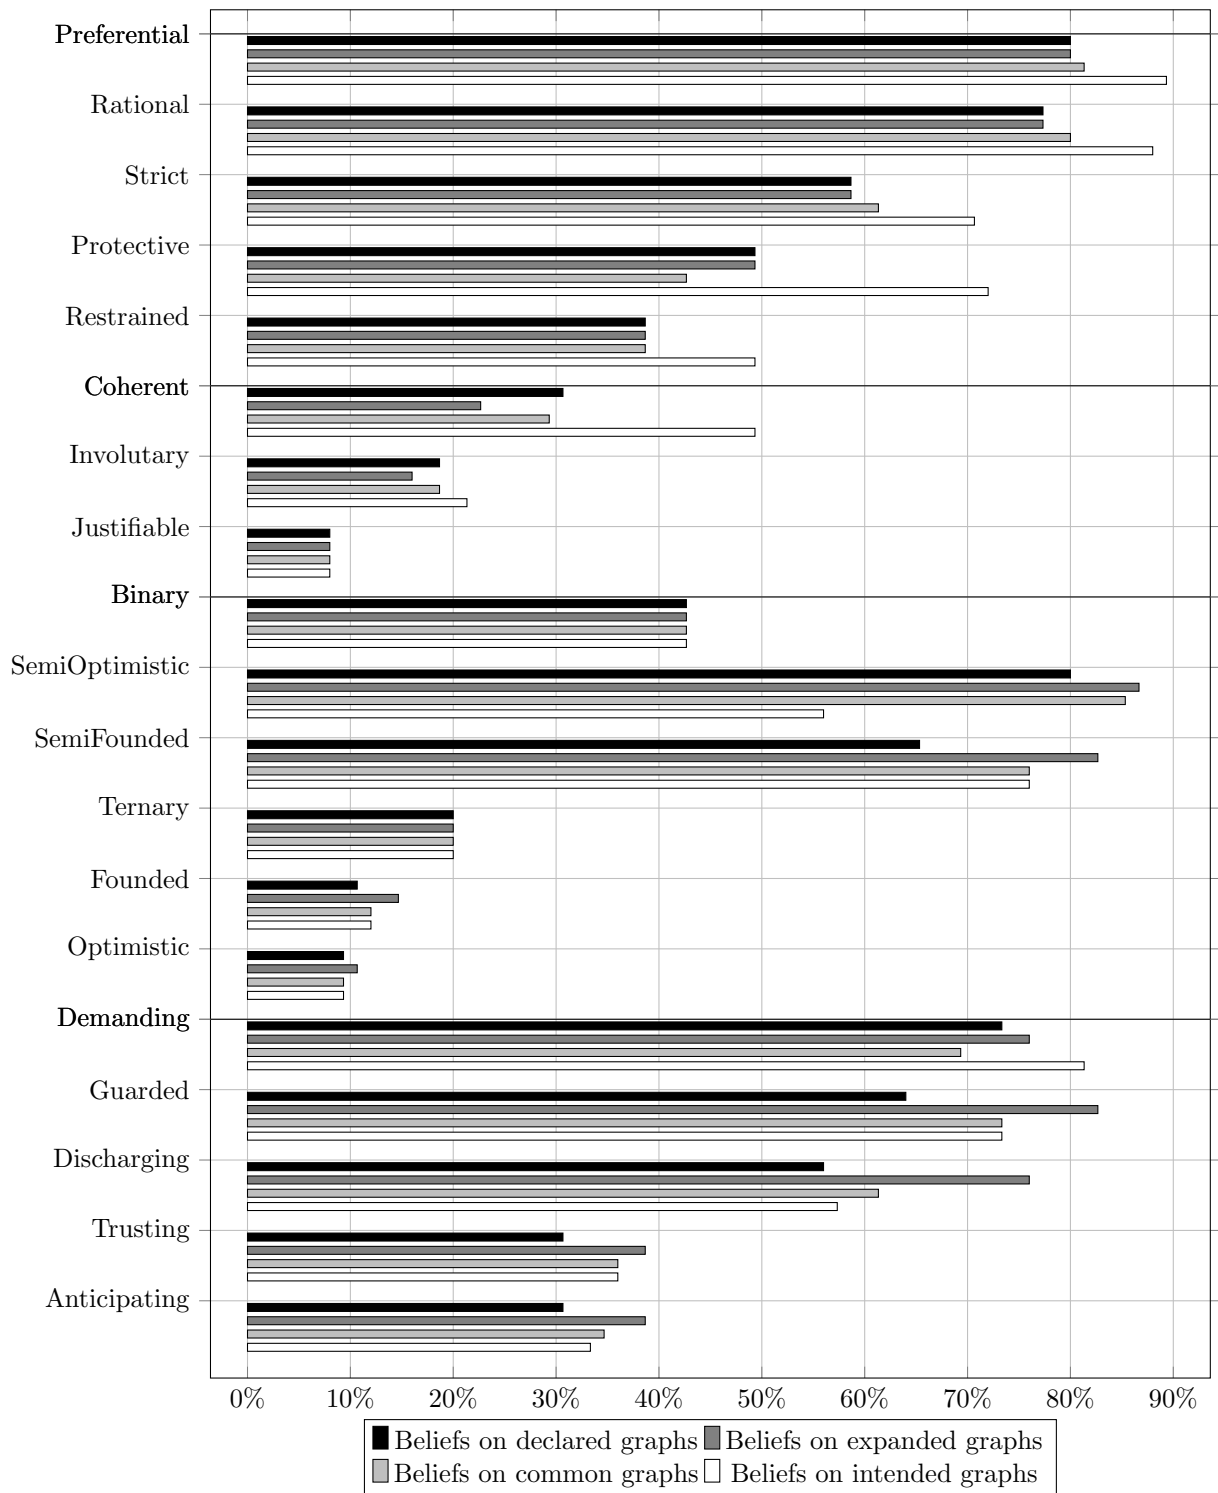


Figure 23: Postulate satisfaction on core sample in Dialogue 2




Figure 24: Usage of different levels of agreement in Dialogue 2

out Dialogue 2 then the probability of them using exactly 7 values throughout the dialogue is 0.876, probability of them using exactly 6 values is 0.121, of using 5 values is 0.003, of using 4 is 0.00002, and of 3 or less is 0.0000000095. Exact multinomial test shows there are significant differences between the obtained and the expected probabilities (obtained p-value was small enough that it was rounded up to 0 by the used software).

6.3. Statements vs Relations

Relations in argumentation frameworks are often defined in terms of the argument structure we are dealing with. Consequently, they are seen as secondary to arguments, and as such their acceptability is rarely considered. One of the exceptions is the attack-based semantics for argumentation frameworks, in which the status assigned to a relation corresponds to the status assigned to its source [53] and vice versa. The same behaviour is replicated by other frameworks allowing the attacks to appear in extensions [54, 55] if we limit ourselves to non-recursive graphs. In this section we would like to verify this view by analyzing the answers to the agreement tasks and relation tasks.

We will perform our analysis in the following manner. For every relation of a given type, by which we understand *A good reason against*, *A somewhat good reason against*, *A good reason for*, *A somewhat good reason for* and *Somehow related, but can't say how*, we create a distribution based on the levels of agreement assigned to the statements carrying them out. Moreover, for every possible level of agreement, we gather the types of relations that the statements assigned a given level carry out. This allows us to observe whether by knowing the value ascribed to a relation we can guess the value assigned to its source and vice versa. In particular, this allows us to verify whether the connection between the values assigned to arguments and relations is as strong as we might think based on some of the research in abstract argumentation. Therefore, what we would expect to obtain is either 1) relations considered “somewhat good” to be carried out primarily by the statements that the participants disagree with and “good” relations to be sourced primarily by the statements that the participants agree with, or 2) relations considered “somewhat good” to be carried out primarily by the statements that have low levels of agreement or disagreement and relations considered “good” to be sourced primarily by the statements that have high levels of agreement or disagreement. In both cases we expect the *Somehow related, but can't say how* edges to be primarily associated with statements of an undecided or unknown status. Our analysis will be carried out on the original as well as pooled data, i.e. one in which certain categories can be merged. Both relation types and agreement levels can be grouped by strength or polarity:

- Agreement grouping:
 - By strength: we group the agreement levels according to how strongly believed or disbelieved they are. We distinguish between *Strong Belief* (created from *Strongly Agree* and *Strongly Disagree*), *Moderate Belief* (combining *Agree* and *Disagree*), *Weak Belief* (grouping *Somewhat Agree* and *Somewhat Disagree*), and *Neither* (created from *Neither Agree nor Disagree* and *Don't Know*) categories.
 - By polarity: we group the agreement levels according to whether they represent belief or disbelief. We distinguish between *Believed* (grouping the *Somewhat Agree*, *Agree* and *Strongly Agree* levels), *Disbelieved* (created from *Somewhat Disagree*, *Disagree* and *Strongly Disagree*), and *Neither* (formed from *Neither Agree nor Disagree* and *Don't Know*) categories.
- Relation grouping:
 - By strength: we group relations according to whether they are good or somewhat good reasons. We distinguish between *Strong Relations* (formed from *A good reason for* and *A good reason against*), *Normal Relations* (grouping *A somewhat good reason for* and *A somewhat good reason against*), and *Dependencies* (representing *Somehow related, but can't say how*).
 - By polarity: we group relations according to whether they are positive or negative. We distinguish between *Attacks* (formed from *A good reason against* and *A somewhat good reason against*), *Support* (grouping *A good for* and *A somewhat good reason for*), and *Dependencies* (representing *Somehow related, but can't say how*).

We consider pooling data in order to be able to more clearly investigate certain trends that may become more apparent with fewer categories.

6.3.1. Dialogue 1

6.3.1.1. Effect of a Relation on its Source

In Figure 25 we can observe how the levels of agreements of the sources of relations of a given kind are distributed. The results show that this agreement distribution is in principle dependent on the nature of the relation in question¹¹. Nevertheless, the distributions visible in Figure 25 do show certain similarities and we carry out additional pairwise-analysis.

We can observe that the distributions of relations marked as being good reasons are relatively similar, as also supported by the results in Table 11 in Section 9.1.2. In both cases, independently of whether we choose the core or total sample, the dominating belief assigned to the source is *Strongly Agree*, with *Agree* and *Strongly Disagree* next in line.

Although some similarities can also be observed between the relations considered as somewhat good reasons for or against a given statement, they are more visible in case of the core sample (Table 11 in Section 9.1.2). In the total sample, the proportions of source agreement levels are in principle distinct. It appears that more participants have chosen disagreement rather than agreement types of sources, though it does not appear that one particular belief is substantially more common than the other in the total sample. In the core sample, the *Somewhat Disagree* option appears to be chosen slightly more often. No particular pattern appears to hold for the *Somewhat related, but can't say how* relation.

Although none of the agreement distributions is random (see Table 12 in Section 9.1.2), we can observe that the “weaker” the relation, the more blurred the picture becomes. Consequently, if we were given a relation, informed of its type and then asked to guess what should the belief in its source be, our best chances of answering correctly are in the case of edges marked as good reasons.

Let us now look at the results obtained for the pooled data. Given the similarity in the behaviour of the relations marked as good reasons for and against in both samples and the somewhat good reasons for and against in the core sample, we will now pool our data according to its strength. We first consider grouping both arguments and relations according to this criterion. We obtain the distributions visible in Figure 26a. The results show that they depend on the chosen relation type, both in the overall¹² and pairwise analysis (see Table 13 in Section 9.1.2), and that they are not random (see Table 14 in Section 9.1.2). We can observe that strong relations are primarily carried out by strong arguments, i.e. those that the participants strongly agree or disagree with. The slightly weaker relations appear to be carried out more by moderately believed arguments, however, the obtained result accounts for less than 50% of the cases. The sources of relations marked as dependencies do not seem to follow any particular pattern. We can only observe that weakly believed arguments are the least common sources.

Let us now consider pooling arguments by polarity and relations by strength. The obtained distributions are depicted in Figure 26b. We can observe that in this case, strong relations are primarily carried out by believed arguments. Normal relations tend to be carried out by disbelieved arguments, though only in the core sample such sources account for more than 50% of the answers. Finally, relations marked as dependencies appear to be carried out by all possible types of arguments, to the point that the results obtained in the case of the core sample are similar to random (see Table 14 in Section 9.1.2). Nevertheless, results in Table 13 in Section 9.1.2 show that despite pooling, the nature of the source is dependent on the type of relation we consider¹³.

Let us now briefly consider pooling relations and arguments according to polarity. Given the similarities in the behaviour of relations marked as good reasons in the original data, it is not surprising that the results obtained for attacks and supports are similar in both core and total samples (see also Table 13 in Section 9.1.2), despite the fact that the overall analysis shows that source distributions depend on the chosen

¹¹G-test for independence yields G-value 317.77 with 28 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library `Deducer` (`likelihood.test` function) in R.

¹²G-test for independence yields G-value 208.01 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library `Deducer` (`likelihood.test` function) in R.

¹³G-test for independence yields G-value 175.33 with 4 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library `Deducer` (`likelihood.test` function) in R.

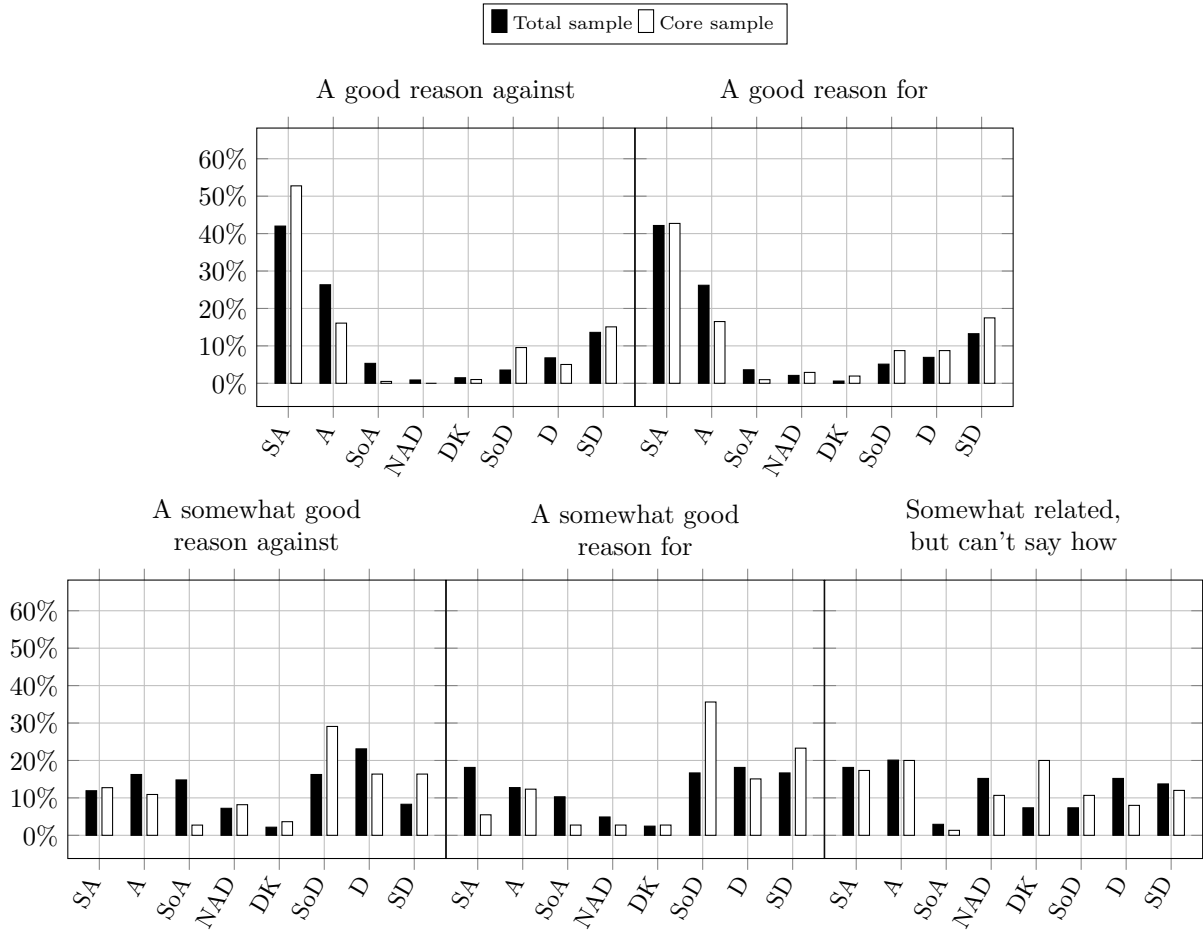



Figure 25: The levels of agreements assigned to the sources of relations of a given type in Dialogue 1. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK)

relation¹⁴. In both cases, the pooled belief marked as *Neither* is the least common option, with *Believed* being the most common choice with the exception of the *Support* relations on the core sample. Again, relations marked as *Dependencies* do not follow any particular pattern, even to the point that the answers can be considered random (see Table 14 in Section 9.1.2). This similarity between attacks and supports persists even if we consider pooling agreement levels according to strength (see Tables 13 and 14 in Section 9.1.2), again despite overall result indicating that all of the relations are distinct¹⁵.


To conclude, we can observe that the dependent relations do not appear to favour sources marked as *Neither* in any particular way, which does not agree with the models we have recalled at the start of the section. Nevertheless, certain interesting patterns can be observed for strong and normal relations. In particular, strong relations tend to be carried out by arguments that are believed (polarity pooling) and strongly believed (strength pooling). Grouping relations by polarity leads to similar behaviour of attack and support relations, which may mean that future experimental data on one of them might be generalized to the other.

¹⁴G-test for independence yields G-value 62.846 with 4 degrees of freedom and p-value of 7.312×10^{-13} . This result was obtained using the library Deducer (likelihood.test function) in R.

¹⁵G-test for independence yields G-value 69.849 with 6 degrees of freedom and p-value of 4.391×10^{-13} . This result was obtained using the library Deducer (likelihood.test function) in R.



(a) Arguments and relations pooled by strength



(b) Arguments pooled by polarity and relations pooled by strength

Figure 26: The levels of agreements assigned to the sources of relations of a given type in Dialogue 1 according to a given type of pooling. We use the following abbreviations: Strong Belief (SB), Normal Belief (NB), Weak Belief (WB), Neither (N), Disbelief (D), Belief (B).

6.3.1.2. Effect of an Argument on the Relations it Carries Out

Let us now consider the dual problem i.e. what relations are carried out by an argument assigned a given agreement level? In Figure 27 we can find the data extracted from the answers of our participants. Although in principle the relation distribution does depend on the strength of the argument carrying it out¹⁶, we perform additional pairwise analysis. We can observe that while in the previous case the charts of relations with the same status but opposite polarity shared certain similarities, the grouping in this case is somewhat different. The graphs of *Strongly Agree*, *Agree* and potentially *Strongly Disagree* do tend to put more focus on the good reason for and against edges, however, they in general lead to different results (see Table 15 in Section 9.1.3). The *Somewhat Agree*, *Somewhat Disagree* and *Disagree* statements on the other hand appear to lead more to relations marked as somewhat good reasons for or against. The similarity in the types of relations these three agreement levels carry out is also reflected by the results in Table 15 in Section 9.1.3. The *Somewhat good reason against* and *Somehow related, but can't say how* are the two most common choices when it comes to the statements that the participants neither agreed nor disagreed with, the latter being the dominant relation resulting from the elements marked as *Don't know*. These two levels of agreement also lead to similar distributions (see Table 15 in Section 9.1.3).

Based on the results in Table 16 in Section 9.1.3, we can also observe that our results are in principle not random. The results for agreement levels marked as *Somewhat* in the core sample are not significant, however, given the distributions in Figure 27, this might be related to the sample size.

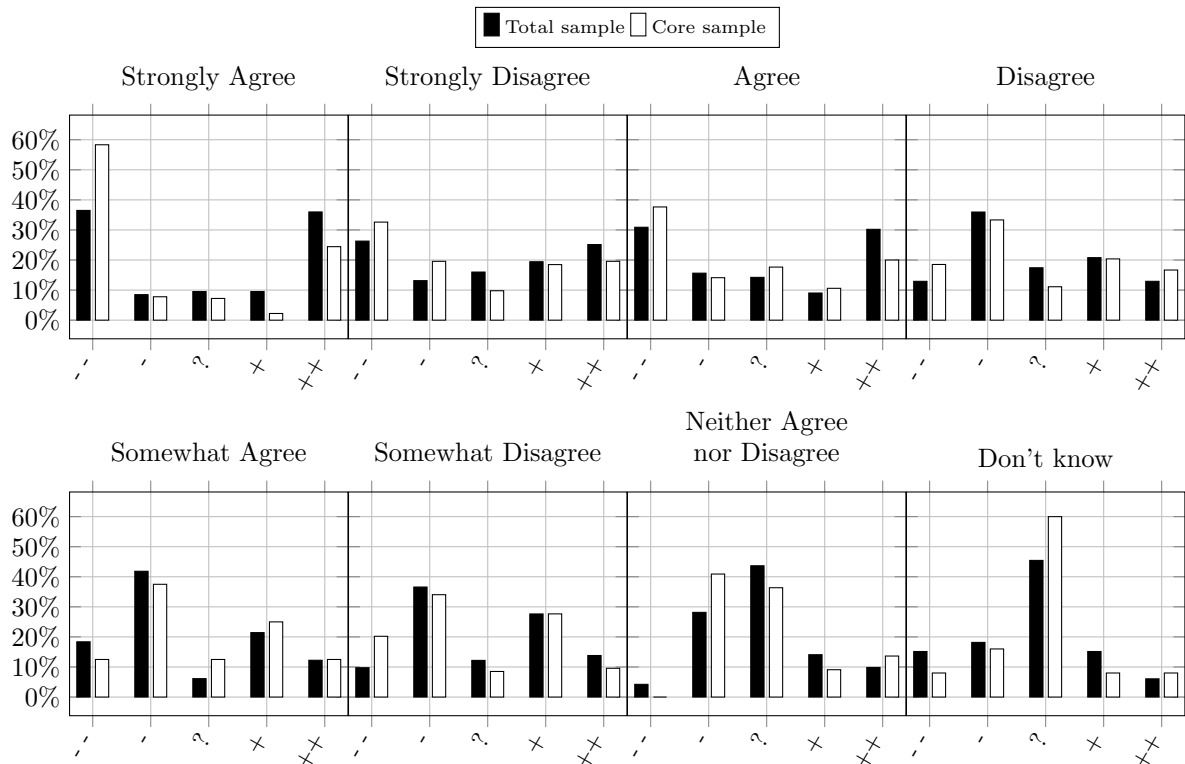


Figure 27: The types of relations carried out by statements of a given level of agreement in Dialogue 1. We use the following abbreviations for the relation types: A good reason against (- -), A somewhat good reason against (-), A good reason for (++), A somewhat good reason for (+), Somewhat related, but can't say how (?)

In many cases, this initial analysis points to the statements with a particular level of agreement favouring relations of a corresponding “power”. Thus, we will now consider pooling our results by strength, as visible

¹⁶G-test for independence yields G-value 317.77 with 28 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

in Figure 28a. All of the belief strengths affect the associated relation distributions in the total sample¹⁷, although the results for the moderate and weak belief appear to be more closely related on the core sample (see Table 17 in Section 9.1.3). Nevertheless, the obtained distributions are not random (see Table 18 in Section 9.1.3) and we can observe certain interesting patterns emerging. Strong belief leads primarily to strong relations and weak belief leads primarily to normal relations (obtained results account for 55% to 72% of the answers, depending on the sample). Relations caused by the arguments with moderate belief are almost evenly split between strong and normal relations. Arguments that are classified as *Neither* often carry out dependencies, however, they do not account for 50% of all of the possible relations.

Let us now consider pooling agreement levels according to belief and relations according to strength. By doing so, we obtain the distributions visible in Figure 28b. We can observe that believed arguments lead primarily to strong relations, disbelieved to normal relations, and neither to dependencies. However, clear dominance is visible only in the case of believed arguments. In all cases, the chosen agreement level leads to a different relation distribution¹⁸ that is not random (see Tables 17 and 18 in Section 9.1.3).

Finally, let us briefly consider agreement levels paired with relations pooled by polarity. Although tests show that the obtained results are in principle distinct, many similarities become apparent in the pairwise analysis¹⁹. If we group levels by polarity, we obtain the result that the proportions of the relation types carried out by believed and disbelieved arguments are not significantly different (see Tables 17 and 18 in Section 9.1.3). This similarity may in part be due to the connections between the *Somewhat Agree*, *Somewhat Disagree* and *Disagree* agreement levels. It might also explain the fact that when we consider merging arguments by strength, the similarities between the weak and moderate or strong levels of agreement become visible (see Tables 17 and 18 in Section 9.1.3).

To conclude, we can observe that the tendencies observable in Figures 28a and 28b seem to agree with what the models we have previously recalled would assume. However, the domination of the desired relations is not as significant in all situations as we would expect. Similarly as in the previous relation analysis, the tendencies for strongly believed (strength pooling) and believed (polarity pooling) arguments are the most visible.

6.3.2. Dialogue 2

6.3.2.1. Effect of a Relation on its Source

In Figure 29 we can observe how the levels of agreements on the sources of relations are distributed. The tests show that the agreement distributions depend on the type of the relation we consider, both in the overall²⁰ and pairwise analysis w.r.t. the total sample and in almost all cases w.r.t. the core sample (see Table 20 in Section 9.2.2). In all cases, the obtained distributions are not random (Table 21 in Section 9.2.2).

As in the previous dialogue, we can observe certain similarities in the distributions of relations marked as good reasons for and against, particularly in the case of the core sample (Table 20 in Section 9.2.2). However, it is easy to see that these results are substantially different from the ones observed in Dialogue 1 (see Figure 25). In this case, the *Agree* and *Neither Agree nor Disagree* options appear to be the two most common choices. Nevertheless, none of them exceeds 25% of the given answers. The options representing agreement appear to be favoured over those representing disagreement, which we will take into consideration when analyzing pooled results.

The *Agree* and *Neither Agree nor Disagree* answers are also common when it comes to relations marked as *A somewhat good reason against*. However, depending on whether we are dealing with the total or core

¹⁷G-test for independence yields G-value 208.01 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

¹⁸G-test for independence yields G-value 175.33 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

¹⁹When pooling arguments by strength, G-test for independence yields G-value 69.849 with 6 degrees of freedom and p-value of 4.391×10^{-13} . When pooling arguments by polarity, G-test for independence yields G-value 62.846 with 4 degrees of freedom and p-value of 7.312×10^{-13} . This result was obtained using the library Deducer (likelihood.test function) in R.

²⁰G-test for independence yields G-value 390.02 with 28 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

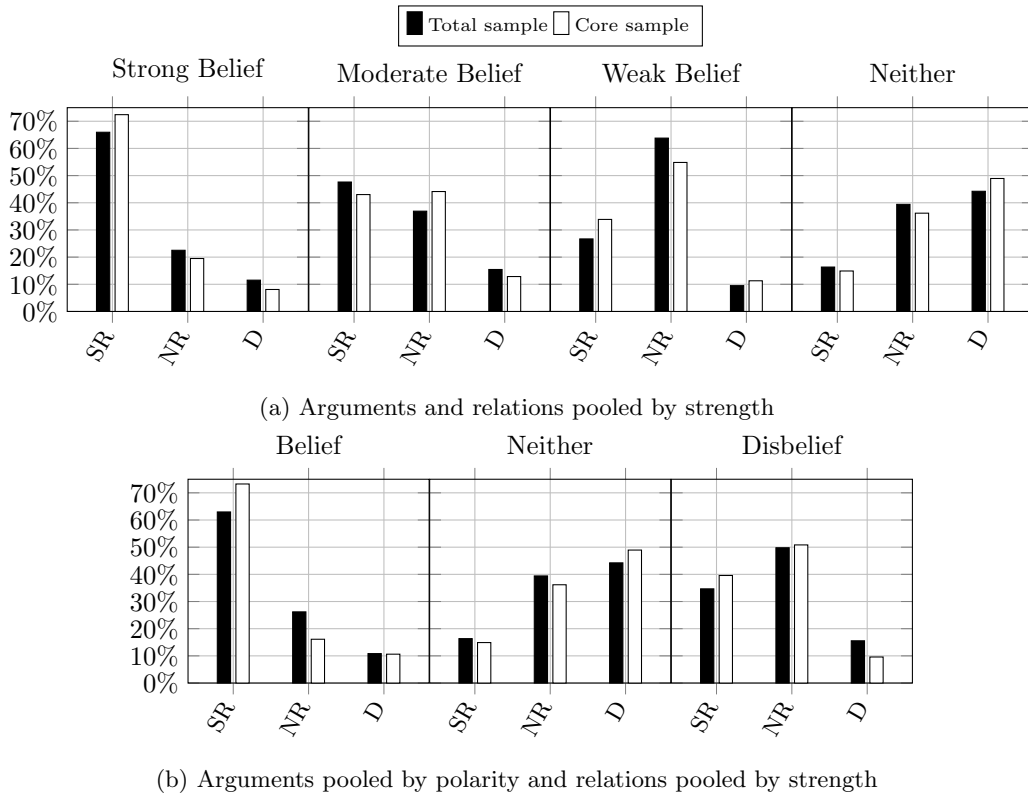


Figure 28: The types of relations carried out by statements of a given level of agreement in Dialogue 1 according to a given pooling. We use the following abbreviations for relation types: Strong Relation (SR), Normal Relation (NR), Dependency (D).

sample, the *Somewhat Agree* or *Somewhat Disagree* value is more frequent. It is worth noticing is that the strongest possible values – *Strongly Agree* and *Strongly Disagree* – are the least common. These values are also rarely chosen in the case of the *A somewhat good reason for* relation. For this edge, the *Neither Agree nor Disagree* level of agreement is the most common value, both in the total and in the core sample.

Finally, we reach the relations marked as *Somewhat related, but can't say how*. This is the only case in which the domination of a single value – *Neither Agree nor Disagree* – is clearly visible, particularly in the case of the core sample. Additionally, in the total sample we can observe that the *Don't know* option may also be of some importance. Thus, we can observe that independently of the chosen relation type, the *Neither Agree nor Disagree* agreement level of the source appears to be quite common, often with *Agree*, *Somewhat Agree* or *Somewhat Disagree* not far behind. The observed behaviour of the relations is also notably different from what we have seen in Dialogue 1.

Let us now analyze pooled data. We start by grouping relations by strength, and pool agreement levels first by strength and afterwards by polarity. We obtain the distributions visible in Figures 30a and 30b. The results show that the way the sources are distributed depends on the chosen relation type, both in the overall²¹ and pairwise analysis (see Table 22 in Section 9.2.2). Furthermore, the results are also not random (Table 23 in Section 9.2.2). Let us now look more closely at the results with strength pooling on both agreement levels and relations. We can observe that unlike in the first dialogue, the distribution associated with the strong relations does not appear to follow any particular pattern. We can only note that weak belief is the

²¹In case both arguments and relations have been pooled by strength, the G-test for independence yields G-value 232.25 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . In pooling arguments by polarity and relations by strength, we obtain G-value 167.86, 4 degrees of freedom and again p-value less than 2.2×10^{-16} . This result was obtained using the library `Deducer` (`likelihood.test` function) in R.

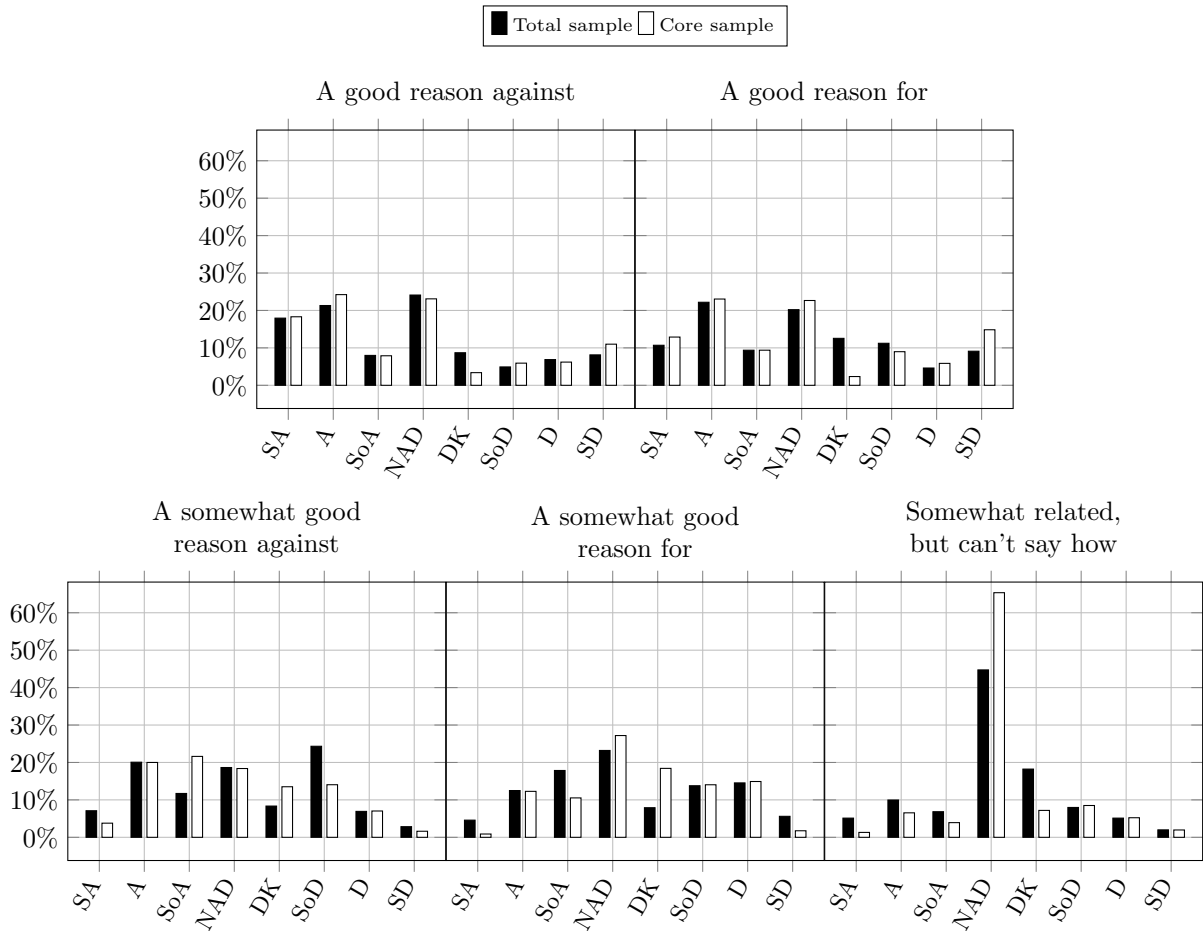


Figure 29: The levels of agreements assigned to sources of relations of a given type in Dialogue 2. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK)

least common type of a source. In a similar fashion, normal relations are carried out by strongly believed arguments the least often, with the remaining types of sources being spread relatively evenly. Finally, while in Dialogue 1 relations marked as *Dependencies* did not appear to follow any patterns, in Dialogue 2 the dominance of sources grouped as *Neither* is clear.

The dominance of *Neither* sources in *Dependencies* is also quite clear if we look at the results of arguments pooled by polarity and relations by strength. We can observe that in this analysis, strong relations tend to be carried out by believed arguments. Although this repeats the results of Dialogue 1, in this case the dominance is not entirely clear and believed sources do not account for 50% of the relations. Additionally, unlike in Dialogue 1, the disbelieved arguments appear to be the least common sources. No particular pattern emerges when we look at normal relations – we can only observe that the sources marked as *Neither* appear to be somewhat less common than other types.

Let us now briefly consider pooling relations and arguments according to polarity. Despite the fact that the overall analysis shows that source distributions depend on the chosen relation²², pairwise analysis gives us a different result. Although the similarities in the behaviour of relations marked as good reasons in the

²²G-test for independence yields G-value 132.84 with 4 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library *Deducer* (likelihood.test function) in R.

original data were only apparent in the core sample, the pooled results obtained for attacks and supports are similar in both core and total samples (see also Table 22 in Section 9.2.2). Similarly as in the first dialogue, the *Believed* level is the most common choice. However, in this dialogue, the relations marked as *Dependencies* do follow a pattern and are carried out primarily by arguments marked as *Neither*. In all cases, our results are not random (see Table 23 in Section 9.2.2). Let us now pool agreement levels according to strength. The overall test shows that the source distribution depends on the chosen relation ²³. However, a detailed analysis shows that agreement distributions associated with attacks and supports still share certain similarities (see Tables 22 and 23 in Section 9.2.2), even though it becomes apparent mostly on the core sample.

To conclude, although there are certain characteristics present in both dialogues, not all of the findings from Dialogue 2 coincide with those from Dialogue 1. This is particularly the case when we consider pooling relations by strength. Results associated with agreements pooled by strength are quite distinct from the results in Dialogue 1. It is worth observing that relations marked as *Dependency*, which previously did not follow any particular pattern, now favour agreement levels marked as *Neither*. While strong relations in both dialogues do tend to favour believed arguments more than disbelieved or neither in polarity pooling of agreement levels, the domination of this value in Dialogue 1 is not replicated in Dialogue 2.

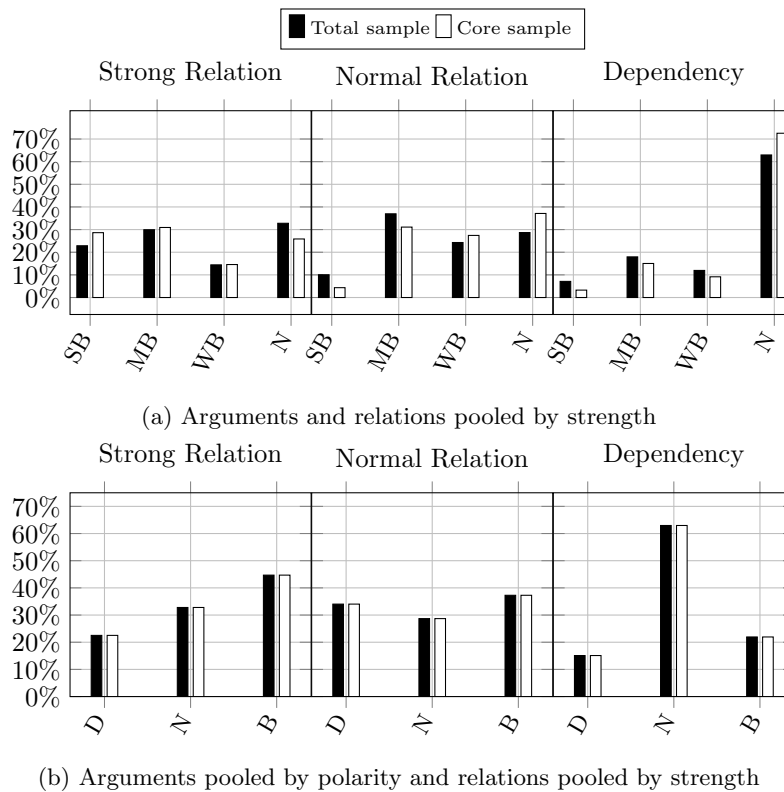


Figure 30: The levels of agreements assigned to the sources of relations of a given type in Dialogue 2 according to a given type of pooling. We use the following abbreviations: Strong Belief (SB), Normal Belief (NB), Weak Belief (WB), Neither (N), Disbelief (D), Belief (B).

²³G-test for independence yields G-value 142.15 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

6.3.2.2. Effect of an Argument on the Relations it Carries Out

In Figure 31 we can find the data showing what sort of relations are carried out by the statements with a particular agreement level. In general, the choice of agreement level affects the resulting relation types²⁴. However, in the pairwise analysis, certain similarities become more visible (see Table 24 in Section 9.2.3).

We may notice that some of the observations made in the case of Dialogue 1 carry over to Dialogue 2. In particular, the charts created for *Strongly Agree*, *Agree* and *Strongly Disagree* tend to put more focus on the good reason for and against edges. Nevertheless, the distributions are technically speaking distinct, and the similarity between the strong agreement and disagreement is detected only in the core sample. The *Disagree* level, depending on the sample we are dealing with, favours *A good reason against* and *Somewhat good reason for* edges, with *A good reason for* and *A somewhat good reason against* next in line.

The *Somewhat Agree* and *Somewhat Disagree* charts are quite similar, as also supported by results in Table 24 in Section 9.2.3. In this case, the *A somewhat good reason against* agreement level is perhaps more common than other options. The *Neither Agree nor Disagree* agreement level appears to favour the *Somewhat related, but can't say how* option the most and the *Somewhat good reason for/against* answers the least. Finally, we can observe that the distributions obtained from the total and core samples in the case of the *Don't know* chart are quite different; one favours the *A good reason for* answer, while the other the *Somewhat good reason for/against* options. We can observe that the distributions associated with the agreement levels are in principle not random (see Table 25 in Section 9.2.3), though there appear to be certain difficulties with the *Disagree* and *Somewhat Disagree* values in the core sample.




Figure 31: The types of relations carried out by statements of a given level of agreement in Dialogue 2. We use the following abbreviations for the relation types: A good reason against (--), A somewhat good reason against (-), A good reason for (++) , A somewhat good reason for (+), Somewhat related, but can't say how (?)

²⁴G-test for independence yields G-value 390.02 with 28 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

Let us now consider pooling our results. We start by grouping arguments and relations by strength, as visible in Figure 32a. All of the belief strengths affect the associated relation distributions, both in the overall²⁵ and pairwise analysis (see Table 26 in Section 9.2.3). Additionally, the answers extracted from the users are not random (see Table 27 in Section 9.2.3). We can observe that the strong relations account for at least 50% of connections carried out by arguments that are strongly or moderately believed. Weak belief appears to be split evenly between the strong and normal relations. Arguments that are neither believed nor disbelieved can lead to all three types of relations, though strong relations are more favoured than others. Consequently, not all of the observations made in the case of Dialogue 1 are repeated in Dialogue 2.

Let us now consider grouping arguments by polarity and relations by strength (see Figure 32b). Again, all of the belief strengths affect the associated relation distributions, both in the overall²⁶ and pairwise analysis (see Table 26 in Section 9.2.3). The answers extracted from the users are also not random (see Table 27 in Section 9.2.3). Nevertheless, there are certain similarities between the obtained distributions. In particular, we can observe that strong relations appear to be the most common relations carried out by arguments marked as *Believed* or *Neither*, independently of the sample, and by *Disbelieved* in the core sample. In particular, in the first and the last case these relations account for over 60% of all edges. What is worth observing is that dependent links are the least common when it comes to arguments that are believed or disbelieved. Nevertheless, it is easy to see that the results obtained for agreement levels marked as *Neither* and *Disbelieved* in the second dialogue are different from the first dialogue.

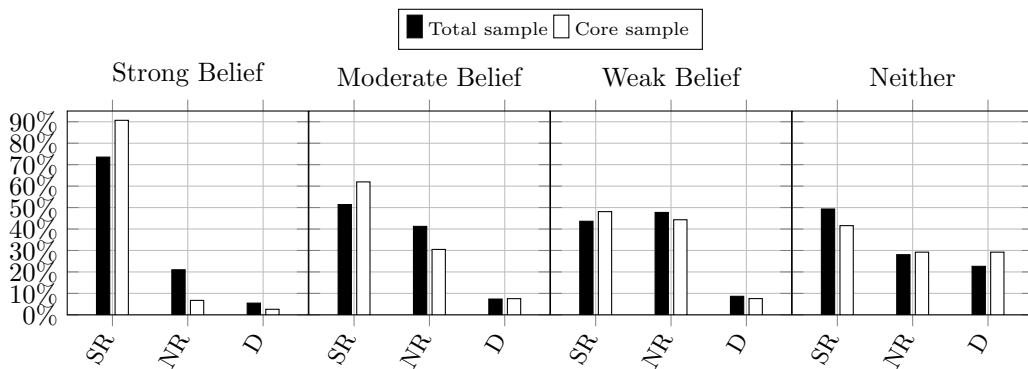
We will now briefly consider agreement levels paired with relations pooled by polarity. Although tests show that the obtained results are in principle distinct, certain similarities become apparent in the pairwise analysis²⁷. If we group agreement levels by polarity, we obtain the result that the proportions of the relation types carried out by the believed and disbelieved arguments are not significantly different (see Table 26 in Section 9.2.3), particularly in the case of the total sample. When we consider merging arguments by strength, certain similarities between the weak, moderate and strong levels of agreement become visible (see Tables 26 and 27 in Section 9.2.3).

To conclude, despite various differences between Dialogue 1 and Dialogue 2, it again appears that analysis in which relations are pooled by polarity leads to situations where relation distributions might not in all cases depend on the agreement level of the source. When pooling relations by strength, we can also observe that similarly to Dialogue 1, the strongly believed (strength pooling) and believed (polarity pooling) arguments lead primarily to strong relations. However, the remaining pooled agreement levels do not appear to follow exactly the same patterns. Even though certain shifts towards weaker relations can be observed in Figure 32a, their progress can be seen as “slower” than in Dialogue 1.


²⁵G-test for independence yields G-value 232.25 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

²⁶G-test for independence yields G-value 167.86 with 4 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.

²⁷When pooling arguments by strength, G-test for independence yields G-value 142.15 with 6 degrees of freedom and p-value less than 2.2×10^{-16} . When pooling arguments by polarity, G-test for independence yields G-value 132.84 with 4 degrees of freedom and p-value less than 2.2×10^{-16} . This result was obtained using the library Deducer (likelihood.test function) in R.



(a) Arguments and relations pooled by strength



(b) Arguments pooled by polarity and relations pooled by strength

Figure 32: The types of relations carried out by statements of a given level of agreement in Dialogue 2 according to a given pooling. We use the following abbreviations for the relation types: Strong Relation (SR), Normal Relation (NR), Dependency (D).

6.4. Changes in Beliefs

In this section we discuss the statement awareness declared by the participants in both of the dialogues and the changes – or rather, the lack of them – in participants’ opinions. In Figure 33 we can observe that the statements we have used in the dialogues were, in general, not common knowledge.

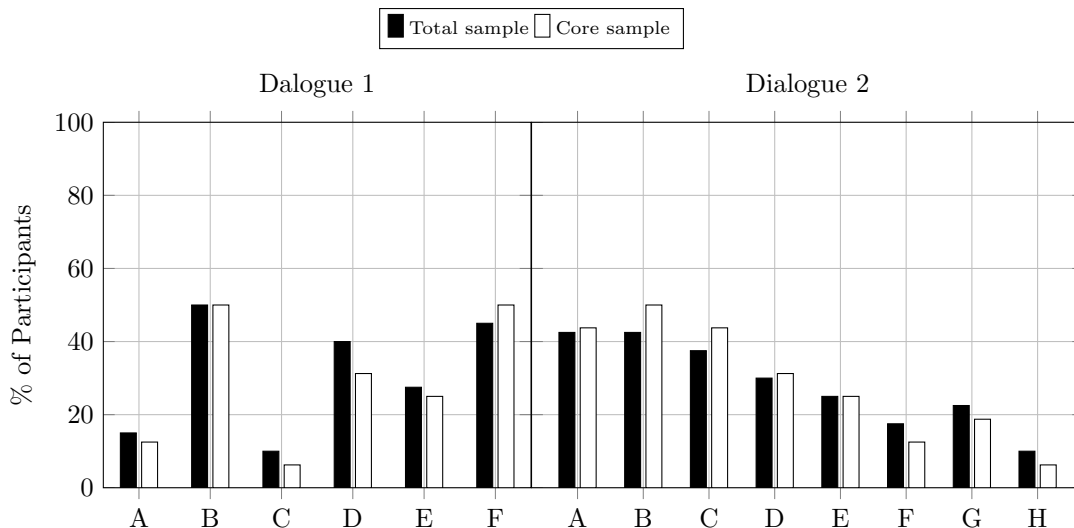


Figure 33: The percentage of participants in the core and total samples in both dialogues who have stated that they were aware of a given argument prior to the experiment

Nevertheless, this lack of awareness of the presented pieces of information did not result in participants changing their opinions easily. For example, in the first dialogue, the majority of participants – resp. 90% and 87.5% in the total and core samples – agreed with *B* on all stages of the dialogue, with the remainder disagreeing with *B* on all stages of the dialogue. The majority of participants – 75% and 68.75% in the total and core samples – disagreed with *A* on all stages of the dialogue, with resp. 17.5% and 25% agreeing with *A* on all stages of the dialogue. This leaves us only with a couple of participants “swapping sides” in the dialogue. More such people can be found in the second dialogue, where we have 17.5% and 26.67% participants changing between agreeing, disagreeing and being undecided at least once in the case of statement *B* in the total and core samples, with 10% and 20% changing their opinions in the case of *A*. With all the opportunities that the participants were given to change their opinion, the number of people who did not at least once alternate between belief and disbelief appears extremely high. Even in a more detailed analysis, where we differentiate between all the declarable levels of agreement rather than just belief–disbelief, the significant changes appear to primarily concern argument *B* in the first dialogue (see Table 28). Hence, even though the dialogues appear to have certain effect on the participants, it is more visible on the fine-grained level.

Despite these results, there are few notable changes that have occurred and which we would like to discuss. In the first dialogue, we have two participants who have changed their position concerning the statement *A*: one who initially disbelieved *A*, but believed it at the end, and one who initially believed *A* and disbelieved it at the end (this particular person was also present in the core sample for in this dialogue). In both cases, the change appeared in the second step of the dialogue. The first person initially disagreed with both *A* and *B* and apparently interpreted the dialogue as discussing whether the hospital staff should be demanded to take the vaccine, not whether it would be best for them to take it. Once statement *C* was presented, which in the opinion of this participant reinforced the fact that the doctors don’t need the shots, he started agreeing with *A* and no further counterarguments fixed this situation. The second participant initially agreed with both *A* and *B* and provided no real explanation for his opinions. However, the appearance of *C*, with which the participant disagreed, led to further disagreement with *A* throughout the rest of the dialogue.

In the second dialogue, we have five participants that we would like to discuss: 1) two for whom initial agreement with *B* turned to disagreement without affecting *A*, 2) one for whom agreement (disagreement) with *B* (with *A*) turned to disagreement (agreement), 3) one for whom disagreement with *A* turned into indecisiveness, and 4) one for whom indecisiveness with *B* turned into disagreement. The last three participants belong to the core sample. In the majority of these cases, the participants, either due to some initial skepticism or lack of knowledge, became swayed by the incorrect statements, namely *C*, *D* and *G*. This was often paired with the inability to verify the correctness of *H*, or with interpreting it as an admission that the contained substances were so harmful that they had to be removed for the sake of the children²⁸.

In Figure 34 we present the average change in beliefs of the participants, which is calculated in the following manner. For every participant, we have summed the differences in their beliefs throughout all the stages in the dialogue on statements they were aware of and those there were not w.r.t. the value assignment explained at the beginning of Section 6. By dividing this sum by the number of statements declared as aware or unaware by a given participant (assuming this number is non-zero), we have obtained an average change in statements of a given type per person. Based on this we calculate the average change in statements of a given type per sample²⁹. As we may observe, the changes are quite modest, and relatively close in both aware and unaware statements. This is particularly visible in Dialogue 1, where the average change is smaller than the single smallest change possible (e.g. the difference between *Strongly Agree* to *Agree*, which are mapped to 6/6 and 5/6 respectively, is approximately 0.167). The variability in Dialogue 2 is somewhat larger, though one also has to bear in mind that more statements were presented to the participants and therefore they were given more opportunities to change their mind per statement. Nevertheless, it is possible that some of the observed changes were noise rather than actual modifications in beliefs caused by the dialogue. This is particularly supported by the fact that various possible explanations for these differences, such as the appearance of an attacker, could have been paired with an increase as well as decrease in the beliefs declared by the participants.




Figure 34: Average change in beliefs in dialogues per sample, calculated using the value assignment from page 21.

It appears that in many cases, the participants were able to evaluate the statements they were not familiar with by using their own knowledge and other facts that they were aware of. Unfortunately, this also happened in the cases where the knowledge of the participants was incorrect, such as the vaccines causing autism or flu shots causing flu. However, the complete inability to verify the pieces of information contained in the statements with their own knowledge had made some participants of Dialogue 2 choose to neither agree nor disagree with the initial statements presented in it. This might mean that depending on the difficulty of the dialogue, too much as well as too little personal knowledge might affect how open to change people are.

We would also like to note that we have found participants who, on at least one occasion, gave answers such as “*I agree/disagree with this statement*”, “*This information is (in)accurate*” or “*I think/believe this is true/false*” in the explanation tasks and refrained from providing a reason for their opinions. By considering only those cases in which the participants claimed they were not aware of the presented pieces of information

²⁸The thimerosal has been removed from the vaccine primarily due to social pressure and overwhelming distrust, not any scientific evidence of its harmfulness.

²⁹Participants who did not have any unaware (resp. aware) statements were excluded from the appropriate calculation.

and the stage of the dialogue they were at should not have given them sufficient reasons for evaluating these statements, we have found respectively 8 and 5 participants for the first and second dialogues resorting to such explanations. These participants also exhibited insignificant, if any, changes in their opinions. Therefore, it might be the case that for various reasons, they could not or did not want to provide us with further details.

7. Related Work

In this section we would like to review the results of other works dealing with empirical evaluation of argumentation approaches and discuss other works dealing with more fine-grained models to argument acceptability. We begin by analyzing the work by Rahwan et al [9], which focuses on the issue of reinstatement. The authors have prepared a total of 10 short dialogues in order to study this phenomenon. A sample dialogue for the case of simple reinstatement is as follows:

A Louis applied the brake and the brake was not faulty. Therefore, the car slowed down.

B The brake fluid was empty. Therefore, the brake was faulty.

C The car had just undergone maintenance service. Therefore, the brake fluid was not empty.

The dialogue was split into three stages, where first **A**, then **A** and **B**, and finally all **A**, **B** and **C** were presented. For each problem, the participants had to choose a value from a 7-point scale ranging from *Certainly false* to *Certainly true* that best describes their confidence in the conclusion of **A**. The obtained results support the notions of defeat and reinstatement, that is, that the belief in **A** decreases once it is defeated and increases when it is defended, though still remaining significantly lower than prior to the defeat.

This study, similarly to ours, lends support to the use of more fine-grained approaches towards describing the beliefs of the participants, such as the epistemic approaches. However, as we can observe, the dialogues used in this study were much simpler and shorter than ours, and unlikely to be affected by any subjective views of the participants. The results show that reinstatement as such does occur and thus support the idea behind semantics such as grounded or preferred. However, at this point it cannot be claimed that they are applicable in a more complex setting and that they are a general proof for the use of argumentation semantics in computational persuasion.

The next study we would like to consider is by Cerutti et al [6], which again studies various forms of reinstatement. The authors have prepared two texts – a base one and an extended one – on four topics (weather forecast, political debate, used car sale and romantic relationship). For example, the texts prepared for the political debate scenario are as follows:

Base Case: In a TV debate, the politician AAA argues that if Region X becomes independent then Xs citizens will be poorer than now. Subsequently, financial expert Dr. BBB presents a document, which scientifically shows that Region X will not be worse off financially if it becomes independent.

Extended Case: In a TV debate, the politician AAA argues that if Region X becomes independent then Xs citizens will be poorer than now. Subsequently, financial expert Dr. BBB presents a document, which scientifically shows that Region X will not be worse off financially if it becomes independent. After that, the moderator of the debate reminds BBB of more recent research by several important economists that disputes the claims in that document.

The participants are given one of the eight scenarios and then asked to determine which of the following positions they think is accurate:

- P_A : I think that AAAs position is correct (e.g. “Xs citizens will be poorer than now”).
- P_B : I think that BBBs position is correct (e.g. “Xs citizens will not be worse off financially”).

- P_U : I cannot determine if either AAAs or BBBs position is correct (e.g. “I cannot conclude anything about Region Xs finances”).

Next, the participants were asked certain questions concerning the text. In particular, concerning the agreement, the participants are asked “How much do you agree with the following statements?” and respond on a 7–point scale from *Disagree* to *Agree* for each statement.

The results show that the majority of participants in the base case scenarios agree with position P_B and with position P_U in the extended scenarios. They suggest a correspondence between the formal theory used by the authors and its representation in natural language. However, additional analyses show that there are some significant deviations, apparently caused by the personal knowledge of the participants. For example, while P_B was the most common choice in the base weather forecast scenario, there were participants who chose P_U and explained their decision with “*All weather forecasts are notoriously inaccurate*”.

Our study, similarly to this one, points to the personal knowledge of the participants affecting their decisions. However, due to a different methodology, it is difficult to compare the other findings. We have presented the participants with an ongoing dialogue and tried to monitor the changes in their beliefs at every step of the way, similarly as in [9]. In this study, a given participant receives only one text to evaluate, i.e. the base and extended scenarios are judged by different people. We therefore cannot claim that if a person was first shown the base case, and then the extended one, the P_U choice would still be more common than P_B , which is what our results would point to.

Another interesting study worth considering is [10], which focuses on exploring the abilities of argumentation theory, relevance heuristics, machine and transfer learning for predicting the argument choices of participants, with a particular focus on the machine learning. Argumentation theory is verified using three experiments, in which the dialogues are used to construct bipolar argumentation frameworks and the sets of arguments selected by the participants are contrasted with grounded, preferred and stable extensions. In the first experiment consisting of 6 scenarios, the authors create bipolar argumentation frameworks which are not known to the participants, present two standpoints from two parties and ask the participants to choose which of the additional four arguments they would use next if they were one of the deliberants in the discussion. In the second experiment, real life data is annotated and structured into a bipolar argumentation framework that is later analyzed. In the third experiment, a chat bot is created, aimed at discussing flu vaccination with the participants. In this case, both the chat and the participant can only use arguments from a predefined list. Finally, in an additional experiment, an argument theory agent was implemented in order to provide suggestions to the participants during a two–person chat.

The authors report that a substantial part of the results (or in some cases, even the majority) do not conform to the results predicted by the semantics. It is worth mentioning that the stated adherence to the conflict–free extension–based semantics is 78% and is similar to our results concerning the rational postulate, which corresponds to this semantics. The created agent was also considered helpful only by a small percentage of the participants. Nevertheless, the causes for such behaviour of the semantics are not investigated, and the participants were not allowed to explain their decisions (the first and the third experiment) or there was no possibility to ask them for further input (the second experiment). Moreover, unlike in our study, the participants were evaluated against the graphs constructed by the authors or annotators. As shown by our results, they do not necessarily reflect how the participants view the relations between the arguments. Finally, there is no discussion concerning whether these particular bipolar argument framework semantics used in this experiment [56] are applicable. The stable and conflict–free semantics in this work are based on direct and supported attacks, and only the direct ones need to be defended from and can be used for defence. Taking into account the fact that this approach has been superseded by a number of different methods since it has been introduced, the presented results indicate that these particular semantics are not useful in modelling of the user behaviour, rather than there exists a deeper issue within the argumentation itself.

The aforementioned studies focused on the empirical evaluation of certain phenomena in argumentation. However, there are also other studies, which focus more on the behavioural methods or computational linguistics perspective rather than the argumentation perspective and can be seen as complementary to ours. These works often use arguments sourced from e.g. social media, and analyze the relation between

the persuasiveness of an argument and its traits, the persuasiveness of an argument or relations between the arguments and the personality or the emotions of the participant, and more. Some of them, such as [57], also point to the importance of the prior knowledge in changing one’s beliefs. These studies are an important line of work which can be harnessed in creating computational persuasion methods that are tailored to the participants and which can provide guidance in transforming logical arguments to natural language arguments. Therefore, they should be considered in the next steps of our work. We refer the readers to [58, 57] for further details.

Although in this work we have mostly focused on the analysis of epistemic probabilities, we would like to note that this is not the only approach allowing a more fine-grained perspective on argument acceptability. There also are other interesting methods that allow score assignments, such as certain forms of preference or weight-based argumentation (see [2, 59] for an overview). However, many of these works share one common problem, namely that the values associated with the arguments are quite abstract and do not have a meaning of their own. Even though using them we can, for example, state that a participant agrees with one argument more than with another, we cannot state whether he agrees with it in the first place. Let us consider a framework containing a single argument and no attacks. Giving this argument a preference 0.3 has no effect, as it will always be accepted in any extension. Giving it the epistemic probability of 0.3, which is interpreted as disbelief, will lead it to it being always rejected. Consequently, despite certain similarities on the structural level, the semantics of the preference and weight-based frameworks are quite different from epistemic probabilities.

8. Conclusions

Our work, through empirically verifying certain prevailing assumptions in argumentation, is about informing the design of formalisms that can model how participants think about arguments arising in a dialogue. We observe that handling the uncertainty concerning both the participants’ opinions about arguments and the structure of the graph is of critical importance, and that a Dung framework equipped only with the standard semantics is insufficient to represent the views of the participants. Our results can be summarized in the following manner:

Observation 1 The data supports the use of the **constellation approach** to probabilistic argumentation for modelling the argument graphs representing the views of dialogue participants. In particular, in Section 6.1 we have seen that people may interpret statements and relations between them differently and without adhering to the intended relations. Furthermore, their personal knowledge can affect their perception and evaluation of the dialogue. Thus, the constellation approach can represent our uncertainty about the argument graphs describing our opponents views.

Observation 2 People may explicitly declare that two given statements are connected, however, they might not be sure of the exact nature of the relation between them. In Section 6.1 we have observed that the portion of the graphs declared by the participants that were not clarified (i.e. contained at least one relation marked as dependent) reached even 70%. We therefore also need to express the uncertainty that a person has about his or her own views. Although for the purpose of the analysis we have introduced the notion of a tripolar framework in order to be represent such situations, this issue can potentially be addressed with the constellation approach to argumentation. By this we understand that associated with a given unclarified framework we can have a number of graphs which interpret the dependent links as attacking or supporting, each with a given probability that could, for example, be obtained from the graph distribution of all participants.

Observation 3 The data supports the use of the **epistemic approach** to probabilistic argumentation. In Section 6.2 we calculated how many values the participants were using throughout the dialogues, and we could have observed that in most of the cases three values were insufficient to represent the participants’ opinions. Another of our important observations concerns the adherence to the epistemic postulates, analyzed in the aforementioned section. While classical semantics tend to represent a

number of properties at the same time, a single postulate tends to focus on a single aspect, as seen in Section 4.1. They therefore allow a more detailed view on the participant behaviour and can allow us to analyze the cases in which classic semantics may fail to explain it. The epistemic postulates can also provide more feedback to argumentation systems, such as for computational persuasion, than normal semantics do. For example, the low performance of the argumentation semantics such as complete does not really inform the system what aspect of the participant reasoning does not meet the semantics requirements. In contrast, the high performance of the preferential postulates and low performance of the explanatory ones, can inform the system that it has insufficient information about the participant’s personal knowledge and that it should proceed with querying him or her for further data. This is particularly important as our data also shows that people use their own personal knowledge in order to make judgments and might not necessarily disclose it.

The extended epistemic approach [27] would also allow us to model situations where the way a given relation is perceived is not necessarily tightly related to the way how much we believe or disbelieve its source, which as seen in Section 6.3, does tend to occur. It is likely that there are more properties of an argument or statement, such as how detailed and informative it is, which affect how the relations carried out by it are seen by people. While our analysis concerned the impact of the source of a relation on the relation and vice versa, it is possible that both source and target can affect the relation and thus approaches from [60, 61] could be verified. Thus, these issues need further investigation.

Observation 4 The data supports the use of **bipolar argumentation frameworks**. In Section 6.1 we have observed that the participants explicitly view certain relations as supporting and that the notion of defence does not account for all of the positive relations that the participants have identified between the presented statements. In particular, we could have observed that there are new support relations arising in the context of the dialogue, such as support coming from statements working towards the same goal.

Observation 5 The data supports the use of **bipolar argumentation in combination with the prudent/careful approaches**. In Section 6.1 we have observed that many additional attacks perceived by the participants can be explained by the existing notions of indirect conflicts in these settings. They can therefore be used to model auxiliary conflicts arising in the context of a dialogue, but not necessarily created on the logical level.

Observation 6 The data shows that **people use their own personal knowledge** in order to make judgments and might not necessarily disclose it. In Section 6.2 we have shown how the differences between the declared and expanded graphs become visible on the postulates highly sensitive to personal knowledge, such as the ones belonging to the explanatory family. Additionally, in Section 6.4 we have noted that a number of participants were able to evaluate the presented statements despite the lack of awareness of certain information. This, combined with them not providing any explanations for their opinions, may mean that argumentation systems need to handle participants not willing to share their knowledge.

Observation 7 The data shows that presenting a new and correct piece of information that a given person was not aware of **does not necessarily lead to changing that person’s beliefs**. In Section 6.4 we have observed that throughout the dialogues, not many participants have changed their opinions in a significant manner. Moreover, quite often those who have, have done so under the influence of incorrect information and presenting them with the correct data has not managed to rectify their opinions.

Our exploratory study shows that the most common approaches to argumentation might be too simplistic in order to adequately grasp human reasoning. However, we do not believe that the argumentation theory as a field is insufficient altogether. In particular, we have highlighted the correspondence of the obtained results to various, less common formalisms, such as probabilistic and bipolar frameworks, and prudent and careful approaches. Consequently, these methods could be merged in the future in order to develop

abstract argumentation tools that can be used in dialogical argumentation with more success. Nevertheless, further and more specialized studies concerning our observations should be carried out. In particular, our experiment could be seen as dynamic, as it concerned two dialogues between different parties. It would be interesting to observe whether our findings could be replicated in a more static setting, where arguments are presented randomly and not in the context of a dialogue. For example, this shift could affect the perception of the indirect attacks, defenses and supports between different statements. We believe further studies will provide more insight into this matter.

References

- [1] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artificial Intelligence* 77 (1995) 321–357.
- [2] G. Brewka, S. Polberg, S. Woltran, Generalizations of Dung frameworks and their role in formal argumentation, *IEEE Intelligent Systems* 29 (1) (2014) 30–38.
- [3] H. Nguyen, J. Masthoff, Designing persuasive dialogue systems: Using argumentation with care, in: H. Oinas-Kukkonen, P. Hasle, M. Harjumaa, K. Segerstahl, P. Øhrstrøm (Eds.), *Proceedings of the 3rd International Conference on Persuasive Technology (PERSUASIVE 2008)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 201–212.
- [4] A. Tversky, D. Kahneman, *The Framing of Decisions and the Psychology of Choice*, Vol. 4 of NATO ASI Series, Springer Berlin Heidelberg, 1985, pp. 107–129.
- [5] A. J. Rothman, P. Salovey, Shaping perceptions to motivate healthy behavior: the role of message framing., *Psychological Bulletin* 121 (1) (1997) 3–19.
- [6] F. Cerutti, N. Tintarev, N. Oren, Formal arguments, preferences, and natural language interfaces to humans: an empirical evaluation, in: T. Schaub, G. Friedrich, B. O’Sullivan (Eds.), *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*, Vol. 263 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2014, pp. 1033–1034.
- [7] J. M. Olson, Psychological barriers to behavior change: how to identify the barriers that inhibit change, *Canadian Family Physician* 38 (1992) 309–319.
- [8] J. Ogden, *Health Psychology: A Textbook*, 5th Edition, Open University Press, 2012.
- [9] I. Rahwan, M. Madakkatel, J. Bonnefon, R. Awan, S. Abdallah, Behavioural experiments for assessing the abstract argumentation semantics of reinstatement, *Cognitive Science* 34 (8) (2010) 1483–1502.
- [10] A. Rosenfeld, S. Kraus, Providing arguments in discussions on the basis of the prediction of human argumentative behavior, *ACM Transactions on Interactive Intelligent Systems* 6 (4) (2016) 30:1–30:33.
- [11] M. Caminada, D. M. Gabbay, A logical account of formal argumentation, *Studia Logica* 93 (2) (2009) 109–145.
- [12] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, *The Knowledge Engineering Review* 26 (4) (2011) 365–410.
- [13] S. Coste-Marquis, C. Devred, P. Marquis, Inference from controversial arguments, in: G. Sutcliffe, A. Voronkov (Eds.), *Proceedings of the 12th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2005)*, Vol. 3835 of *Lecture Notes in Computer Science*, Springer-Verlag Berlin Heidelberg, 2005, pp. 606–620.
- [14] S. Coste-Marquis, C. Devred, P. Marquis, Prudent semantics for argumentation frameworks, in: A. Lim (Ed.), *Proceedings of the 17th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2005)*, IEEE Computer Society, Washington, DC, USA, 2005, pp. 568–572.
- [15] C. Cayrol, M.-C. Lagasque-Schiex, Bipolar abstract argumentation systems, in: I. Rahwan, G. R. Simari (Eds.), *Argumentation in Artificial Intelligence*, Springer, 2009, pp. 65–84.
- [16] C. Cayrol, M.-C. Lagasque-Schiex, Coalitions of arguments: A tool for handling bipolar argumentation frameworks, *International Journal of Intelligent Systems* 25 (1) (2010) 83–109.
- [17] C. Cayrol, M.-C. Lagasque-Schiex, Bipolarity in argumentation graphs: Towards a better understanding, *International Journal of Approximate Reasoning* 54 (7) (2013) 876–899.
- [18] G. Boella, D. Gabbay, L. van der Torre, S. Villata, Support in abstract argumentation, in: P. Baroni, F. Cerutti, M. Giacomin, G. R. Simari (Eds.), *Proceedings of the 3rd International Conference on Computational Models of Argument (COMMA 2010)*, Vol. 216 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, Amsterdam, The Netherlands, The Netherlands, 2010, pp. 111–122.
- [19] F. Nouioua, V. Risch, Argumentation frameworks with necessities, in: S. Benferhat, J. Grant (Eds.), *Proceedings of the 5th International Conference on Scalable Uncertainty Management (SUM 2011)*, Vol. 6929 of *Lecture Notes in Artificial Intelligence*, Springer-Verlag Berlin Heidelberg, 2011, pp. 163–176.
- [20] F. Nouioua, AFs with necessities: Further semantics and labelling characterization, in: W. Liu, V. Subrahmanian, J. Wijsen (Eds.), *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, Vol. 8078 of *Lecture Notes in Computer Science*, Springer-Verlag Berlin Heidelberg, 2013, pp. 120–133.
- [21] N. Oren, T. J. Norman, Semantics for evidence-based argumentation, in: P. Besnard, S. Doutre, A. Hunter (Eds.), *Proceedings of the 2nd International Conference on Computational Models of Argument (COMMA 2008)*, Vol. 172 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2008, pp. 276–284.
- [22] N. Oren, C. Reed, M. Luck, Moving between argumentation frameworks, in: P. Baroni, F. Cerutti, M. Giacomin, G. R. Simari (Eds.), *Proceedings of the 3rd International Conference on Computational Models of Argument (COMMA 2010)*, Vol. 216 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, Amsterdam, The Netherlands, The Netherlands, 2010, pp. 379–390.
- [23] S. Polberg, N. Oren, Revisiting support in abstract argumentation systems, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014)*, Vol. 266 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2014, pp. 369–376.
- [24] S. Polberg, Intertranslatability of abstract argumentation frameworks, *Tech. Rep. DBAI-TR-2017-104*, Institute for Information Systems, Technical University of Vienna (2017).
- [25] C. Cayrol, M.-C. Lagasque-Schiex, An axiomatic approach to support in argumentation, in: E. Black, S. Modgil, N. Oren (Eds.), *Proceedings of the 3rd International Workshop on Theory and Applications of Formal Argumentation (TAFA 2015)*, Vol. 9524 of *Lecture Notes in Computer Science*, Springer International Publishing, 2015, pp. 74–91.
- [26] A. Hunter, A probabilistic approach to modelling uncertain logical arguments, *International Journal of Approximate Reasoning* 54 (1) (2013) 47–81.

- [27] S. Polberg, A. Hunter, M. Thimm, Belief in attacks in epistemic probabilistic argumentation, in: Proceedings of the 11th International Conference on Scalable Uncertainty Management (SUM 2017), Lecture Notes in Artificial Intelligence, Springer-Verlag Berlin Heidelberg, 2017, forthcoming.
- [28] M. Thimm, A probabilistic semantics for abstract argumentation, in: L. D. Raedt, C. Bessiere, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, P. Lucas (Eds.), Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012), Vol. 242 of Frontiers in Artificial Intelligence, IOS Press, 2012, pp. 750–755.
- [29] A. Hunter, M. Thimm, Probabilistic argument graphs for argumentation lotteries, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014), Vol. 266 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2014, pp. 313–324.
- [30] P. Baroni, M. Giacomin, P. Vicig, On rationality conditions for epistemic probabilities in abstract argumentation, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014), Vol. 266 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2014, pp. 121–132.
- [31] A. Hunter, M. Thimm, Probabilistic reasoning with abstract argumentation frameworks, *Journal of Artificial Intelligence Research* forthcoming.
- [32] A. Hunter, M. Thimm, Probabilistic argumentation with epistemic extensions and incomplete information, Tech. rep., ArXiv (May 2014).
- [33] A. Hunter, M. Thimm, On partial information and contradictions in probabilistic abstract argumentation, in: C. Baral, J. Delgrande, F. Wolter (Eds.), Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR 2016), AAAI Press, 2016, pp. 53–62.
- [34] A. Hunter, Some foundations for probabilistic abstract argumentation, in: B. Verheij, S. Szeider, S. Woltran (Eds.), Proceedings of the 4th International Conference on Computational Models of Argument (COMMA 2012), Vol. 245 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2012, pp. 117–128.
- [35] P. M. Dung, P. M. Thang, Towards (probabilistic) argumentation for jury-based dispute resolution, in: P. Baroni, F. Cerutti, M. Giacomin, G. R. Simari (Eds.), Proceedings of the 3rd International Conference on Computational Models of Argument (COMMA 2010), Vol. 216 of Frontiers in Artificial Intelligence and Applications, IOS Press, Amsterdam, The Netherlands, The Netherlands, 2010, pp. 171–182.
- [36] H. Li, N. Oren, T. J. Norman, Probabilistic argumentation frameworks, in: S. Modgil, N. Oren, F. Toni (Eds.), Proceedings of the 1st International Workshop on the Theory and Applications of Formal Argumentation (TFAFA 2011), Vol. 7132 of Lecture Notes in Computer Science, Springer-Verlag Berlin Heidelberg, 2011, pp. 1–16.
- [37] P. Dondio, Multi-valued and probabilistic argumentation frameworks, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014), Vol. 266 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2014, pp. 253–260.
- [38] P. Dondio, Toward a computational analysis of probabilistic argumentation frameworks, *Cybernetics and Systems* 45 (3) (2014) 254–278.
- [39] B. Fazzinga, S. Flesca, F. Parisi, On the complexity of probabilistic abstract argumentation, in: F. Rossi (Ed.), Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013), AAAI Press, 2013, pp. 898–904.
- [40] B. Fazzinga, S. Flesca, F. Parisi, On the complexity of probabilistic abstract argumentation frameworks, *ACM Transactions on Computational Logic* 16 (3) (2015) 1–39.
- [41] H. Li, N. Oren, T. J. Norman, Relaxing independence assumptions in probabilistic argumentation, in: Proceedings of the 10th International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2013), 2013.
- [42] S. Polberg, D. Doder, Probabilistic abstract dialectical frameworks, in: E. Fermé, J. Leite (Eds.), Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA 2014), Vol. 8761 of Lecture Notes in Computer Science, Springer International Publishing, Cham, 2014, pp. 591–599.
- [43] D. Doder, S. Woltran, Probabilistic argumentation frameworks – a logical approach, in: U. Straccia, A. Calì (Eds.), Proceedings of the 8th International Conference on Scalable Uncertainty Management (SUM 2014), Vol. 8720 of Lecture Notes in Computer Science, Springer International Publishing, 2014, pp. 134–147.
- [44] T. Bosc, E. Cabrio, S. Villata, Tweets squabbling: Positive and negative results in applying argument mining on social media, in: P. Baroni, T. F. Gordon, T. Scheffler, M. Stede (Eds.), Proceedings of the 6th International Conference on Computational Models of Argument (COMMA 2016), Vol. 287 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2016, pp. 21–32.
- [45] E. Cabrio, S. Villata, Node: A benchmark of natural language arguments, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014), Vol. 266 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2014, pp. 449–450.
- [46] S. Polberg, A. Hunter, Data Appendix, http://www0.cs.ucl.ac.uk/staff/a.hunter/papers/dialogue_study_appendix.zip (2017).
- [47] S. Polberg, Understanding the Abstract Dialectical Framework, in: L. Michael, A. Kakas (Eds.), Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016), Vol. 10021 of Lecture Notes in Artificial Intelligence, Springer International Publishing, 2016, pp. 430–446.
- [48] F. Cerutti, A. Palmer, A. Rosenfeld, J. Snajder, F. Toni, A pilot study in using argumentation frameworks for online debates, in: Proceedings of the 1st International Workshop on Systems and Algorithms for Formal Argumentation (SAFA 2016), 2016, pp. 63–74.
- [49] S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasque-Schiex, P. Marquis, On the merging of Dung’s argumentation systems, *Artificial Intelligence* 171 (10) (2007) 730–753.
- [50] R. Booth, M. Caminada, M. Podlaskowski, I. Rahwan, Quantifying disagreement in argument-based reasoning, in:

- W. van der Hoek, L. Padgham, V. Conitzer, M. Winikoff (Eds.), Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012), IFAAMAS, 2012, pp. 493–500.
- [51] C. Cayrol, M.-C. Lagasquie-Schiex, Weighted argumentation systems: A tool for merging argumentation systems, in: T. M. Khoshgoftaar, X. Zhu (Eds.), Proceedings of the 23rd IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2011), IEEE Computer Society, 2011, pp. 629–632.
- [52] H. Prakken, On support relations in abstract argumentation as abstractions of inferential relations, in: T. Schaub, G. Friedrich, B. O’Sullivan (Eds.), Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014), Vol. 263 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2014, pp. 735–740.
- [53] S. Villata, G. Boella, L. Van Der Torre, Attack semantics for abstract argumentation, in: T. Walsh (Ed.), Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), AAAI Press, 2011, pp. 406–413.
- [54] P. Baroni, F. Cerutti, M. Giacomin, G. Guida, AFRA: Argumentation framework with recursive attacks, International Journal of Approximate Reasoning 52 (1) (2011) 19–37.
- [55] A. Cohen, S. Gottifredi, A. J. García, G. R. Simari, On the acceptability semantics of argumentation frameworks with recursive attack and support, in: P. Baroni, T. F. Gordon, T. Scheffler, M. Stede (Eds.), Proceedings of the 6th International Conference on Computational Models of Argument (COMMA 2016), Vol. 287 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2016, pp. 231–242.
- [56] L. Amgoud, C. Cayrol, M. C. Lagasquie-Schiex, P. Livet, On bipolarity in argumentation frameworks, International Journal of Intelligent Systems 23 (10) (2008) 1062–1093.
- [57] S. Lukin, P. Anand, M. Walker, S. Whittaker, Argument strength is in the eye of the beholder: Audience effects in persuasion, in: M. Lapata, P. Blunsom, A. Koller (Eds.), Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017): Volume 1, Long Papers, Association for Computational Linguistics, 2017, pp. 742–753.
- [58] S. Villata, E. Cabrio, S. Jraidi, Imène; Benlamine, M. Chaouachi, C. Frasson, F. Gandon, Emotions and personality traits in argumentation: An empirical evaluation, Argument & Computation 8 (1) (2017) 61–87.
- [59] E. Bonzon, J. Delobelle, S. Konieczny, N. Maudet, A comparative study of ranking-based semantics for abstract argumentation, in: D. Schuurmans, M. Wellman (Eds.), Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI 2016), AAAI Press, 2016, pp. 914–920.
- [60] C. Cayrol, M.-C. Lagasquie-Schiex, From preferences over arguments to preferences over attacks in abstract argumentation: A comparative study, in: N. Bourbakis (Ed.), Proceedings of the 25th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2013), IEEE Computer Society, 2013, pp. 588–595.
- [61] S. Kaci, C. Labreuche, Arguing with valued preference relations, in: W. Liu (Ed.), Proceedings of the 11th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2011), Vol. 6717 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2011, pp. 62–73.

9. Statistics Appendix

9.1. Dialogue 1

9.1.1. Postulate Satisfaction

Total Sample	Intended vs Declared		Intended vs Expanded		Intended vs Common		Declared vs Expanded		Declared vs Common		Expanded vs Common	
	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value
Preferential	1.428	0.149	1.472	0.143	1.999	0.125	1.000	1.000	-0.221	0.819	-0.544	0.601
Rational	2.417	0.019	2.426	0.018	1.732	0.250	1.000	1.000	-2.017	0.061	-2.199	0.036
Strict	1.609	0.115	1.795	0.079	1.732	0.250	1.000	1.000	-1.079	0.306	-1.364	0.203
Protective	1.102	0.289	1.102	0.289	1.732	0.250	NA	NA	-0.608	0.577	-0.608	0.577
Restrained	0.880	0.411	0.880	0.411	1.414	0.500	NA	NA	-0.608	0.577	-0.608	0.577
Coherent	-1.783	0.075	-0.252	0.804	NA	NA	2.447	0.031	1.783	0.075	0.252	0.804
Involuntary	-1.344	0.183	-0.461	0.650	NA	NA	2.235	0.063	1.344	0.183	0.461	0.650
Justifiable	2.917	0.004	2.917	0.004	NA	NA	NA	NA	-2.917	0.004	-2.917	0.004
Semi Optimistic	1.799	0.064	1.420	0.144	-0.296	1.000	-1.414	0.500	-1.462	0.146	-0.990	0.331
Semi Founded	4.107	0.000	1.847	0.066	-1.732	0.250	-3.304	0.001	-4.682	0.000	-2.225	0.025
Founded	3.353	0.001	0.956	0.351	-2.449	0.031	-3.153	0.002	-4.412	0.000	-1.870	0.062
Optimistic	3.213	0.001	1.150	0.260	-2.000	0.125	-2.823	0.008	-4.084	0.000	-1.735	0.083
Demanding	-1.083	0.354	-2.383	0.016	1.999	0.125	-2.447	0.031	2.299	0.026	3.025	0.002
Guarded	4.220	0.000	2.109	0.034	-1.414	0.500	-3.153	0.002	-4.682	0.000	-2.356	0.017
Discharging	4.220	0.000	2.133	0.032	-1.414	0.500	-3.153	0.002	-4.682	0.000	-2.363	0.017
Trusting	4.224	0.000	2.203	0.026	-1.732	0.250	-3.153	0.002	-4.772	0.000	-2.559	0.009
Anticipating	4.224	0.000	2.203	0.026	-1.732	0.250	-3.153	0.002	-4.772	0.000	-2.559	0.009
Core Sample	Intended vs Declared		Intended vs Expanded		Intended vs Common		Declared vs Expanded		Declared vs Common		Expanded vs Common	
	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value
Preferential	1.732	0.250	1.730	0.250	1.000	1.000	1.000	1.000	-0.500	1.000	-1.414	0.500
Rational	1.730	0.250	1.730	0.250	1.000	1.000	1.000	1.000	-1.414	0.500	-1.732	0.250
Strict	1.414	0.500	1.732	0.250	1.000	1.000	1.000	1.000	-0.577	1.000	-1.414	0.500
Protective	0.318	1.000	0.318	1.000	1.000	1.000	NA	NA	0.138	0.750	0.138	0.750
Restrained	-0.138	0.750	-0.138	0.750	NA	NA	NA	NA	0.138	0.750	0.138	0.750
Coherent	-0.138	0.750	1.208	0.375	NA	NA	1.413	0.500	0.138	0.750	-1.208	0.375
Involuntary	0.615	0.750	1.033	0.500	NA	NA	1.000	1.000	-0.615	0.750	-1.033	0.500
Justifiable	1.730	0.250	1.730	0.250	NA	NA	NA	NA	-1.730	0.250	-1.730	0.250
SemiOptimistic	0.615	0.750	0.034	1.000	-1.000	1.000	-1.000	1.000	-1.033	0.500	-0.034	1.000
SemiFounded	0.841	0.500	-1.587	0.124	-1.413	0.500	-2.617	0.016	-1.995	0.125	1.313	0.227
Founded	0.382	1.000	-2.169	0.031	-1.732	0.250	-2.617	0.016	-1.996	0.125	1.671	0.094
Optimistic	0.382	1.000	-1.932	0.055	-1.732	0.250	-2.431	0.031	-1.996	0.125	1.361	0.191
Demanding	1.000	1.000	-0.046	1.000	NA	NA	-1.000	1.000	-1.000	1.000	0.046	1.000
Guarded	0.841	0.500	-1.390	0.176	-1.413	0.500	-2.431	0.031	-1.995	0.125	1.089	0.313
Discharging	0.841	0.500	-1.389	0.176	-1.413	0.500	-2.431	0.031	-1.995	0.125	1.115	0.266
Trusting	0.841	0.500	-1.442	0.152	-1.730	0.250	-2.431	0.031	-2.226	0.063	0.855	0.426
Anticipating	0.841	0.500	-1.442	0.152	-1.730	0.250	-2.431	0.031	-2.226	0.063	0.855	0.426

Table 10: The results of the Wilcoxon signed rank test with Pratt adjustment for a given postulate evaluated on two separate graphs on the total and core samples in Dialogue 1. The results have been obtained using R library coin. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. In case the adherence rate for a given postulate was identical for two graphs for all participants, the algorithm has returned NA.

9.1.2. Statements vs. Relations: Effect of a Relation on its Source

Total Sample	-			++			+			?		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
--	153.455	7	0.000	5.076	7	0.651	89.014	7	0.000	94.698	7	0.000
-				145.156	7	0.000	15.290	7	0.032	52.819	7	0.000
++							82.477	7	0.000	84.490	7	0.000
+										37.405	7	0.000

Core Sample	-			++			+			?		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
--	85.098	7	0.000	10.231	7	0.176	77.735	7	0.000	69.700	7	0.000
-				38.538	7	0.000	6.837	7	0.446	26.354	7	0.000
++							45.756	7	0.000	30.842	7	0.000
+										35.980	7	0.000

Table 11: Results of G-test for independence between relations of given type on total and core samples. We use the following abbreviations for the relation types: A good reason against (--), A somewhat good reason against (-), A good reason for (++) , A somewhat good reason for (+), Somewhat related, but can't say how (?). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

	Total Sample					Core Sample				
	--	-	++	+	?	--	-	++	+	?
χ^2	401.219	66.134	393.783	42.745	40.784	340.297	44.036	108.651	57.959	17.267
DF	7	7	7	7	7	7	7	7	7	7
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.016

Table 12: Results of chi-squared goodness of fit test on levels of agreement of sources of a relation of a given type in Dialogue 1. We use the following abbreviations for the relation types: A good reason against (--), A somewhat good reason against (-), A good reason for (++) , A somewhat good reason for (+), Somewhat related, but can't say how (?). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

Relations pooled according to strength												
Agreement pooled according to strength	Total Sample						Core Sample					
	Normal Relation			Dependency			Normal Relation			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Strong Relation	145.715	3	0.000	89.373	3	0.000	66.202	3	0.000	61.424	3	0.000
Normal Relation				47.221	3	0.000				19.351	3	0.000

Relations pooled according to polarity												
Agreement pooled according to polarity	Total Sample						Core Sample					
	Normal Relation			Dependency			Normal Relation			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Strong Relation	113.155	2	0.000	103.253	2	0.000	86.140	2	0.000	53.169	2	0.000
Normal Relation				25.713	2	0.000				31.549	2	0.000

Relations pooled according to polarity												
Agreement pooled according to polarity	Total Sample						Core Sample					
	Support			Dependency			Support			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Attack	0.673	2	0.714	48.684	2	0.000	4.906	2	0.086	35.441	2	0.000
Support				54.445	2	0.000				29.528	2	0.000

Relations pooled according to strength												
Agreement pooled according to strength	Total Sample						Core Sample					
	Support			Dependency			Support			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Attack	7.553	3	0.056	48.315	3	0.000	2.302	3	0.512	39.198	3	0.000
Support				56.617	3	0.000				29.107	3	0.000

Table 13: Results of G-test for independence between relations of given type in Dialogue 1 on total and core samples according to a given pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

Relations pooled according to strength						
Agreement pooled according to strength	Total Sample			Core Sample		
	Strong Relation	Normal Relation	Dependency	Strong Relation	Normal Relation	Dependency
χ^2	472.913	78.459	30.627	297.046	46.399	9.853
DF	3	3	3	3	3	3
p-value	0.000	0.000	0.000	0.000	0.000	0.020

Relations pooled according to polarity						
Agreement pooled according to polarity	Total Sample			Core Sample		
	Strong Relation	Normal Relation	Dependency	Strong Relation	Normal Relation	Dependency
χ^2	519.513	136.832	11.412	185.490	97.475	0.960
DF	2	2	2	2	2	2
p-value	0.000	0.000	0.003	0.000	0.000	0.619

Relations pooled according to polarity						
Agreement pooled according to polarity	Total Sample			Core Sample		
	Attack	Support	Dependency	Attack	Support	Dependency
χ^2	272.556	251.086	11.412	120.544	64.511	0.960
DF	2	2	2	2	2	2
p-value	0.000	0.000	0.003	0.000	0.000	0.619

Relations pooled according to strength						
Agreement pooled according to strength	Total Sample			Core Sample		
	Attack	Support	Dependency	Attack	Support	Dependency
χ^2	184.929	229.567	30.627	185.019	79.000	9.853
DF	3	3	3	3	3	3
p-value	0.000	0.000	0.000	0.000	0.000	0.020

Table 14: Results of chi-squared goodness of fit test on levels of agreement of sources of a relation of a given type in Dialogue 1 on total and core samples according to a given pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

9.1.3. Statements vs. Relations: Effect of an Argument on the Relations it Carries Out

Total sample	D		SoD		NAD		DK		SoA		A		SA	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
SD	34.873	0.000	35.761	0.000	44.077	0.000	17.742	0.001	33.778	0.000	11.239	0.024	23.965	0.000
D			3.452	0.485	19.668	0.001	13.085	0.011	8.647	0.071	61.659	0.000	118.932	0.000
SoD					25.352	0.000	19.199	0.001	6.513	0.164	63.386	0.000	105.564	0.000
NAD							4.571	0.334	38.749	0.000	59.589	0.000	97.628	0.000
DK									25.826	0.000	25.004	0.000	40.261	0.000
SoA											47.747	0.000	80.135	0.000
A													13.623	0.009

Core sample	D		SoD		NAD		DK		SoA		A		SA	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
SD	11.457	0.022	5.418	0.247	23.947	0.000	28.026	0.000	2.570	0.632	5.042	0.283	36.282	0.000
D			2.413	0.660	19.445	0.001	29.463	0.000	0.477	0.976	25.615	0.000	89.884	0.000
SoD					13.304	0.010	20.222	0.000	0.351	0.986	13.042	0.011	51.616	0.000
NAD							6.920	0.140	4.976	0.290	22.987	0.000	50.419	0.000
DK									6.329	0.176	19.897	0.001	50.310	0.000
SoA											4.931	0.294	14.571	0.006
A													20.806	0.000

Table 15: Results of G-test for independence between different argument acceptance levels in Dialogue 1 on total and core samples. In all cases we have obtained 4 degrees of freedom. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

	Total Sample								Core Sample							
	SD	D	SoD	NAD	DK	SoA	A	SA	SD	D	SoD	NAD	DK	SoA	A	SA
χ^2	11.314	32.225	33.057	35.972	14.727	35.980	57.000	171.296	12.239	23.340	7.296	13.909	25.600	2.000	18.706	190.611
DF	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
p-value	0.023	0.000	0.000	0.000	0.005	0.000	0.000	0.000	0.016	0.000	0.121	0.008	0.000	0.736	0.001	0.000

Table 16: Results of chi-squared goodness of fit test on relations carried out by arguments of a given acceptance level in Dialogue 1. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

Agreement pooled according to strength												
Relations pooled according to strength	Total Sample						Core Sample					
	Moderate Belief		Weak Belief		Neither		Moderate Belief		Weak Belief		Neither	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
Strong Belief	36.110	0.000	122.613	0.000	100.750	0.000	40.240	0.000	33.854	0.000	65.784	0.000
Moderate Belief			44.053	0.000	52.513	0.000			2.165	0.339	29.473	0.000
Weak Belief					49.153	0.000					20.024	0.000
Relations pooled according to polarity	Total Sample						Core Sample					
	Moderate Belief		Weak Belief		Neither		Moderate Belief		Weak Belief		Neither	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
Strong Belief	7.965	0.019	5.437	0.066	56.972	0.000	4.124	0.127	2.060	0.357	41.822	0.000
Moderate Belief			4.981	0.083	37.544	0.000			0.229	0.892	25.974	0.000
Weak Belief					49.160	0.000					19.549	0.000
Agreement pooled according to polarity												
Relations pooled according to strength	Total Sample				Core Sample							
	Believed		Neither		Believed		Neither					
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value				
Disbelieved	98.123	0.000	40.289	0.000	74.909	0.000	37.864	0.000				
Believed			100.516	0.000			62.282	0.000				
Relations pooled according to polarity	Total Sample				Core Sample							
	Believed		Neither		Believed		Neither					
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value				
Disbelieved	5.811	0.055	37.922	0.000	5.042	0.080	36.045	0.000				
Believed			62.659	0.000			34.388	0.000				

Table 17: Results of G-test for independence between different argument acceptance levels in Dialogue 1 on total and core samples according to a given pooling. In all cases we have obtained 2 degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

Agreement pooled according to strength								
Relations pooled according to strength	Total Sample				Core Sample			
	Strong Belief	Moderate Belief	Weak Belief	Neither	Strong Belief	Moderate Belief	Weak Belief	Neither
χ^2	280.351	75.107	102.118	13.865	192.360	33.832	17.645	8.340
DF	2	2	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.015
Relations pooled according to polarity	Total Sample				Core Sample			
	Strong Belief	Moderate Belief	Weak Belief	Neither	Strong Belief	Moderate Belief	Weak Belief	Neither
χ^2	121.032	74.476	63.430	7.000	116.919	43.486	15.516	6.298
DF	2	2	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.030	0.000	0.000	0.000	0.043
Agreement pooled according to polarity								
Relations pooled according to strength	Total Sample			Core Sample				
	Disbelieved	Believed	Neither	Disbelieved	Believed	Neither		
χ^2	84.105	333.680	13.865	65.475	197.077	8.340		
DF	2	2	2	2	2	2		
p-value	0.000	0.000	0.001	0.000	0.000	0.015		
Relations pooled according to polarity	Total Sample			Core Sample				
	Disbelieved	Believed	Neither	Disbelieved	Believed	Neither		
χ^2	69.584	180.390	7.000	69.475	107.868	6.298		
DF	2	2	2	2	2	2		
p-value	0.000	0.000	0.030	0.000	0.000	0.043		

Table 18: Results of chi-squared goodness of fit test on relations carried out by arguments of a given acceptance level in Dialogue 1 on total and core samples according to a given pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

9.2. Dialogue 2

9.2.1. Postulate Satisfaction

Total Sample	Intended vs Declared		Intended vs Expanded		Intended vs Common		Declared vs Expanded		Declared vs Common		Expanded vs Common	
	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value
Preferential	3.231	0.001	3.256	0.001	2.995	0.004	1.414	0.500	-1.632	0.132	-1.849	0.080
Rational	3.021	0.002	3.035	0.002	2.642	0.016	1.000	1.000	-2.029	0.050	-2.044	0.048
Strict	1.957	0.050	2.550	0.010	3.585	0.000	1.999	0.125	-0.143	0.892	-0.941	0.351
Protective	2.410	0.015	2.485	0.012	4.295	0.000	1.000	1.000	0.378	0.712	0.307	0.766
Restrained	2.019	0.043	2.026	0.042	3.718	0.000	1.000	1.000	-0.131	0.905	-0.182	0.863
Coherent	2.019	0.043	3.794	0.000	4.190	0.000	3.845	0.000	0.472	0.640	-2.421	0.015
Involuntary	0.465	0.641	2.425	0.015	2.235	0.063	3.154	0.002	0.497	0.613	-1.914	0.063
Justifiable	2.236	0.063	2.000	0.125	NA	NA	-1.000	1.000	-2.236	0.063	-2.000	0.125
SemiOptimistic	-3.196	0.001	-3.834	0.000	-4.498	0.000	-2.643	0.016	-1.062	0.308	0.416	0.757
SemiFounded	4.291	0.000	0.389	0.704	NA	NA	-3.963	0.000	-4.291	0.000	-0.389	0.704
Founded	2.825	0.008	-0.971	0.315	NA	NA	-3.156	0.002	-2.825	0.008	-0.971	0.315
Optimistic	2.827	0.008	0.284	0.918	NA	NA	-2.644	0.016	-2.827	0.008	-0.284	0.918
Demanding	1.363	0.250	-0.508	0.684	1.999	0.125	-1.867	0.094	-0.255	1.000	1.275	0.173
Guarded	4.072	0.000	0.194	0.850	-1.732	0.250	-3.842	0.000	-4.298	0.000	-0.631	0.535
Discharging	3.365	0.000	-0.157	0.880	-2.824	0.008	-3.841	0.000	-4.500	0.000	-0.888	0.381
Trusting	4.208	0.000	0.916	0.363	-1.414	0.500	-3.588	0.000	-4.330	0.000	-1.259	0.213
Anticipating	3.636	0.000	0.346	0.735	-1.999	0.125	-3.305	0.001	-4.104	0.000	-0.986	0.331

Core Sample	Intended vs Declared		Intended vs Expanded		Intended vs Common		Declared vs Expanded		Declared vs Common		Expanded vs Common	
	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value	Z	p-value
Preferential	2.226	0.063	2.226	0.063	1.994	0.125	NA	NA	-1.000	1.000	-1.000	1.000
Rational	2.224	0.063	2.224	0.063	1.729	0.250	NA	NA	-1.413	0.500	-1.413	0.500
Strict	2.025	0.063	2.025	0.063	1.993	0.125	NA	NA	-0.961	0.625	-0.961	0.625
Protective	2.124	0.035	2.124	0.035	2.924	0.004	NA	NA	0.435	0.703	0.435	0.703
Restrained	1.983	0.051	1.983	0.051	2.229	0.063	NA	NA	-0.411	0.813	-0.411	0.813
Coherent	2.215	0.027	2.647	0.006	2.777	0.008	2.229	0.063	-0.494	1.000	-1.679	0.125
Involuntary	0.961	0.625	1.594	0.219	1.414	0.500	1.414	0.500	0.049	1.000	-0.961	0.625
Justifiable	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
SemiOptimistic	-2.647	0.008	-3.056	0.002	-3.172	0.001	-1.994	0.125	-0.960	0.500	0.535	1.000
SemiFounded	2.224	0.063	-1.061	0.359	NA	NA	-2.433	0.031	-2.224	0.063	1.061	0.359
Founded	1.000	1.000	-0.659	0.500	NA	NA	-1.413	0.500	-1.000	1.000	0.659	0.500
Optimistic	NA	NA	-1.000	1.000	NA	NA	-1.000	1.000	NA	NA	1.000	1.000
Demanding	1.729	0.250	0.618	0.750	1.730	0.250	-1.000	1.000	0.049	1.000	0.659	0.500
Guarded	1.993	0.125	-1.393	0.188	NA	NA	-2.429	0.031	-1.993	0.125	1.393	0.188
Discharging	0.745	0.594	-1.687	0.102	-1.413	0.500	-2.427	0.031	-1.994	0.125	1.359	0.191
Trusting	1.730	0.250	-0.130	0.906	NA	NA	-1.729	0.250	-1.730	0.250	0.130	0.906
Anticipating	0.659	0.500	-0.809	0.531	-1.000	1.000	-1.729	0.250	-1.413	0.500	0.513	0.688

Table 19: The results of the Wilcoxon signed rank test with Pratt adjustment for a given postulate evaluated on two separate graphs on the total and core samples in Dialogue 2. The results have been obtained using R library coin. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. In case the adherence rate for a given postulate was identical for two graphs for all participants, the algorithm has returned NA.

9.2.2. Statements vs. Relations: Effect of a Relation on its Source

Total Sample	-			++			+			?		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
--	146.771	7	0.000	44.558	7	0.000	112.838	7	0.000	122.577	7	0.000
-				70.229	7	0.000	49.829	7	0.000	125.574	7	0.000
++							82.165	7	0.000	110.601	7	0.000
+										90.918	7	0.000

Core Sample	-			++			+			?		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
--	86.269	7	0.000	7.506	7	0.378	83.920	7	0.000	123.81499	7	0.000
-				70.549	7	0.000	18.330	7	0.011	89.734	7	0.000
++							74.209	7	0.000	110.2263	7	0.000
+										42.580	7	0.000

Table 20: Results of G-test for independence between relations of given type in Dialogue 2 on total and core samples. We use the following abbreviations for the relation types: A good reason against (--), A somewhat good reason against (-), A good reason for (++), A somewhat good reason for (+), Somewhat related, but can't say how (?). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

	Total Sample					Core Sample				
	--	-	++	+	?	--	-	++	+	?
χ^2	220.015	183.075	146.260	87.918	378.960	134.442	59.454	79.750	46.842	396.177
DF	7	7	7	7	7	7	7	7	7	7
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 21: Results of chi-squared goodness of fit test on levels of agreement of sources of a relation of a given type in Dialogue 2. We use the following abbreviations for the relation types: A good reason against (--), A somewhat good reason against (-), A good reason for (++), A somewhat good reason for (+), Somewhat related, but can't say how (?). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

Relations pooled according to strength												
Agreement pooled according to strength	Total Sample						Core Sample					
	Normal Relation			Dependency			Normal Relation			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Strong Relation	99.442	3	0.000	120.905	3	0.000	98.797	3	0.000	127.352	3	0.000
Normal Relation				127.914	3	0.000				54.266	3	0.000

Relations pooled according to polarity												
Agreement pooled according to polarity	Total Sample						Core Sample					
	Normal Relation			Dependency			Normal Relation			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Strong Relation	38.622	2	0.000	108.870	2	0.000	14.090	2	0.001	120.824	2	0.000
Normal Relation				127.904	2	0.000				55.808	2	0.000

Relations pooled according to polarity												
Agreement pooled according to polarity	Total Sample						Core Sample					
	Support			Dependency			Support			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Attack	3.595	2	0.166	122.975	2	0.000	9.883	2	0.007	107.440	2	0.000
Support				104.779	2	0.000				78.490	2	0.000

Relations pooled according to strength												
Agreement pooled according to strength	Total Sample						Core Sample					
	Support			Dependency			Support			Dependency		
	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value	G-value	DF	p-value
Attack	9.338	3	0.025	128.599	3	0.000	0.965	3	0.810	104.209	3	0.000
Support				106.480	3	0.000				82.158	3	0.000

Table 22: Results of G-test for independence between relations of given type in Dialogue 2 on total and core samples according to a given pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

Relations pooled according to strength						
Agreement pooled according to strength	Total Sample			Core Sample		
	Strong Relation	Normal Relation	Dependency	Strong Relation	Normal Relation	Dependency
χ^2	118.474	145.419	278.048	38.630	73.749	188.726
DF	3	3	3	3	3	3
p-value	0.000	0.000	0.000	0.000	0.000	0.000

Relations pooled according to polarity						
Agreement pooled according to polarity	Total Sample			Core Sample		
	Strong Relation	Normal Relation	Dependency	Strong Relation	Normal Relation	Dependency
χ^2	108.641	10.771	141.128	61.437	7.732	106.235
DF	2	2	2	2	2	2
p-value	0.000	0.005	0.000	0.000	0.021	0.000

Relations pooled according to polarity						
Agreement pooled according to polarity	Total Sample			Core Sample		
	Attack	Support	Dependency	Attack	Support	Dependency
χ^2	63.392	24.454	141.128	59.744	4.805	106.235
DF	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.000	0.090	0.000

Relations pooled according to strength						
Agreement pooled according to strength	Total Sample			Core Sample		
	Attack	Support	Dependency	Attack	Support	Dependency
χ^2	112.872	83.463	278.048	22.326	20.270	188.726
DF	3	3	3	3	3	3
p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table 23: Results of chi-squared goodness of fit test on levels of agreement of sources of a relation of a given type in Dialogue 2 on total and core samples according to a given type of pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

9.2.3. Statements vs. Relations: Effect of an Argument on the Relations it Carries Out

Total sample	D		SoD		NAD		DK		SoA		A		SA	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
SD	88.802	0.000	41.240	0.000	57.479	0.000	39.781	0.000	39.381	0.000	18.469	0.001	15.769	0.003
D			46.449	0.000	116.328	0.000	72.405	0.000	30.161	0.000	72.721	0.000	129.046	0.000
SoD					40.291	0.000	45.535	0.000	5.529	0.237	46.450	0.000	60.854	0.000
NAD							10.459	0.033	50.145	0.000	77.095	0.000	74.098	0.000
DK									41.020	0.000	41.392	0.000	55.609	0.000
SoA											37.557	0.000	68.731	0.000
A													21.771	0.000
Core sample	D		SoD		NAD		DK		SoA		A		SA	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
SD	48.190	0.000	37.146	0.000	64.903	0.000	83.274	0.000	48.583	0.000	19.065	0.001	6.361	0.174
D			4.049	0.399	26.022	0.000	10.662	0.031	6.652	0.155	22.830	0.000	61.635	0.000
SoD					21.170	0.000	11.723	0.020	11.731	0.019	17.529	0.002	46.062	0.000
NAD							43.599	0.000	56.900	0.000	72.986	0.000	87.561	0.000
DK									18.957	0.001	53.371	0.000	94.664	0.000
SoA											17.817	0.001	55.203	0.000
A													20.983	0.000

Table 24: Results of G-test for independence between different argument acceptance levels in Dialogue 2 on total and core samples. In all cases we have obtained 4 degrees of freedom. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

	Total Sample								Core Sample							
	SD	D	SoD	NAD	DK	SoA	A	SA	SD	D	SoD	NAD	DK	SoA	A	SA
χ^2	87.128	117.032	22.202	36.956	37.706	26.618	137.961	156.983	90.706	5.596	7.067	59.016	16.133	32.727	98.417	140.519
DF	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.231	0.132	0.000	0.003	0.000	0.000	0.000

Table 25: Results of chi-squared goodness of fit test on relations carried out by arguments of a given acceptance level in Dialogue 2. We use the following abbreviations: Strongly Agree (SA), Agree (A), Somewhat Agree (SoA), Neither Agree nor Disagree (NAD), Somewhat Disagree (SoD), Disagree (D), Strongly Disagree (SD), Don't Know (DK). DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

Agreement pooled according to strength												
Relations pooled according to strength	Total Sample						Core Sample					
	Moderate Belief		Weak Belief		Neither		Moderate Belief		Weak Belief		Neither	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
Strong Belief	64.174	0.000	89.900	0.000	101.892	0.000	55.723	0.000	88.655	0.000	146.970	0.000
Moderate Belief			7.532	0.023	97.014	0.000			10.005	0.007	59.026	0.000
Weak Belief					76.925	0.000					40.821	0.000

Agreement pooled according to polarity												
Relations pooled according to polarity	Total Sample						Core Sample					
	Moderate Belief		Weak Belief		Neither		Moderate Belief		Weak Belief		Neither	
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value
Strong Belief	1.847	0.397	9.971	0.007	78.619	0.000	6.068	0.048	5.081	0.079	72.506	0.000
Moderate Belief			7.261	0.026	88.552	0.000			0.000	1.000	55.628	0.000
Weak Belief					49.026	0.000					39.409	0.000

Agreement pooled according to polarity									
Relations pooled according to strength	Total Sample				Core Sample				
	Believed		Neither		Believed		Neither		
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	
Disbelieved	34.173	0.000	100.464	0.000	8.874	0.012	43.919	0.000	
Believed			104.099	0.000			114.249	0.000	

Agreement pooled according to polarity									
Relations pooled according to polarity	Total Sample				Core Sample				
	Believed		Neither		Believed		Neither		
	G-value	p-value	G-value	p-value	G-value	p-value	G-value	p-value	
Disbelieved	2.630	0.268	74.999	0.000	15.702	0.000	40.888	0.000	
Believed			106.419	0.000			104.415	0.000	

Table 26: Results of G-test for independence between different argument acceptance levels in Dialogue 2 on total and core samples according to a given type of pooling. In all cases we have obtained 2 degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using library Deducer (likelihood.test function) in R.

Agreement pooled according to strength								
Relations pooled according to strength	Total Sample				Core Sample			
	Strong Belief	Moderate Belief	Weak Belief	Neither	Strong Belief	Moderate Belief	Weak Belief	Neither
χ^2	348.713	273.129	134.568	116.882	286.052	136.630	55.665	11.626
DF	2	2	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.003

Agreement pooled according to polarity								
Relations pooled according to polarity	Total Sample				Core Sample			
	Strong Belief	Moderate Belief	Weak Belief	Neither	Strong Belief	Moderate Belief	Weak Belief	Neither
χ^2	168.530	271.362	134.827	50.852	94.518	107.849	65.200	8.311
DF	2	2	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.016

Agreement pooled according to polarity						
Relations pooled according to strength	Total Sample			Core Sample		
	Disbelieved	Believed	Neither	Disbelieved	Believed	Neither
χ^2	213.405	463.158	116.882	105.506	281.212	11.626
DF	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.000	0.000	0.003

Agreement pooled according to polarity						
Relations pooled according to polarity	Total Sample			Core Sample		
	Disbelieved	Believed	Neither	Disbelieved	Believed	Neither
χ^2	213.633	351.965	50.852	68.486	212.382	8.311
DF	2	2	2	2	2	2
p-value	0.000	0.000	0.000	0.000	0.000	0.016

Table 27: Results of chi-squared goodness of fit test on relations carried out by arguments of a given acceptance level in Dialogue 2 on total and core samples according to a given pooling. DF stands for degrees of freedom. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. These results were obtained using R.

9.3. Dialogues 1& 2: Changes in Beliefs

		Argument	A		B		C		D		E		F		G		
		Between steps	Z	pvalue	Z	pvalue	Z	pvalue	Z	pvalue	Z	pvalue	Z	pvalue	Z	pvalue	
Dialogue 1	Total	1 and 2	-0.551	0.716	-5.997	0.000											
		2 and 3	1.890	0.125	5.846	0.000	1.783	0.085									
		3 and 4	-1.986	0.063	-5.826	0.000	0.780	0.544	-5.351	0.000							
		4 and 5	2.157	0.039	5.773	0.000	1.964	0.051	-1.134	0.453	2.875	0.005					
		1 and 2	0.069	1.000	-3.726	0.000											
	Core	2 and 3	1.732	0.250	3.792	0.000	0.000	1.000									
		3 and 4	-0.577	1.000	-3.873	0.000	2.000	0.125	-0.069	1.000							
		4 and 5	1.000	1.000	3.873	0.000	0.046	1.000	-3.553	0.000	1.513	0.188					
		1 and 2	-0.378	1.000	-4.126	0.000											
		2 and 3	-0.794	0.656	-1.102	0.406	5.449	0.000	0.622	0.627							
Dialogue 2	Total	3 and 4	-2.448	0.031	0.740	0.688	-4.781	0.000	-0.839	0.440	1.635	0.156	0.349	0.699			
		4 and 5	-0.013	1.000	0.764	0.445	1.297	0.227	1.154	0.359	-0.952	0.434	-1.393	0.227	0.421	0.591	
		1 and 2	0.000	1.000	0.577	1.000											
		2 and 3	1.413	0.500	-3.011	0.003	3.316	0.000	1.413	0.500							
		3 and 4	-1.995	0.125	3.372	0.000	-1.717	0.125	-1.466	0.234	-2.653	0.003	-2.232	0.033			
	Core	4 and 5	0.494	1.000	1.000	1.000	-0.647	0.564	1.996	0.125	-0.960	0.500	-0.576	0.750	1.308	0.219	

Table 28: The results of the Wilcoxon signed rank test with Pratt adjustment between the acceptance levels declared by the participants at steps i and $i + 1$ for a given argument, in a given dialogue and sample. The results have been obtained using R library coin. We have highlighted the fields with p-value greater than 0.05, i.e. those for which we cannot reject the null hypothesis. Due to the fact that arguments F and H appear only at the last stages of Dialogues 1 and 2 respectively, they possess only a single distribution and the change analysis cannot be performed for them.