
Argumentation Semantics as Formal Discussion

MARTIN CAMINADA

ABSTRACT. In the current chapter, we interpret a number of mainstream argumentation semantics by means of structured discussion. The idea is that an argument is justified according to a particular argumentation semantics iff it is possible to win a discussion of a particular type. Hence, different argumentation semantics correspond to different types of discussion. Our aim is to provide an overview of what these discussions look like, and their formal correspondence to argumentation semantics.

1 Introduction

The term “argumentation”, when used in an informal way, calls upon intuitions of arguments being exchanged in some kind of interactive discussion. Yet, the notion of discussion plays a relatively limited role in abstract argumentation theory, which mainly focuses on various principles (called “argumentation semantics”) for selecting nodes from a graph. As such, there seems to be quite a gap between (abstract) argumentation theory as described in much of the literature,¹ as it occurs in everyday life.²

In order to address this gap, attempts have been made to express argumentation semantics in terms of structured discussion. More precisely, the idea is that an argument is accepted w.r.t. a particular argumentation semantics iff it is possible to successfully defend the argument using a particular kind of discussion. In the current chapter we provide an overview of what the different kinds of discussion are, and how they formally relate to their associated argumentation semantics.

Although the discussion protocols (which we will often refer to as “discussion games”) can serve as proof procedures of their associated argumentation semantics, their potential application is much wider than that. One could for instance use the discussion games for the purpose of human computer interaction. Suppose a knowledge-based system has determined that a particular argument (say, about how to treat a patient) should be accepted, and communicates this to its user (say, a doctor). When the user asks why this is the

¹As for instance described in chapters 4 (Abstract argumentation frameworks and their semantics), 5 (Abstract dialectical frameworks), 6 (Abstract rule-based argumentation) 9 (A review of argumentation based on deductive arguments) and 19 (Argumentation, nonmonotonic reasoning and argumentation and logic) of this volume.

²As for instance described in Chapter 12 (Processing natural language argumentation) of this volume.

case, what should probably be avoided is a highly technical answer of the form “because the argument is in the minimal fixpoint of monotonic function F ”.³ Instead, one would like the user to critically question the answer,⁴ and be able to utter counter arguments to see whether these are properly addressed (by the system providing counter counter arguments). As an example of such a human-computer discussion, consider the following dialogue:

System: The patient is best off with medicine X, because this is the most effective.

User: But the patient is diabetic, for which medicine X could have side effects.

System: Recent studies have shown that these side effects are relatively minor. So instead of the system immediately providing the full justification for its answer (say, by providing the entire grounded extension) it engages in a discussion with its user. Ideally, such a discussion should be “natural” in the sense that the human-computer interaction looks as much as possible as human-human interaction (say, if the doctor were to discuss the case with a more senior colleague).

Apart from being natural, the discussion should also be sound and complete. That is, the ability to win the discussion for a particular argument (that is, to have a winning strategy for the argument in the discussion game) should coincide with the argument being justified according to a pre-defined argumentation semantics. Soundness and completeness imply that if the system provides an answer (“argument A is (or is not) justified according to a particular argumentation semantics”) the system can successfully defend itself in the discussion with the user. When this discussion is also perceived as natural by the user, this will hopefully increase the user’s confidence in the system’s answer.

Soundness and completeness also imply that what we are looking for are essentially proof procedures for particular argumentation semantics. Several of these have been stated in the literature. Inclusion in the current book chapter is done based on two criteria:

- (1) does the discussion game have any link with natural discussion concepts, like described in philosophy or linguistics?
- (2) is the discussion game such that it guarantees the absence of any exponential blowups, in either time or space?

Criterion (1) is the reason why for instance we have not included any discussion games for sceptical preferred semantics (like those of Doutre and Mengin [2004] and Dung and Thang [2007]). Criterion (2) is the reason why we did not include a detailed treatment of tree-based discussion games (like those of Prakken and Sartor [1997], Caminada [2004], Modgil and Caminada [2009] and Dung *et al.*

³which basically says the argument is in the grounded extension.

⁴See for instance Chapter 11 (Argumentation schemes) of this volume.

[2007]).⁵

The remaining part of this chapter is structured as follows. First, in Section 2 we briefly recall some basic definitions and results from abstract argumentation theory. Then, in Section 3 we describe a discussion game for (credulous) preferred semantics [Caminada *et al.*, 2014], and explain that it contains aspects of Socratic discussion. Then, in Section 4 we briefly state how this discussion game can be reapplied in the context of ideal semantics [Caminada *et al.*, 2014]. In Section 5 we subsequently describe a discussion game for stable semantics [Caminada and Wu, 2009], basically by making minor modifications to the earlier described discussion game for (credulous) preferred semantics. In Section 6 we then describe a different discussion game in the context of grounded semantics [Caminada, 2015a] and explain its relationship with persuasion dialogue. Then, in Section 7 we briefly examine tree-based discussion games and explain one of their main disadvantages: the possibility of an exponential blowup in time or space. We round off with a discussion in Section 8.

2 Formal Preliminaries

In the current section, we briefly recall some basic definitions from abstract argumentation theory.⁶ For current purposes, we restrict ourselves to finite argumentation frameworks.

Definition 1 (argumentation framework) *An argumentation framework is a pair (Ar, att) where Ar is a finite set of entities called arguments and att is a binary relation on Ar .*

Given an argumentation framework (Ar, att) , $A, A' \in Ar$ and $Args, Args' \subseteq Ar$, we say that (1) A attacks A' iff $(A, A') \in att$, (2) A attacks $Args$ iff A attacks some argument in $Args$, (3) $Args$ attacks A iff some argument in $Args$ attacks A , and (4) $Args$ attacks $Args'$ iff some argument in $Args$ attacks some argument in $Args'$.

Definition 2 (preliminaries, extension-based) *Let (Ar, att) be an argumentation framework. A set $Args \subseteq Ar$ is conflict-free iff $Args$ does not attack itself. A set $Args \subseteq Ar$ defends $A \in Ar$ iff for each $B \in Ar$ that attacks A , it holds that $Args$ attacks B .*

Definition 3 (admissibility, extension-based) *Let (Ar, att) be an argumentation framework. A set $Args \subseteq Ar$ is admissible iff $Args$ is conflict-free and each $A \in Args$ is defended by $Args$.*

Definition 4 (strong admissibility, extension-based) *Let (Ar, att) be an argumentation framework. A set $Args \subseteq Ar$ is strongly admissible iff each*

⁵How tree-based discussion games can lead to an exponential blowup is explained in Section 7.

⁶We refer Chapter 4 (Abstract argumentation frameworks and their semantics) of this volume for a more thorough discussion.




Figure 1. An argumentation framework to illustrate strong admissibility.

$A \in \text{Args}$ is defended by some $\text{Args}' \subseteq \text{Args} \setminus \{A\}$ which in its turn is again strongly admissible.

It has been proved that each strongly admissible set is conflict-free as well as admissible [Baroni and Giacomin, 2007; Caminada, 2014].

As an example, consider the argumentation framework of Figure 1. Here, the set $\{A, C\}$ is strongly admissible as A is defended by $\emptyset \subseteq \{A, C\} \setminus \{A\}$ which is trivially strongly admissible, and C is defended by $\{A\} \subseteq \{A, C\} \setminus \{C\}$ which is strongly admissible (as A is defended by $\emptyset \subseteq \{A\} \setminus \{A\}$). The set $\{G\}$, however, is admissible but not strongly admissible as G is not defended by any subset of $\{G\} \setminus \{G\}$.

Definition 5 (completeness, extension-based) Let (Ar, att) be an argumentation framework. A set $\text{Args} \subseteq Ar$ is a complete extension iff Args is conflict-free and the set of arguments defended by Args is equal to Args .

Definition 6 (semantics, extension-based) Let (Ar, att) be an argumentation framework. A set $\text{Args} \subseteq Ar$ is called

1. a grounded extension iff Args is the minimal (w.r.t. \subseteq) complete extension
2. a preferred extension iff Args is a maximal (w.r.t. \subseteq) complete extension
3. a stable extension iff Args is a complete extension that attacks each argument in $Ar \setminus \text{Args}$
4. an ideal extension iff Args is the maximal (w.r.t. \subseteq) complete extension that is not attacked by any complete extension

We recall that each argumentation framework has precisely one grounded extension, precisely one ideal extension, one or more preferred extensions and zero or more stable extensions.

The above definition describes grounded, preferred, stable and ideal semantics uniformly in terms of complete semantics. However, for our purposes it is sometimes useful to describe these semantics in terms of (strong) admissibility.

Theorem 1 (semantics, extension-based) *Let (Ar, att) be an argumentation framework. A set $Args \subseteq Ar$ is*

1. *a preferred extension iff $Args$ is a maximal (w.r.t. \subseteq) admissible set*
2. *a grounded extension iff $Args$ is the maximal (w.r.t. \subseteq) strongly admissible set*
3. *a stable extension iff $Args$ is an admissible set that attacks each argument in $Ar \setminus Args$*
4. *an ideal extension iff $Args$ is the maximal (w.r.t. \subseteq) admissible set that is not attacked by any admissible set*

Apart from the extension-based view on argumentation semantics, there is also the labelling-based view [Caminada, 2006; Caminada and Gabbay, 2009; Caminada, 2011; Baroni *et al.*, 2011] of which we now provide a brief overview.

Definition 7 (preliminaries, labelling-based) *Let (Ar, att) be an argumentation framework. An argument labelling is a function $\mathcal{L}ab : Ar \rightarrow \{\text{in}, \text{out}, \text{undec}\}$. We define $\text{in}(\mathcal{L}ab)$ as $\{A \in Ar \mid \mathcal{L}ab(A) = \text{in}\}$, $\text{out}(\mathcal{L}ab)$ as $\{A \in Ar \mid \mathcal{L}ab(A) = \text{out}\}$ and $\text{undec}(\mathcal{L}ab)$ as $\{A \in Ar \mid \mathcal{L}ab(A) = \text{undec}\}$. We sometimes write a labelling as a triple $(\text{in}(\mathcal{L}ab), \text{out}(\mathcal{L}ab), \text{undec}(\mathcal{L}ab))$. If $\mathcal{L}ab_1$ and $\mathcal{L}ab_2$ are labellings, we write $\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$ when $\text{in}(\mathcal{L}ab_1) \subseteq \text{in}(\mathcal{L}ab_2)$ and $\text{out}(\mathcal{L}ab_1) \subseteq \text{out}(\mathcal{L}ab_2)$. Moreover, we write $\mathcal{L}ab_1 \approx \mathcal{L}ab_2$ when $\text{in}(\mathcal{L}ab_1) \cap \text{out}(\mathcal{L}ab_2) = \emptyset$ and $\text{out}(\mathcal{L}ab_1) \cap \text{in}(\mathcal{L}ab_2) = \emptyset$.*

Definition 8 (admissibility, labelling-based) *Let $\mathcal{L}ab$ be a labelling of argumentation framework (Ar, att) . $\mathcal{L}ab$ is called an admissible labelling iff for each $A \in Ar$ it holds that*

1. *if $\mathcal{L}ab(A) = \text{in}$ then for each $B \in Ar$ that attacks A it holds that $\mathcal{L}ab(B) = \text{out}$*
2. *if $\mathcal{L}ab(A) = \text{out}$ then there exists a $B \in Ar$ that attacks A such that $\mathcal{L}ab(B) = \text{in}$*

In order to define strong admissibility in the context of argument labellings, we first need to introduce the concept of a min-max numbering.

Definition 9 (min-max numbering) *Given an admissible labelling $\mathcal{L}ab$ of argumentation framework (Ar, att) , a min-max numbering is a function $\mathcal{M}\mathcal{M}_{\mathcal{L}ab} : \text{in}(\mathcal{L}ab) \cup \text{out}(\mathcal{L}ab) \rightarrow \mathbb{N} \cup \{\infty\}$ such that for each $A \in \text{in}(\mathcal{L}ab) \cup \text{out}(\mathcal{L}ab)$ it holds that*

- *if $\mathcal{L}ab(A) = \text{in}$ then $\mathcal{M}\mathcal{M}_{\mathcal{L}ab}(A) = \max(\{\mathcal{M}\mathcal{M}_{\mathcal{L}ab}(B) \mid B \text{ attacks } A \text{ and } \mathcal{L}ab(B) = \text{out}\}) + 1$ (with $\max(\emptyset)$ defined as 0)*

- if $\mathcal{L}ab(A) = \text{out}$ then $\mathcal{M}\mathcal{M}_{\mathcal{L}ab}(A) = \min(\{\mathcal{M}\mathcal{M}_{\mathcal{L}ab}(B) \mid B \text{ attacks } A \text{ and } \mathcal{L}ab(B) = \text{in}\}) + 1$ (with $\min(\emptyset)$ defined as ∞)

It can be proved that each admissible labelling has a unique min-max numbering [Caminada, 2014].⁷

Definition 10 (strong admissibility, labelling-based) *Let $\mathcal{L}ab$ be a labelling of argumentation framework (Ar, att) . $\mathcal{L}ab$ is called a strongly admissible labelling iff it is an admissible labelling whose associated min-max numbering yields natural numbers only (so no argument is numbered ∞).*

From Definition 10 it trivially follows that each strongly admissible labelling is also an admissible labelling.

As an example, consider the argumentation framework shown in Figure 1. Here $\mathcal{L}ab_1 = (\{A, C, E, G\}, \{B, D, H\}, \{F\})$ is an admissible labelling with associated min-max numbering $\mathcal{M}\mathcal{M}_{\mathcal{L}ab_1} = \{(A: 1), (B: 2), (C: 3), (D: 4), (E: 5), (G: \infty), (H: \infty)\}$, which implies that $\mathcal{L}ab_1$ is not strongly admissible. Furthermore, $\mathcal{L}ab_2 = (\{A, C, E\}, \{B, D, F\}, \{G, H\})$ is an admissible labelling with associated min-max numbering $\mathcal{M}\mathcal{M}_{\mathcal{L}ab_2} = \{(A: 1), (B: 2), (C: 3), (D: 4), (E: 5), (F: 2)\}$, which implies that $\mathcal{L}ab_2$ is indeed a strongly admissible labelling.

Definition 11 (completeness, labelling-based) *Let $\mathcal{L}ab$ be a labelling of argumentation framework (Ar, att) . $\mathcal{L}ab$ is called a complete labelling iff for each $A \in Ar$ it holds that*

1. if $\mathcal{L}ab(A) = \text{in}$ then for each $B \in Ar$ that attacks A it holds that $\mathcal{L}ab(B) = \text{out}$
2. if $\mathcal{L}ab(A) = \text{out}$ then there exists a $B \in Ar$ that attacks A such that $\mathcal{L}ab(B) = \text{in}$
3. if $\mathcal{L}ab(A) = \text{undec}$ then not for each $B \in Ar$ that attacks A it holds that $\mathcal{L}ab(B) = \text{out}$ and there does not exist a $B \in Ar$ that attacks A such that $\mathcal{L}ab(B) = \text{in}$

Definition 12 (semantics, labelling-based) *Let (Ar, att) be an argumentation framework. A labelling $\mathcal{L}ab$ is called*

1. a grounded labelling iff it is the minimal (w.r.t. \sqsubseteq) complete labelling
2. a preferred labelling iff it is a maximal (w.r.t. \sqsubseteq) complete labelling

⁷The min-max numbering can be constructed in an iterative way, starting from the unnumbered in-labelled arguments without attackers (these are numbered 1), then the unnumbered out-labelled arguments that are attacked by these (these are numbered 2), etc. When a particular iteration provides no new argument numbers, the remaining unnumbered in and out-labelled arguments are numbered ∞ . See the work of Caminada [2014] for details.

3. a stable labelling iff it is a complete labelling with $\text{undec}(\mathcal{L}ab) = \emptyset$
4. an ideal labelling iff it is the maximal (w.r.t. \sqsubseteq) complete labelling that is compatible (\approx) with every complete labelling

We recall that each argumentation framework has precisely one grounded labelling, precisely one ideal labelling, one or more preferred labellings and zero or more stable labellings.

The above definition describes grounded, preferred, stable and ideal semantics in terms of complete labellings. However, it is sometimes useful to be able to describe these semantics in terms of (strong) admissibility, similar to what was done earlier for the extension-based semantics.

Theorem 2 (semantics, labelling-based) *Let (Ar, att) be an argumentation framework. A labelling $\mathcal{L}ab$ is called*

1. a preferred labelling iff it is a maximal (w.r.t. \sqsubseteq) admissible labelling
2. a grounded labelling iff it is the maximal (w.r.t. \sqsubseteq) strongly admissible labelling
3. a stable labelling iff it is an admissible labelling with $\text{undec}(\mathcal{L}ab) = \emptyset$
4. an ideal labelling iff it is the maximal (w.r.t. \sqsubseteq) admissible labelling that is compatible (\approx) with every admissible labelling

To be able to easily switch between the labelling-based approach and the extension-based approach, we introduce two functions Lab2Ext and Ext2Lab , such that for an admissible labelling $\mathcal{L}ab$, $\text{Lab2Ext}(\mathcal{L}ab)$ is defined as $\text{in}(\mathcal{L}ab)$, and for an admissible set $\mathcal{A}rgs$, $\text{Ext2Lab}(\mathcal{A}rgs)$ is defined as $(\mathcal{A}rgs, \{A \in Ar \mid A \text{ is attacked by } \mathcal{A}rgs\}, \{A \in Ar \mid A \notin \mathcal{A}rgs \text{ and } A \text{ is not attacked by } \mathcal{A}rgs\})$ where Ar is the set of all arguments in the argumentation framework. It holds that if $\mathcal{L}ab$ is a (strongly) admissible labelling (resp. a complete, grounded, preferred, stable or ideal labelling) then $\text{Lab2Ext}(\mathcal{L}ab)$ is a (strongly) admissible set (resp. a complete, grounded, preferred, stable or ideal extension). It also holds that if $\mathcal{A}rgs$ is a (strongly) admissible set (resp. a complete, grounded, preferred, stable or ideal extension) then $\text{Ext2Lab}(\mathcal{A}rgs)$ is a (strongly) admissible labelling (resp. complete, grounded, preferred, stable or ideal labelling). Moreover, when restricted to complete (or resp. grounded, preferred, stable or ideal) extensions and labellings, the functions Lab2Ext and Ext2Lab become bijections that are each other's inverses [Caminada, 2006; Caminada and Gabbay, 2009].

The above results imply that:

- in order to determine whether an argument is in a preferred extension, it suffices to determine whether the argument is labelled in by an admissible labelling

- in order to determine whether an argument is in the grounded extension, it suffices to determine whether the argument is labelled **in** by a strongly admissible labelling
- in order to determine whether an argument is in a stable extension, it suffices to determine whether the argument is labelled **in** by an admissible labelling without **undec**
- in order to determine whether an argument is in the ideal extension, it suffices to determine whether the argument is labelled **in** by an admissible labelling that is compatible with every admissible labelling

In the sections that follow, we will apply the above observations to provide discussion games for preferred, grounded, stable and ideal semantics.

3 Preferred Semantics

In the current section, we describe the discussion game for preferred semantics as stated by Caminada *et al.* [2014].⁸ The idea of the preferred discussion game is to show membership of a preferred extension by constructing an admissible labelling where the argument in question is labelled **in**.

The preferred discussion game has two players which we will refer to as M and S. Player M starts; his task is to defend the fact that he has a reasonable position (admissible labelling) in which a particular argument is accepted (labelled **in**). Player S then tries to confront M with the consequences of M's own position, and asks for these consequences to be resolved. Player M is successful if he is able to address all the issues pointed out by player S, without being led to a contradiction.

As an example of how such a discussion can take place, consider the argumentation framework of Figure 2.




Figure 2. An argumentation framework

Here, the player M can win the discussion game for argument *D* in the following way.

⁸The discussion game of Caminada *et al.* [2014] consists of a labelling-based reinterpretation of the work of Vreeswijk and Prakken [2000].

Example 1

- M*: in(*D*)
 “I have an admissible labelling in which *D* is labelled in.”
- S*: out(*C*)
 “But then in your labelling it must also be the case that *D*’s attacker *C* is labelled out. Based on which grounds?”
- M*: in(*B*)
 “*C* is labelled out because *B* is labelled in.”
- S*: out(*A*)
 “But then in your labelling it must also be the case that *B*’s attacker *A* is labelled out. Based on which grounds?”
- M*: in(*B*)
 “*A* is labelled out because *B* is labelled in.”

As is shown in the above example, the discussion moves of player *M* are statements that particular arguments are labelled *in* in *M*’s labelling. The moves of player *S*, on the other hand, are meant to confront *M* with the consequences of his own position: “if you think that argument *X* is labelled *in* then you must also hold that *X*’s attacker *Y* is labelled *out* in your labelling.” That is, by uttering out(*Y*), player *S* points out that player *M* is implicitly committed to the fact that *Y* should be rejected. This means that player *M* has to explain *why* *Y* should be rejected. That is, the moves of player *S* can be seen as *questions* about why a particular argument *Y* should be labelled out. The moves of player *M* (except his first move) can then be interpreted as the *answers* to the questions of player *S*. Each answer follows directly to the question raised by player *S*. That is:

Each move of M (except the first) contains an attacker of the argument in the directly preceding move of S. (1)

Every time player *M* claims that an argument is labelled *in*, player *S* should be given the opportunity to state that as a consequence of this, player *M* is implicitly committed that *all* attackers of the argument are labelled *out*. The problem, however, is that each move of player *S* is a statement about just *one* argument. In order to deal with this problem, player *S* should be given the opportunity to react on the same *in*-labelled argument several times, each time confronting player *M* with a different *out*-labelled argument. This means that player *S* should be allowed to react not just on the immediately preceding move of player *M*, but on *any* previous move of player *M*.

Each move of player S contains an attacker of an argument contained in some (not necessarily the directly preceding) move of player M. (2)

Another issue is whether player *S* should be allowed to repeat his own moves. Recall that each move essentially contains a question (“Based on which grounds

is argument Y labelled **out**?”). At the moment player S repeats one of his moves, this question has already been answered by player M , so there is no good reason to ask again. In order to avoid the discussion from going round in circles, it does not make sense to allow player S to repeat his moves.

Player S is not allowed to repeat his moves. (3)

On the other hand, Example 1 does illustrate the need for player M to be able to repeat his moves (like **in**(B)). This is because some of the questions of S (like “why is argument C **out**” and “why is argument A **out**”) can have the same answer (“because argument B is **in**”).

Player M is allowed to repeat his moves. (4)

The argumentation framework of Figure 2 can also be used for an example of a game won by player S :

Example 2

M : **in**(E)
 “I have an admissible labelling in which E is labelled **in**.”
 S : **out**(D)
 “But then in your labelling it must be the case that E ’s attacker D is labelled **out**. Based on which grounds?”
 M : **in**(C)
 “ D is labelled **out** because C is labelled **in**.”
 S : **out**(E)
 “But then in your labelling it must be the case that C ’s attacker E is labelled **out**. This contradicts with your earlier claim that E is labelled **in**.”

The above example illustrates that when player S manages to use an argument uttered previously by player M , player S has won the game. After all, if player M claims an argument to be **in** and player S subsequently manages to confront player M with the fact that in M ’s own position, the same argument should be labelled **out**, then player S has successfully pointed out a contradiction in M ’s position.

If player S uses an argument previously used by player M , then player S wins the discussion game. (5)

One can ask a similar question regarding what happens when player M uses one of the arguments previously used by player S . The fact that player S performed an **out** move means that the argument must be labelled **out** in the labelling of player M . If player M then subsequently claims that the same argument is labelled **in**, then he has directly contradicted himself.

If player M uses an argument previously used by player S, then player S wins the discussion game. (6)

There also exists a third condition under which player S wins the game. This is when player M is unable to answer one of the questions of S. This can be the case when there exists no attacker against an argument uttered by player S. Hence, player S asks why a particular argument is labelled **out** but player M is unable to come up with any attacker to be labelled **in**. In that case, player M has lost the game, for not being able to answer the critical questions of player S.

If player M cannot make a move any more, player S wins the discussion game. (7)

Similarly, one might examine what happens when it is player S who cannot make a move any more. This essentially means that player S has run out of questions. All possible relevant questions have already been asked; all relevant issues have already been raised. Moreover, player M has managed to answer all questions in a satisfactory way. Therefore, player M has survived the process of critical questioning, hence winning the discussion.

If player S cannot make a move any more, player M wins the discussion game. (8)




Figure 3. An argumentation framework with floating attack

As a last illustration of how the discussion game functions, consider the argumentation framework of Figure 3. Argument *C* is not in any admissible set. It is illustrative to see what happens if player M tries to defend *C*.

Example 3

- M*: $\text{in}(C)$
 “I have an admissible labelling in which C is labelled **in**.”
- S*: $\text{out}(A)$
 “But then in your labelling C ’s attacker A must be labelled **out**. Based on which grounds?”
- M*: $\text{in}(B)$
 “ A is labelled **out** because B is labelled **in**.”
- S*: $\text{out}(B)$
 “But from the fact that you hold C to be **in**, it follows that C ’s attacker B must be labelled **out**. This contradicts with your earlier claim that B is labelled **in**.”

The above example illustrates the need for player S to be able to respond not only to the immediately preceding move, but to any past move of player M ; in the example, $\text{out}(B)$ is a response to $\text{in}(C)$. This is because, as we have mentioned before, for an argument to be labelled **in**, all its attackers have to be **out**, so player S may need to respond to a move of player M with more than one countermove.

When putting observations (1) to (8) together, we obtain the following description of the discussion game

Definition 13 *Let (Ar, att) be an argumentation framework. A preferred discussion is a sequence of moves $[\Delta_1, \Delta_2, \dots, \Delta_n]$ ($n \geq 0$) such that:*

- each move Δ_i ($1 \leq i \leq n$) where i is odd is called an *M-move* and is of the form $\text{in}(A)$, where $A \in Ar$
- each move Δ_i ($1 \leq i \leq n$) where i is even is called an *S-move* and is of the form $\text{out}(A)$, where $A \in Ar$
- for each *S-move* $\Delta_i = \text{out}(A)$ ($2 \leq i \leq n$) there exists an *M-move* $\Delta_j = \text{in}(B)$ ($j < i$) such that A attacks B
- for each *M-move* $\Delta_i = \text{in}(A)$ ($3 \leq i \leq n$) it holds that Δ_{i-1} is of the form $\text{out}(B)$, where A attacks B
- there exist no two *S-moves* Δ_i and Δ_j with $i \neq j$ and $\Delta_i = \Delta_j$

A preferred discussion $[\Delta_1, \Delta_2, \dots, \Delta_n]$ is said to be finished iff (1) there exists no Δ_{n+1} such that $[\Delta_1, \Delta_2, \dots, \Delta_n, \Delta_{n+1}]$ is a preferred discussion, or there exists an *M-move* and an *S-move* containing the same argument, and (2) no subsequence $[\Delta_1, \dots, \Delta_m]$ ($m < n$) is finished. A finished preferred discussion is won by player S if there exist an *M-move* and an *S-move* containing the same argument. Otherwise, it is won by the player making the last move (Δ_n).

The soundness and completeness of the game described above is stated in the following theorem.

Theorem 3 ([Caminada and Wu, 2009; Caminada *et al.*, 2014])

Let (Ar, att) be an argumentation framework and $A \in Ar$.

1. If there exists a preferred discussion for A that is won by player M , then there exists a preferred extension that contains A .
2. If there exists a preferred extension that contains A then player M has a winning strategy⁹ for the preferred discussion game.

The correctness of Theorem 3 can be seen as follows. As for point 1, it has to be observed that what the game essentially does is to build an admissible labelling of which the **in**-labelled arguments coincide with the M-moves and the **out**-labelled arguments coincide with the S-moves (all the other arguments are labelled **undec**). The resulting labelling is well-defined in the sense that no argument is labelled both **in** and **out** (otherwise there would be an argument that is subject to both an M-move and an S-move, in which case player S would have won the discussion). Moreover, the fact that player M wins the discussion also means that he made the last move, which implies that (i) each **out**-labelled argument has an **in**-labelled attacker. Also, the fact that player S cannot move anymore implies that (ii) each **in**-labelled argument has all its attackers labelled **out**. From (i) and (ii) it follows that the labelling yielded by the game is indeed an admissible one, satisfying the conditions of Definition 8. In this admissible labelling, argument A is labelled **in** (since A was the subject of the first M-move). This implies that A is element of an admissible set, and therefore also element of a preferred extension.

As for point 2, it should be mentioned that the fact that A is in a preferred extension by definition implies that A is in an admissible set ($Args$), which then implies that A is labelled **in** by an admissible labelling $Lab = \text{Ext2Lab}(Args)$. This makes it possible for player M to win the game simply by staying within the borders of admissible labelling Lab . That is, as long as player M only plays arguments that are labelled **in** by Lab , each move of player S has to be an argument that is labelled **out** by Lab , which then implies that player M can always react with an argument that is labelled **in** by Lab , etc. If player M follows such a strategy, there will never be an M-move and an S-move for the same argument (this is because Lab is a well-defined labelling, meaning that no argument is labelled both **in** and **out**). Moreover, the fact that player S cannot repeat himself means that the game has to finish in a finite number of moves. As player M can *always* react on a move of player S, this means that the last move has to be an M-move. Hence, player M wins the game.

From points 1 and 2 together, it follows that if there is at least one preferred discussion that is won by player M, then M has a winning strategy for the preferred discussion game. This is not the case in alternative discussion games

⁹Winning strategy in the sense of [Caminada *et al.*, 2014, Definition 5.6]. Informally this means that player M has a way of winning the discussion, regardless of what moves player S decides to play.

for preferred semantics, like the one described by Modgil and Caminada [2009]. In their approach, a single discussion game does not prove membership (for this, the presence of a winning strategy is really necessary). From informal perspective, this is rather odd, as in everyday life the aim of a (persuasion) discussion is to convince the other party in a single discussion. This means that at the end of the discussion, the other party has to have heard sufficient evidence to accept the main claim. This is the case in the above described preferred discussion game, but not in the alternative discussion game of Modgil and Caminada [2009].

As we have observed, an admissible labelling can serve as a “roadmap” for winning the preferred discussion game.¹⁰ However, an argument can be labelled **in** by more than one admissible labelling, which raises the question of which admissible labelling to choose as a basis to play the game. It can be verified that given an admissible labelling $\mathcal{L}ab$ (with $\mathcal{L}ab(A) = \mathbf{in}$ and $\mathbf{out}(\mathcal{L}ab)$ being minimal w.r.t. set inclusion) the number of moves required in the game for main argument A is $2 \cdot |\mathbf{out}(\mathcal{L}ab)| + 1$ (see [Caminada *et al.*, 2014] for details). Hence, in order to be able to finish the game in as few moves as possible (which could be desirable from the perspective of human-computer interaction if the aim of the game is to convince a human user) one should try to find an admissible labelling $\mathcal{L}ab$ where $|\mathbf{out}(\mathcal{L}ab)|$ is minimal. This is a computationally hard problem, as even verifying whether a particular admissible labelling has this property is coNP complete [Caminada *et al.*, 2014].

The essential nature of the preferred discussion game is that of critically questioning a particular position, and to see whether the proponent of this position (player M) can avoid being led to a contradiction (by player S). As such, the preferred discussion game bears a close resemblance to the concept of Socratic discussion, as well as to its modern variants like critical interviews or cross-examinations in court.¹¹ The general idea is to have somebody take a position and then iteratively confront him (through questioning) with what appears to be the consequences of this position, in the hope of ultimately leading him to a contradiction. We refer to the work of Caminada *et al.* [2014] for a details.

4 Ideal Semantics

An ideal set of arguments, as was originally defined by Dung *et al.* [2007], is an admissible set that is a subset of each preferred extension. It can be proved that the maximal ideal set (commonly known as the *ideal extension*) is unique and is a complete extension as well.

An alternative but equivalent way of characterising the ideal extension is as the maximal admissible set that is not attacked by any admissible set (like is done in Theorem 1) or as the maximal complete extension that is not attacked

¹⁰For details, we refer to the work of Caminada *et al.* [2014].

¹¹In fact, in the work of Caminada *et al.* [2014] player S stands for Socrates and player M stands for Menexenus, which is one of Socrates’s historic discussion partners.

by any complete extension (like is done in Definition 6). It can be proved that for each admissible sets $\mathcal{A}rgs_1$ and $\mathcal{A}rgs_2$ it holds that $\mathcal{A}rgs_1$ attacks $\mathcal{A}rgs_2$ iff $\mathcal{A}rgs_2$ attacks $\mathcal{A}rgs_1$. This gives rise to the labelling-based descriptions of ideal semantics of Theorem 2 and Definition 12.¹²

For current purposes, our characterisation of the ideal extension is as the maximal admissible set that is not attacked by any admissible set. To determine membership of the ideal extension, one then needs to find an admissible set (although not necessarily the maximal one) that contains the argument in question and is not attacked by any admissible set. This makes it possible to express ideal semantics using the preferred discussion game. Basically, the discussion whether an argument is in an ideal extension consists of two phases. In the first phase, one runs the preferred discussion game, as is described in the previous section. This is to determine whether the argument is in an admissible set. Then, in the second phase of the discussion, one needs to determine whether this set is attacked by another admissible set. This is done by again running the preferred discussion game for each of the arguments that were rejected (labelled **out**) during the first phase of the discussion, this time trying to defend (label **in**) the argument.

As an example, consider again the argumentation framework of Figure 2. Now consider the question of whether argument D is in an ideal set. The first phase of the discussion would be like Example 1 (page 495). Then, in the second phase of the discussion, one has to try to find an argument that was labelled out during the first phase¹³ (say A) and can be defended in a new preferred discussion game. Such a game would be as follows.

- M: **in**(A)
*“I have a reasonable position (admissible labelling) in which A is accepted (labelled **in**).”*
- S: **out**(B)
*“Then in your position, argument B must be rejected (labelled **out**). Based on which grounds?”*
- M: **in**(A)
*“ B is rejected (labelled **out**) because A is accepted (labelled **in**).”*

Hence, we have an admissible set $\{A\}$ that attacks the admissible set $\{B, D\}$ found during the first phase, so the admissible set $\{B, D\}$ of the first phase is not an ideal set.¹⁴

The overall procedure for ideal semantics puts an extra burden on the proponent of the argument. Not only does he have to win the preferred discussion game in the first phase, but he has to win it in such a way¹⁵ that the resulting

¹²Recall that each complete extension (labelling) is also an admissible set (labelling).

¹³Recall that the preferred game is such that the **out**-labelled arguments are the attackers of the **in**-labelled arguments (which is not necessarily the case for admissible labellings in general).

¹⁴In fact, for the argumentation framework of Figure 2, the only ideal set is the empty set.

¹⁵Since an argument can be element of more than one admissible set, there can be different

position (labelling) cannot be argued against in the second phase.

5 Stable Semantics

In the current section, we describe a discussion game for credulous stable semantics based on the work of Caminada and Wu [2009]. Before doing so, it may be illustrative to see why the preferred discussion game does not work for stable semantics. Consider again the argumentation framework of Figure 2. Even though A is in an admissible set and in a preferred extension ($\{A\}$), A is not in a stable extension. To see why A is in an admissible set, consider the following discussion:

M: $\text{in}(A)$ “I have an admissible labelling where A is labelled **in**.”

S: $\text{out}(B)$ “Then in your labelling, argument B must be labelled **out**. Based on which grounds?”

M: $\text{in}(A)$ “ B is labelled **out** because A is labelled **in**.”

The point is, however, that once it has been decided that A is labelled **in** and B is labelled **out**, it is not possible anymore to label the remaining arguments such that final result will be a stable labelling. This can be seen as follows. Suppose C is labelled **in**. Then E must be labelled **out**, so D should be labelled **in**, which means that C would be labelled **out**. Contradiction. Similarly, suppose that C is labelled **out**. Then E must be labelled **in**, so D should be labelled **out**, so C should be labelled **in**. Again, contradiction.

In general, there are many ways to characterize a stable extension [Caminada and Gabbay, 2009]. For our purposes, the most useful characterization is that of an admissible set which attacks every argument that is not in it (Theorem 1). When one translates this to labellings, one obtains an admissible labelling where each argument is labelled either **in** or **out** (that is, no argument is labelled **undec**, Theorem 2).

It appears that a discussion game for stable semantics requires an additional type of move: **question**. To illustrate the role of this new move, imagine a politician being interviewed for TV. At first the discussion may be about financial matters (say, whether the banking system should be nationalized). Then, the discussion may be about the consequences of the politician’s opinion (“If you accept to nationalize the banks, then you must reject the possibility to improve healthcare, because there will not be enough money left to do so.”). However, at some moment, the interviewer could choose to totally change topic (“By the way, what are your opinions about abortion?”). It is this change of topic that is enabled by the **question** move.¹⁶

For the discussion game for stable semantics, we use the **question** move to

ways to win the preferred discussion game.

¹⁶One of the reasons the **question** move is needed is because stable semantics does not satisfy the property of *directionality* [Baroni and Giacomin, 2007]. This means that for determining the status of an argument, not just the ancestors (the attackers, the attackers of these attackers, etc) are relevant but also the offspring (the attacked, the attacked of the attacked, etc) as well as arguments from unconnected parts of the graph. See also Chapter 16 (The principle-based approach to abstract argumentation semantics) of this volume.

involve those arguments that have never been uttered before so that we are able to label all the arguments in Ar . By questioning an argument ($\text{question}(A)$), player S (the opponent) asks player M (the proponent) to give an explicit opinion on whether A should be labelled **in** or **out**. If player M thinks that A should be labelled **in** then he should respond with $\text{in}(A)$. If, on the other hand, player M thinks that A should be labelled **out** then he should respond with $\text{in}(B)$ where B is a attacker of A . The discussion game for stable semantics can thus be described as follows:

- Player M (the proponent) and player S (the opponent) take turns. Player M starts.
- Each move of player S is either of the form $\text{out}(A)$, where A is a attacker of some (not necessarily the directly preceding) move of player M, or of the form $\text{question}(A)$, where A is an argument that has not been uttered in the discussion before (by either player M or player S).
- The first move of player M is of the form $\text{in}(A)$, where A is the main argument of the discussion. The following moves of player M are also of the form $\text{in}(A)$ although A no longer needs to be the main claim. If the directly preceding move of player S is of the form $\text{out}(B)$ then A is a attacker of B . If the directly preceding move of player S is of the form $\text{question}(B)$ then A is either equal to B or a attacker of B .
- Player S is not allowed to repeat any of his **out** moves.
- Player M is allowed to repeat his own **in** moves.

Player S wins if there is an argument A that has been subject to both an **in** move (by player M) and an **out** move (by player S). Otherwise, the discussion continues until one of the players cannot move anymore, in which case the discussion is won by the player making the last move.

To illustrate the use of the discussion game, consider the argumentation framework depicted in Figure 4.




Figure 4. Another argumentation framework

Suppose player M would like to start a discussion about A .

M: $\text{in}(A)$ "I have a stable labelling in which A is labelled **in**."

S: $\text{out}(B)$ "Then in your labelling, A 's attacker B must be labelled **out**. Based on which grounds?"

M: $\text{in}(A)$ " B is labelled **out** because A is labelled **in**."

S: **question**(C) “What about C ?”

M: **in**(C) “ C is labelled **in**.”

S: **out**(D) “Then C ’s attacker D must be labelled **out**. Based on which grounds?”

M: **in**(C) “ D is labelled **out** because C is labelled **in**.”

Player M wins the discussion, since player S cannot move anymore.

The above example also shows that the outcome of a discussion may depend on player M’s response to a question move. For instance, if player M would have replied to **question**(C) with **in**(D), then he would have lost the discussion, since player S would then move **out**(D).

As an example of a discussion that cannot be won by player M, consider the discussion for argument B . This discussion has to be lost by player M since the argumentation framework of Figure 4 has only one stable extension: $\{A, C\}$, which does not include B .

M: **in**(B) “I have a stable labelling in which B is labelled **in**.”

S: **out**(A) “Then in your labelling, B ’s attacker A must be labelled **out**. Based on which grounds?”

M: **in**(B) “ A is labelled **out** because B is labelled **in**.”

S: **question**(C) “What about C ?”

M: **in**(D) “ C is labelled **out** because its attacker D is labelled **in**.”

S: **out**(D) “Then D ’s attacker D (itself) must be labelled **out**. Contradiction.”

Player M would still have lost the discussion if he had responded to **question**(C) with **in**(C) instead of with **in**(D). This is because then player S would have reacted with **out**(B) and would therefore still have won the discussion.

Formally, the stable discussion game can be described as follows.

Definition 14 *Let (Ar, att) be an argumentation framework. A stable discussion is a sequence of moves $[\Delta_1, \Delta_2, \dots, \Delta_n]$ ($n \geq 0$) such that:*

- each Δ_i ($1 \leq i \leq n$) where i is odd (which is called an M-move) is of the form **in**(A), where $A \in Ar$.
- each Δ_i ($1 \leq i \leq n$) where i is even (which is called an S-move) is of the form **out**(A) where $A \in Ar$, or of the form **question**(A) where $A \in Ar$.
- For each S-move $\Delta_i = \mathbf{out}(A)$ ($2 \leq i \leq n$) there exists an M-move $\Delta_j = \mathbf{in}(B)$ ($j < i$) where A attacks B .
- For each M-move $\Delta_i = \mathbf{in}(A)$ ($3 \leq i \leq n$) it either holds that (1) $\Delta_{i-1} = \mathbf{out}(B)$ where A attacks B , or (2) $\Delta_{i-1} = \mathbf{question}(B)$ where either $A = B$ or A attacks B .
- For each S-move $\Delta_i = \mathbf{out}(A)$ ($1 \leq i \leq n$) there does not exist an S-move $\Delta_j = \mathbf{out}(A)$ with $j < i$.
- For each S-move $\Delta_i = \mathbf{question}(A)$ ($1 \leq i \leq n$) there does not exist any move Δ_j ($j < i$) of the form **in**(A), **out**(A) or **question**(A).

- For each M-move $\Delta_i = \text{in}(A)$ ($1 \leq i \leq n$) there does not exist an S-move $\Delta_j = \text{out}(A)$ with $j < i$.

A stable discussion $[\Delta_1, \Delta_2, \dots, \Delta_n]$ is said to be finished iff (1) there exists no Δ_{n+1} such that $[\Delta_1, \Delta_2, \dots, \Delta_n, \Delta_{n+1}]$ is a stable discussion, or there exists an M-move $\text{in}(A)$ and an S-move $\text{out}(A)$ for the same argument A , and (2) no subsequence $[\Delta_1, \dots, \Delta_m]$ ($m < n$) is finished. A finished stable discussion is won by player S if there exists an M-move $\text{in}(A)$ and an S-move $\text{out}(A)$ for the same argument A . Otherwise it is won by the payer making the last move Δ_n .

It turns out that an argument is in at least one stable extension iff the proponent can win the stable discussion game for it.

Theorem 4 Let (Ar, att) be an argumentation framework and $A \in Ar$.

1. If there exists a stable discussion for A that is won by player M, then A is in a stable extension.
2. If A is in a stable extension, then player M has a winning strategy for the stable discussion game.

As for point 1, it can be observed that what the discussion game essentially does is to build a stable labelling \mathcal{Lab} with $\text{in}(\mathcal{Lab}) = \{A \mid \text{there exists an M-move } \text{in}(A)\}$ and $\text{out}(\mathcal{Lab}) = \{A \mid \text{there exists an S-move } \text{out}(A)\} \cup \{A \mid \text{there exists an S-move } \text{question}(A) \text{ that was responded to with } \text{in}(B) \text{ where } B \text{ attacks } A\}$. It can be verified that \mathcal{Lab} is an admissible labelling without any argument being labelled **undec**. Hence, \mathcal{Lab} is a stable labelling in the sense of Theorem 2. As A is labelled **in** by \mathcal{Lab} (since A is the subject of the first M-move) it holds that A is in $\text{Lab2Ext}(\mathcal{Lab})$. Hence, A is in a stable extension.

As for point 2, it should be mentioned that player M can win the game simply by staying within the borders of the stable labelling $\mathcal{Lab} = \text{Ext2Lab}(\text{Args})$ (with Args being the stable extension that contains A , the argument that the discussion will start with). That is, as long as player M only plays arguments that are labelled **in** by \mathcal{Lab} , each **out** move of player S will be labelled **out** by \mathcal{Lab} , which then implies that player M can always react with an argument that is labelled **in** by \mathcal{Lab} , etc. Moreover, when player S does a **question**(A) move, either A itself or an attacker of A is labelled **in** by \mathcal{Lab} , which again means that player M can always respond with an argument that is labelled **in** by \mathcal{Lab} . As the argumentation framework is finite and player S cannot repeat himself, it follows that the game will finish in a finite number of moves. As player M can *always* react to the moves of player S, this means that the last move has to be an M-move. Hence, player M wins the game.¹⁷

Definition 14 describes the discussion game for credulous stable semantics (that is, it can be used to determine whether an argument is in at least one stable

¹⁷A more elaborate proof can be found in [Caminada and Wu, 2009].

extension). It is, however, relatively straightforward to re-apply this game in the context of sceptical stable semantics (that is, to determine whether an argument is in every stable extension). The idea is that an argument A is in each stable extension iff no attacker of A is in any stable extension. So in order to determine whether A is in every stable extension, one could try to play the stable discussion game for each attacker of A . If for none of these attackers the discussion game can be won, argument A is in each stable extension.

6 Grounded Semantics

So far, we have mainly focussed on the preferred discussion game and its slightly modified variants for ideal and stable semantics. In the current section we will focus on a fundamentally different type of discussion game, in the context of grounded semantics.

One of the main differences between the preferred discussion game and the grounded discussion game to be introduced in the current section is a conceptual one. To properly understand this difference, it is useful to take the perspective of complete labellings. We recall that a complete labelling (Definition 11) is a labelling where one has reasons for each argument one accepts (because all its attackers are rejected), reasons for each argument one rejects (because it has an attacker that is accepted), and reasons for each argument one abstains from having an explicit opinion about (because there are insufficient grounds to accept it and insufficient grounds to reject it). As such, a complete labelling can be seen as a reasonable position on how to evaluate the conflicting information represented in the argumentation framework. The preferred discussion game determines whether an argument is accepted (labelled *in*) by *at least one* such reasonable position.¹⁸ The grounded discussion game, to be introduced in the current section, determines whether an argument is accepted (labelled *in*) by *every* such reasonable position.¹⁹ That is, from the perspective of complete labellings, the preferred discussion game is about whether an argument *can be* accepted, whereas the grounded discussion game is about whether an argument *has to be* accepted.

The difference between determining whether an argument can be accepted and whether an argument has to be accepted is reflected in the nature of the associated discussion game. If the discussion is merely about whether an argument can be accepted (that is, about whether there exists a reasonable position in which the argument is accepted) then arguing against this means pointing out that any position in which the argument is accepted is somehow not reasonable. That is, the opponent tries to lead the proponent of such a position towards a contradiction.²⁰ Hence, the admissible discus-

¹⁸This is because an argument is labelled *in* by some admissible labelling i it is labelled *in* by some complete labelling.

¹⁹This is because an argument is labelled *in* by the grounded labelling i it is labelled *in* by every complete labelling.

²⁰like saying, if you think that argument X is labelled *in*, then it follows that X 's attacker Y should be labelled *out*, but previously you claimed that Y should be labelled *in*.

sion game has at least some properties of Socratic discussion [Caminada, 2008; Caminada *et al.*, 2014]. If, on the other hand, the discussion is about whether an argument has to be accepted (that is, about whether the argument is accepted in each reasonable position) then the discussion gets a totally different nature. If an argument is accepted in each reasonable position, then in particular one's discussion partner, by being reasonable, should accept the argument. So the discussion becomes one of trying to *convince* the discussion partner that he has to accept a particular argument. That is, the discussion partner should be shown that by being reasonable, he cannot avoid having to accept the argument in question. As such, the nature of the discussion becomes that of persuasion dialogue [Walton and Krabbe, 1995].

Now that the conceptual difference between the preferred discussion game and the grounded discussion game has been explained, we will take a closer look at the technical differences. Although the preferred discussion game is used to determine membership of a preferred extension, it does so by determining membership of an admissible set (labelling).²¹ This has the advantage of not having to construct the entire preferred extension (labelling), as constructing an admissible set (labelling) will be sufficient. Similarly, although the grounded discussion game is used to determine membership of the grounded extension, it does so by determining membership of a strongly admissible set (labelling) [Baroni and Giacomin, 2007; Caminada, 2014].²² This has the advantage of not having to construct the entire grounded extension (labelling) as constructing a strongly admissible set (labelling) will be sufficient.

The grounded discussion game [Caminada, 2015a; Caminada, 2015b] that we will describe in the current section has two players (proponent and opponent) and is based on four different moves, each of which has an argument as a parameter.

HTB(A) (“*A* has to be the case”)

With this move, the proponent claims that *A* has to be labelled **in** by every complete labelling, and hence also has to be labelled **in** by the grounded labelling.

CB(B) (“*B* can be the case, or at least cannot be ruled out”)

With this move, the opponent claims that *B* does not have to be labelled **out** by every complete labelling. That is, the opponent claims there exists a complete labelling where *B* is labelled **in** or **undec**, and that *B* is therefore not labelled **out** by the grounded labelling.

CONCEDE(A) (“I agree that *A* has to be the case”)

With this move, the opponent indicates that he now agrees with the

²¹Recall that an admissible set (labelling) can always be extended to a preferred extension (labelling), as a preferred extension (labelling) is a maximal admissible set (labelling).

²²Recall that a strongly admissible set (labelling) can always be extended to the grounded extension (labelling), as the grounded extension (labelling) is the maximal strongly admissible set (labelling) (see Theorem 2 and the work of Baroni and Giacomin [2007] and Caminada [2014]).

proponent (who previously did an $HTB(A)$ move) that A has to be the case (labelled **in** by every complete labelling, including the grounded).

$RETRACT(B)$ (“I give up that B can be the case”)

With this move, the opponent indicates that he no longer believes that B can be **in** or **undec**. That is, the opponent acknowledges that B has to be labelled **out** by every complete labelling, including the grounded.

One of the key ideas of the discussion game is that the proponent has burden of proof. He has to establish the acceptance of the main argument and make sure the discussion does not go around in circles. The opponent merely has to cast sufficient doubts.

The game starts with the proponent uttering an HTB statement. After each HTB statement (either the first one or a subsequent one) the opponent utters a sequence of one or more CB , $CONCEDE$ and $RETRACT$ statements, after which the proponent again utters an HTB statement, etc. In the argumentation framework of Figure 1 the discussion could go as follows.

- | | |
|-----------------|---------------------|
| (1) P: $HTB(C)$ | (4) O: $CONCEDE(A)$ |
| (2) O: $CB(B)$ | (5) O: $RETRACT(B)$ |
| (3) P: $HTB(A)$ | (6) O: $CONCEDE(C)$ |

In the above discussion, C is called *the main argument* (the argument the discussion starts with). The discussion above ends with the main argument being conceded by the opponent, so we say that the proponent wins the discussion.

As an example of a discussion that is lost by the proponent, it can be illustrative to examine what happens if the proponent claims that B has to be the case.

- | | |
|-----------------|----------------|
| (1) P: $HTB(B)$ | (2) O: $CB(A)$ |
|-----------------|----------------|

After the second move, the discussion is terminated, as the proponent cannot make any further move, since A does not have any attackers. This brings us to the precise preconditions of the discussion moves.

$HTB(A)$ Either this is the first move, or the previous move was $CB(B)$, where A attacks B , and no $CONCEDE$ or $RETRACT$ move is applicable.

$CB(A)$ A is an attacker of the last $HTB(B)$ statement that is not yet conceded, the directly preceding move was not a CB statement, argument A has not yet been retracted, and no $CONCEDE$ or $RETRACT$ move is applicable.

$CONCEDE(A)$ There has been an $HTB(A)$ statement in the past, of which every attacker has been retracted, and $CONCEDE(A)$ has not yet been moved.

$RETRACT(A)$ There has been a $CB(A)$ statement in the past, of which there exists an attacker that has been conceded, and $RETRACT(A)$ has not yet been moved.

Apart from the preconditions mentioned above, all four statements also have the additional precondition that no *HTB-CB* repeats have occurred. That is, there should be no argument for which *HTB* has been uttered more than once, *CB* has been uttered more than once, or both *HTB* and *CB* have been uttered. In the first and second case, the discussion is going around in circles, which the proponent has to prevent as he has burden of proof. In the third case, the proponent has been contradicting himself, as his statements are not conflict-free. In each of these three cases, the discussion comes to an end with no move being applicable anymore. The above conditions are made formal as follows.

Definition 15 *Let $AF = (Ar, att)$ be an argumentation framework. A grounded discussion is a sequence of discussion moves constructed by applying the following principles.*

BASIS (HTB) *If $A \in Ar$ then $[HTB(A)]$ is a grounded discussion.*

STEP (HTB) *If $[M_1, \dots, M_n]$ ($n \geq 1$) is a grounded discussion without *HTB-CB* repeats,²³ and no *CONCEDE* or *RETRACT* move is applicable,²⁴ and $M_n = CB(A)$ and B is an attacker of A then $[M_1, \dots, M_n, HTB(B)]$ is also a grounded discussion.*

STEP (CB) *If $[M_1, \dots, M_n]$ ($n \geq 1$) is a grounded discussion without *HTB-CB* repeats, and no *CONCEDE* or *RETRACT* move is applicable, and M_n is not a *CB* move, and there is a move $M_i = HTB(A)$ ($i \in \{1 \dots n\}$) such that the discussion does not contain *CONCEDE(A)*, and for each move $M_j = HTB(A')$ ($j > i$) the discussion contains a move *CONCEDE(A')*, and B is an attacker of A such that the discussion does not contain a move *RETRACT(B)*, then $[M_1, \dots, M_n, CB(B)]$ is a grounded discussion.*

STEP (CONCEDE) *If $[M_1, \dots, M_n]$ ($n \geq 1$) is a grounded discussion without *HTB-CB* repeats, and *CONCEDE(B)* is applicable then $[M_1, \dots, M_n, CONCEDE(B)]$ is a grounded discussion.*

STEP (RETRACT) *If $[M_1, \dots, M_n]$ ($n \geq 1$) is a grounded discussion without *HTB-CB* repeats, and *RETRACT(B)* is applicable then $[M_1, \dots, M_n, RETRACT(B)]$ is a grounded discussion.*

It can be observed that the preconditions of the moves are such that a proponent move (*HTB*) can never be applicable at the same moment as an opponent move

²³We say that there is a *HTB-CB* repeat i $\exists i, j \in \{1 \dots n\} \exists A \in Ar : (M_i = HTB(A) \vee M_i = CB(A)) \wedge (M_j = HTB(A) \vee M_j = CB(A)) \wedge i \neq j$.

²⁴A move *CONCEDE(B)* is applicable i the discussion contains a move *HTB(A)* and for every attacker A of B the discussion contains a move *RETRACT(B)*, and the discussion does not already contain a move *CONCEDE(B)*. A move *RETRACT(B)* is applicable i the discussion contains a move *CB(B)* and there is an attacker A of B such that the discussion contains a move *CONCEDE(A)*, and the discussion does not already contain a move *RETRACT(B)*.

(*CB*, *CONCEDE* or *RETRACT*). That is, proponent and opponent essentially take turns in which each proponent turn consists of a single *HTB* statement, and every opponent turn consists of a sequence of *CONCEDE*, *RETRACT* and *CB* moves.

Definition 16 *A grounded discussion $[M_1, \dots, M_n]$ is called terminated iff there exists no move M_{n+1} such that $[M_1, \dots, M_n, M_{n+1}]$ is a grounded discussion. A terminated grounded discussion (with A being the main argument) is won by the proponent iff the discussion contains $CONCEDE(A)$, otherwise it is won by the opponent.*

To illustrate why the discussion has to be terminated after the occurrence of an *HTB-CB* repeat, consider the following discussion in the argumentation framework of Figure 1.

- | | |
|---------------------------|---------------------------|
| (1) P: <i>HTB</i> (G) | (3) P: <i>HTB</i> (G) |
| (2) O: <i>CB</i> (H) | |

At the third move, an *HTB-CB* repeat occurs and the discussion is terminated (opponent wins). Hence, termination after an *HTB-CB* repeat is necessary to prevent the discussion from going on perpetually.

Theorem 5 *Every discussion will terminate after a finite number of steps.*

From the fact that a discussion terminates after an *HTB-CB* repeat, the following result follows.

Lemma 1 *No discussion can contain a $CONCEDE$ and $RETRACT$ move for the same argument.*

The soundness and completeness of the game described above is stated in the following theorem.

Theorem 6 ([Caminada, 2015a]) *Let (Ar, att) be an argumentation framework and let $A \in Ar$.*

1. *If there exists a grounded discussion for A that is won by player P , then A is labelled **in** by the grounded labelling.*
2. *If A is labelled **in** by the grounded labelling, then player P has a winning strategy for A in the grounded discussion game.*

The correctness of Theorem 6 can be seen as follows. As for point 1, it can be observed that what the discussion game actually does is to construct a strongly admissible labelling of which the **in**-labelled arguments coincide with the *CONCEDE* moves, and the **out**-labelled arguments coincide with the *RETRACT* moves. In fact, it can be proved by induction that at each state of the discussion, the labelling where each *CONCEDE* move is labelled **in** and

each retract move is labelled *out* is strongly admissible [Caminada, 2015b]. The fact that the discussion is won by player *P* implies that the main argument (*A*) has been conceded. So at the end of the discussion, we have a strongly admissible labelling where argument *A* is labelled *in*. Hence, by Theorem 2, *A* is labelled *in* by the grounded labelling.

As for point 2, it should be mentioned that a strongly admissible labelling (for instance the grounded labelling) with its associated min-max numbering can serve as a roadmap for winning the discussion. The proponent will be able to win if, whenever he has to do an *HTB* move, he prefers to use an *in* argument with the lowest min-max number that attacks the directly preceding *CB* move. We refer to this as a *lowest number strategy*.²⁵

It turns out that when applying such a strategy, the game stays within the boundaries of the strongly admissible labelling (that is, within its *in* and *out* labelled part). As long as each *HTB* move of the proponent is related to an *in*-labelled argument, it follows that all the attackers are labelled *out* (Definition 8, first bullet) so each *CB* move the opponent utters in response will be related to an *out*-labelled argument. This *out*-labelled argument will then have at least one *in*-labelled attacker (Definition 8, second bullet) as a candidate for the proponent's subsequent *HTB* move.

The next thing to be observed is that when the proponent applies a lowest number strategy, the game will not terminate due to any *HTB-CB* repeats. This is due to the facts that (1) after a move *HTB(A)* is played (for some argument *A*) all subsequent *CB* and *HTB* moves will be related to arguments with lower min-max numbers than *A* until a move *CONCEDE(A)* is played, and (2) after a move *CB(A)* is played (for some argument *A*), all subsequent *HTB* and *CB* moves will be related to arguments with lower min-max numbers than *A*, until a move *RETRACT(A)* is played. We refer to [Caminada, 2015b] for details.

7 Tree-Based Discussion Games

The discussion games that were described in the previous sections are not the only ones that have been stated for preferred, stable, ideal and grounded semantics. In fact, various alternative dialectical proof procedures can be found in the literature, many of them are based on the concept of dialectical trees [Dung *et al.*, 2007; Modgil and Caminada, 2009; Thang *et al.*, 2009]. In the current section, we aim to provide an impression of these tree-based discussion games, and explain some of their disadvantages compared to the discussion games described in the previous sections. Rather than giving an overview of all tree-based discussion games that have been stated in the literature, we will focus our attention on one of them: the Standard Grounded Game [Prakken and Sartor, 1997; Caminada, 2004; Modgil and Caminada, 2009].

²⁵We write *a* lowest number strategy instead of *the* lowest number strategy as a lowest number strategy might not be unique due to different lowest numbered *in*-labelled arguments being applicable at a specific point. In that case it is sufficient to pick an arbitrary one.

The Standard Grounded Game (SGG) [Prakken and Sartor, 1997; Caminada, 2004; Modgil and Caminada, 2009] is one of the earliest dialectical proof procedures for grounded semantics. Each game²⁶ consists of a sequence $[A_1, \dots, A_n]$ ($n \geq 1$) of arguments, moved by the proponent and opponent taking turns, with the proponent starting. That is, a move A_i ($i \in \{1 \dots n\}$) is a proponent move iff i is odd, and an opponent move iff i is even. Each move, except the first one, is an attacker of the previous move. In order to ensure termination even in the presence of cycles, the proponent is not allowed to repeat any of his moves. A game is terminated iff no next move is possible; the player making the last move wins. Formally, the Standard Grounded Game can be defined as follows.

Definition 17 *A discussion in the Standard Grounded Game is a finite sequence $[A_1, \dots, A_n]$ ($n \geq 1$) of arguments (sometimes called moves), of which the odd moves are called P-moves (Proponent moves) and the even moves are called O-moves (Opponent moves), such that:*

1. *every O-move is an attacker of the preceding P-move (that is, every A_i where i is even and $2 \leq i \leq n$ attacks A_{i-1})*
2. *every P-move except the first one is an attacker of the preceding O-move (that is, every A_i where i is odd and $3 \leq i \leq n$ attacks A_{i-1})*
3. *P-moves are not repeated (that is, for every odd $i, j \in \{1, \dots, n\}$ it holds that if $i \neq j$ then $A_i \neq A_j$)*

A discussion is called terminated iff there is no A_{n+1} such that $[A_1, \dots, A_n, A_{n+1}]$ is a discussion. A terminated discussion is said to be won by the player making the last move.

As an example, in the argumentation framework of Figure 1 $[C, B, A]$ is terminated and won by the proponent (as A has no attackers, the opponent cannot move anymore) whereas $[G, H]$ is terminated and won by the opponent (as the only attacker of H is G , which the proponent is not allowed to repeat). It is sometimes possible for the proponent to win a game even if the main argument is not in the grounded extension. An example would be $[F, B, A]$. This illustrates that in order to show that an argument is in the grounded extension, a single game won by the proponent is not sufficient. Instead, what is needed is a *winning strategy*. This is essentially a tree in which each node is associated with an argument such that (1) each path from the root to a leaf constitutes a terminated discussion won by the proponent, (2) the children of each proponent node (a node corresponding with a proponent move) coincide with all attackers of the associated argument, and (3) each opponent node (a node corresponding with an opponent move) has precisely one child, whose argument attacks the argument of the opponent node.

²⁶What we call an SGG game is called a *line of dispute* in [Modgil and Caminada, 2009].

Formally, *argument tree* is a tree of which each node (\mathbf{n}) is labelled with an argument ($Arg(\mathbf{n})$). The *level* of a node is the number of nodes in the path to the root. This leads to the following formal definition of a winning strategy in the context of the Standard Grounded Game.

Definition 18 *A winning strategy of the Standard Grounded Game for argument A is an argument tree, where the root is labelled with A , such that*

1. *for each path from the root (\mathbf{n}_{root}) to a leaf node (\mathbf{n}_{leaf}) it holds that the arguments on this path form a terminated discussion won by P*
2. *for each node at odd level \mathbf{n}_P it holds that $\{Arg(\mathbf{n}_{child}) \mid \mathbf{n}_{child} \text{ is a child of } \mathbf{n}_P\} = \{B \mid B \text{ attacks } Arg(\mathbf{n}_P)\}$ and the number of children of \mathbf{n}_P is equal to the number of attackers of $Arg(\mathbf{n}_P)$*
3. *each node of even level \mathbf{n}_O has precisely one child \mathbf{n}_{child} , and $Arg(\mathbf{n}_{child})$ attacks $Arg(\mathbf{n}_O)$*

It has been proved that an argument is in the grounded extension iff the proponent has a winning strategy for it in the SGG [Prakken and Sartor, 1997; Caminada, 2004]. Moreover, it has also been shown that an SGG winning strategy defines a strongly admissible labelling, when each argument of a proponent node is labelled **in**, each argument of an opponent node is labelled **out** and all remaining arguments are labelled **undec** [Caminada, 2014].

As an example, in the argumentation framework of Figure 1 the winning strategy for argument E would be the tree consisting of the two branches $E - B - A$ and $E - D - C - B - A$, thus proving its membership of the grounded extension by yielding the strongly admissible labelling $(\{A, C, E\}, \{B, D\}, \{F, G, H\})$.

As can be observed from this example, a winning strategy of the SGG can contain some redundancy when it comes to multiple occurrences of the same arguments in different branches. In the current example, the redundancy is relatively mild (consisting of just the two arguments A and B) but other cases exist where the SGG requires a number of moves in the winning strategy that is *exponential* w.r.t. the size of the strongly admissible labelling the winning strategy is defining. As an example, consider the argumentation framework of Figure 5 (top left). The winning strategy of the SGG is in the same figure (top right). Now consider what would happen if one would start to extend the argumentation framework by duplicating the middle part. That is, suppose we have arguments B_1, \dots, B_n and C_1, \dots, C_n (with n being an odd number), as well as arguments A and D . Suppose that for every $i \in \{1, \dots, n-1\}$ B_{i+1} attacks B_i , and C_{i+1} attacks C_i , and that for each *even* $i \in \{2, \dots, n-1\}$ B_{i+1} attacks C_i , and C_{i+1} attacks B_i , and that B_1 and C_1 attack A , and that D attacks B_n and C_n . In that case, the branches in the SGG winning strategy would split at every O-move. So for $n = 3$ (as is the case in Figure 5) the number of branches is four, for $n = 5$ it is eight, etc. In general, the number of

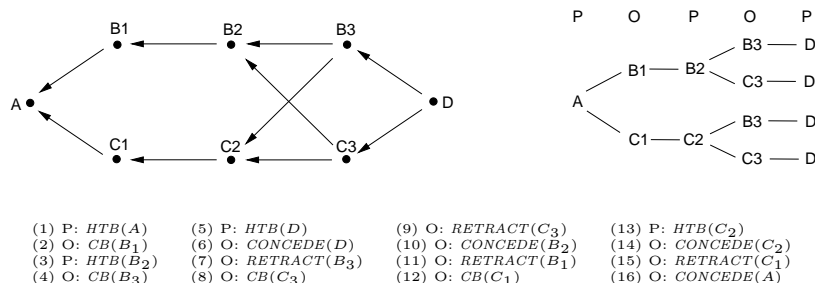


Figure 5. The Standard Grounded Game (SGG) versus the Grounded Discussion Game (GDG).

branches in the SGG winning strategy is $2^{(n+1)/2}$, with the number of nodes in the SGG winning strategy being $1 + 2\sum_{i=1}^{(n+1)/2} 2^i$. Hence, the number of steps needed in a winning strategy of the SGG can be *exponential* in relation to the size (number of *in* and *out* labelled arguments) of the strongly admissible labelling that the SGG winning strategy is constructing.²⁷

As for the Grounded Discussion Game (GDG) as described in Section 6, the situation is different. As was mentioned in Section 6, what the GDG essentially does is to construct a strongly admissible labelling of which the *in* labelled arguments coincide with the *CONCEDE* moves and the *out* labelled arguments coincide with the *RETRACT* moves. It can be observed that no argument occurs in both a *CONCEDE* and *RETRACT* move (otherwise the argument would also have occurred in both an *HTB* and *CB* move, and the discussion would have terminated before reaching the *CONCEDE* and *RETRACT* moves) and that for each argument there exists at most one *CONCEDE* move and at most one *RETRACT* move. As we assume the game is won by the proponent, who is playing a lowest number strategy, there will be no *HTB-CB* repeats. This implies that for each *CONCEDE* move, there exists precisely one *HTB* move, and for each *RETRACT* move, there exists precisely one *CB* move. This means that the total number of moves (in a game won by the proponent, who is applying a lowest number strategy) is two times the number of *in* labelled arguments (which accounts for the *HTB* and *CONCEDE* moves) plus two times the number of *out* labelled arguments (which accounts for the *CB* and *RETRACT* moves). Hence, the number of moves in the game is *linear* in relation to the size (number of *in* and *out* labelled arguments) of the strongly admissible labelling the GDG is constructing.²⁸

Hence, whereas for the Grounded Discussion Game, constructing a strongly admissible labelling (which is needed to show membership of the grounded extension) requires a linear number of moves, for the Standard Grounded Game this requires a potentially exponential number of moves. This makes the GDG

²⁷We thank Mikolaj Podlaszewski for this example.

²⁸See [Caminada, 2015a] for details.

a better choice for purposes of human-computer interaction, assuming that the human user’s time is precious.

It should be mentioned that the possibility of an exponential blowup in the number of moves is not restricted to the SGG, but is a feature of tree-based discussion games in general. For instance, the above sketched example also leads to an exponential number of moves in the preferred semantics game of Modgil and Caminada [2009] and in the ideal semantics game of Dung *et al.* [2007]. The key feature of these approaches is that they require a winning strategy to show membership of a (grounded, preferred or ideal) extension. It is this winning strategy that is responsible for the exponential blowup. In the discussion games described in sections 3, 5 and 6, however, no winning strategy is required, as just a single game won by the proponent is sufficient to prove membership of a (preferred, stable or grounded) extension.²⁹

8 Discussion

What the above described discussion games for preferred semantics (Section 3), stable semantics (Section 5) and grounded semantics (Section 6) have in common is that (1) a single game won by the proponent is sufficient to prove membership of a (preferred, stable or grounded) extension, and (2) if an argument is member of a (preferred, stable or grounded) extension then the proponent has a winning strategy for it. This is evidenced by theorems 3, 4 and 6. In tree-based discussion games, like those of Dung *et al.* [2007], Modgil and Caminada [2009] and Thang *et al.* [2009] point (1) is altered such that a single game won by the proponent is *not* sufficient to prove membership of an extension; for this a winning strategy is needed. Having to provide such a winning strategy in a dialectical way can be troublesome for two reasons. First of all, the tree of the winning strategy would need to be “linearized” as discussions take place not in branching time but in linear time. But even if linearization takes place, one still has to deal with the fact that the original (tree-based) winning strategy could have a size that is exponentially related to the (strongly) admissible labelling it is based on. The discussion games presented in sections 3, 5 and 6 have the advantage that they are not tree-based and hence do not have these problems.

One can ask the question whether it is always possible (for any argumentation semantics) to define a discussion game that satisfies the points (1) and (2) mentioned above. For instance, the procedure sketched in Section 4 (ideal semantics) does not satisfy point (1). This is because in the second phase of the discussion, when trying to find an admissible set that attacks the admissible set obtained in the first phase of the discussion, not finding such a set could be due to the proponent making the “wrong” choices during the second phase, rather than due to the actual absence of such a set. It would be a challenge

²⁹It can be proved that the preferred discussion game (Section 3) is linear in the number of moves required. See [Caminada *et al.*, 2014] for details. Using similar techniques one can also prove that the stable discussion game (Section 5) requires only a linear number of moves.

to change the discussion procedure for ideal semantics such that both points (1) and (2) are satisfied. An even greater challenge would be to formulate discussion games (still satisfying points (1) and (2)) for semi-stable, stage or even CF2 semantics.

As the tree-based discussion games of Dung *et al.* [2007], Modgil and Caminada [2009] and Thang *et al.* [2009] violate point (1) but satisfies point (2), one can ask the question of whether there also exists a discussion game that satisfies point (1) but violates point (2). The answer is affirmative, as is evidenced by the work of Caminada and Podlaskowski [2012a; 2012b]. Here, the ability to win the discussion game might depend on cooperation of the opponent. So even though an argument being in the grounded extension implies the existence of a discussion for it that is won by the proponent, it does not imply that the proponent also has a winning strategy.³⁰ For the purpose of human-computer interaction, this property is undesirable, as the computer should be able to win the discussion (for an argument that is actually in the grounded extension) regardless of how the human user chooses to utter the possible counterarguments.

The discussion games presented in the current chapter have been stated in the context of *abstract* argumentation theory. This raises the question of whether these discussion games are also suitable in the context of *instantiated* argumentation, like ASPIC+ [Modgil and Prakken, 2014]³¹, ABA [Toni, 2014]³² or logic-based argumentation [Gorogiannis and Hunter, 2011]³³. Technically, this should not be a problem, as each of these formalisms provides an instantiation of Dung’s abstract argumentation theory. That is, each of these formalisms specifies what arguments can be constructed and how these attack each other, starting from a particular knowledge base. Although applying the discussion games in the context of instantiated argumentation is technically straightforward, there is a catch. The question is whether the notion of attack of the instantiated argumentation formalism is defined in such a way that it allows for moves that can be considered as intuitive during the course of the discussion. For instance, in ASPIC+ it can be the case that a discussion partner utters an argument with conclusion c , which cannot be replied to with an argument for conclusion $\neg c$ (even though such an argument is well-formed and perhaps even justified) because the definition of attack is such that it does not attack the argument with conclusion c . This is like having your discussion partner uttering an argument for a claim (c) which you know is not the case, but you’re not allowed to reply with an argument that directly rebuts this claim. We refer to Chapter 15 (“Rationality postulates: applying argumentation theory for non-monotonic reasoning”) of this volume for details.

As mentioned in the introduction, one of the possible applications of the

³⁰We refer to [Caminada, 2015a] for a specific example.

³¹See also Chapter 6 (Abstract rule-based argumentation) of this volume.

³²See also Chapter 7 (assumption-based argumentation: disputes, explanations, preferences) of this volume.

³³See also Chapter 9 (A review of argumentation based on deductive arguments) of this volume.

discussion games is for the purpose of human-computer interaction. The context here is that of a shared knowledge base³⁴ (say, of medical research and clinical evidence) that allows for the construction of arguments (say, regarding to how to treat a particular patient). As the knowledge base can be complex and huge, it is not always directly obvious what the justified arguments are. Although a software implementation of (instantiated) argumentation theory can help to provide an answer, the correctness of this answer might need to be explained to a human user. Our hypothesis is that human-computer discussion can contribute to acceptance of argument-based entailment. In order to test this hypothesis, one would need to perform experiments in which the user's confidence in the argument-based entailment is tested, before and after performing the discussion game. Experiments like these is what we would like to perform in the near future.

BIBLIOGRAPHY

- [Baroni and Giacomin, 2007] P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.
- [Baroni *et al.*, 2011] P. Baroni, M.W.A. Caminada, and M. Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.
- [Caminada and Gabbay, 2009] M.W.A. Caminada and D.M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2-3):109–145, 2009. Special issue: new ideas in argumentation theory.
- [Caminada and Podlaszewski, 2012a] M.W.A. Caminada and M. Podlaszewski. Grounded semantics as persuasion dialogue. In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Computational Models of Argument - Proceedings of COMMA 2012*, pages 478–485, 2012.
- [Caminada and Podlaszewski, 2012b] M.W.A. Caminada and M. Podlaszewski. User-computer persuasion dialogue for grounded semantics. In Jos W.H.M. Uiterwijk, Nico Roos, and Mark H.M. Winands, editors, *Proceedings of BNAIC 2012; The 24th Benelux Conference on Artificial Intelligence*, pages 343–344, 2012.
- [Caminada and Sakama, 2015] M.W.A. Caminada and Ch. Sakama. On the issue of argumentation and informedness. In *2nd International Workshop on Argument for Agreement and Assurance*, 2015.
- [Caminada and Wu, 2009] M.W.A. Caminada and Y. Wu. An argument game of stable semantics. *Logic Journal of IGPL*, 17(1):77–90, 2009.
- [Caminada *et al.*, 2014] M.W.A. Caminada, W. Dvorak, and S. Vesic. Preferred semantics as socratic discussion. *Journal of Logic and Computation*, 2014. (in print).
- [Caminada, 2004] M.W.A. Caminada. *For the sake of the Argument. Explorations into argument-based reasoning*. PhD thesis, Vrije Universiteit Amsterdam, 2004.
- [Caminada, 2006] M.W.A. Caminada. On the issue of reinstatement in argumentation. In M. Fischer, W. van der Hoek, B. Konev, and A. Lisitsa, editors, *Logics in Artificial Intelligence; 10th European Conference, JELIA 2006*, pages 111–123. Springer, 2006. LNAI 4160.
- [Caminada, 2008] M.W.A. Caminada. A formal account of socratic-style argumentation. *Journal of Applied Logic*, 6(1):109–132, 2008.
- [Caminada, 2011] M.W.A. Caminada. A labelling approach for ideal and stage semantics. *Argument & Computation*, 2:1–21, 2011.

³⁴A particularly interesting situation is where such a shared knowledge base is absent, that is, where proponent and opponent each have their own private knowledge base and associated argumentation framework. In that case, both proponent and opponent learn new information from each other during the course of the discussion. This puts additional constraints on the discussion protocol. We refer to [Caminada and Sakama, 2015] for details.

- [Caminada, 2014] M.W.A. Caminada. Strong admissibility revisited. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Computational Models of Argument; Proceedings of COMMA 2014*, pages 197–208. IOS Press, 2014.
- [Caminada, 2015a] M.W.A. Caminada. A discussion game for grounded semantics. In Elizabeth Black, Sanjay Modgil, and Nir Oren, editors, *Theory and Applications of Formal Argumentation (proceedings TAFA 2015)*, pages 59–73. Springer, 2015.
- [Caminada, 2015b] M.W.A. Caminada. A discussion protocol for grounded semantics (proofs). Technical report, University of Aberdeen, 2015.
- [Doutre and Mengin, 2004] S. Doutre and J. Mengin. On sceptical versus credulous acceptance for abstract argument systems. In *Proceedings of the 9th European Conference on Logics in Artificial Intelligence (JELIA-2004)*, pages 462–473, 2004.
- [Dung and Thang, 2007] P.M. Dung and P.M. Thang. A sound and complete dialectical proof procedure for sceptical preferred argumentation. In *Proc. of the LPNMR Workshop on Argumentation and Nonmonotonic Reasoning (ArgNMR07)*, pages 49–63, 2007.
- [Dung et al., 2007] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10-15):642–674, 2007.
- [Gorogiannis and Hunter, 2011] N. Gorogiannis and A. Hunter. Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, 175(9-10):1479–1497, 2011.
- [Modgil and Caminada, 2009] S. Modgil and M.W.A. Caminada. Proof theories and algorithms for abstract argumentation frameworks. In I. Rahwan and G.R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 105–129. Springer, 2009.
- [Modgil and Prakken, 2014] S. Modgil and H. Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument & Computation*, 5:31–62, 2014. Special Issue: Tutorials on Structured Argumentation.
- [Prakken and Sartor, 1997] H. Prakken and G. Sartor. Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-Classical Logics*, 7:25–75, 1997.
- [Thang et al., 2009] Phan Minh Thang, Phan Minh Dung, and Nguyen Duy Hung. Towards a common framework for dialectical proof procedures in abstract argumentation. *Journal of Logic and Computation*, 19(6):1071–1109, 2009.
- [Toni, 2014] F. Toni. A tutorial on assumption-based argumentation. *Argument & Computation*, 5:89–117, 2014. Special Issue: Tutorials on Structured Argumentation.
- [Vreeswijk and Prakken, 2000] G.A.W. Vreeswijk and H. Prakken. Credulous and sceptical argument games for preferred semantics. In *Proceedings of the 7th European Workshop on Logic for Artificial Intelligence (JELIA-00)*, number 1919 in Springer Lecture Notes in AI, pages 239–253, Berlin, 2000. Springer Verlag.
- [Walton and Krabbe, 1995] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Series in Logic and Language. State University of New York Press, Albany, NY, USA, 1995.

Martin Caminada
 Cardiff University
 Email: CaminadaM@cardiff.ac.uk