An overview of formal models of argumentation and their application in philosophy

Henry Prakken
Department of Information and Computing Sciences, Utrecht University
& *Faculty of Law, University of Groningen, The Netherlands*

April 4, 2011

Abstract

Argumentation is the process of supporting claims with grounds and defending them against attack. In the last decades argumentation has become an important topic in philosophy and artificial intelligence. In philosophy, the criticisms of Toulmin and Perelman of formal logic in the 1950s and 1960s gave rise to the field of informal logic, which studies informal models of reasoning and argumentation. In artificial intelligence, formal models of argumentation have been proposed as models of commonsense reasoning and multi-agent conflict resolution. This paper discusses how the formal models resulting from this research can clarify philosophical problems and issues, including those raised in the field of informal logic. An important point will be that while formal logic in the days of Toulmin and Perelman only focused on mathematical reasoning, non-mathematical forms of reasoning can still be formalised.

1 Introduction

Introductions to logic often portray logically valid inference as 'foolproof' reasoning: an argument is valid if the truth of its premises guarantees the truth of its conclusion. However, we all construct arguments from time to time that are not foolproof in this sense but that merely make their conclusion plausible when their premises are true. For example, if we are told that John, a professor in economics, says that reducing taxes increases productivity, we conclude that reducing taxes increases productivity since we know that experts are usually right within their domain of expertise. Sometimes such arguments are defeated by counterarguments. For example, if we are also told that John has political ambitions, we have to retract our previous conclusion that he is right about the effect of taxes if we also believe that people with political ambitions are often biased when it comes to taxes. Or, to use an example of practical instead of epistemic reasoning, if we accept that reducing taxes increases productivity and that increasing productivity is good, then we conclude that the taxes should be reduced, unless we also accept that reducing taxes increases inequality, that this is bad and that equality is more important than productivity. However, as long as such counterarguments are not available, we are happy to live with the conclusions of our fallible arguments. The question is: are we then reasoning fallaciously or is there still logic in our reasoning?

An answer to this question has been given in the development of argumentation logics. In a nutshell, the answer is that there is such logic but that it is inherently dialectic: an argument only warrants its conclusion if it is acceptable, and an argument is acceptable if, firstly, it is properly constructed and, secondly, it can be defended against counterarguments. Thus argumentation logics must define three things: how arguments can be constructed, how they can be defeated by counterarguments and how they can be defended against such defeats.

Argumentation logics are a form of nonmonotonic logic (see e.g. Antoniou (1997)), since their notion of warrant is nonmonotonic: new information may give rise to new counterarguments defeating arguments that were originally acceptable. One attractive feature of argumentation logics as a model of nonmonotonic reasoning is that they are close to concepts like 'argument', 'rebuttal' and 'defeat' that are used in ordinary discourse, in philosophy and in professions such as the law. Another attractive feature is that argumentation has a dialogical side: notions like argument, attack and defence naturally apply when (human or artificial) agents try to persuade each other to adopt or give up a certain point of view. There is thus a natural relation between argumentation logics (which define what conclusions can be drawn from a given body of information) and dialogue systems for argumentation (which regulate how such a body of information can evolve during a dialogue).

This paper aims to show how formal models of argumentation can clarify philosophical problems and issues. Some of these arise in the field of epistemology. Pollock (1974) argued that the principles by which knowledge can be acquired are defeasible. Later he made this precise in a formal system (Pollock; 1995), which became a source of inspiration for the development of argumentation logics in artificial intelligence (AI). Rescher (1977) also stressed the dialectical nature of theories of knowledge and presented a disputational model of scientific inquiry.

Other issues and problems originate from the fields of informal logic and argumentation theory. In 1958, Stephen Toulmin launched his influential attack on logic research of those days, accusing it of only studying mathematical reasoning while ignoring other forms of reasoning, such as commonsense reasoning and legal reasoning (Toulmin; 1958). He argued that outside mathematics the standards for the validity of arguments are context-dependent and procedural: according to him an argument is valid if it has been properly defended in a dispute, and different fields can have different rules for when this is the case. Moreover, in his famous argument scheme he drew attention to the fact that different premises can have different roles in an argument (data, warrant or backing) and he noted the possibility of exceptions to rules (rebuttals). Perelman argued that arguments in ordinary discourse should not be evaluated in terms of their syntactic form but on their rhetorical potential to persuade an audience (Perelman and Olbrechts-Tyteca; 1969). These criticisms gave rise to the fields of informal logic and argumentation theory, which developed notions like argument schemes with critical questions and dialogue systems for argumentation. Many scholars in these fields distrusted or even rejected formal methods, but one point of this paper is that formal methods can also clarify these aspects of reasoning. Another claim often made in these fields is that arguments can only be evaluated in the context of a dialogue or procedure. A second point of this paper is that this can be respected by embedding logical in dialogical accounts of argumentation.

The problems to be discussed in this paper then are:

- How can formal standards for argumentation-based inference be developed?
- How can reasoning with argument schemes be formalised?
- Can the use of arguments to persuade be formalised?
- How can a procedural and context-dependent account of argument evaluation be given?

These questions will be answered in the following way. First in Section 2 a fully abstract framework for argument evaluation will be presented, which in Section 3 will be supplemented with a framework for accounts of argument construction and the nature of defeat. In Section 4 it is then explained how reasoning with argument schemes can be formalised in the resulting formal framework. In Section 5 the idea of dialogue systems for argumentation is introduced and used to clarify the remaining problems.

The present paper focuses on the use of formal methods to analyse these problems and cannot give a comprehensive survey of the formal study of argumentation. A systematic (although somewhat outdated) introduction to argumentation logics is Prakken and Vreeswijk (2002) while a recent collection of survey papers on argumentation in AI is Rahwan and Simari (2009).

2 Dung's abstract argumentation frameworks

In 1995 Phan Minh Dung introduced an abstract formalism for argumentation-based inference (Dung; 1995), which assumes as input nothing but a set (of arguments) ordered by a binary relation (by Dung called 'attack' but in this paper the term 'defeat' will be used).

Definition 2.1 [Abstract argumentation framework] An *abstract argumentation framework* (AF) is a pair $\langle \mathcal{A}, \mathit{Def} \rangle$. $\mathcal{A}$ is a set arguments and $\mathit{Def} \subseteq \mathcal{A} \times \mathcal{A}$ is a binary relation of defeat. We say that an argument A defeats an argument B iff $(A, B) \in \mathit{Def}$, and that A strictly defeats B if A defeats B while B does not defeat A. A set S of arguments is said to defeat an argument A iff some argument in S defeats A.

All further definitions in this section are relative to an implicitly assumed AF.

Definition 2.2 [Conflict-free, Defence] Let $\mathcal{B} \subseteq \mathcal{A}$.

- A set $\mathcal{B}$ is *conflict-free* iff there exist no A_i, A_j in $\mathcal{B}$ such that A_i defeats A_j.
- A set $\mathcal{B}$ *defends* an argument A_i iff for each argument $A_j \in \mathcal{A}$, if A_j defeats A_i, then there exists A_k in $\mathcal{B}$ such that A_k defeats A_j.

Definition 2.3 [Acceptability semantics] Let $\mathcal{B}$ be a set of arguments.

- $\mathcal{B}$ is an *admissible set* iff $\mathcal{B}$ is conflict-free and $\mathcal{B}$ defends all its members.
- $\mathcal{B}$ is a *preferred extension* iff $\mathcal{B}$ is a maximal (w.r.t. set-inclusion) admissible set.
- $\mathcal{B}$ is a *stable extension* iff $\mathcal{B}$ is conflict-free and $\mathcal{B}$ attacks all arguments in $\mathcal{A} \backslash \mathcal{B}$.
- $\mathcal{B}$ is a *complete extension* iff $\mathcal{B}$ is admissible and contains all arguments it defends.
- $\mathcal{B}$ is a *grounded extension* iff $\mathcal{B}$ is the least (wrt set inclusion) complete extension.

These definitions formalise so-called preferred, stable, grounded and complete semantics for abstract argumentation frameworks. Some known facts are that

- each grounded, preferred or stable extension of an AF is also a complete extension of that AF;
- the grounded extension is unique but all other semantics allow for multiple extensions of an AF;
- each AF has a grounded and at least one preferred and complete extension, but there are AFs without stable extensions;
- the grounded extension of an AF is contained in all other extensions of that AF.
- If all arguments have at most a finite number of defeaters, then the grounded extension can be obtained by iterating the function $\mathcal{F}$ on the empty set. More precisely, the grounded extension is then $\mathcal{B} = \mathcal{B}_0 \cup \ldots \cup \mathcal{B}_n$ where
  - $\mathcal{B}_0 = \varnothing$;
  - $\mathcal{B}_{i+1} = \mathcal{B} \cup \{A \mid A \text{ is defended by } \mathcal{B}_i\}$

  Otherwise, thus a subset of the grounded extension can be obtained.

Argument extensions can also be characterised in terms of so-called status assignments or labellings (Verheij; 1996; Jakobovits and Vermeir; 1999; Caminada; 2006).

Definition 2.4 A *status assignment* assigns to zero or more members of $\mathcal{A}$ either the status *in* or *out* (but not both) such that:

1. an argument is *in* iff all arguments defeating it are *out*.
2. an argument is *out* iff it is defeated by an argument that is *in*.

Let $In = \{A \mid A$ is *in*$\}$ and $Out = \{A \mid A$ is *out*$\}$ and *Undecided* $= \mathcal{A} \setminus (In \cup Out)$. Then

1. A status assignment is *stable* if *Undecided* $= \varnothing$.
2. A status assignment is *preferred* if *Undecided* is minimal (wrt set inclusion)
3. A status assignment is *grounded* if *Undecided* is maximal (wrt set inclusion)

4. Any status assignment is *complete*.

These notions coincide with those of Definition 2.3 as follows. Let $S \in \{$stable, preferred, grounded, complete$\}$. Then (In, Out) is an S-status assignment iff In is an S-extension.

To obtain a definition of the acceptability status of arguments a further refinement is necessary (here given for argument labellings):

Definition 2.5 [acceptability status of arguments] For grounded semantics an argument A is

1. *justified* iff A is *in* in the grounded status assignment;
2. *overruled* iff A is *out* in the grounded status assignment;
3. *defensible* iff A is undecided in the grounded status assignment;

For stable and preferred semantics an argument A is

1. *justified* iff A is *in* in all stable/preferred status assignments;
2. *overruled* iff A is *out* or undecided in all stable/preferred status assignments;
3. *defensible* iff A is *in* in some but not all stable/preferred status assignments.

Let us illustrate the definitions with some examples, where defeat relations are graphically depicted as arrows[1].

Example 2.6 (Reinstatement)

$$A \longleftarrow B \longleftarrow C$$

All semantics produce the same, unique extension, namely, $\{A, C\}$. Hence in all semantics A and C are justified while B is overruled. It is sometimes said that C *reinstates* A by defeating its defeater B.

Example 2.7 (Even defeat loop)

$$\begin{tikzcd} A \arrow[r, bend left] & B \arrow[l, bend left] \end{tikzcd}$$

The grounded extension is empty while the preferred-and-stable extensions are $\{A\}$ and $\{B\}$. All these extensions are also complete. Hence in all semantics both A and B are defensible.

The next example shows a difference between stable and preferred semantics.

[1]The pictures are copied from Prakken and Vreeswijk (2002).

Example 2.8 (Odd defeat loop)

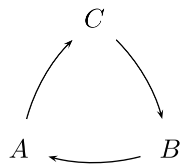

This example has no stable extensions while there is a unique grounded, preferred and complete extension, which is empty. Note that if a fourth argument D is added with no defeat relations with the other three arguments, there is still no stable extension while the unique grounded, preferred and complete extension is $\{D\}$.

The following example shows a difference between grounded and preferred semantics.

Example 2.9 Consider the arguments A, B, C and D such that A and B defeat each other, both A and B defeat C and C defeats D. The grounded extension is empty while the two preferred (and stable) extensions are $\{A, D\}$ and $\{B, D\}$. Thus while in grounded semantics all arguments are defensible, in preferred and stable semantics A and B are defensible, D is justified and C is overruled. The two corresponding preferred-and-stable status assignments are shown in the following figure:

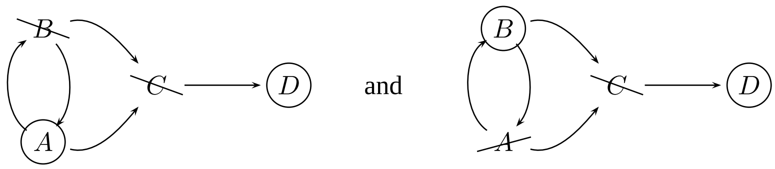

While some researchers give reasons why one semantics would be better than another, others argue that the choice of semantics may depend on the reasoning context and the nature of the knowledge involved. We will return to this issue in subsequent sections.

Dung's abstract approach has been further developed in various ways. To start with, other semantics have been proposed and investigated. For an overview see Baroni and Giacomin (2009). Furthermore, Amgoud and Cayrol (2002) have decomposed the defeat relation into a more basic *attack* relation, standing just for notions of syntactic conflict, and a binary *preference* relation on arguments. Argument A is then said to *defeat* argument B if A attacks B and B is not preferred to A. Modgil (2009) takes this a step further in allowing *attacks on attacks* in addition to attacks on arguments. Intuitively, if argument C claims that argument B is preferred to argument A, and A attacks B, then C undermines the success of A's attack on B (i.e., A does not *defeat* B) by pref-attacking A's attack on B. Since arguments attacking attacks can themselves be attacked, as can these attacks, and so on, Modgil's *extended argumentation frameworks* can fully model argumentation about whether an argument defeats another.

In the semantics of argumentation-based inference the main focus is on characterising properties of *sets* of arguments, without specifying procedures for determining whether a given argument is a member of the set. The proof theory of argumentation-based inference amounts to specifying such procedures. An elegant form of argumentation-based proof procedures is that of an *argument game* between two players, a proponent and an opponent of an argument. The exact rules of the game depend on the semantics the game is meant to capture. The rules should be chosen such that the existence of a winning strategy (in the usual game-theoretic sense) for the proponent of an argument corresponds to the investigated semantic status of the argument, for example, 'justified in grounded semantics' or 'defensible in preferred semantics'.

To give an example argument game, the following game is sound and complete for grounded semantics in that proponent has a winning strategy for argument A just in case A is in the grounded extension (Dung; 1994; Prakken and Sartor; 1997).

Definition 2.10 [argument game for grounded semantics] An argument game for grounded semantics between a proponent P and opponent O of an argument A_1 is a finite nonempty sequence of moves $move_i = (Player_i, A_i)$ $(i > 0)$, such that:

1. $Player_i = P$ iff i is odd; and $Player_i = O$ iff i is even;
2. $move_1 = (P, A_1)$;
3. If $Player_i = Player_j = P$ and $i \neq j$, then $A_i \neq A_j$;
4. If $Player_i = P$ and $i \neq 0$ then A_i defeats A_{i-1} while A_{i-1} does not defeat A_i;
5. If $Player_i = O$, then A_i defeats A_{i-1} .

A game is *terminated* if it cannot be extended with further moves. The player who moves last in a terminated game *wins* the game.

Informally, the proponent starts a game with an argument and then the players take turns, trying to defeat the previous move of the other player. In doing so, the proponent must strictly defeat the opponent's arguments while he is not allowed to repeat his own arguments. The winning rule of this game in fact says that the proponent has a winning strategy if he has a way to make the opponent run out of moves (from the implicitly assumed AF) whatever choice the opponent makes.

As remarked in the introduction, argumentation logics must define three things: how arguments can be constructed, how they can be defeated and how they can be defended against defeating counterarguments. Dung's abstract formalism only answers the third question. To answer the first two questions, accounts are needed of argument construction and the nature of defeat. We next discuss a general framework for formulating such accounts.

3 An abstract framework for structured argumentation

In the European ASPIC project (2004-2006) an abstract account was developed of how Dungean AFs can be generated from more basic information, building on earlier work of Vreeswijk (1993) and Pollock (1995) on the structure of arguments and of Pollock (1974; 1995) and others on the nature of defeat. The ASPIC framework assumes an unspecified logical language and defines arguments as inference trees formed by applying strict or defeasible inference rules, the nature of which is also unspecified. The notion of an argument as an inference tree leads to three ways of attacking an argument: attacking an inference, attacking a conclusion and attacking a premise. To resolve such conflicts, preferences may be used, which leads to three corresponding kinds of defeat. To characterise them, some minimal assumptions on the logical object language are made. First, a contrariness function is assumed on the object language, generalising classical negation as in Bondarenko et al. (1997). Second, defeasible inference rules are assumed to be named in the object language. Apart from these assumptions the framework is still abstract: it applies to any set of inference rules divided into strict and defeasible rules, and to any logical language that satisfies the above assumptions.

Below the latest version of the ASPIC framework is summarised, as described by Prakken (2010). The basic notion is that of an argumentation system.

Definition 3.1 [Argumentation system] An *argumentation system* is a tuple $AS = (\mathcal{L}, {}^{-}, \mathcal{R}, \leq)$ where

- $\mathcal{L}$ is a logical language.
- ${}^{-}$ is a contrariness function from $\mathcal{L}$ to $2^{\mathcal{L}}$, such that if $\varphi \in \overline{\psi}$ then if $\psi \notin \overline{\varphi}$ then φ is called a *contrary* of ψ, otherwise φ and ψ are called *contradictory*. The latter case is denoted by $\varphi = -\psi$ (i.e., $\varphi \in \overline{\psi}$ and $\psi \in \overline{\varphi}$).
- $\mathcal{R} = \mathcal{R}_s \cup \mathcal{R}_d$ is a set of strict ($\mathcal{R}_s$) and defeasible ($\mathcal{R}_d$) inference rules such that $\mathcal{R}_s \cap \mathcal{R}_d = \varnothing$.
- $\leq$ is a partial preorder on $\mathcal{R}_d$.

Henceforth, a set $S \subseteq \mathcal{L}$ is said to be *directly consistent* iff $\nexists\, \psi, \varphi \in S$ such that $\psi \in \overline{\varphi}$, otherwise it is *directly inconsistent*. And S is said to be *indirectly (in)consistent* if its closure under application of strict inference rules is directly (in)consistent.

Arguments are built by applying inference rules to one or more elements of $\mathcal{L}$. Strict rules are of the form $\varphi_1, \ldots, \varphi_n \rightarrow \varphi$, interpreted as 'if the antecedents $\varphi_1, \ldots, \varphi_n$ are acceptable, then the consequent φ must be accepted, no matter what'. Defeasible rules are written as $\varphi_1, \ldots, \varphi_n \Rightarrow \varphi$, meaning 'if the antecedents are acceptable, then the consequent must be accepted if there is no good reason not to accept it'. As is usual in logic, inference rules can be specified by schemes in which a rule's antecedents and consequent are metavariables ranging over $\mathcal{L}$.

Arguments are constructed from a knowledge base, which is assumed to contain three kinds of formulas.

Definition 3.2 [Knowledge bases] A *knowledge base* in an argumentation system $(\mathcal{L}, ^{-}, \mathcal{R}, \leq)$ is a pair $(\mathcal{K}, \leq')$ where $\mathcal{K} \subseteq \mathcal{L}$ and $\leq'$ is a partial preorder on $\mathcal{K} \setminus K_n$. Here, $\mathcal{K} = \mathcal{K}_n \cup \mathcal{K}_p \cup \mathcal{K}_a$, the *necessary*, *ordinary* and *assumption* premises, where these subsets of $\mathcal{K}$ are disjoint .

Intuitively, arguments can only be attacked on their ordinary and assumption premises. Attacks on assumption premises always result in defeat while attacks on ordinary premises are resolved by using preferences.

Arguments can be constructed step-by-step from knowledge bases by chaining inference rules into trees. Arguments thus contain subarguments, which are the structures that support intermediate conclusions (plus the argument itself and its premises as limiting cases). In what follows, for a given argument the function `Prem` returns all its premises, `Conc` returns its conclusion, `Sub` returns all its sub-arguments and `DefRules` returns all defeasible rules of the argument.

Definition 3.3 [Argument] An *argument* A on the basis of a knowledge base $(\mathcal{K}, \leq')$ in an argumentation system $(\mathcal{L}, ^{-}, \mathcal{R}, \leq)$ is:

1. φ if $\varphi \in \mathcal{K}$ with: $\texttt{Prem}(A) = \{\varphi\}$; $\texttt{Conc}(A) = \varphi$; $\texttt{Sub}(A) = \{\varphi\}$; $\texttt{DefRules}(A)$ = $\varnothing$.
2. $A_1, \ldots A_n \rightarrow/\Rightarrow \psi$ if $A_1, \ldots, A_n$ are arguments such that there exists a strict/defeasible rule $\texttt{Conc}(A_1), \ldots, \texttt{Conc}(A_n) \rightarrow/\Rightarrow \psi$ in $\mathcal{R}_s/\mathcal{R}_d$.
   $\texttt{Prem}(A) = \texttt{Prem}(A_1) \cup \ldots \cup \texttt{Prem}(A_n)$,
   $\texttt{Conc}(A) = \psi$,
   $\texttt{Sub}(A) = \texttt{Sub}(A_1) \cup \ldots \cup \texttt{Sub}(A_n) \cup \{A\}$.
   $\texttt{DefRules}(A) = \texttt{DefRules}(A_1) \cup \ldots \cup \texttt{DefRules}(A_n)$.

Then A is: *strict* if $\texttt{DefRules}(A) = \varnothing$; *defeasible* if $\texttt{DefRules}(A) \neq \varnothing$; *firm* if $\texttt{Prem}(A) \subseteq \mathcal{K}_n$; *plausible* if $\texttt{Prem}(A) \not\subseteq \mathcal{K}_n$.

Example 3.4 Consider a knowledge base in an argumentation system with

$$\mathcal{R}_s = \{p, q \rightarrow s;\ u, v \rightarrow w\}; \mathcal{R}_d = \{p \Rightarrow t;\ s, r, t \Rightarrow v\}$$
$$\mathcal{K}_n = \{q\}; \mathcal{K}_p = \{p, u\}; \mathcal{K}_a = \{r\}$$

An argument for w is displayed in Figure 1. The type of a premise is indicated with a superscript and defeasible inferences are displayed with dotted lines. Formally the argument and its subarguments are written as follows:

A_1: p	A_5: $A_1 \Rightarrow t$
A_2: q	A_6: $A_1, A_2 \rightarrow s$
A_3: r	A_7: $A_5, A_3, A_6 \Rightarrow v$
A_4: u	A_8: $A_7, A_4 \rightarrow w$

We have that

$\texttt{Prem}(A_8) =$	$\{p, q, r, u\}$
$\texttt{Conc}(A_8) =$	w
$\texttt{Sub}(A_8) =$	$\{A_1, A_2, A_3, A_4, A_5, A_6, A_7, A_8\}$
$\texttt{DefRules}(A_8) =$	$\{p \Rightarrow t;\ s, r, t \Rightarrow v\}$

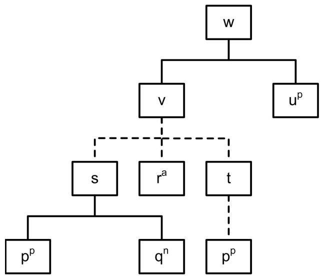

Figure 1: An argument

Combining an argumentation system and a knowledge base with an *argument ordering* results in an *argumentation theory*. The argument ordering is a partial preorder $\preceq$ on arguments (with its strict counterpart $\prec$ defined in the usual way), and is assumed to be 'admissible', i.e., firm-and-strict arguments are strictly better than all other arguments, and a strict inference cannot make an argument strictly better or worse than its weakest proper subargument. The argument ordering can but needs not be defined in terms of the orderings on $\mathcal{R}_d$ and $\mathcal{K} \setminus K_n$.

Definition 3.5 [Argumentation theories] An *argumentation theory* is a triple $AT = (AS, KB, \preceq)$ where AS is an argumentation system, KB is a knowledge base in AS and $\preceq$, a partial preorder, is an admissible ordering on the set of all arguments that can be constructed from KB in AS.

Informally, an argument ordering is admissible if it makes all strict-and-firm arguments strictly preferred over all other arguments and if strict inferences cannot make an argument weaker or stronger.

As indicated above, when arguments are inference trees, three syntactic forms of attack are possible: attacking a premise, a conclusion, or an inference.

Definition 3.6 [Attacks]

- Argument A *undercuts* argument B (on B') iff $\texttt{Conc}(A) \in \overline{B'}$ for some $B' \in \texttt{Sub}(B)$ of the form $B''_1, \ldots, B''_n \Rightarrow \psi$.[2]
- Argument A *rebuts* argument B on (B') iff $\texttt{Conc}(A) \in \overline{\varphi}$ for some $B' \in \texttt{Sub}(B)$ of the form $B''_1, \ldots, B''_n \Rightarrow \varphi$. In such a case A *contrary-rebuts* B iff $\texttt{Conc}(A)$ is a contrary of φ.
- Argument A *undermines* B (on φ) iff $\texttt{Conc}(A) \in \overline{\varphi}$ for some $\varphi \in \texttt{Prem}(B) \setminus \mathcal{K}_n$. In such a case A *contrary-undermines* B iff $\texttt{Conc}(A)$ is a contrary of φ or if $\varphi \in \mathcal{K}_a$.

[2]Here an unspecified method is assumed to name defeasible inferences in the object language.

In Example 3.4 argument A_8 can be undercut in two ways: by an argument with conclusion $\overline{A_5}$, which undercuts A_8 on A_5, and by an argument with conclusion $\overline{A_7}$, which undercuts A_8 on A_7. Moreover, argument A_8 can be rebutted on A_5 with an argument for $\bar{t}$ and on A_7 with an argument for $\overline{v}$. Moreover, if $\bar{t} = -t$ and the rebuttal has a defeasible top rule, then A_5 in turn rebuts the argument for $\bar{t}$. However, A_8 itself does not rebut that argument, except in the special case where $w \in \bar{\bar{t}}$. Finally, argument A_8 can be undermined with an argument that has conclusion $\overline{p}$, $\overline{r}$ or $\overline{u}$.

Attacks combined with the preferences defined by an argument ordering yield three kinds of defeat. For undercutting attack no preferences are needed to make it succeed, since otherwise a weaker undercutter and its stronger target might be in the same extension. The same holds for the other two ways of attack as far as they involve contraries (i.e., non-symmetric conflict relations between formulas).

Definition 3.7 [Successful rebuttal, undermining and defeat]

- A *successfully rebuts* B if A rebuts B on B' and either A contrary-rebuts B' or $A \nprec B'$.
- A *successfully undermines* B if A undermines B on φ and either A contrary-undermines B or $A \nprec \varphi$.
- A *defeats* B iff A undercuts or successfully rebuts or successfully undermines B.

The success of rebutting and undermining attacks thus involves comparing the conflicting arguments at the points where they conflict. The definition of successful undermining exploits the fact that an argument premise is also a subargument.

Recall that argumentation logics must define three things: how arguments can be constructed, how they can be defeated and how they can be defended against defeating counterarguments. While Dung's abstract argumentation semantics addresses the last issue, we can now combine it with the ASPIC framework to address the first two issues and obtain a general framework for the definition of argumentation logics. More precisely, argumentation theories generate Dungean AFs as follows:

Definition 3.8 An *abstract argumentation framework* AF_{AT} *corresponding to an argumentation theory* AT is a pair $\langle \mathcal{A}, Def \rangle$ such that $\mathcal{A}$ is the set of arguments defined by AT as in Definition 3.3, and Def is the relation on $\mathcal{A}$ given by Definition 3.7.

Then any semantics for Dung frameworks can be used to define the acceptability status of arguments. This in turn enables a definition of a consequence notion for well-formed formulas. Several definitions are possible. One is:

Definition 3.9 [Acceptability of conclusions] For any semantics S and for any argumentation theory AT and formula $\varphi \in \mathcal{L}_{AT}$:

1. φ is *S-justified* in AT if and only if all S-extensions of AT contain an argument with conclusion φ;
2. φ is *S-defensible* in AT if and only if there exists an S-extension of AT that contains an argument with conclusion φ.

An alternative definition of S-justification is

1. φ is *S-justified* in AT if and only if there exists an argument with conclusion φ that is contained in all S-extensions of AT.

While the original definition allows that different extensions contain different arguments for a justified conclusion, the alternative definition requires that there is one argument for it that is in all extensions. The significance of this difference is illustrated by the following example (which is a structured counterpart of Example 2.9).

Example 3.10 Assume that people who are born in the Netherlands are usually Dutch, people with a Norwegian name are usually Norwegian and that both Dutch and Norwegians like ice skating. Assume furthermore that nobody can be both Dutch and Norwegian and that Brygt Rykkje was born in the Netherlands and has a Norwegian name. The following argumentation theory formalises this example. Here $\mathcal{R}_s$ consists of all classically valid inferences while $\mathcal{R}_d$ contains a modus ponens rule for a connective $\rightsquigarrow$ for defeasible conditionals in $\mathcal{L}$. Next, $\mathcal{K}_p$ consists of:

$$\forall x(\texttt{BornInNL}(x) \rightsquigarrow \texttt{Dutch}(x))$$
$$\forall x(\texttt{NorwegianName}(x) \rightsquigarrow \texttt{Norwegian}(x))$$
$$\forall x((\texttt{Dutch}(x) \vee \texttt{Norwegian}(x)) \rightsquigarrow \texttt{LikesIceSkating}(x))$$
$$\texttt{BorninNL}(b)$$
$$\texttt{NorwegianName}(b)$$
$$\forall x \neg(\texttt{Dutch}(x) \wedge \texttt{Norwegian}(x))$$

We leave it to the reader to verify the following analysis. First, distinct arguments can be constructed for the following conclusions

A_1 for $\texttt{Dutch}(b)$	A_2 for $\texttt{LikesIceSkating}(b)$
B_1 for $\texttt{Norwegian}(b)$	B_2 for $\texttt{LikesIceSkating}(b)$

such that A_1 is a subargument of A_2 and B_1 is a subargument of B_2. Second, if all arguments are of equal strength then the grounded extension contains neither of these arguments while there are two preferred-and-stable extensions, one which contains A_1 and A_2 but not B_1 or B_2 and another which contains B_1 and B_2 but not A_1 or A_2. Hence in grounded semantics the conclusion $\texttt{LikesIceSkating}(b)$ is defensible while in preferred and stable semantics it is defensible according to the first definition of S-justification but justified according to the alternative definition.

One possible analysis of such examples is that some definitions of justification are better than others. Another analysis is that different semantic definitions capture different senses or strengths of justification, which each may have their use in certain contexts.

4 The nature of inference rules

While we now have a general framework for the definition of argumentation logics, much more can be said. To start with, the framework can be instantiated in many ways, so there is a need for principles that can be used in assessing the quality of instantiations. Caminada and Amgoud (2007) formulated several so-called rationality postulates, namely, that each extension should be closed under subarguments and under strict rule application, and directly and indirectly consistent. Prakken (2010) identifies some broad classes of instantiations of the ASPIC framework that satisfy these postulates.

The next question is, what are 'good' collections of strict and defeasible inference rules? In AI there is a tradition to let inference rules express domain-specific information, for example, Reiter's (1980) default logic, Prakken and Sartor's (1997) system based on extended logic programming and many applications of Bondarenko et al.'s (1997) assumption-based argumentation. This runs counter to the usual practice in logic, in which inference rules express general patterns of reasoning, such as modus ponens, universal instantiation and so on. More in line with this practice are logics for so-called classical argumentation, studied by e.g. Besnard and Hunter (2008). These logics are in fact a special case of the ASPIC framework with $\mathcal{L}$ being the language of standard propositional or first-order logic (or some other deductive logic), the contrary relation conforming to classical negation, the strict rules being all valid propositional or first-order inferences (or of some other deductive logic), and with the additional requirement that the premises of an argument are classically consistent. In these logics arguments can thus only be attacked on their premises.

The last observation indicates that within the ASPIC framework deductive logics (in the Tarskian sense) model the special case of arguments that can only be attacked on their premises. This also illuminates a distinction that is sometimes made between *plausible* and *defeasible* reasoning; cf. Rescher (1976) and Vreeswijk (1993, Ch. 8). Vreeswijk describes plausible reasoning as sound (i.e, deductive) reasoning on an uncertain basis and defeasible reasoning as unsound (but still rational) reasoning on a solid basis. We now see that ASPIC argumentation theories with only strict inference rules formalise plausible reasoning while theories that include defeasible inference rules and only have necessary premises formalise defeasible reasoning. The full ASPIC framework gives a unified account of these two kinds of reasoning.

While this answers what are 'good' strict inference rules, an answer to the same question for defeasible rules can be given by combining the full ASPIC framework with the idea that inference rules should express general patterns of reasoning. This can clarify Pollock's (1974; 1995) notion of *prima facie reasons* and argumentation-theory's notion of *argument schemes* (Walton et al.; 2008). Pollock's prima facie reasons are general patterns of epistemic defeasible reasoning. He formalised reasons for perception, memory, induction, temporal persistence and the statistical syllogism, as well as undercutters for these reasons. In the ASPIC framework prima facie reasons can be expressed as schemes (in the logical sense, with metavariables ranging over $\mathcal{L}$) for defeasible inference rules.

The difference between domain-specific and general defeasible inference rules is illustrated with the following example. Consider the information that all Frisians are Dutch, that the Dutch are usually tall and that Wiebe is Frisian. With domain-specific inference rules this can in a propositional language be represented as follows:

$$
\begin{array}{lcl}
\mathcal{R}_s & = & \{Frisian \rightarrow Dutch\} \\
\mathcal{R}_d & = & \{Dutch \Rightarrow Tall\} \\
\mathcal{K}_p & = & \{Frisian\}
\end{array}
$$

The argument that Wiebe is tall then has the form as displayed on the left in Figure 2.

With general inference rules the two rules must instead be represented in the object language $\mathcal{L}$. The first one can be represented with the material implication but for the second one a connective for defeasible conditionals must be added to $\mathcal{L}$ and a defeasible modus-ponens inference rule must be added for this connective. For example:

$$
\begin{array}{lcl}
\mathcal{R}_s & = & \{\varphi, \varphi \supset \psi \rightarrow \psi \text{ (for all } \varphi, \psi \in \mathcal{L}), \ldots\} \\
\mathcal{R}_d & = & \{\varphi, \varphi \rightsquigarrow \psi \Rightarrow \psi \text{ (for all } \varphi, \psi \in \mathcal{L}), \ldots\} \\
\mathcal{K}_p & = & \{Frisian \supset Dutch, Dutch \rightsquigarrow Tall, Frisian\}
\end{array}
$$

Then the argument that Wiebe is tall has the form as displayed on the right in Figure 2.

$$
\dfrac{\dfrac{\textit{Frisian}}{\textit{Dutch}}}{\textit{Tall}}
\qquad\qquad
\dfrac{\dfrac{\textit{Frisian} \qquad \textit{Frisian} \supset \textit{Dutch}}{\textit{Dutch}} \qquad \textit{Dutch} \sim> \textit{Tall}}{\textit{Tall}}
$$

Figure 2: Domain-specific vs. general inference rules

The same analysis applies to argument schemes, which are taken to be stereotypical non-deductive patterns of reasoning. Uses of argument schemes are evaluated in terms of critical questions specific to the scheme. In the literature on argumentation theory many collections of argument schemes have been proposed, both for epistemic and for practical reasoning. An example of an epistemic argument scheme is the scheme from expert opinion (Walton et al.; 2008, p. 310):

E is an expert in domain D
E asserts that P is true
P is within D

P is true

This scheme has six critical questions:

1. How credible is E as an expert source?
2. Is E an expert in domain D?
3. What did E assert that implies P?
4. Is E personally reliable as a source?
5. Is P consistent with what other experts assert?
6. Is E's assertion of P based on evidence?

A practical argument scheme is the scheme from good (bad) consequences (here in a formulation that deviates from Walton et al. (2008) to stress its abductive nature):

$$\frac{\begin{array}{l}\text{Action } A \text{ results in } P\\ P \text{ is good (bad)}\end{array}}{A \text{ should (not) be done.}}$$

This scheme has three critical questions:

1. Does A result in P?
2. Does A also result in something which is bad (good)?
3. Is there another way to realise P?

Argument schemes can also be formalised as schemes for defeasible inference rules; then critical questions can be regarded as pointers to counterarguments. Some critical questions challenge an argument's premise and therefore point to undermining attacks, others point to undercutting attacks, while again other questions point to rebutting attacks. In the scheme from expert opinion questions (2) and (3) point to underminers (of, respectively, the first and second premise), questions (4), (1) and (6) point to undercutters (the exceptions that the expert is biased or incredible for other reasons and that he makes scientifically unfounded statements) while question (5) points to rebutting applications of the expert opinion scheme. In the scheme from good (bad) consequences question (1) points to underminers of the first premise, question (2) points to rebuttals using the opposite version of the scheme while question (3) points to undercutters. Thus we also see that Pollock's prima facie reasons are examples of epistemic argument schemes and that his undercutters are negative answers to critical questions.

This account of argument schemes can also clarify Toulmin's (1958) distinction between warrants (rule-like premises) and backings of warrants. For example, a warrant can be that smoking causes cancer while its backing can be an expert opinion or a scientific study. In fact, several argument schemes studied in the literature are for source-based reasoning (such as the above scheme for expert opinion and the witness testimony scheme), and the account of argument schemes proposed here formalises such reasoning about the backing of warrants.

The distinction between epistemic and practical reasoning can shed some light on the issue of which consequence notion is the best. If, for instance, the scheme from good consequences can be applied to two incompatible actions (say reducing and increasing taxes) for two different good consequences (say increasing productivity and increasing equality) and there is no reason to prefer one consequence over the other, then an arbitrary choice is rational. If, on the other hand, two experts disagree about whether reducing taxes increases productivity, then an arbitrary choice for one of them seems irrational. So it might be argued that in practical reasoning a defensible argument can be good enough while in epistemic reasoning we should aim for justified arguments or conclusions.

One question about the Wiebe example remains: what is the 'logic' of the $\rightsquigarrow$ connective, that is, which inference rules other than defeasible modus ponens apply to it? The ASPIC framework abstracts from this issue: it may be that a suitable model-theoretic semantics of $\leadsto$ (Kraus et al.; 1990; Pearl; 1992) generates suitable sets of strict and defeasible inference rules. However, the literature on argument schemes reveals that often another way of reasoning about defaults is more relevant, namely, whether a default is based on an adequate epistemic or authoritive source (such as an expert, a witness, a scientific study, a statute). Toulmin (1958) was perhaps the first to highlight this difference, with his notion of backings for warrants. In fact, many argument schemes studied in the literature are for source-based reasoning, and the account of argument schemes proposed here formalises such reasoning about the sources of defaults.

5 Argumentation as a form of dialogue

As stated in the introduction, argumentation-theorists often claim that arguments can only be evaluated in the context of a dialogue or procedure. More specifically, Walton (1996) regards argument schemes as dialogical devices, determining dialectical obligations and burdens of proof. An argument is a move in a dialogue and the scheme that it instantiates determines the allowed and required responses to that move. More precisely (and in present terms), asking whether a premise is true creates a burden on the other side to back the premise with further grounds, while asking questions that point to rebuttals or undercutters does not shift the burden back to the other side: instead, the one who asks such a question must back it up with some evidence as to why the exception would be true. Only if such evidence is provided, the burden of proof shifts back to the proponent of the original argument. At first sight, our account of argument schemes as defeasible inference rules would seem to be incompatible with Walton's dialogical account. However, these two accounts can be reconciled by embedding argumentation logics in dialogue systems for argumentation.

While argumentation logics define notions of consequence from a given body of information, dialogue systems for argumentation (Walton and Krabbe; 1995) regulate disputes between real agents, who each have their own body of information, and who may be willing to learn from each other so that their information state may change. Moreover, during the dialogue they construct a joint theory on the issue in dispute, which also evolves over time. Essentially, dialogue systems define a communication language (the well-formed utterances) and a protocol (when a well-formed utterance may be made and when the dialogue terminates). The logical argument games described in Section 2 cannot be such a dialogue system, for two reasons. First, these games assume a single and fixed body of information, so they do not apply to contexts with distributed and possibly changing bodies of information. Moreover, in argumentation dialogues other utterances can be made than just stating arguments.

Consider the following simple example, with a dialogue system that allows players to move arguments and to challenge, concede or retract premises and conclusions of these arguments. Each challenge must be answered with a ground for the challenged statement or else the statement must be retracted. The two agents have their own ASPIC argumentation theory in a shared ASPIC argumentation system with a propositional language and three defeasible inference rules: $p \Rightarrow q$, $r \Rightarrow p$ and $s \Rightarrow \neg r$. Paul's and

Olga's knowledge bases contain, respectively, single ordinary premises p and r. Let us assume that all arguments are of equal preference. Paul wants to persuade Olga that q is the case. He can internally construct the following argument for q:

A_1: r $\qquad$ A_2: $A_1 \Rightarrow p$ $\qquad$ A_3: $A_2 \Rightarrow q$

However, a well-known argumentation heuristic says that arguments in dialogue should be made as sparse as possible in order to avoid attacks. Therefore, Paul only utters the last step in the argument, hoping that Olga will accept p so that Paul does not have to defend r. This leads to the following dialogue.

P_1:	q *since* p	O_1:	*why* p
P_2:	p *since* r	O_2:	$\neg r$ *since* s
P_3:	*retract* r		
P_4:	*retract* q		

What has happened here? If Olga had been a trusting person who concedes a statement if she cannot construct an argument for the opposite, then she would have conceded p and q after P_1. However, q is not a justified conclusion from the joint knowledge bases, so this outcome is undesirable. In fact, Olga was less trusting and first asked Paul for his reasons for p. Since Paul was honest, he gave his true reasons, which allowed Olga to discover that she could attack Paul with an undermining counterargument. Paul could not defend himself against this attack, so he realised that he cannot persuade Olga that q is true; he therefore retracted r and q, after which the dialogue terminated.

Why moves are also relevant in legal contexts. For example, Dutch civil procedure combines a silence-implies-consent principle with burdens of proof: normally, plaintiff must prove his claims but this proof burden only becomes effective after defendant has challenged this claim, otherwise the judge must accept it as true.

Argumentation logic applies here in several ways. It can model the agents' internal reasoning but it can also be applied at each dialogue stage to the joint theory that the agents have created at that stage. For example, after O_2 the logic says that q is overruled on the basis of $\mathcal{K} = \mathcal{K}_p = \{p, r, s\}$ while after P_4 the logic says that no argument for q can be constructed on the basis of $\mathcal{K} = \mathcal{K}_p = \{p, s\}$. Argumentation logic can also be used as a component of notions of soundness and completeness of protocols, such as:

- A protocol is *sound* if whenever at termination p is accepted, p is justified by the participants' joint knowledge bases.
- A protocol is *weakly* complete if whenever p is justified by the participants' joint knowledge bases, there is a legal dialogue at which at termination p is accepted.
- A protocol is *strongly* complete if whenever p is justified by the participants' joint knowledge bases, all legal dialogues terminate with acceptance of p.

Similar notions can be defined relative to the joint theory constructed during a dialogue, while the notions can also be made conditional on particular agent strategies and heuristics (for example, a protocol could be sound and complete on the condition that all agents are honest but not trusting). For an overview of current research on these issues and other desirable properties of dialogue systems see Prakken (2006) and several chapters in Rahwan and Simari (2009).

We can now without giving up the idea of an argumentation logic make sense of the argumentation-theorists' claim that arguments should be evaluated in the context of a dialogue or procedure. The dialogue provides the relevant statements and arguments at each stage of the dialogue. The logic then determines the justified arguments at that stage. The logic also points at the importance of investigation. Since arguments can be defeated by counterarguments, the process of searching for information that gives rise to counterarguments is an essential part of testing an argument's viability: the harder and more thorough this search has been, the more confident we can be that an argument is justified if we cannot find defeaters. The ultimate justification of an argument is then determined by applying the logic to the final information state. Thus the ultimate justification of an argument depends on both logic and dialogue, or more generally on both logic and investigation.

On this account the critical questions of argument schemes have a dual role. On the one hand they define possible counterarguments to arguments constructed with the scheme (logic) while on the other hand they point at investigations that could be done to find such counterarguments (dialogue and procedure).

The combined logical/dialogical account of argumentation can also clarify notions of burden of proof, especially as they are used in the law. For space limitations the reader is referred to Prakken and Sartor (2009) for more on this issue.

This account also gives a second explanation why argument evaluation is context-dependent, besides the fact that different domains may have different sets of accepted argument schemes. The second explanation is that different contexts may require different protocols for dialogue: when a decision has to be reached in reasonable time (as in legal proceedings or a business meeting), a protocol may be more restrictive than in settings like academic debate. For example, the possibility to give alternative replies to a move may be restricted so that agents are forced to think what is their best reply.

Finally, on this account persuasiveness of arguments can be modelled as follows. Each dialogical agent has an internal argumentation theory and evaluates incoming arguments in terms of how they fit with the AF that it can internally generate. Given an *acceptance attitude* the agent will either accept the argument's premises and/or conclusion, or attack it with a counterargument, or ask for further grounds for a premise. Personality models can help modelling which types of arguments an agent of a certain type tends to accept. This gives a third way in which argument evaluation is context-dependent: the persuasive force of an argument depends on the listener. Current work of this kind is still preliminary but fascinating and promising (see e.g. the proceedings of the annual *ArgMas* workshops on argumentation in multi-agent systems). In fact this work provides a formal or even computational account of Perelman's New Rhetoric (Perelman and Olbrechts-Tyteca; 1969).

6 Conclusion

In this paper we discussed four philosophical problems concerning argumentation, with the aim to show how formal methods can be used to clarify them. We first showed how formal standards for argumentation-based inference can be developed, by presenting an abstract framework for argument evaluation given a set of arguments and their defeat relations, and by supplementing it with accounts of argument construction and the nature of defeat. We then clarified how a dialogical account of argument evaluation can be given in formal terms, by discussing the embedding of argumentation logics in dialogue systems for argumentation. This embedding also clarified how reasoning with argument schemes can be formalised: argument schemes are defeasible inference rules and their critical questions point at counterarguments. We also clarified how the use of arguments to persuade can be formalised, by adding the notions of argumentation strategies and heuristics and suggesting the use of personality models of argumentative agents. We then gave three reasons why argument evaluation is context-dependent: different domains may have different sets of argument schemes, different contexts may require more or less strict protocols for dialogue and the persuasive force of arguments depends on the listener.

References

Amgoud, L. and Cayrol, C. (2002). Inferring from inconsistency in preference-based argumentation frameworks, *International Journal of Automated Reasoning* **29**: 125–169.

Antoniou, G. (1997). *Nonmonotonic Reasoning*, MIT Press, Cambridge, MA.

Baroni, P. and Giacomin, M. (2009). Semantics of abstract argument systems, *in* I. Rahwan and G. Simari (eds), *Argumentation in Artificial Intelligence*, Springer, Berlin, pp. 25–44.

Besnard, P. and Hunter, A. (2008). *Elements of Argumentation*, MIT Press, Cambridge, MA.

Bondarenko, A., Dung, P., Kowalski, R. and Toni, F. (1997). An abstract, argumentation-theoretic approach to default reasoning, *Artificial Intelligence* **93**: 63–101.

Caminada, M. (2006). On the issue of reinstatement in argumentation, *Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA 2006)*, number 4160 in *Springer Lecture Notes in AI*, Springer Verlag, Berlin, pp. 111–123.

Caminada, M. and Amgoud, L. (2007). On the evaluation of argumentation formalisms, *Artificial Intelligence* **171**: 286–310.

Dung, P. (1994). Logic programming as dialog game. Unpublished paper, Division of Computer Science, Asian Institute of Technology, Bangkok.

Dung, P. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and *n*–person games, *Artificial Intelligence* **77**: 321–357.

Jakobovits, H. and Vermeir, D. (1999). Robust semantics for argumentation frameworks, *Journal of Logic and Computation* **9**: 215–261.

Kraus, S., Lehmann, D. and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence* **44**: 167–207.

Modgil, S. (2009). Reasoning about preferences in argumentation frameworks, *Artificial Intelligence* **173**: 901–934.

Pearl, J. (1992). Epsilon-semantics, *in* S. Shapiro (ed.), *Encyclopedia of Artificial Intelligence*, Wiley, pp. 468–475.

Perelman, C. and Olbrechts-Tyteca, L. (1969). *The New Rhetoric. A Treatise on Argumentation*, University of Notre Dame Press, Notre Dame, Indiana.

Pollock, J. (1974). *Knowledge and Justification*, Princeton University Press, Princeton.

Pollock, J. (1995). *Cognitive Carpentry. A Blueprint for How to Build a Person*, MIT Press, Cambridge, MA.

Prakken, H. (2006). Formal systems for persuasion dialogue, *The Knowledge Engineering Review* **21**: 163–188.

Prakken, H. (2010). An abstract framework for argumentation with structured arguments, *Argument and Computation* **1**: 93–124.

Prakken, H. and Sartor, G. (1997). Argument-based extended logic programming with defeasible priorities, *Journal of Applied Non-classical Logics* **7**: 25–75.

Prakken, H. and Sartor, G. (2009). A logical analysis of burdens of proof, *in* H. Kaptein, H. Prakken and B. Verheij (eds), *Legal Evidence and Proof: Statistics, Stories, Logic*, Ashgate Publishing, Farnham, pp. 223–253.

Prakken, H. and Vreeswijk, G. (2002). Logics for defeasible argumentation, *in* D. Gabbay and F. Günthner (eds), *Handbook of Philosophical Logic*, second edn, Vol. 4, Kluwer Academic Publishers, Dordrecht/Boston/London, pp. 219–318.

Rahwan, I. and Simari, G. (eds) (2009). *Argumentation in Artificial Intelligence*, Springer, Berlin.

Reiter, R. (1980). A logic for default reasoning, *Artificial Intelligence* **13**: 81–132.

Rescher, N. (1976). *Plausible Reasoning*, Van Gorcum, Assen.

Rescher, N. (1977). *Dialectics: a Controversy-oriented Approach to the Theory of Knowledge*, State University of New York Press, Albany, N.Y.

Toulmin, S. (1958). *The Uses of Argument*, Cambridge University Press, Cambridge.

Verheij, B. (1996). Two approaches to dialectical argumentation: admissible sets and argumentation stages, *Proceedings of the Eighth Dutch Conference on Artificial Intelligence (NAIC-96)*, Utrecht, The Netherlands, pp. 357–368.

Vreeswijk, G. (1993). *Studies in Defeasible Argumentation*, Doctoral dissertation Free University Amsterdam.

Walton, D. (1996). *Argumentation Schemes for Presumptive Reasoning*, Lawrence Erlbaum Associates, Mahwah, NJ.

Walton, D. and Krabbe, E. (1995). *Commitment in Dialogue. Basic Concepts of Interpersonal Reasoning*, State University of New York Press, Albany, NY.

Walton, D., Reed, C. and Macagno, F. (2008). *Argumentation Schemes*, Cambridge University Press, Cambridge.