Chapter 1

Introduction

The subject of this thesis is argumentation. We consider argumentation as a process in which arguments supporting a conclusion are taken into account. During the process of argumentation, a conclusion originally justified by some argument can become unjustified. This is the result of the *defeasibility* of arguments.[1] Our central theme is how argumentation and the defeasibility of arguments can be formally modeled.

In this chapter, we first introduce the central concepts used throughout the thesis: argumentation (section 1), arguments (sections 2 and 3) and defeasibility (section 4). Then the research of the thesis is introduced. We put our research in perspective by giving a brief survey of recent related research (section 5), and by explaining our general aims and biases (section 6). The introductory chapter concludes with the research goals and method (section 7), and an outline of the thesis (section 8).

1 The process of argumentation

Argumentation is a process.[2] Its purpose is to justify conclusions (see, e.g., Pollock, 1987). Which conclusions are justified changes during the argumentation process. For instance, let us consider a story about John. It starts as follows.

> John is going to work and notices that it is still freezing. He sees some people skating on the lake that he passes each day, and he realizes that the ice is finally thick enough. After his arrival at the office, he notices that his colleague Mary is not there, and wonders whether she has taken a day off. Later that morning, he meets Harry at the coffee machine. Harry tells John that whenever the ice is

---

1 The term 'defeasibility' was introduced by Hart in 1948 (cf. Loui, 1995a).

2 This is an old idea in philosophy and can already be found in the work of Aristotle (cf. Rescher, 1977). Recently, the importance of process for argumentation has been reemphasized, e.g., by Loui (1992), Vreeswijk (1993) and Lodder (1996).

> thick enough, Mary takes a day off to go skating. John concludes that Mary has indeed taken a day off.

This first part of our story about John shows the process character of argumentation. Which conclusions are justified depends on the stage of the argumentation. For instance, only at the end of the story does John consider the conclusion that Mary has taken a day off to be justified. First, John notices that the ice is thick enough. Later this turns out to support his conclusion. Then he considers the question whether Mary has taken a day off. However, at that moment he does not connect this with the state of the ice. He makes the connection after Harry tells him that whenever the ice is thick enough, Mary takes a day off to go skating. Only then does he conclude that Mary has taken a day off.

Justified conclusions can be used in argumentation to support new conclusions, as we can see in the following continuation of our story about John.

> John returns to his desk unhappily. He knows that each time Mary takes a day off, he cannot finish his work. An hour later Mary's boss, Anne, passes by. She asks how things are going. John reluctantly tells her that he will not be able to finish his work, and explains why.

The conclusion that Mary has taken a day off is used to support the conclusion that John cannot finish his work.

The story shows that during the process of argumentation new conclusions become justified. New conclusions are justified by already available information, or by new information. But this is not all: not only can new conclusions become justified, but also old conclusions can become unjustified. This is shown in our story, which ends as follows.

> After John has told Anne that he will not be able to finish his work, she laughs. Anne says that she has forbidden Mary to take a day off. Mary is not at her desk because of a meeting.

Now that John knows that Anne has forbidden Mary to take a day off, his conclusions that Mary has taken a day off, and that he cannot finish his work, are no longer justified.

The story about John shows that justified conclusions are not necessarily true conclusions. John's conclusion that Mary had taken a day off was justified, but false. Falsity of justified conclusions can have two causes. First, it can be due to false premises, and second it can result from a lack of relevant information. One example is the following argument:

> Amsterdam is the capital of Denmark.
> So, the Danish government resides in Amsterdam.

Even though the government of a country usually resides in the capital of the country, this argument has a false conclusion, because its premise is false. Another example is the argument

> Amsterdam is the capital of the Netherlands.
> So, the Dutch government resides in Amsterdam.

This second argument has a false conclusion because the information is lacking that in the Netherlands the government of the country does not reside in the capital of the country, in contrast with the usual situation.

Which conclusions are justified depends on the information taken into account at the stage of the argumentation process. As a result, justified conclusions are at best an approximation of true conclusions. Normally, the approximation becomes better, as the justified conclusions are at successive stages based on more information.

2 Arguments as reconstructions

Central in argumentation are arguments which are used to justify conclusions. Conclusions become justified if they are supported by arguments. An example of an argument is the following:

> The ice is thick enough for skating. Whenever the ice is thick enough for skating, Mary takes a day off to go skating.
> So, Mary takes a day off to go skating.

In our story, John does not explicitly give an argument that Mary has taken a day off. However, if he were asked why he thought that Mary had taken a day off, he would give an argument similar to the one above. The argument is a reconstruction of how John arrived at the conclusion that Mary has taken a day off.

The structure of the example argument is depicted in Figure 1. The conclusion is supported by an argument that consists of one step from the premises to the conclusion. The arrow indicates that the premises support the conclusion.

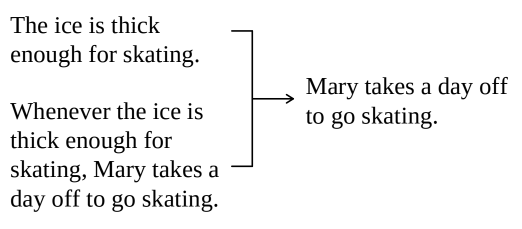

Figure 1: The structure of an argument

An argument from premises to conclusion consists of one or more steps. For instance, the following argument that John cannot finish his work consists of two steps:

> The ice is thick enough for skating.
> Whenever the ice is thick enough for skating, Mary takes a day off to go skating.
> So, Mary takes a day off to go skating.
> Whenever Mary takes a day off, John cannot finish his work.
> So, John cannot finish his work.

This argument has three premises. Two of the premises are used in the first step of the argument, and support the intermediate conclusion that Mary takes a day off. The third premise is used in the second step of the argument to support the conclusion that John cannot finish his work. The structure of the argument is shown in Figure 2. It shows the premises, the intermediate conclusion, and the conclusion of the argument.

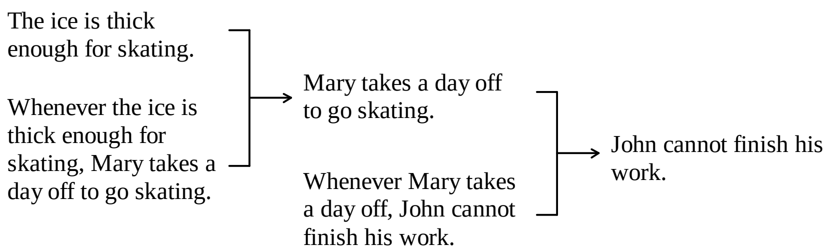

Figure 2: A two step argument

In the story, it turns out that John's conclusions that he cannot finish his work and that Mary takes a day off are in the end not justified. The arguments are defeated because of Anne's prohibition. The defeasibility of arguments is our central theme.

Summarizing, we treat arguments as reconstructions of how conclusions are supported. We regard argumentation as the process of collecting arguments in order to justify conclusions. A property of the process of argumentation is that whether arguments justify their conclusions can change during the process. They can become defeated. Arguments that at an early stage in the argumentation process justify their conclusion do not necessarily justify it at a later stage.

3 Arguments and proofs

In this thesis, we deal with formal models of argumentation. A well-known formal model of argumentation is the proof theory of classical deductive logic (in its various guises: Propositional Logic, First-Order Predicate Logic, Modal Logic).[3] Proof theory deals with proofs, that in some ways resemble arguments.

For instance, a proof that resembles the argument in Figure 1 is

$$\frac{\text{Thick-ice} \qquad \text{Thick-ice} \rightarrow \text{Day-off}}{\text{Day-off}}$$

Informally, a proof is a series of proof steps starting from given premises. Proof steps are instances of deduction rules. The example consists of one proof step that is an instance of the deduction rule known as Modus Ponens:

$$\frac{\textit{Sentence-1} \qquad \textit{Sentence-1} \rightarrow \textit{Sentence-2}}{\textit{Sentence-2}}$$

Here *Sentence-1* and *Sentence-2* are any two sentences of the logical language.

The similarity of proofs and arguments is clear. For instance, the structure of proofs is closely related to the structure of arguments. Like an argument, a proof supports its conclusion. Like an argument, a proof can consist of several steps from premises to conclusion.

In one respect, however, proofs differ from arguments: arguments are defeasible. Additional information may have the effect that an argument does no longer justify its conclusion and becomes defeated. This does not hold for proofs in deductive logic. Additional proofs never make other proofs unacceptable. Therefore, the proof theory of deductive logic is inappropriate as a model of argumentation. In this thesis, formal models of argumentation are discussed that can deal with the defeasibility of arguments.

4 The defeasibility of arguments

In this section, we informally discuss four cases in which arguments may become defeated. These are meant as illustrations of the defeasibility of arguments, and not as a taxonomy of types of defeasibility.

4.1 Exceptions to rules

So far, we have seen examples of arguments, but we have not yet investigated how the steps in an argument arise. We reconsider our example.

---

[3] Lukaszewicz (1990) and Gabbay *et al.* (1993) give overviews.

The ice is thick enough for skating. Whenever the ice is thick enough for skating, Mary takes a day off to go skating.
So, Mary takes a day off to go skating.

This is an argument that consists of a single step. The argument above is an instance of the following scheme:

Situation-1. Whenever *Situation-1*, *Situation-2*.
So, *Situation-2*.

Not only the argument above, but all instances of this scheme are arguments. There is some kind of relation between the premises and the conclusion of the argument. This relation is called a *rule*.[4] If a rule gives rise to acceptable arguments, it is valid.

The rule behind the argument scheme above is closely related to the deduction rule Modus Ponens of classical deductive logic (see section 3). If an instance of a (valid) rule is used as a step in an argument, we say that the rule is applied.

A characteristic of rules is that they can have exceptions: the conclusion of a rule does not always follow if its condition is satisfied. In the case of an exception to a rule, arguments that contain a step warranted by that rule are defeated.

We have already seen an exception to the rule above, namely the case that Mary's boss prohibited her to take a day off. In such a case the rule is not applied. Exceptions to the rule can exist, since even if Mary normally goes skating when the ice is thick enough, there can be other reasons why she does not go.

4.2 Conflicting arguments

If arguments have incompatible conclusions, we speak of conflicting arguments. For instance, Mary can have a reason to go to work, and at the same time a reason to take a day off. Not going to work may cause problems at the office, but not taking the day off means that she misses one of the few opportunities to go skating. So, Mary might consider the following two arguments :

There will be problems at the office, if I take a day off.
So, I go to work.

---

[4] One might think that 'Whenever the ice is thick enough for skating, Mary takes a day off to go skating' is a rule. In the argument in the text it is however a premise of the argument. If it is considered to be a valid rule, it gives rise to the following argument:

The ice is thick enough for skating.
So, Mary takes a day off to go skating.

We come back to this difference in chapter 4, section 1, where we discuss syllogistic and enthymematic arguments.

I miss one of the few opportunities to go skating, if I go to work.
So, I take a day off.

However, it is impossible to go to work and to take a day off, so the conclusions of Mary's two arguments are incompatible, and the arguments are conflicting.

In the arguments above, 'There will be problems at the office, if I take a day off' is a reason for 'I go to work', and 'I miss one of the few opportunities to go skating, if I go to work' is a reason for 'I take a day off'. In general, we call the direct predecessor of a conclusion in an argument a *reason* for that conclusion. In a case such as in this example, where reasons support incompatible conclusions, we say that the reasons are conflicting.

Conflicting arguments must be distinguished from contradictory proofs in deductive logic. If two proofs are contradictory, anything can be proven, and there must be a false premise. If two arguments conflict, there is not necessarily a false premise. It can also be the case that one (or both) of the arguments should be considered defeated. In the example above, probably both arguments are defeated, and replaced by an argument in which Mary takes her preferences into account:

There will be problems at the office, if I take a day off.
I miss one of the few opportunities to go skating, if I go to work.
Opportunities to go skating are extremely rare, and the problems can be solved tomorrow.
So, I take a day off.

4.3 Conclusive force

Not all arguments support their conclusion equally well; arguments have different degrees of conclusive force. Some arguments make their conclusion more plausible than others ('If it was Mary who told you John is nice, I believe he is. If Anne told you, I don't know'). If an argument uses statistical evidence, one conclusion can be more probable than another ('John's boss is probably male'). If the conclusive force of an argument is too weak, it is defeated.

The depth of an argument influences its conclusive force: a series of argument steps is often less cogent than one step. For instance, the argument that there will be problems at the office is less cogent than the shorter argument that Mary takes a day off. The conclusive force becomes less because the larger argument can be defeated by exceptions to *both* argument steps.

In the story there is an exception to the first argument step. The conclusion that Mary has taken a day off is not supported, because Anne prohibited it. As a result, the conclusion that there will be problems at the office is then also no longer justified. But if it was justified to believe that Mary took a day off, there could be still be an exception to the second step. For instance, if a temporary employee is hired, it is not justified to believe that there will be problems at the office. We call this the *sequential weakening* of an argument.

Also the number of arguments that support a conclusion or intermediate conclusion influences the conclusive force of an argument. An argument that contains more (independent) reasons for some conclusion can become more cogent. For instance,

Harry told you John is nice; Pat told you John is nice.
So, I believe he is.

can justify its conclusion, while

Harry told you John is nice.
So, I believe he is.

may not. We call this the *parallel strengthening* of an argument.

4.4 Other arguments taken into account

Whether an argument is defeated is influenced by the other arguments taken into account. We have already seen an example: the argument that Mary takes a day off to go skating is defeated as soon as there is another argument that justifies the conclusion that Mary's boss forbids her to. In our story the latter argument was actually a statement: it did not contain an argument step.[5]

In the example, there is an exception to the rule that Mary takes a day off to go skating if the ice is thick enough. The argument that Mary takes a day off can also be defeated by an argument that explicitly takes the exception into account:

The ice is thick enough for skating.
If the ice is thick enough for skating and Mary's boss does not forbid her to take a day off, Mary takes a day off to go skating.
If the ice is thick enough for skating and Mary's boss forbids her to take a day off, Mary does not take a day off to go skating.
Mary's boss forbids her to take a day off.
So, Mary does not take a day off to go skating.

Another example of the influence of arguments on each other is that arguments can challenge each other. We say that one argument challenges another argument if the challenged argument is defeated in case the challenging argument is not. For instance, the argument

John dislikes Mary.
So, I think that Mary is not nice.

---

[5] By convention, we treat statements as arguments with trivial structure (cf. chapter 5, section 2.1).

might be challenged by the statement:

John and Mary had a relationship, and Mary finished it.

5 Related research

In the previous sections, we introduced the central concepts of the thesis. We continue with an introduction to the research in the thesis, and start with a brief survey of related research.

Recently there has been a revival of research on argumentation and defeat. This revival has been motivated by several cross-disciplinary interests. For instance, the following - not necessarily disjoint - disciplines have stimulated the research on argumentation and defeat:

- *Logic*: The research on nonmonotonic logics remains popular (Gabbay *et al.* (1994b) give an overview) and now encompasses the defeasibility of arguments as a special topic (cf. Nute, 1994).
- *Computer science*: The computational complexity of nonmonotonicity attracted the attention of the logic programming community (e.g., Dung, 1993, 1995; Bondarenko *et al.*, 1993).
- *Artificial intelligence*: Since reasoning with defeasible arguments seems to lead to successful behavior of people, artificial intelligence researchers try to capture its essence (e.g., Nute, 1988; Geffner and Pearl, 1992; Simari and Loui, 1992).
- *Epistemology*: Questions about the justification and support of beliefs have resulted in formal epistemological theories (e.g., Loui, 1987, 1991; Pollock, 1987-1995).
- *Argumentation theory*: Notions such as counterargument and reinstatement have been formally studied (e.g., Vreeswijk, 1991, 1993; Verheij, 1995a, b, c).
- *Dialectics*: Several game-like formalisms have been proposed in which two parties are disputing an issue (e.g., Loui, 1992; Gordon, 1993a, 1993b, 1995; Vreeswijk, 1993; Brewka, 1994; Leenes *et al.*, 1994; Hage *et al.*, 1994; Lodder and Herczog, 1995).
- *Legal theory*: The pragmatic solutions in legal reasoning to deal with exceptions and conflicts have inspired researchers and have been formally analyzed (e.g., Hage, 1993, 1995; Prakken, 1993a, b, 1995; Sartor, 1994).

This brief survey contains only a selection of recent research to give an idea of the current activity and the diversity of perspectives. Overviews of research on argumentation and defeat have recently been given by Bench-Capon (1995), who focuses on artificial intelligence and law, and Loui (1995b), who focuses on computational dialectics.

6 General aims and biases of research

Research on the modeling of argumentation can have rather different aims, for instance:

- to describe and evaluate actual human argumentation by means of empirical investigation, e.g., in cognitive science or psychology;
- to apply an argumentation model in order to build intelligent computers and programs, e.g., in computer science and artificial intelligence;
- to investigate and enhance our conceptualizations of argumentation in order to better understand its nature, e.g., in philosophical and mathematical logic.

Of course, doing research on the modeling of argumentation, one does not normally have only one of the aims above: even if one is mostly interested in intelligent computer programs, one can be inspired by actual human argumentation, and be led to the enhancement of one's initial model of argumentation. Nevertheless, research is often biased towards one or more of the mentioned aims of study.

In Figure 3, we have visualized the biases of some research on the modeling of argumentation in a triangular diagram. The three corners of the triangle correspond to the three aims mentioned, and are suggestively labeled 'Minds and humans', 'Machines and programs', and 'Theories and models'. Researchers or subjects of research are indicated by a labeled dot.[6] The closer a dot is to one of the corners, the more the corresponding researcher or subject of research is biased towards the aim of study of that corner.[7]

Some research is mostly biased to one of the three aims. We give examples of each. First, we mention First-Order Predicate Logic.[8] As a model of argumentation, it is most appropriate as a theoretical model, due to its nice mathematical properties, but less as an empirical model or as a computational model. It is therefore indicated in the upper corner. Second, Van Eemeren and Grootendorst (Van Eemeren *et al.*, 1981, 1987) have provided a model of argumentation explicitly meant to analyze argumentation as it occurs in argumentative texts, and are therefore indicated in the lower-left corner. Third, the research on logic programming is clearly mostly aimed at building intelligent machines, notwithstanding its theoretical achievements, and is therefore indicated in the lower-right corner.

---

[6] Several of the indicated researchers or subjects of research are extensively discussed later on.

[7] The triangle has *barycentric coordinates*. One can think of the triangle as the set of points (x, y, z) in the plane $x + y + z = 1$, such that $0 \leq x \leq 1$, $0 \leq y \leq 1$ and $0 \leq z \leq 1$. For instance, the corners of the triangle are the points where one of the coordinates is equal to 1. The sides of the triangle are the points where one of the coordinates is equal to 0. The values of each of the three coordinates represent the bias level towards one of the corners.

[8] Van Dalen (1983) and Davis (1993) give introductions to First-Order Predicate Logic.

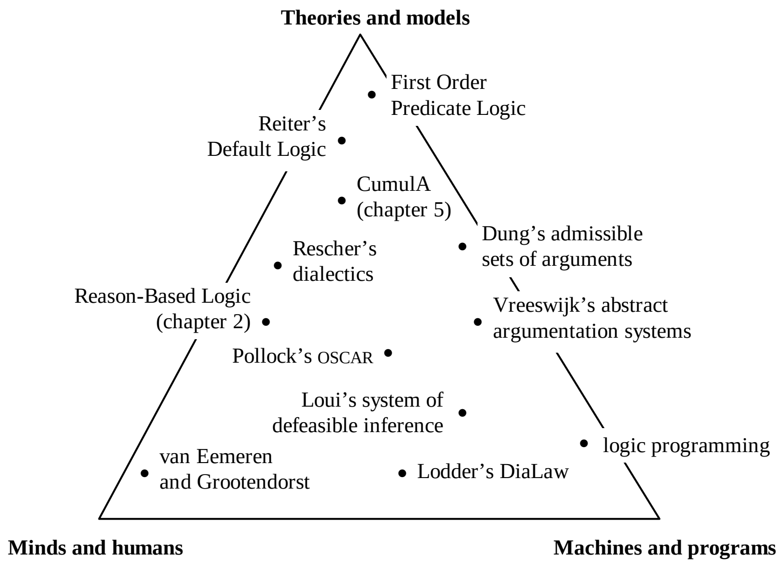

Figure 3: Biases diagram of some research on modeling argumentation

Some research is equally biased towards two of the aims, and is therefore indicated near the middle of one of the sides of the triangle. For instance, Vreeswijk’s abstract argumentation systems (1991, 1993) were meant both as a model for theoretical study and for computational application.[9] Pollock’s (1995) research on OSCAR is indicated in the middle of the triangle, since it equally contains elements of all three aims: Pollock has applied his philosophical theories on epistemology in the computer program OSCAR that is designed to argue as people do (or should do).

In order to show our aims of research, we have included our two main topics, Reason-Based Logic and CumulA, in the triangle. The first, Reason-Based Logic (see chapter 2), is indicated near the middle of the left-side of the triangle. It was inspired by actual human argumentation, especially in the field of law (see chapter 4), but it was also developed in order to compare it with other models. The second main topic of research, CumulA (see chapter 5), is indicated near the upper corner, since it was mainly designed as an abstract model of argumentation, that can be used to analyze different approaches towards modeling argumentation.

Although the diagram is merely tentative and we make no claim about its ‘truth’, we nevertheless hope that the biases diagram illustrates how differently biased research on modeling argumentation can be, and what the biases of our own research are.

---

9 Vreeswijk (1995) describes the program IACAS, which was written to demonstrate his abstract argumentation systems.

7 Research goals and method

Our starting point of research is that the currently available models of argumentation are not fully satisfactory. This starting point, although certainly not new, remains valuable, despite the abundance of newly presented models (see section 5). As Haack (1978) put it, when discussing the paradigmatic example of a rule in argumentation, the material conditional of First-Order Predicate Logic,

> ‘(...) the significance of the discrepancies between ‘if’ and ‘→’ will depend on the answers to at least two (...) questions: for what purpose(s) is the formalisation intended? and, does that purpose require something stronger than the material conditional? Both (...) are deep and difficult questions.’
> (Haack, 1978, p. 38)

The recent revival of research (cf. section 5) is partly due to a new answer to the first of these questions: recent research often is concerned with defeasible arguments, leading to other formalizations of rules. As for rules, new purposes of formalizing argumentation, such as capturing the role of counterarguments and of the process character of argumentation, lead to new models.

The purpose of our research is to find answers to two groups of research questions.

- What is the role of rules and reasons in argumentation with defeasible arguments? What properties of rules and reasons are relevant for argumentation and defeat? How do these properties relate?
- What is the role of process in argumentation with defeasible arguments? How is the defeat of an argument determined by its structure, counterarguments and the argumentation stage?

Trying to answer these groups of questions, we study argumentation and defeat from two angles, resulting in formalisms of different nature, Reason-Based Logic and CumulA.

Reason-Based Logic is a model of the nature of rules and reasons, which are at the basis of argumentation. We investigate the properties of rules and reasons that are relevant for the argumentation and defeat, and how these properties relate to each other. This part of the research is joint work with Hage, who initiated the development of Reason-Based Logic (see chapter 2).

CumulA is a model of argumentation in stages. We investigate how the structure of an argument is related to defeat, when arguments are defeated by counterarguments, and how the status of arguments is affected by the argumentation stage.

The thesis has five goals:

- Providing a model of rules and reasons, Reason-Based Logic, focusing on properties that are relevant for the defeasibility of arguments.
- Demonstrating the usefulness of the model by providing examples in the field of law.
- Discussing how Reason-Based Logic relates to previously proposed models.
- Providing a model of argumentation, CumulA, that focuses on the process of taking arguments into account, and shows how the status of an argument is determined by the structure of the argument, the counterarguments and the stage of the argumentation process.
- Demonstrating how CumulA can be used to analyze other models of argumentation.

Our method of research can be summarized, as follows:

> Developing formal models of argumentation on the basis of informal examples.

The advantage of formal models is that they are clear and precise, which is necessary to show the intentions of the model and is useful for revealing errors and shortcomings. A drawback of formal models, as put forward by Van Eemeren *et al.*, discussing the attraction of Toulmin's less formal model (Toulmin, 1958), is that:

> 'Studying formal logic systems requires quite a lot of effort, its relevance for practical purposes is not immediately apparent and the return on the effort spent is slight.' (Van Eemeren *et al.*, 1987, p. 206)

This is felt so by many people, and indeed the feeling seems to be justified by the research on nonmonotonic logics, which has become a mathematically inclined subject, even though it was initially inspired by intuitive examples.

The drawback can partly be circumvented by providing informal examples. We not only do this to make the text more legible, but also as an essential ingredient of our method: without informal examples, a formalism remains uninterpreted, and therefore much less useful. We are backed by Haack (1978), who in her 'Philosophy of logics' stresses the importance of informal interpretation and extra-systematic judgments (p. 32ff.) for devising and evaluating a formal model.

As a result, in this thesis, we stick to the precision and rigor of formal models, but precede all formal definitions by informal examples, needed to interpret the formalism.

8 Outline of the thesis

The structure of this thesis follows the research goals discussed in the previous section.

In chapter 2, we describe Reason-Based Logic. We determine types of facts concerning rules and reasons that are relevant for the defeasibility of arguments, and show their relations. Using this semantics of rules and reasons, we determine some intuitively attractive modes of reasoning. However, these lead to the difficulties of nonmonotonic reasoning. We show how the ideas of Reiter (1980, 1987) can be used to define rigorously which conclusions nonmonotonically follow from a given set of premises.

Chapter 3 contains a series of examples of Reason-Based Logic, taken from the field of law. We give applications of Reason-Based Logic to the theory of legal reasoning: we describe three different ways of reconstructing reasoning by analogy, and provide an integrated view on rules and principles, which seem fundamentally different (cf. Dworkin, 1978, p. 22ff. and 71ff.).

In chapter 4, we survey other models of rules, and compare them to Reason-Based Logic. We do this by treating a number of issues concerning the formalization of rules, and discussing various approaches to deal with these issues.

In chapter 5, the second part of the thesis starts with a discussion of CumulA. It is a formal model of argumentation with defeasible arguments, focusing on the process of taking arguments into account. The main ingredients of the formalism are arguments, defeaters, argumentation stages and lines of argumentation.

In chapter 6, we show how CumulA can be used to analyze models of argumentation. We investigate types of argument structure and of defeat, the role of inconsistency and counterarguments for defeat, and directions of argumentation. As a result, we are able to distinguish a number of argumentation theories (based on existing argumentation models) on formal grounds.

The thesis ends with the results and conclusions of the research (chapter 7). We also give some suggestions for future research.