A THEORY OF LEGAL REASONING AND A LOGIC TO MATCH[1]

JAAP HAGE
University of Limburg, Department of Metajuridica,
email: jaap.hage@metajur.rulimburg.nl

Abstract.
This paper describes a model of legal reasoning and a logic for reasoning with rules, principles and goals that is especially suited to this model of legal reasoning. The paper consists of three parts.
The first part describes a model of legal reasoning based on a two-layered view of the law. The first layer consists of principles and goals that express fundamental ideas of a legal system. The second layer contains legal rules which in a sense summarise the outcome of the interaction of the principles and goals for a number of case types. Both principles, goals and rules can be used in legal arguments, but their logical roles are different. One characteristic of the model of legal reasoning described in the first part of the paper is that it takes these logical differences into account. Another characteristic is that it pays serious attention to the phenomena of reasoning about the validity and acceptance of rules, respectively principles and goals, and about the application of legal rules, and the implications of these arguments for the use of rules, principles and goals in deriving legal conclusions for concrete cases.
The second part of the paper first describes a logic (Reason-Based Logic) that is especially suited to deal with legal arguments as described in terms of the previously discussed model. The facilities of the logic are illustrated by means of examples that correspond to the several aspects of the model.
The third part of the paper deals with a number of logico-philosophical reflections on Reason-Based Logic. The occasion is also used to compare these presuppositions with theories of defeasible reasoning based on the comparison of arguments.

Keywords: legal reasoning, reasoning with principles and rules, nonmonotonic logic

1. Introduction for non-logicians: Why bother about nonmonotonic logics?

Many persons with an interest in the field of AI and Law, especially those with a legal background, sometimes wonder what is the use of all the research in complicated logics and in

[1] The ideas developed in this paper are based on the draft of my book 'Reasoning with rules' which will be published by Kluwer Academic Publishers in the Law and Philosophy Series. The book offers not only more elaborate and sometimes slightly different treatments of the topics of this paper, but also pays more attention to the philosophical background of this work.

particular so-called nonmonotonic ones.[2] In the first place, one might argue, because the essence of legal reasoning has to do with making the right decisions, that is with matters of evaluation, rather than logic. And in the second place because the logic that seems to be used is rather simple. We have a rule with a number of conditions, and if these conditions are satisfied, the rule conclusion is drawn, and otherwise not. There may be discussions about whether the rule conditions are satisfied, but that is the evaluative issue that is at stake, not the logical one. And to the extent that the evaluative discussion makes use of logic, it is the same simple logic of rule application again. In other words, all the logic that is needed is the logic that accounts for the application of rules. Technically this means that no other logic is necessary than that which accounts for the validity of Modus Ponens arguments.

This is a simple argument against the usefulness of logics, and although many will see that it is too simple, many will also have the feeling that for *most* of their legal work they hardly use any complex reasoning, let alone logic. One of the purposes of this paper is to show that ordinary legal reasoning is much more complex than it might seem at first sight. By means of examples which are not too unrealistic, I will try to illustrate that legal reasoning involves more than simple rule application, and that a logic that takes this complexity into account will be nonmonotonic. In other words, I will try to show that the intuitions of many lawyers about what they are logically doing are wrong, and that - often unknowingly - lawyers argue in ways that are logically rather complex.

However, even if it is granted that legal reasoning is more complex than just applying applicable rules, it has not yet been shown that AI and Law needs nonmonotonic logics. In particular not, because nonmonotonic logics tend to be computationally inefficient, and are therefore not so suitable as the basis of legal inference mechanisms.

This objection overlooks that logics need not fulfil a role in the inference mechanism of (legal) knowledge systems, but that they can also be an intermediary between on the one hand a jurisprudential account of legal reasoning, and on the other hand the specification of an inference engine for legal knowledge systems. Jurisprudence is a domain theory which is too much interwoven with characteristics of legal systems to provide the specifications for an inference engine. A logic that gives a satisfactory account of the reasoning involved in legal decision making, abstracts from the legal background that will dominate the jurisprudential theory. Still it does not commit to a particular inference mechanism for knowledge systems. A logic rather provides a list of desiderata, demands that we would like to make on an inference mechanism, but which may have to be compromised against demands of implementation.[3]

This paper is meant for non-logicians who are prepared to take logic seriously. It is, however, not meant as an introduction to nonmonotonic reasoning for lawyers. Such an introduction can be found in [Prakken 1993]. This paper rather wants to show how particular characteristics of legal

---

[2] A logic is monotonic if the sentences that can be derived from a superset of a theory are a superset of the sentences that can be derived from the theory. So if $\Rightarrow$ stands for derivability, the monotonicity of a logic means that if $T \Rightarrow S$, then $T, P \Rightarrow S$. A logic that is not monotonic is nonmonotonic.

Informally, the nonmonotonicity of a logic means that the addition of new information to a theory can make sentences underivable which used to be derivable on the basis of the smaller theory. A legal example would be that new information about an exception makes a rule inapplicable and the conclusion of this rule underivable.

[3] By the way, this is also a reason why a logical account of legal reasoning should not be compromised by the aim to end up with something which is computational tractable.

reasoning ask for complex, even nonmonotonic logics, and how a logic which takes these characteristics seriously may look like. Moreover, it hopes to demonstrate by means of extensive examples, how a particular nonmonotonic logic, Reason-Based Logic, can be used to model these relatively complex forms of legal reasoning.

2. Introduction for logicians

The recent proliferation of logical systems has made it impossible to assume that there is something like THE LOGIC, a universal theory about valid reasoning, which holds true for all fields in which arguments play a role. It is rather the case that we must choose a logic that is most adequate for the kind of reasoning task that we are dealing with. In the field of Artificial Intelligence and Law, a number of logics has been proposed (e.g. [Gordon 1994], [Prakken and Sartor 1995], [Hage 1995] to mention some of the most recent versions). To make a reasoned choice between these logics, or to prefer a combination of them or still some other logic, we need a standard to measure these logics against.

Because these logics are meant to deal with legal reasoning, the proper standard is whether they give an adequate account of valid legal reasoning. The problem in this connection is, however, that we lack an independent account of valid legal reasoning. On the one hand we need a theory of valid legal reasoning to measure logics against, while on the other hand we need a logical theory to distinguish between valid and invalid legal reasoning. The 'solution' to this dilemma is to develop simultaneously a theory of valid legal reasoning and a matching logic for legal reasoning. The jurisprudential part of this dual theory must be acceptable as an idealised reconstruction of legal practice, while the logical part must match our intuitive understanding of valid reasoning. In the end, we need three matches: the account of valid legal reasoning must match legal practice, while the legal-logical theory must match both general logical intuitions and the legal-theoretical account of legal reasoning. Cf. figure 1.

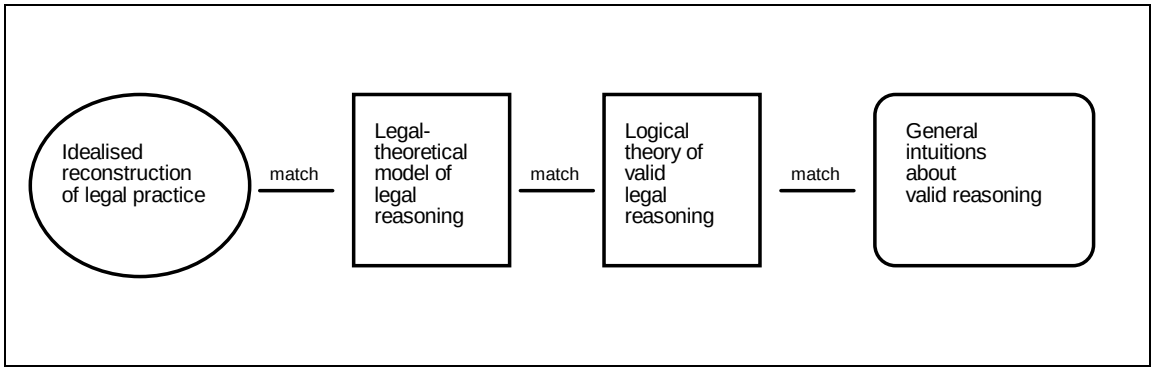

Figure 1: Matches between legal practice, legal theory, legal logic and logical intuitions

These three matches are reflected in the three main parts of which this paper consists. Part A (sections 3 to 18) of the paper contains a jurisprudential account of the legal reasoning which is involved in 'solving' particular cases. This account is based on a two-layered model of the law. The first layer consists of principles and goals that express fundamental ideas of a legal system. The second layer contains legal rules which in a sense summarise the outcome of the interaction of the principles and goals for a number of case types. Both principles, goals and rules can be used in legal arguments, but their logical roles are different. One characteristic of the proposed model

of legal reasoning is that the model takes these logical differences into account. Another characteristic is that it pays serious attention to the phenomena of reasoning about the validity and acceptance of rules, respectively principles and goals, and about the application of legal rules, and the implications of these arguments for the use of rules, principles and goals in deriving legal conclusions for concrete cases.

Part B (sections 19 to 35) of the paper describes Reason-Based Logic (further: RBL), a nonmonotonic logic that mirrors a number of aspects of legal reasoning that are part of the previously described model. By means of examples that illustrate the several parts of the model of legal reasoning, it is shown how RBL matches this model.

To show that RBL is not merely a formal description of the model, but a logic on its own right, part C (sections 36 to 39) of the paper contains some logico-philosophical reflections on RBL, pictures RBL against the background of a general theory of valid reasoning, and finally compares RBL to logics based on the comparison of arguments.

A. The two layers of the law

3. Two kinds of legal reasoning

At first sight it seems that at least two different kinds of reasoning occur within the law.[4] The first kind of reasoning is illustrated by the application of a legal rule whose conditions are satisfied. The result of this application is that the rule conclusion is drawn by means of an argument which is very similar to a Modus Ponens argument. In fact, several authors have described the application of a rule as a case of deduction [e.g. MacCormick 1978, and Alexy 1978]. The following is an example of this type of legal reasoning:

Rule: *Thieves are punishable.*

Fact: *John is a thief.*

Conclusion: *John is punishable.*

The second kind of reasoning is maybe best exemplified when a particular case is to be classified (or not) by an evaluative term, where there is no antecedent legal basis in the form of rules or case law for making this classification. For instance, what amounts to *creating danger* in the case of making genetic modifications to some bacterium? Arguments in such a case will consist of adducing reasons that plead for and against the dangerous nature of the modifications, where none of these reasons will be decisive by itself and the final decision depends on the 'weighing' of all available reasons. This type of reasoning cannot without distortion be modelled as a Modus Ponens argument, nor as any other kind of deductive argument.[5]

These two kinds of reasoning are reminiscent of Dworkin's distinction between legal rules and legal principles. According to Dworkin [1978, pp. 22f.], rules apply in an all-or-nothing fashion. If

---

[4] I do not claim that the two kinds of reasoning described here are exhaustive as a categorisation of kinds of legal reasoning. Another obvious distinction is the one between rule-based reasoning and case-based reasoning.

[5] With distortion it is possible to model any finite argument as a Modus Ponens argument. Suppose we have an argument consisting of the set of premises S and the conclusion C. This argument can be modelled as a Modus Ponens argument by adding the 'hidden' premise that the conjunction of the elements of S materially implies C. Seemingly valid arguments can be made invalid by adding 'hidden' elements to the antecedent of an implication.

the conditions of a valid legal rule are satisfied, the conclusion of this rule must follow. A principle, on the contrary, 'states a reason that argues in one direction, but does not necessitate a particular decision' (p. 26). As a consequence of this first difference, Dworkin mentions a second difference, namely that other than rules, principles have a dimension of weight or importance. According to Dworkin 'When principles intersect ... one who must resolve the conflict has to take into account the relative weight of each.' Rules, on the contrary, cannot be weighed: 'If two rules conflict, one of them cannot be a valid rule. (p. 27)'

To facilitate the discussion of these two types of reasoning, I will call them *rule-based reasoning* and *principle-based reasoning*, where the latter includes reasoning on the basis of goals (cf. section 8).

4. Principle-based reasoning

A discussion of principle-based reasoning can very well start with the notion of a reason. A reason is an elementary or compound fact which is significant for the presence of some other fact, or for believing that some other fact is present. For instance, the compound fact that John took away Mary's car and that Mary did not allow John to take her car away, is a reason why John is a thief. The compound fact that somebody saw John climbing out of a window of Mary's house at three at night and that the following morning Mary found her antique clock missing in a (for her) unexplainable way, is a reason to believe that John is a thief (of the clock). The former fact makes John into a thief; it is a *constitutive* reason. The latter fact is evidence for John's being a thief; it is an *epistemic* reason. In both cases the reason forms a basis for an argument to the conclusion that John is a thief, although the significance of the reason for the conclusion is somewhat different.

One important characteristic of reasons is that they are amenable to 'weighing' . There can both be reasons for and against a particular conclusion, and if this is the case, the judgement on the conclusion depends on all the available reasons.[6] Let us consider another example. John sells the clock he has stolen from Mary to Gerald. The fact that John was not the owner of the clock is a reason why the sale does not make Gerald the owner of the clock. However, the fact that Gerald could not know that John had stolen the clock and that he paid a reasonable price for it is a reason why Gerald should become the owner. These reasons why Gerald has not, respectively has become the owner of the clock have to be weighed in order to determine the final legal consequence.

Another important characteristic of reasons is that they are - as Hare [1963] called the phenomenon - *universalisable*. If a fact of a type T is significant for a conclusion of the type C, then all facts of the type T are significant for conclusions of the type C. This presupposes, of course, that facts of the type T are significant *because* they belong to type T and the conclusion belongs to type C. So, if the fact that he is a thief is a reason why John is punishable, the fact that somebody is a thief is in general a reason why this person is punishable. The latter formulation can be considered as the expression of a *principle*: Thieves are punishable.

If any fact of the type R is a reason for a conclusion of the type C, this can be universalised into the formulation of the principle 'If there is a case of R, then there is a case of C'. Notice, however, that the formulation expresses no more than a principle. It does not state that all cases of R are cases of C. In fact it does not *state* anything. The formulation of the principle only *denotes* a

---

[6] The need for weighing reasons is also stressed by, amongst others, Raz [1975 and 1979] and Peczenik [1989].

connection between types of facts, that is the connection that facts of the type R are a reason for facts of the type C.

5. Directions of fit

The distinction between stating and denoting which I made in connection with formulations of principles, can be clarified by another distinction, namely the distinction popularised by Searle [1983] between directions of fit. The distinction between directions of fit is explained by Searle by borrowing an example of Anscombe's [Anscombe 1957]. I will use the same example. Suppose I make a shopping list which I use in the supermarket to put items in my trolley. A detective follows me and makes a list of everything that I put in my trolley. After I am finished, the list of the detective will be identical to my shopping list. However, the lists had different functions. If I use the list correctly, I place exactly those items in my trolley that are indicated on the list. My behaviour is to be adapted to what is on my list. In the case of the detective it is just the other way round; the list should reflect my shopping behaviour. If we consider my behaviour as (part of) the world, we can say that my shopping list has the world to word direction of fit, because my behaviour must fit the words on the list. The detective's list, on the contrary, has the word to world direction of fit, because his list must fit the world (my behaviour).

In this example, the world to word direction of fit is illustrated by means of a directive. The shopping list directs my behaviour and in that way makes the world fit the words. The word to world direction of fit is illustrated by means of a description. The list of the detective describes what is put in my trolley. However, directives are not the only case of world to word fit. Baptising is another example of a case in which the world (having a particular name) fits words (those used for baptising).

My interest in the world to word fit focuses on still a different case, namely that of *constitutive* principles (or rules). Suppose that we adopt a particular standard for determining which behaviour counts as coward. This standard runs that a soldier who flees at the approach of the enemy is coward. Notice that the standard is not a descriptive sentence which is true or false depending on the facts of the world. On the contrary, the adoption of this standard *makes* that certain facts occur in the world. Without the standard, there is only a soldier who flees at the approach of the enemy. Assuming the standard, there is also a coward. Adoption of the standard makes that certain facts also count as other facts. The fleeing of the soldier becomes to count as an act of cowardice. This is a clear case of the world to word direction of fit: thanks to the standard the world counts an additional fact (a case of cowardice) which would not be present without the standard.[7]

Do not misunderstand this. Physically the world would be the same with or without the standard. The difference created by the adoption of the standard is in the layer of facts built on top of the purely physical ones. In this connection Hare writes about 'supervenience' [Hare 1952, p. 80f]: Some facts supervene on other ones; they are present to the extent that these other facts are also present. A logical example is that the truth of the sentence p & q is supervenient on the truth of the sentences p and q.

---

[7] Notice the clear connection with the phenomenon of institutional facts as discussed in several papers in [MacCormick and Weinberger 1986]. Cf. also [Searle 1995].

The same thing can also be approached from a slightly different angle: The standard (a kind of principle) that soldiers who flee at the approach of the enemy are coward makes the fact that a particular soldier fled at the approach of the enemy into a reason why this soldier is coward. The principle makes a particular fact into a reason, which is a kind of constitution.

Given the distinction between directions of fit, we can also make another distinction. On the one hand we have descriptive sentences which are statements, which have the word to world direction of fit, and which are consequently true or false. On the other hand we have (amongst others) principles, which have the world to word direction of fit, and therefore do not state but rather express a connection between types of facts.

Although the distinction between directions of fit does not dictate a different logic for on the one hand reasoning with statements, and on the other hand reasoning with principles, rules and goals, it provides at least the beginning of an explanation of the differences in logical behaviour between these two categories. We will see that the distinction plays an important role in the model of legal reasoning that will be developed and in the logic that corresponds to this model.[8]

6. The narrow view of principles

The phenomenon that reasons (insofar as they are based on principles) have to be weighed against each other is a consequence of the narrow view of principles. Consider the following example: Gerald bought an antique clock from John, who stole the clock. Gerald does not become the owner of the clock, because his predecessor, John, was not the owner either. This reason is based on the principle that nobody can transfer a right that he has not got himself, the so-called nemo plus-principle.

Suppose moreover that Gerald thought that John was the owner of the clock, and that he paid a reasonable price for the clock. These facts form a reason why Gerald should become the owner of the clock, but they are not taken into consideration by the nemo plus-principle. This principle only has an eye for one particular aspect of the case, namely that John was not the owner of the clock. All other aspects, whether they are legally relevant for the outcome or not, are not taken into consideration by this principle.

That is what I mean by the narrow view of principles. Principles only have attention for the aspect of a case that corresponds to the condition part of the principle, and leave everything else out of consideration. Although the latter also counts for rules, rules aim - as I will argue in section 9 - to take all legally relevant aspects of a case into account. A typical legal principle, such as the nemo plus-principle, does not even attempt to avoid narrow-mindedness.

The narrow view of principles is compensated for by the possibility that a reason generated by a principle is outweighed by one or more reasons that plead in a different direction. In our example, the fact that John was not the owner of the clock as a reason against Gerald becoming the new owner, may be outweighed by the reason that Gerald acted in good faith and paid a reasonable price for the clock. In other words, the narrow view of principles is counteracted by the possibility of colliding principles (and goals) which underlie reasons that plead in a different direction. By basing the final conclusion on the interaction of all available reasons, the narrow view of principles does not become a major disadvantage.

---

[8] Another approach to the same fundamental distinction can be found in [Toulmin 1958], where warrants are strictly distinguished from data and from their backings.

7. Reasoning with principles

One principle of Penal Law might be that if somebody committed a crime, he ought to be punished. Since this is only a principle, it does not imply anything more than that having committed a crime is a reason for being punished. In particular it does not imply that everybody who committed a crime ought to be punished.

Suppose that John committed a crime. On the basis of the principle that if somebody committed a crime, he ought to be punished, this fact becomes a reason why John ought to be punished. If there are no reasons why John ought not to be punished, or if these reasons are outweighed by the reason that John committed a crime, this reason leads to the conclusion that John ought to be punished.

It turns out that the argument to the conclusion that John ought to be punished consists of two steps. The first step leads from the fact that John committed a crime to the intermediate conclusion that this fact is a reason why John ought to be punished. The second step leads from this reason, possibly in combination with other reasons, to the conclusion that John ought to be punished. Notice that both steps are cases of constitution, and are on the *ontological* level. The principle *makes* a fact into a reason, and this reason *makes* it the case that (or - if there are also other reasons - contributes to) the final conclusion that John ought to be punished. The argument is the mental (verbal) reproduction of this constitution. Cf. figure 2, which on purpose makes use of the structure proposed by Toulmin.

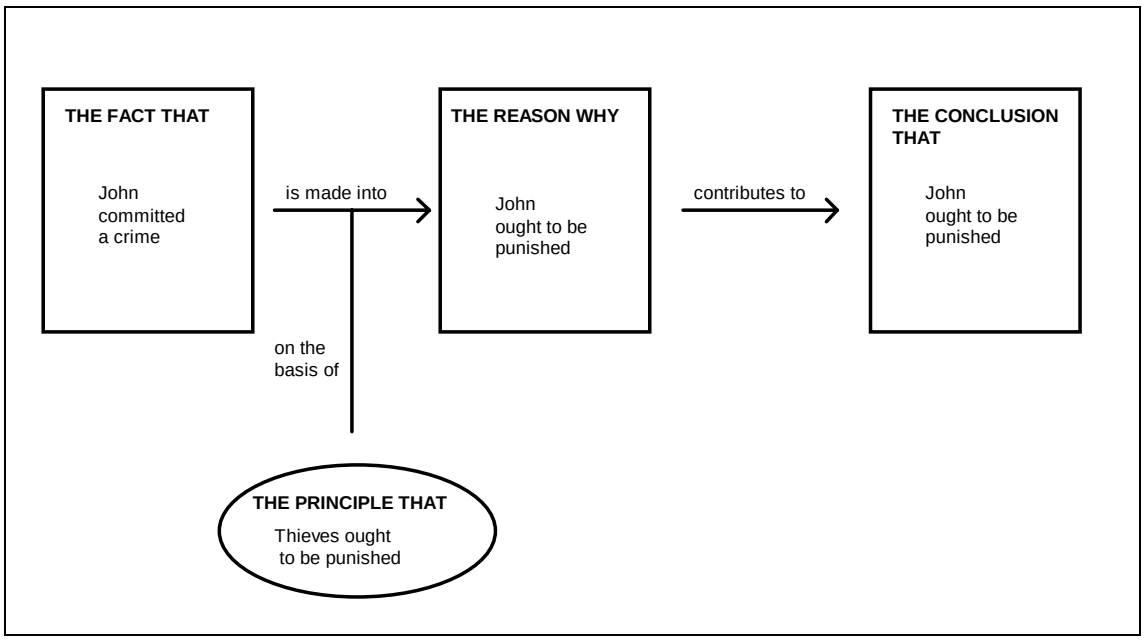

Figure 2: Reasoning with principles

8. Goals

Not only principles, but also goals can underlie reasons. In this connection, goals are to be taken in a very broad sense, which includes governmental policies, values, interests, and - in some sense of the word - rights [cf. Alexy 1985]. In contrast to principles, which can underlie all kinds of

reasons such as reasons for classification, reasons to believe, reasons in the sense of causes, and reasons for acting, goals can only underlie deontic reasons, that is reasons for acting, reasons why something ought to be done (is forbidden / permitted), or reasons why something ought (not) to be the case, or may (not) be the case.[9]

Suppose that we have the goal to protect juveniles from harmful environments. Suppose, moreover, that John is a juvenile thief, and that punishing him (imprisonment) would bring John into a harmful environment. The goal to protect juveniles from harmful environments makes the fact that punishing John detracts from this goal into a reason against punishing John. Cf. figure 3.

This example can be generalised to the following informal rules:

- The goal to promote state S makes the fact that some action A contributes to S into a reason why A ought to be performed, or into a reason to perform A.
- The goal to promote state S makes the fact that some action A detracts from S into a reason why A ought not to be performed, or into a reason against performing A.
- The goal to promote state S makes the fact that some state T contributes to S into a reason why T ought to be the case.
- The goal to promote state S makes the fact that some state T detracts from S into a reason why T ought not to be the case.

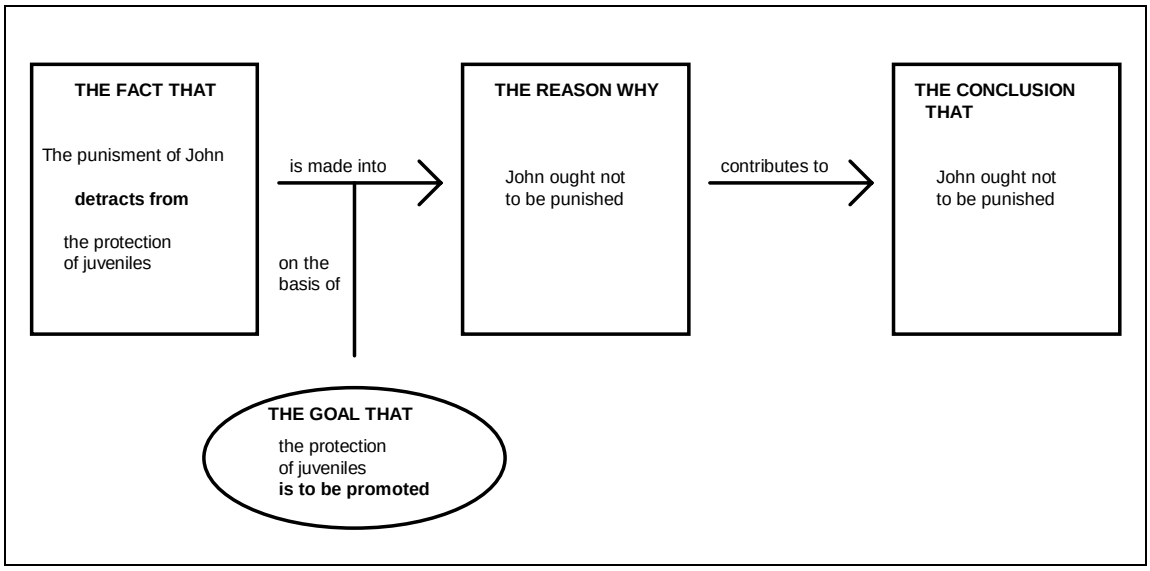

Figure 3: A goal-based argument

Goals suffer from the same narrowness of view that also characterises principles. Take for instance the goal to protect parties in good faith. The fact that if Gerald becomes the owner of the antique clock which he bought, this contributes to the protection of parties in good faith, is made into a reason for the conclusion that Gerald ought to become the owner of the clock. The goal only has an eye for one aspect of this case. The fact that John, who sold the clock to Gerald, was

9 The limitation that goals can only underlie deontic reasons is not as serious as it might seem, because of the phenomenon of *deontic collapse*. Cf. [Hage 1995 and Reasoning].

not the owner of the clock is not taken into account. That is why the reason based on the goal to protect parties in good faith has still to be weighed against the reason based on the nemo plus-principle.

9. Rules

Both (legal) principles and goals reflect that some aspects of a case pull the legal solution of this case in a certain direction. Since both will usually be unwritten, so that their precise contents can be subject to dispute, and since both give only rise to reasons which need to be weighed against other reasons with uncertain outcome, legal reasoning solely on the basis of goals and principles suffers from legal insecurity. The use of legal rules is one way to minimise this evil. In the first place because rules are often written down so that, except for the uncertainty of interpretation, the contents of the rules are known. In the second place because rules determine the outcome of the cases to which they are applied. It is this second characteristic which requires some more attention.

Consider again the example of Gerald who, in good faith, bought an antique clock from John who had stolen the clock. There is both a reason why Gerald ought to become the owner of the clock, based on the goal to protect parties who acted in good faith, and a reason why Gerald does not become the owner, based on the nemo plus-principle. Suppose that the weighing of these reasons leads to the conclusion that Gerald becomes the owner.[10] To prevent that a next time in a comparable case these same reasons have to be weighed again, with the risk that the outcome will be different, a rule can be adopted to the effect that if somebody in good faith obtained a right from somebody who did not have the capacity to transfer this right, this party in good faith still obtains this right. Given this rule, the next time it is not necessary to identify and weigh the original reasons anymore. Instead it is possible to apply the rule and let the rule determine the legal consequences of the case.

A *stereotypical* rule differs fundamentally from a *stereotypical* principle. A principle identifies a particular *aspect* that may be part of many divergent types of cases and attaches to the presence of this aspect in a case the tendency to solve the case in a particular way. Because the case may have other relevant aspects, there may also be tendencies to solve the case in a different way.

A rule identifies a particular *type of case* and indicates the consequences of this type of case. Types of cases are characterised by sets of characteristics which will be mentioned in the rule conditions. Each of these characteristics may represent an aspect of the case that is made relevant by a legal principle. The rule has the pretension to take all relevant aspects of the case into account and to indicate the outcome of the case *based on all relevant aspects.* Because of this pretension, there is no need to weigh the outcome of the rule against other possible outcomes, because the reasons for these alternatives are (at least in theory) taken into account in the rule, and need not to be taken into account again.

In our example, the outcome of the rule needs not to be reconsidered in the light of the reason based on the nemo plus-principle why Gerald does not become the owner. This principle has been

[10] There appears to be a problem with weighing the reasons, because they seemingly do not conflict. One reason is a reason why Gerald *ought* to be come the owner, while the other reason is a reason why Gerald *actually* has become the owner. Because of the phenomenon of *deontic collapse* [Hage Reasoning], this problem disappears.

taken into account in making the rule and needs not to be considered anymore if the rule is applied.

In the terminology of Raz [1975], the facts that make a rule apply in a particular case are an *exclusionary reason* for the rule conclusion. This means that they are both a reason for the conclusion of the rule, and a reason not to apply competing goals and principles.[11] If these competing goals and principles are not applied, they do not generate reasons, and as a consequence the reason based on the rule needs not to be weighed against other reasons anymore. The rule applies in an all-or-nothing fashion.

However, although a principle or goal that is replaced by a rule does not generate a reason anymore that can compete with the reason generated by the rule, this does not mean that they have no role in legal reasoning anymore. Just because a rule replaces particular goals and principles, these goals and principles are relevant for the *interpretation* of the rule [Soeteman 1991]. The exclusion of goals and principles by applicable rules only means that these principles and goals loose their *independent* function in solving the cases in which their application is excluded.

The same phenomenon can also be described in terms of *replacing reasons*. The reason based on the rule comes instead of, replaces, the original reasons based on the principle and the goal. These latter reasons, which have been replaced, do not occur anymore [cf. Hage 1995].

Despite the distinction between stereotypical rules and principles, many principles (including goals) and rules will not be stereotypical and take an intermediate position on a scale the extremes of which are formed by the stereotypical principle and the stereotypical rule. Take for instance the moral principle not to hurt other people except in cases of self-defence. This principle is the result of taking the exception of self-defence into account in reformulating the one-sided principle that one ought not hurt other people. Originally there would be two reasons for and against hurting some other person. One reason is that the hurting was necessary to defend oneself; the other that it was a human being who suffered the hurting. If the self-defence reason outweighs the reason that a human being is hurt, and the result is universalised, the outcome is the principle not to hurt except in case of self-defence.

Because there are so many more issues involved in the question whether one may hurt another person (agreement, for instance), it would be overdone to call the more complex principle a rule. Still, this principle has in common with a rule that it replaces the original principles (not to hurt human beings, and self-defence is allowed) and that the application of these latter is excluded by the application of the complex principle. There can be principles and goals which are the result of combining other principles and/or goals and which replace these latter principles and goals.

There can also be rules which do not replace *all* principles and goals which might be relevant for the type of case with which the rule deals. For instance, the rule formulated above that if a party in good faith obtained a right from a person who did not have the capacity to transfer the right, this party still obtained the right, does not take the factor into account whether the original owner was to be blamed for loosing the power over his property. It may be argued that an owner who lost his good through theft deserves a better protection against parties in good faith than the

---

[11] My use of the notion of an exclusionary reason is similar but not identical to Raz's use. According to Raz, an exclusionary reason is a second order reason not to act for some reasons, while in my view an exclusionary reason is a reason not to apply a rule, goal or principle, with the effect that some reasons do not even arise. For the present purposes this difference has only theoretical significance.

owner who lend out his property, thereby causing the appearance that somebody else had the capacity to transfer the ownership. A factor that was not taken into account in formulating a rule is arguably not replaced as a reason by the application of this rule. A rule need not exclude the application of all possibly relevant goals and principles. There may be principles and goals that remain applicable in addition to an applicable rule [cf. Raz 1979, p. 22]. Such a rule is less of a rule than the stereotypical one.

There seems to be no clear boundary between, on the one hand, complex goals and principles which are the result of combining other principles and goals, and, on the other hand, less than perfect rules, which do not take into account all factors that are relevant for the type of case they deal with. One extreme is the stereotypical principle, which deals only with one aspect of a case, and which does not exclude any other principle or goal. The other extreme is the prototypical rule which takes into account all factors that may be relevant for the type of case governed by this rule, and which excludes the application of all principles and goals that deal with factors which occur in cases of this type. Most principles and rules will be somewhere in between these extremes. In fact, because it is practically impossible to determine in advance all aspects which may be relevant for all cases which belong to a particular type, stereotypical rules probably do not even exist.

10. Reasoning with rules

Still there remains an important difference between on the one hand rules, and on the other hand principles and goals, and this difference is that most rules are the product of some kind of decision making, while principles and goals generally are not [Raz 1975, p. 71]. This difference is especially clear in the contrast between rules which are the product of legislation and unwritten principles and goals, but is also present if rules are abstracted from case law. Where rules are the product of decision making, there are additional reasons to apply them if their conditions are satisfied, based on authority, economy of deliberation and - in the case of legal rules - on legal security.

Because rules generate exclusionary reasons, reasoning with rules is somewhat different from reasoning with principles. If a rule is applied to a case, there is a reason for the conclusion of the rule for this case. This reason replaces all other reasons which might be relevant for the case, and the application of the goals and principles which might underlie these reasons would be excluded. The reason generated by the rule is the only available one, and the rule governs the case on its own. The rule applies in an all-or-nothing fashion, just like Dworkin described the application of rules.

Because ideally a rule takes all relevant factors in a case into account, for all possible cases there should only be one rule that deals with that case. In other words, it should not be possible that rules conflict. This unrealistic ideal is to a certain extent upheld by the presence of conflict rules that dictate which rule is applicable in case of a rule conflict. The application of the other rule is then excluded. Notice that in the case of principles and goals, conflict rules are superfluous. Collisions are so to speak built into goals and reasons, because goals and principles only have an eye for one or a limited number of the relevant aspects of a case. These collisions are dealt with by weighing the colliding reasons, not by excluding some of them.

Because sometimes a rule does not take all relevant aspects of a case into account, it may be the case that the rule is not applied, even though it is applicable. The reasons which were 'overlooked'

by the rule may become reasons not to apply the rule to this particular case. Moreover, there may also be reasons to apply a rule to a case although the conditions of the rule are not (all) fulfilled in this case. Most of these aspects of reasoning with rules are dealt with more elaborately in the sections 14 and 15.

11. The two step-model of reasoning with principles and rules

Principles contribute to the conclusions of arguments in that they generate reasons for their conclusions. A model of reasoning with principles should incorporate a step that determines whether the reasons generated by a principle are decisive. That is why an adequate model of reasoning with principles contains a step that deals with the interaction of the reasons generated by the principles.

Therefore there are two steps involved in the model of reasoning with principles that is presented here. The first step consists of applying all relevant and valid principles, in order to gather all reasons which plead for or against a conclusion. The second step consists of a 'weighing' process, in which it is established which conclusion follows from the collected reasons.

The first step can again be subdivided in three steps. For each possible principle which deals with the conclusion of the argument, it must first be established whether it is valid. Second, it must be decided whether the principle's conditions are satisfied by the case to which the principle is to be applied. And third, it is to be determined whether the principle, given the (non)satisfaction of its conditions, is to be applied to the case at hand.

The reasons which were found by applying principles are divided in a set of reasons which plead for the conclusion and a set of reasons which plead against the conclusion. The weighing of reasons can be dealt with in a rational manner if one of these sets can be argued to outweigh the other set. If this is not possible, there is no rational solution to the weighing problem.

The main elements of the two step-model are depicted in figure 4. In the following sections, I discuss all of the steps involved in the model.

Because many legal rules will not be perfect in the sense that they do not incorporate all factors which might be relevant for the type of case they are dealing with, reasoning with legal rules can be treated as a special kind of reasoning with legal principles, with the difference that the application of a rule excludes the application of the principles and goals which it replaces. Even this difference is not absolute, because there will also be complex legal principles, which replace more elementary principles and goals, and the application of which excludes the application of these more elementary principles and goals.[12]

[12] In [Hage Reasoning] I make a more strict distinction between principles and rules. There I assume that if a rule applies to a case, the rule's conclusion invariably follows.

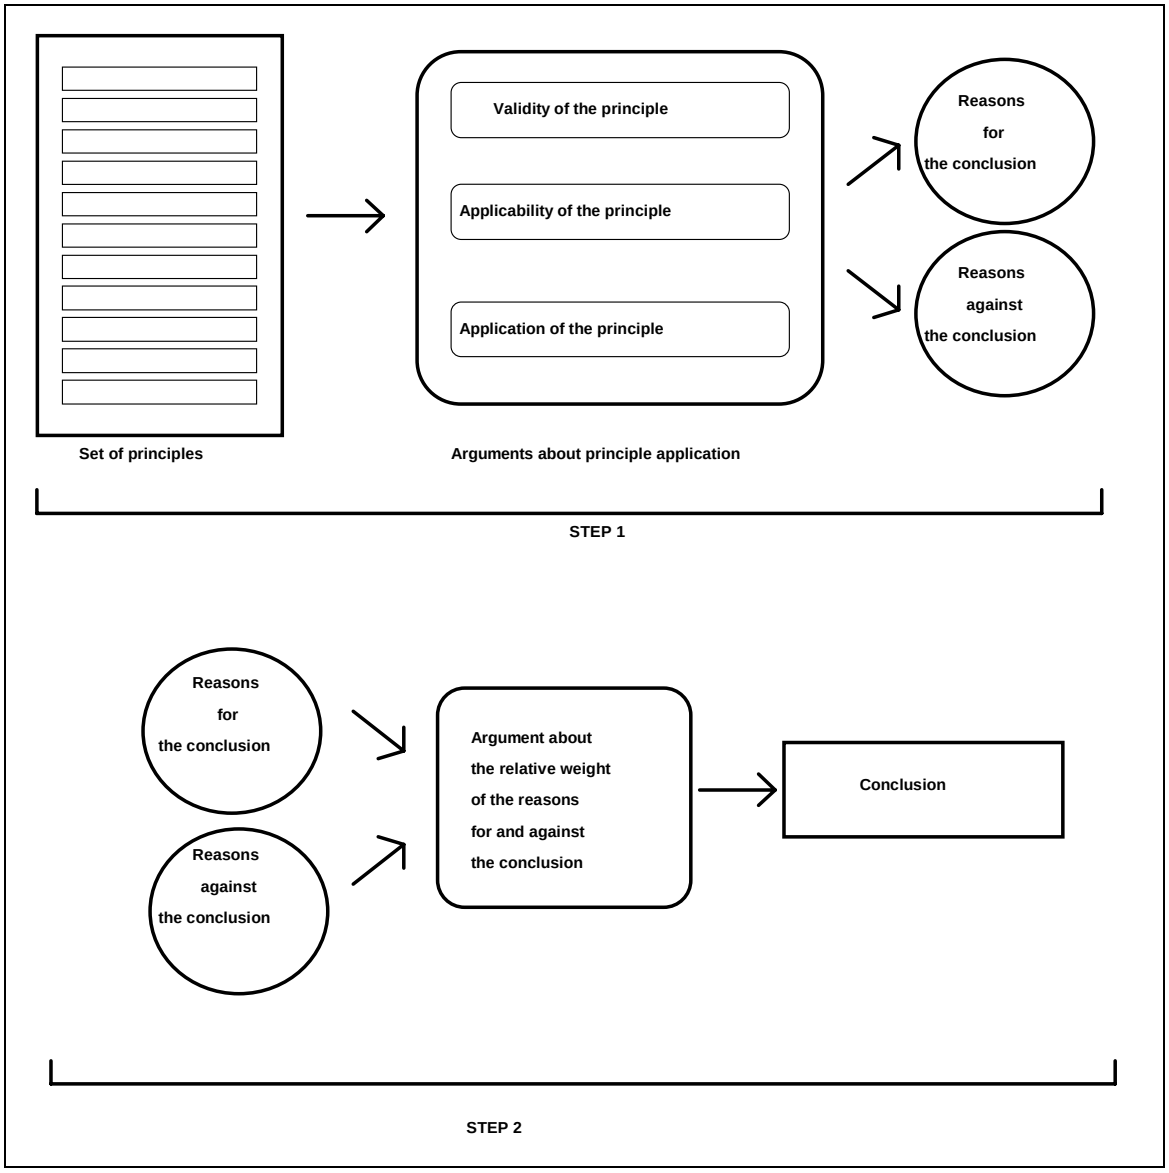

Figure 4: The two step-model of principle application

For this reason, I will write primarily about principles. The discussion of principles also applies to rules, unless the contrary follows from the text. Goals differ from principles in the way they generate reasons. The reasons generated by goals function just like reasons generated by principles.

12. The validity of rules, principles, and goals

A rule, principle or goal can only generate reasons if it is a valid one. Whether a goal, principle or rule is valid or not often depends on the point of view. Legal definitions, for instance, are only valid for the legal point of view. If the validity of rules is attached to a particular point of view, this attachment also holds for the reasons based on these rules, and for the conclusions based on the reasons. The facts that somebody took away somebody else's property with the intention of

appropriation are a *legal* reason, and the conclusion based on that reason is that this person is, *legally speaking*, a thief. The legal classification may or may not coincide with the same classification from other points of view, but this does neither add nor subtract from the relativization of the reasons and the conclusions to the point of view of the rule on which they are based.

Practically this means that if we are engaged in a legal discussion and we collect rules to base reasons on, these rules must be *legally* valid. Legal conclusions must be based on legal reasons, which in their turn must be based on legal rules (or legal principles, or legal goals, etc.).

This is a heavy demand, which is relaxed somewhat because intermediate conclusions, most often classifications, need not be based on legal reasons. Sometimes there are legal rules which deal with classification. More often, however, the law borrows the conditions of applicability for the terms it uses from common parlance [cf. Valente 1995, p. 57].

The issue of the validity of legal rules has also some impact on legal interpretation. Rules can be considered as consisting of a conclusion and a number of conditions. Legal rules are often based on legislation. The results of legislation are texts which are meant to contain the contents of the rules. Most often they do not exhibit a clear structure of conditions and conclusions.

If the law is to be applied to a concrete case, we must on the one hand have a legal rule and on the other hand a case description, such that the conditions of the rule match the description of the case. For instance, if a rule prescribes that vehicles are to drive on the right hand side of the road, it must be applied to a case described in terms of vehicles, rather than, for instance, cars or bicycles. In other words, the conditions of the rule and the facts of the case should be on the same level of abstraction.

Interpretation and classification deal with respectively the rule side and the case side of obtaining this match. The rule formulation is obtained by the interpretation of *legal sources*, such as statutes, social practice, and case law. The case description is obtained by classifying the most 'brute' description of the case.

To a certain extent, it is immaterial which part of the way from the texts of the legal sources to the text of the brute case description is travelled by means of interpretation, and which part by classification. It does not matter whether the rule is given a rather concrete interpretation, or whether the facts are classified rather abstractly.[13] For the present discussion I propose to cut the knot as follows: Issues concerning the precise conditions and conclusion of a rule are the topic of interpretation; issues concerning the denotation of these conditions and conclusion are a matter of classification. On this proposal, issues concerning syntactic ambiguity are dealt with by interpretation, while issues concerning semantic ambiguity (vagueness, open texture) are dealt with by classification.

Given this convention, different interpretations of a statute lead to different rules. Arguments for a particular interpretation of a statute can then be considered as arguments why a particular rule, rather than another one is valid on the basis of the statutory text.

Often the sources of law lend themselves to more than one interpretation. This has given rise to extensive methodological discussions, especially concerning the interpretation of statutory texts.

---

[13] The distinction between interpretation and classification is legally relevant where courts are only allowed to deal with issues of the law, and not with factual issues. Classification is easily (and - in my opinion - wrongly) seen as merely a factual issue.

In my opinion, the criteria to answer the question which interpretation of a text is to be preferred over another interpretation, are part of the law itself.[14] Sometimes they are even the object of legislation, which means that there are rules that deal with the interpretation of legal sources. Most often, however, the *principles of interpretation* - because that is what we are talking about, will be social rules. These rules or principles can be applied to generate reasons concerning the validity of (other) legal rules. In that connection, there can again be a discussion about the validity (acceptance) of these principles in turn. Cf. figure 5.

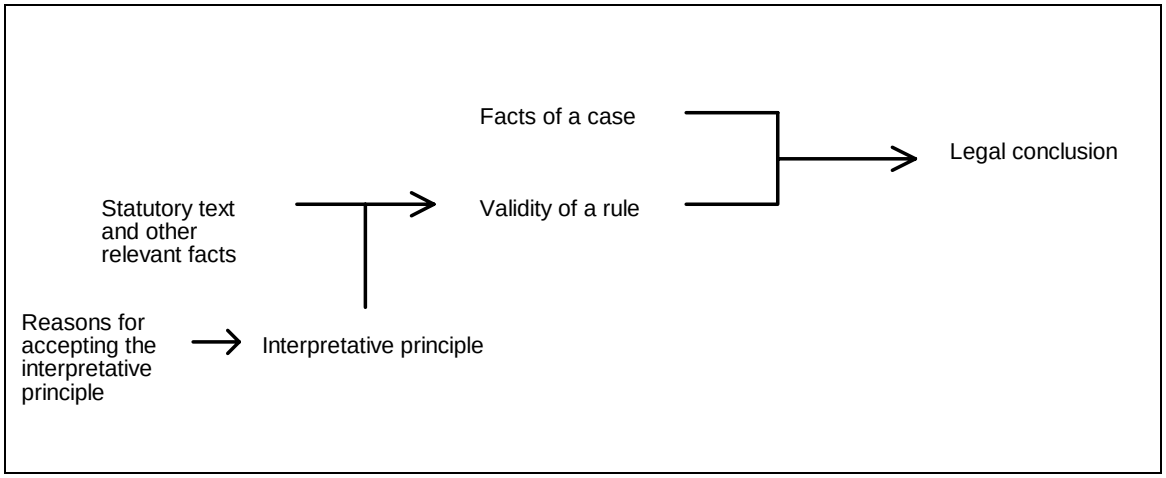

Figure 5: Interpretative principles

13. The constitution of reasons

The most elementary case of the constitution of a reason is where a principle or rule has one or more conditions which must all be satisfied in order for it to apply. Suppose that John took away the car of Ellen, with the intention to appropriate this car. In that case, the facts that 1. John took away the car, 2. The car belonged to Ellen, and 3. John intended to appropriate the car, *together* become a reason why John is a thief. Neither one of these facts by itself is a reason.

If a rule has alternative conditions, it can generate different reasons. Take for instance the rule 'If a person P was either physically incapable to perform his duty, or if P had a justified mental inhibition to do his duty, P acted under force majeure'. This rule has two conditions, but they are alternatives for each other. One can act under force majeure either because of physical impossibility, or because of justified mental inhibition. Both physical impossibility and justified mental inhibition constitute a reason for being under force majeure.

In addition to reasons that plead for a particular conclusion, there are also reasons that plead against a conclusion. If a rule has a negative conclusion, the facts that satisfy the rule's conditions are potential reasons against the fact that is denied in the rule conclusion.

A more subtle case is where the conclusion of a rule is not the denial of a particular thesis, but is incompatible with it. If we have the rule that somebody is a coward if, as a soldier, he flees at the approach of the enemy, the fact that soldier George's flees at the approach of the enemy is a

[14] Cf. [Peczenik 1989, pp. 375 f.] on reasoning norms.

reason against George's being a hero. It is assumed that being a hero and being a coward are incompatible, and therefore a reason for the one fact is ipso facto a reason against the other fact.

The basic mechanism by which goals generate deontic reasons is that if an action contributes to the goal, this is a reason why this action ought to be performed. Analogously, if an action detracts from a goal, this is a reason not to perform that action.

This basic mechanism can sometimes be extended from actions to states. For instance, the fact that continuation of a rent contract in case the rented house changes owners, contributes to the protection of the lessees, is a reason why the rent contract ought to be continued with the new owner.

14. The exclusion of rules, principles and goals

14.1 Scope limitations

Suppose that in the Netherlands there is a valid legal rule which says that those who murder somebody else are punishable and that in the United Kingdom a murder takes place. Suppose moreover, that both the murderer and the victim were inhabitants of the UK and that the case exhibits no connections with the Netherlands or with Dutch law at all. It is clear that the rule from the Dutch criminal law is not applicable to this case, although its conditions (which do not mention the nationality of the murderer or his victim) are satisfied.

This example illustrates that the satisfaction of a rule's conditions does not guarantee the applicability of the rule. The reason for the non-applicability of the rule in this example is that the case falls outside the scope of the rule. In addition to their conditions for applicability, rules have scope limitations which limit the application of the rule in, amongst others, time and space.

In general we must distinguish between the span of time in which a rule of law is valid, and the span of time to which it may be applicable. Normally these two time spans (more or less) coincide, but there may be exceptions. Retroactive laws are a case in point, because they are applicable to cases at a time in which the rules were not yet valid.

Usually scope limitations are not considered as part of the conditions of the rule whose scope they concern. There are several good reasons for that. First because scope limitations in the law often apply to sets of rules. It is more economical to state them in general for the whole set of rules to which they apply, rather than to add the limitations to the conditions of all concerning rules.

Second, and more importantly, the logical role of scope limitations is different from the other conditions. Often the satisfaction of the scope limitation is assumed and the non-satisfaction must be proven. The division between rule conditions on the one hand, and scope limitations on the other hand, has a corollary in the division of the burden of proof.[15]

Finally, if a legal rule is not applicable because its scope limitations are not satisfied, this leads to a different legal verdict than in case its normal conditions are not satisfied. In a Dutch criminal case, for instance, non-satisfaction of scope limitations leads to the dismissal of the public prosecutor or dismissal of the suspect, while non-satisfaction of conditions leads to acquittal of the suspect.

---

[15] Cf. the distinction made in [Sartor 1991] between principal and secondary facts in legal language, in particular the distinction between facts to prove and facts to which contrary prove must not be given. Cf. also Baker's views of defeasibility as described in [Baker 1977].

14.2 Exceptions and rule conflicts

Scope limitations are not the only factor that can make a rule non-applicable, even though its conditions are satisfied. Exceptions to a rule have a similar effect. Sometimes, exceptions to legal rules are just negative rule conditions. It must be shown that the exception does not apply to prove that the rule is applicable. Exceptions of this type are normally included in the formulation of the rule in the source from which the rule is derived.

Very often, however, exceptions to rules are formulated outside the provisions on which the rules are based. In these cases exceptions usually are like scope limitations in that they describe special circumstances under which the rule should not be applied. I will call these exceptions *proper exceptions*. Typical cases in point are situations of force majeure in civil law, and grounds for justification or exculpation in criminal law. Proper exceptions can both be explicitly formulated in rules of law, and be implicit.[16]

For proper exceptions often holds the same as for (other) scope limitations. Their absence is assumed by default, and if they are present, they lead to other verdicts than the mere non-satisfaction of the rule conditions. If a crime is committed under force majeure, for instance, the verdict will be dismissal. Exceptions can apply to categories of rules (e.g. all rules of criminal law, or all laws of contract), but they can also apply to individual rules in particular cases.

Scope limitations that limit the application of a rule in time and place and proper exceptions can make a rule inapplicable although the conditions of the rule are satisfied. Their effects are absolute: if a case falls outside the scope of a rule, or if there is a proper exception to the rule, the rule is not applicable anymore with the consequence that the rule will not be applied. In the terminology of reasons, this means that scope limitations and exceptions give rise to *exclusionary reasons*. If a scope limitation is not satisfied, or if an exception applies, the application of the concerning rule is *excluded*.

In the law, a particularly important category of scope limitations derives from the phenomenon of rule conflicts. Legal ideology will have it that rules of law do not conflict. If two rules seem to conflict, at least one of them is not applicable. The scopes of conflicting rules are disjoint.

The exclusion of a rule is to be distinguished form the invalidity of the rule. If a legal rule is invalid, this means that legally it does not even exist, and as a consequence it cannot generate any reasons at all. Exclusion, on the contrary, is case-related. A rule can only be excluded if it is valid; exclusion means that the valid rule cannot be applied in the particular case(s) in which its application is excluded.

The exclusion of goals is similar to the exclusion of rules and principles. It means that even if some action or state contributes to some adopted goal, there may nevertheless be no reason for performing that action, or for the presence of this state, because the goal is irrelevant *in this particular case*.

15. Reasoning about rule application

The possibility that the application of rules and principles is excluded, and that goals are irrelevant illustrates the instrumental nature of goals, principles, and rules. They must be used to generate

[16] Prakken pointed out to me that sometimes proper exceptions are mentioned in the rule to which they form an exception, as in ' ...unless the law indicates otherwise'.

reasons. This opens the possibility to argue about whether they should be used. In the following subsections, we will study the reasons for and against rule application.

15.1 Reasons not to apply an applicable rule

In the context of rule application, the term 'applicable' has been used as a technical term. Applicability of a rule is distinguished from the satisfaction of the rule's conditions. The latter is a necessary, but not a sufficient condition for the former. A rule is applicable to a case, if and only if its conditions are satisfied, and its application is not excluded.

In addition to the distinction between the *applicability* of a rule and the *satisfaction* of its conditions, I will distinguish between the applicability and the *application* of a rule. Applicability is a reason to apply a rule. Unless there are other reasons not to apply the rule, an applicable rule should be applied. Moreover, applicability is the major reason why a rule should be applied. The normal situation is that a rule should be applied if it is applicable, and that it should not be applied if it is not applicable.

If application of a rule would be *against the purpose* of the rule, this is a reason not to apply the rule. Fuller gave an example of a prohibition to sleep in the railway station, which was motivated by the desire to retain tramps from spending their night on the station [Fuller 1958]. It would be against the purpose of this rule to apply it to the traveller who dozed away a few minutes while waiting on a retarded late train.

If there is a 'normal' reason against the application of a rule, this reason must still be weighed against the reasons for application, the most important of which will be that the rule is applicable. This situation differs from the one in which there is an exclusionary reason against application of the rule, because exclusionary reasons against rule application are not weighed against reasons to apply the rule. The fact that the conditions of a rule are satisfied is not a reason to apply American Federal law to an ordinary Dutch case. The fact that the rule is applicable to the late traveller from Fuller's example still counts as a reason for applying the rule, although this reason may be outweighed by the reason that such application would be against the rule's purpose.

15.2 Analogous application of a rule

A rule is applied analogously if it is applied to a case in which the rule conditions are not satisfied, but that sufficiently resembles the cases in which the rule conditions are satisfied to justify the application of the rule.

Some authors argue that analogous rule application does not actually occur, but that sometimes a more abstract rule is applied to a case that does not satisfy the conditions of the original rule.[17] For example, a Dutch rule of law says that if a house (some good) is sold, the existing rent contract is continued with the new owner. This rule appears to be applied analogously to the case where the house is not sold, but the property is transferred on the basis of some other title, say donation. According to the authors that reject analogous rule application, the seeming application of the rule about the sale of a house to other cases in which the property of the house was

[17] This view of analogous rule application is defended in [Tammelo and Schreiner 1977, p. 112 f.], and [Prakken 1993, p. 22]. Prakken would say that the analogy is not involved in the application of the rule, but in the suggestion of the rule's new formulation. In this he follows [Soeteman 1989, p. 239].

transferred only shows that the applied rule is actually about transfer of houses (or at least also about the donation of houses), and not only about the sale of houses. Selling and donating a house are then special cases to which the more abstract rule (or principle) can be applied in a normal fashion.

The inclination to treat analogous rule application as the application of a more abstract rule is explainable from the tendency to consider rules from a logical point of view as a kind of statements. Clearly statements cannot be applied analogously. The only way to deal with analogy in case of statements is to assume that the statements were actually more abstract than they originally seemed to be.

However, as soon as it is recognised that rules are not statements, and that they can obey a different logic than statements, analogous rule application becomes easy to explain. Rules should be considered as a kind of tools, applied by humans to structure the (legal) world. For instance, we use the rule that thieves are punishable to create a connection between the facts that somebody is a thief and that this person is punishable. This rule is normally used in the cases that somebody is a thief; these cases are the normal conditions in which the tool is used. Sometimes, however, a tool may even be used if the normal circumstances are not present, but some action is nevertheless necessary. For instance, we encounter a case in which somebody illegally copies software. This is not a real case of theft, but we need a legal tool to deal with this kind of case. Analogous rule application might be considered as such a case in which a rule is used under non-ideal circumstances. If it is applied analogously, the rule generates a reason for its conclusion, even though its conditions are not satisfied. (However, there are serious objections against applying criminal laws analogously.)

The account of the relation between legal rules and legal principles and goals that was given above makes it relatively easy to deal with analogous rule application. Let us reconsider the example about the rule that a rent contract is continued with the new owner if a rented house is sold. This rule is applied analogously to cases in which the property of the house is transferred on the basis of another title than sale.

To obtain an account of what happens, we should consider the legal goal that underlies the rule. This goal is that the persons that rent a house should be protected in their right to live in the house. This goal is elaborated in the conclusion of the rule that if the house is sold, the rent contract is continued with the new owner. Notice that the elaboration is much more specific than the goal itself. Other elaborations with a comparable effect are imaginable, such as that in order to be able to sell a rented house, the owner should offer the inhabitants an equivalent alternative.

If a rented house is sold, the legal rule generates a reason why the rent contract is continued with the new owner and excludes the application of underlying legal principles and goals such as the principle that contracts only are in force between the contracting parties, and the goal that the inhabitants of a rented house should be protected in their right of inhabitation. This principle and this goal were both taken into account in making the rule, and if the rule is applicable, the application of the principle and the goal are excluded.

What happens if a case does not satisfy the conditions of the rule? One possibility is that the rule is not applied to the case. In that case the rule does not generate a reason for the solution of the case, but neither does it generate exclusionary reasons that prevent the application of the mentioned legal goal and principle. If, for instance, the house is donated, and the rule about sale of the house is not applied, the case must be decided on the basis of the principle that contracts

only are in force between the contracting parties, and the goal that the inhabitants of a rented house should be protected in their right of inhabitation. Both the principle and the goal would generate a reason, a reason why the inhabitants should be able to remain in their house, and a reason why they have no right to remain in the house at all. These reasons must be weighed to obtain a solution for the case. Notice, however, that neither reason explicitly deals with the continuation of the rent contract with the new owner. The principle and the goal are much too abstract to generate reasons against, respectively for such a specific solution.

Another possibility is that the principle and the goal generate a reason against, respectively for, application of the rule about sale of the house. The normal reason for applying a rule is formed by the facts that make the rule applicable. In the present case these facts would be that a house is sold and that this house is rented. It is not impossible, however, that there are other reasons for applying a rule, because applying a rule is a kind of acting, and there may be all sorts of reasons to act in a particular way. One reason to apply the rule about selling a house in case the house is donated is that the legal goal of protecting the rights of inhabitants of rented houses is served by applying the rule. A reason not to apply the rule is that application would be in conflict with the principle that contracts only are in force between the contracting parties. Another reason would be based on the demand for legal security.

Reasons to apply a rule even if its conditions are not satisfied, will usually be based on principles or goals that led the legislator to make the rule in the first place. Therefore the cases to which the rule is applied, although the rule conditions are not satisfied, will normally resemble cases to which the rule is applicable. That is why we normally speak of 'analogous application' of a rule, if a rule is applied to a case in which it is not applicable.

In the first way of dealing with an inapplicable rule we have direct reasons for and against measures that give the inhabitants the possibility to continue their inhabitation of the house. The conclusion of these reasons is still rather vague. In the second way, on the contrary, we obtain reasons that indirectly (namely via the intermediary of rule application) plead for and against a much more specific solution of the case, namely continuation of the rent contract. Both situations, application of a principle underlying a rule, and real application of a rule although its conditions are not satisfied, will occur in legal practice.

The reasons to apply a rule even if its conditions are not satisfied can also be interpreted as reasons to apply a more abstract rule which is directly applicable to the case at hand. This option might be preferred by the authors mentioned above, who argue that analogy is not an inference scheme, but rather a heuristic to suggest more abstract rules which can be applied 'normally'. The advantage of this option, that no deviant inference scheme is needed, is balanced by the disadvantage that the more abstract rule will not be legally valid.[18]

15.3 E contrario 'application' of a rule

If a rule is not applicable, it may still be applied. Analogous application leads to a reason for the rule's conclusion even if the conditions of the rule are not satisfied. The opposite can, however, also happen. Sometimes it is possible to obtain a reason *against* the rule's conclusion if the rule's conditions are not satisfied. In such cases the rule is said to be used in an argument *e contrario*.

---

[18] A related third option is discussed in [Verheij and Hage 1995].

According to the traditional view, an argument e contrario presupposes that the conditions of the rule not only state a sufficient condition for the rule's conclusion, but also a necessary one (e.g. [Soeteman 1989, p. 39]). The rule would express an equivalence between its conditions and its conclusion. Such a view would go too far, however. First because it leaves unexplained why of two rules which superficially seem not to differ, one states only a sufficient condition for its conclusion, while the other one states an equivalency between its conditions and its conclusion.

Second, because the conclusion of a rule that is used e contrario might follow on the basis of some other rule. For instance, the Dutch law has a rule that forbids widows to remarry within 306 days after the dead of their husbands. The ratio of this rule is to avoid confusion about who is the father of a child of which the woman might be pregnant. Clearly this prohibition is not applicable to widowers, but just as clearly it would be wrong to interpret the condition that a widow is involved as a necessary condition for the prohibition of a marriage. A woman may also be prohibited to marry because she is already married. The conditions of a rule can be a necessary condition for the *application* of the rule, but not for the *conclusion* of the rule.

E contrario 'application' of a rule can be explained by looking at the goals and principles that underlie the rule. If in a particular kind of case the conditions of a rule are not satisfied, this can have one of three causes. First, this may be because the rule has nothing to do with this kind of case. For instance, a tax rule has nothing to do with an insurance contract. In such a case the rule should completely be discarded.

A second possible cause is that the legislator overlooked this kind of case. In this situation, the case is to be solved on the basis of the underlying goals and principles, eventually leading to analogous application of the rule.

The third situation is that the kind of case we are dealing with has been considered by the legislator in drafting the rule, and that this consideration has led to a formulation of the rule conditions that made the rule inapplicable to this kind of case. Then we may assume that the legislator did not want the legal consequences of the rule to obtain in this kind of case, and this is a reason to conclude to the opposite of the rule conclusion. Of course there may still be reasons, on the basis of other rules, which plead for the conclusion of the rule.[19]

An example of this third case would be that the legislator made a law that says that tramps are forbidden to sleep in the railway station. By explicitly referring to tramps, the legislator made it clear that he did not want a general prohibition against sleeping in the railway station. The traveller who dozed away while waiting on a retarded late train is probably not covered by this rule. This means that such travellers are allowed to doze in the railway station, unless their behaviour falls under another prohibition.

Does this mean that in the last case the rule conditions stated both necessary and sufficient conditions for the rule conclusion? The answer is negative. The legislator did not deal with this kind of case *by means of legislation*. But we still have the underlying principles and goals, and the knowledge that the legislator decided that these principles and goals were no reason to let this kind of case fall under the rule. In combination this may lead us to the conclusion that for this kind of case we should solve the case by denying it the consequence of the rule. In other words, we 'apply' the rule e contrario. In fact, however, we do not apply the rule, but its underlying goals and principles.

---

[19] This may even have been intended by the legislator who decided to make the rule inapplicable.

16. Defeasibility of reasoning with rules, goals and principles

Rules and principles differ from statements not only in their direction of fit, but also in that they need to be applied. Application of premises does not play a role in traditional logical in determining whether inferences are valid. If all murderers may be punished and John is a thief, it logically follows that John may be punished, in the sense that this conclusion must be true if the premises are true. There is no need to 'apply' the premise that all murderers may be punished. The premise has a certain information content, and given this content it must be the case that John may be punished. In fact, it is not clear at all what application of a factual statement would amount to.[20]

Rules and principles, on the contrary, need to be applied. The principle that murderers may be punished, for example, must be applied if it is to generate a reason to punish John. Mere validity of this principle is not sufficient.

Most often, the application of this principle will be 'automatic'; it is not given much thought and as a consequence we may overlook the necessity of application. However, it is always possible to adduce reasons why a principle or rule should not be applied in a particular case, and this possibility only makes sense if application of the principle is an issue.

The fact that rules and principles need to be applied explains two characteristics of reasoning with rules that cannot be explained if rules are treated as statements. The first characteristic is that reasoning with rules (and principles) is defeasible[21]; the second one is that rules can be applied analogously.

Analogous rule application was already dealt with in section 15.2. Reasoning with rules, goals and principles is also defeasible. In connection with rules and principles, defeasibility can be defined as the situation where the conditions of a valid rule or principle are satisfied, but the conclusion nevertheless does not follow.

The first reason for defeasibility is that rules and principles need to be applied in order to contribute to their conclusions. If there are decisive reasons not to apply them, they do not generate reasons for their conclusions, and these conclusions will most probably not follow.[22] More or less the same holds if a goal is left out of consideration because of its irrelevance in a particular case.

The second reason is that goals and principles[23], even if they are applied, only generate reasons for their conclusions. If there are reasons against the conclusion that balance or even outweigh the reasons generated by the applied goal or principle, the conclusion does not follow, even though the goal or principle was applied.

---

[20] One might argue that application of a statement is the use of this statement in an argument. However, this is only relevant for the *proof* of the conclusion, not for the issue whether the conclusion logically *follows* from the premises.

[21] It seems that reasoning with incomplete knowledge is also defeasible, even if this knowledge consists solely of statements. This overlooks however that reasoning with incomplete knowledge involves the use of inference rules, and it is precisely this use of inference rules which causes the defeasibility.

[22] At least theoretically it is possible that the conclusions follow for other reasons than those based on the rule or principle involved.

[23] I have some doubts whether this also holds for rules. Cf. also note 12.

17. Weighing reasons

Every principle or rule with a particular fact in its conclusion is a potential generator of reasons for that fact. Conversely, every principle or rule with the negation of a particular fact in its conclusion[24] is a potential generator of reasons against that fact. If each of these principles and rules is checked for its application in a case, a set of zero or more reasons for, and a set of zero or more reasons against the fact obtain. The reasons from these two sets must be 'weighed', to determine whether the fact, its negation, or nothing at all follows as a conclusion.

If we have no reasons why one set of reasons outweighs another set of reasons, we cannot make a rational decision. Sometimes, however, we are able to argue which sets of reasons outweighs the other set. In a legal context we can, for instance, refer to precedent: If we already had an earlier case where the same reasons obtained, and where one set of reasons was assumed (for whatever reason or cause) to outweigh the other set, we can call to the principle that similar cases are to be treated similarly, and make the same weighing as the previous time.[25]

In a model for reasoning with rules we should ask for reasons why a set of reasons outweighs another set. If such 'meta'-reasons are lacking, there are no rational means to obtain a conclusion. If such meta-reasons are available, it should be possible to draw the conclusion which is supported by the strongest set of reasons. One meta-reason is very obvious in this connection: If we only have reasons that plead for a conclusion, and no reasons against it, the reasons for outweigh the reasons against. Analogously, if there are only reasons against, they 'outweigh' the (non-existent) reasons for the same conclusion.

18. Summary and recommendations

Reasoning with principles, goals and rules turns out to be a lot more complex than simple rule application according to the model of Modus Ponens. On the one hand, rules can be applied (analogously) even though their conditions are not satisfied. On the other hand, even if a rule's conditions are satisfied, it is not guaranteed that the rule will be applied: a case that satisfies the rule's conditions may fall outside the rule's scope or purpose. And finally, in the case principles, even if a principle is applied to a case and generates a reason, there is no guarantee that its conclusion follows. Weighing reasons may lead to the result that the principle's conclusion is still not drawn. On the basis of these characteristics of reasoning with rules, goals and principles, we can draw up a list of properties we would want in a logic that deals with rules. These properties both concern the expressive power of the logical language, and the nature of the rules of inference.

Although weighing reasons will only play a role in those arguments in which there are both reasons for and against a particular conclusion, a logic for goals and principles should offer facilities to deal with the weighing of reasons. It should contain a rule of inference to the effect that a conclusion follows if the reasons that plead for it override the reasons against that conclusion.

In this connection, the logic should also offer the possibility to conclude that a particular set of facts forms a reason for or against a conclusion. As a consequence, it must be possible to refer to

---

[24] Incompatible conclusions will also do the job; contradiction is not necessary. Cf. section 13.

[25] This approach to weighing reasons is extensively discussed in section 34.

(sets of) facts and reasons. Since it is also necessary to reason about the facts that make up reasons, and to build on these results while arguing with the reasons, the logical connection should be maintained between the sentences that express certain facts, and the sentences that refer to these same facts as reasons. For instance, we want to be able to argue that John is a thief, to refer to the fact that John is a thief, and to maintain the logical relationship between the sentence 'John is a thief' and the fact that John is a thief.

It should also be possible to reason about the relative weight of sets of reasons.

A rule or principle can only be applied if it is valid (accepted). An important part of legal reasoning deals with the question as to the precise conditions of legal rules, a discussion which can be seen as dealing with the validity or acceptance of rules and principles. A logic for rules should make it possible to deal with arguments about the validity of rules. As a consequence, it should allow reference to rules. Reference to rules is also necessary for arguments about the exclusion, the applicability, and the application of a rule. At the same time a logic for rules should be able to use the internal structure of rules to model the application of rules in generating reasons. All of this also holds for the adoption and the relevance of goals.

The rules of inference that deal with rule application should make the application of a rule dependent on the reasons that plead for and against application. The applicability of a rule should be an important reason for application of this rule, but there must be a possibility to have other reasons for (analogous) and against application of a rule.

B. Reason-based logic

In the previous sections I have offered a theory of legal reasoning as based on rules, principles and goals. This theory gives an account of reasoning with rules, etc that does not fit well in, for instance, predicate logic. In this chapter, I will introduce a logic that was especially developed to deal with the characteristics of rules, principles, goals, and reasons that we encountered above. Because reasons play a central role in this logic, the logic is called Reason-Based Logic (RBL).[26]

In the following sections, the essentials of RBL are exposed. In a semi-formal introduction, which takes the sections 19 to 28, I describe some characteristics of the logic that result from the demands posed by the nature of reasoning with rules, principles and goals. Moreover, I introduce the elements of the language of RBL. Then follows a formal characterisation of the derivability relation of RBL. (sections 29). The exposition of RBL is concluded with a number of applications of RBL to not too unrealistic examples (sections 30 to 36).

19. Ontological presuppositions of Reason-Based Logic

RBL uses a rich ontology, that includes principles, goals, rules, cases, states of affairs, facts, and reasons.

A *state of affairs* is that part of reality that is expressed by a sentence. For instance, the sentence

[26] RBL is the result of gradual development. This development has been laid down in a number of papers that describe different versions of this logic. Cf. [Hage 1991, 1993, and 1995; Hage and Verheij 1995]. The version described here is very similar, but not completely identical to the versions of the more recent papers. There are some differences with the version that is described in [Hage Reasoning] .

```
Thief(john)
```

expresses the state of affairs that John is a thief.

A *fact* is a state of affairs that is expressed by a *true* sentence. Facts are states of affairs that actually *obtain*. False sentences express states of affairs that do not actually obtain.

A *reason* is a fact that has a particular significance for a conclusion, in that it either pleads for or against it.

From a logical point of view, all states of affairs, including facts and reasons, are individuals, which means that they can be the subject of sentences. Note that this does not involve a mix up of object and meta-language, because states of affairs are not linguistic entities, although their presence presupposes language.[27]

RBL provides means to reason with sentences and also to reason about the states of affairs (facts, reasons) that are expressed by those sentences. Moreover, the logical connections between a sentence and the state of affairs expressed by it are maintained. To achieve this, RBL makes use of a special convention to achieve the logical connection between a sentence and the state of affairs that is expressed by it. All RBL-sentences begin with a capital and all terms (or function expressions) that denote individuals begin with an small character. Moreover, if a term is identical to a sentence, with the only difference that the term begins with an small character, this term *denotes* the state of affairs that is *expressed* by that sentence. Cf. figure 6.

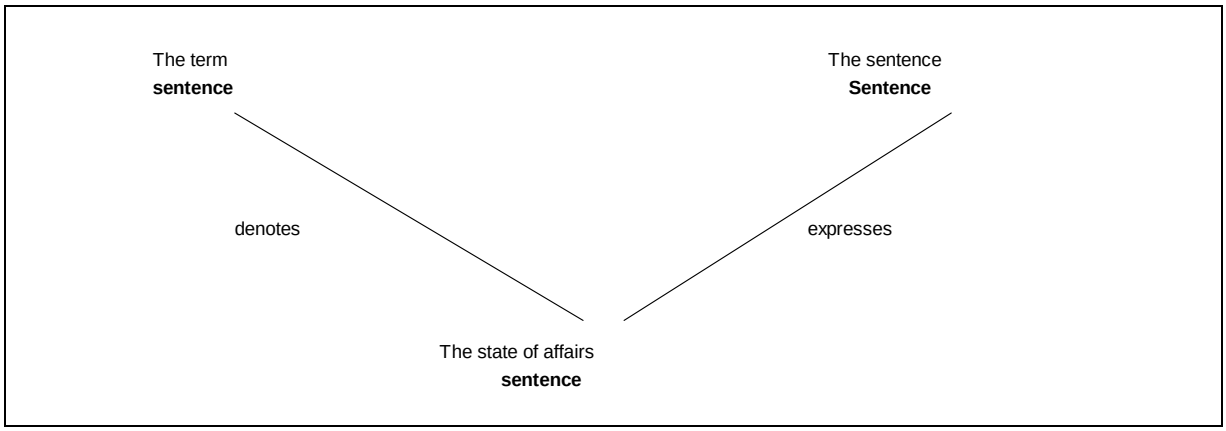

Figure 6: The relation between sentences and terms that deal with the same state of affairs

For instance, the term

```
thief(john)
```

denotes the state of affairs that John is a thief, a state of affairs that is expressed by the sentence

```
Thief(john)
```

Since facts and reasons are also states of affairs, the convention to denote states of affairs also applies to facts and reasons. This makes it possible to state about a particular fact that it is a reason for the presence of some other state of affairs. For instance, the following RBL-sentence expresses that the fact that John is a thief is a reason why John is punishable:

[27] Cf. in this connection [Strawson 1971, pp. 193f.] on the language-dependence of facts.

```
Reason(thief(john), punishable(john), pro)
```

The convention concerning the relationship between sentences and their corresponding terms also applies to logically compound sentences. This means that the term

```
contract(john, geraldine)& defaults(john)& damage(3000))
```

denotes the compound state of affairs that there is a contract between John and Geraldine, that John defaulted in this contract, and that the resulting damage is 3000 dollar. Moreover, it also holds for sentences that contain variables. The term corresponding to a sentence with free variables is a function expression.[28] For instance, the function expression in which *x* is a variable

```
killer_of(x, marcella)
```

denotes the person who killed Marcella.

20. Case-related facts

If we want to argue that a rule is applicable in a particular case, or that it applies or is excluded in a case, we must often refer to facts that obtain in that case or not. For instance, the fact that John is a thief in a particular case is the reason why the rule that thieves are punishable is applicable *to that case*. That is the reason why RBL needs a way to say that a particular state of affairs obtains in some case.

The RBL-sentences that say that a state of affairs obtains in a particular case have the following form:

`Obtains(`*`state_of_affairs`*`, case(`*`c`*`))`[29]

In this sentence, *`state_of_affairs`* stands for the state of affairs that obtains in the case denoted by `case(`*`c`*`).` Suppose that we have a particular case, called `case(john's_theft)`, in which John is a thief. The fact that John is a thief in this case is expressed by the sentence:

```
Obtains(thief(john), case(john's_theft))
```

Very often when a case is described, it will be clear to which case the facts belong. In such cases, the description of the facts as belonging to a particular case is rather cumbersome. That is why the case-related notation is only used when necessary. Normally facts are described by means of the sentences that express these facts without mentioning the case. For instance,

```
Obtains(thief(john), case(john's_theft))
```

will normally be replaced by

```
Thief(john).
```

---

[28] This convention does not hold for variables in the conditions and the conclusion of rules.

[29] Italicised terms are variables. Sometimes variables are used to indicate that only the form of the sentence in which they occur is relevant. The present sentence is an example of this use of variables.

21. Rules as structured logical individuals

The distinction between rules and principles is reflected only indirectly in RBL. The language of RBL only has rules, which are used to represent both rules and principles.

There are many kinds of arguments in which (RBL-)rules are involved, and rules play several roles in them. Most often the rule is just applied, and generates a reason for its conclusion. However, there are also arguments about the validity of rules, and about the application of rules to particular cases. In these arguments, a rule is the topic of the argument. To formalise arguments in which rules are the topic, there should be a possibility to refer to rules.

This need to mention rules leads to a characteristic of RBL that makes it different from most other logics, namely that rules are treated as logical individuals. As a consequence, it is not possible to use a rule as a complete sentence in RBL. Rules can only be introduced into an RBL-theory by making statements about them, e.g. the statement that a rule is valid.[30]

Traditionally, individuals are represented in formal logics by means of an unstructured character or string. The representation of rules as logical individuals in RBL is different, however. The rules of inference of RBL must be able to distinguish between the conditions and the conclusion of an RBL-rule, and as a consequence, the representation of these logical individuals exhibits a structure.

A rule contains three elements, that are kept together by the function symbol 'rule'. These elements are an identifier, which is a kind of name, the condition-part of the rule, and the conclusion of the rule. An example of a rule denoting expression would be:

`rule(theft1, thief(`*x*`), punishable(`*x*`))`

Because the identifier of a rule is unique, it is also possible to refer to the rule with the identifier *r* by means of the expression

`rule(`*r*`)`

Technically, the symbol `rule` can be seen as a function symbol that maps rule identifiers into rules.

If the rule identifier `r` is a constant, the term `rule(r)` is a function expression that denotes the rule `r`. If *r* is a variable, however, the expression `rule(`*r*`)` is a logical variable.

The condition-part of a rule consists of one RBL-term which denotes a state of affairs. If the natural language version of the rule has more than one condition, the term that denotes the conditions will be logically compound.

The conclusion of a rule must always be an elementary fact; it may not contain conjunctions or disjunctions.

Both the condition and the conclusion-part of a rule may contain variables, quantifiers and/or negations.

An example of a not too complex rule would be:

---

[30] Rules are, to a certain extent, also treated as individuals in the logics proposed by [Gordon 1994], [Sartor 1994], and [Prakken and Sartor 1995].

```
rule(rule_of_recognition,
 made_by_legislator(rule(r)) &
 ~abolished(rule(r)),
 valid(rule(r)))
```

This rule says that a rule is valid if it has been made by the legislator and has not been abolished. The name of the rule is `rule_of_recognition`. The use of variables in the rule conditions and the conclusion indicates that this rule of recognition can be applied to any rule.

Remember that a rule expression is not a well-formed RBL sentence. Well-formed sentences are for instance statements about rules. In theory, any kind of statement can be made about a rule. For instance:

```
Made_by_legislator(rule(theft1, thief(x), punishable(x)))
```

or, shorter,

```
Made_by_legislator(rule(theft1))
```

The statements about rules which have a special logical role in RBL are the statements that a rule is valid, that a rule is applicable, that a rule applies, and that a rule is excluded. These statements are discussed next.

VALIDITY

If a rule or principle is to be used in making a derivation, it should be valid. The statement that a rule or principle is valid has one of the following forms:

```
Valid(rule(r, rule_conditions, rule_conclusion))
```

or

```
Valid(rule(r))
```

The validity of a rule is expressed in a statement which can be derived in RBL like any other 'normal' statement. It can also be a premise of an argument.

APPLICABILITY

If the conditions of a rule are satisfied in a particular case, and if the rule is not excluded in that case, the rule is said to be *applicable* to that case. To state this, RBL has a special predicate `Applicable,` that has four parameters. The first parameter is meant for the rule that is applicable, the second parameter indicates the case to which the rule is applicable, while the third and fourth parameter stand for respectively the instantiated rule conditions and the instantiated rule-conclusion, which indicate respectively the facts that will be a reason on the basis of the rule and the conclusion for which these facts will be a reason.

Suppose that we have a case that consists of the facts that John is engaged in a contract with Geraldine, that John defaulted, and that the damage is $ 3000. The facts of the case called `contract`, are represented by:

```
Obtains(contract(john, geraldine), case(contract))
Obtains(defaults(john), case(contract))
Obtains(damage(3000), case(contract))
```

Moreover, we also have

```
Valid(rule(contractual_default,
 contract(x, y) & defaults(x) & damage(z),
 owes(x, y, z)))
```

(If x has a contract with y, x defaults, and the resulting damage amounts to z, then x owes y the amount of z.)

In this case, the following statement would normally - that is, if the rule is not excluded - be true:

```
Applicable(rule(contractual_default), case(contract),
 contract(john, geraldine) & defaults(john) & damage(3000),
 owes(john, geraldine, 3000))
```

The statement that a rule is applicable can only be derived on the basis of a valid rule, the conditions of which are satisfied, and which is (assumed to be) not excluded. It cannot be the premise of an argument.

APPLICATION

In RBL a distinction is made between the applicability of a rule and the application of a rule. There is a special predicate `Applies`, that has the same parameters as the `Applicable` predicate, and which is used to state that a rule applies in a particular case.

The following is an example of the use of the `Applies`-predicate:

```
Applies(rule(contractual_default),
 case(contract),
 contract(john, geraldine)& defaults(john) & damage(3000),
 owes(john, geraldine, 3000))
```

The statement that a rule applies can only be based on the weighing of reasons and cannot be a premise of an argument.

EXCLUSION

If there is an exclusionary reason not to apply a rule in a particular case, this can be stated by saying that this rule is excluded in that case:

```
Excluded(rule(contractual_default), case(contract))
```

The statement that a rule is excluded can only be based on the weighing of reasons and cannot be a premise of an argument.

22. Goals

In RBL, goals are structured individuals, just like rules. The function symbol `goal` has two parameters. The first parameter is the identifier of the goal, the second one denotes the state of affairs that is to be pursued. This second parameter must denote a literal state of affairs. The only allowed logical operator in it is the negation.

For instance, the term

```
goal(youth_protection, youth_protected)
```

denotes the goal with the identifier `youth_protection`, and this goal-state is that the youth is protected.
Because identifiers are unique, it is possible to refer to the goal `youth_protection` by the term

```
goal(youth_protection)
```

VALIDITY OF GOALS

I will use the predicate `Valid` not only to represent the validity of rules, but also to represent the validity of goals. So, the fact that the goal to protect the youth is recognised can be represented in the language of RBL by

```
Valid(goal(youth_protection))
```

or, longer:

```
Valid(goal(youth_protection, youth_protected))
```

The statement that a goal is valid can be derived in RBL like any other 'normal' statement. It can also be a premise of an argument.

CONTRIBUTION TO AND DETRACTION FROM A GOAL

A goal generates a reason for a conclusion if there is some action or state that contributes to, or detracts from the goal. Suppose that the action[31] that John is withheld from using drugs contributes to the goal `youth_protection`. This is expressed by the predicate `Contributes_to`, which goes with two parameters. The first denotes the action that serves the goal. The second parameter denotes the particular instance of the goal that is being served by means of an instantiation of the goal state. For instance:

```
Obtains(
 contributes_to(withhold(john, using_drugs), youth_protected),
 case(john's_drug_addiction)
 )
```

It can also be expressed that some action detracts from a goal. To keep the language of RBL relatively simple, the fact that some action detracts from a goal is represented as the fact that this action contributes to the absence (negation) of the goal state:

```
Contributes_to(encourage(john, using_drugs), ~youth_protected)
```

Not only actions, but also states of affairs (if they obtain) can contribute to, or detract from a goal state. The sentences that express this have the same form as the sentences that express that action types contribute to or detract from goal states, with the difference that reference to an action type is replaced by reference to a state of affairs:

```
Contributes_to(owner(gerald), parties_in_good_faith_protected)
```

---

[31] In this connection, actions are always action types, such as stealing, and not action tokens, such as this particular instance of John's stealing.

(If Gerald becomes the new owner, this contributes to the protection of parties who acted in good faith.)
The statement that something contributes to a goal can be derived in RBL like any other 'normal' statement. It can also be a premise of an argument.

EXCLUSION OF GOALS

A goal only generates a reason in a case, if it is not excluded in that case. The exclusion of goals is represented by the same `Excluded` predicate that is also used in the case of rules and principles. For instance, the sentence

```
Excluded(goal(privacy_protection), case(royal_divorce))
```

expresses the state of affairs that the goal `privacy_protection` is irrelevant in a particular case of a royal divorce.
The statement that a goal is excluded can only be based on the weighing of reasons and cannot be a premise of an argument.

23. From the validity to the applicability of a rule

The derivation of a conclusion on the basis of one or more rules takes two steps. In the first step, the sets of all reasons for, respectively against the hypothetical conclusion are collected. In the second step, the reasons in these sets are weighed against each other.
A reason for a conclusion is based on a valid rule with this conclusion as its conclusion-part. If the conditions of this rule are satisfied, and if it cannot be derived that the rule is excluded in the relevant case, the rule is said to be applicable in that case. The instantiated conditions of the rule are the facts that make the rule applicable. If the rule is applicable, this is a reason for its application.
In theory it is possible that there are more reasons to apply the rule, or that there are also reasons against application, but normally the applicability of the rule will be the only reason that is relevant for the decision to apply the rule. If the rule is actually applied, it generates a reason for its conclusion. The facts that are the reason usually are the instantiated rule conditions that make the rule applicable. Only if the rule is applied analogously, the reason will consist of other facts.

Let us have a closer look at these steps. Suppose that `facts` is an instantiation of the rule conditions, and that `iconcl` is the corresponding instantiation of the rule-conclusion. In that case, the following schematic sentences indicate the first step on the way to the generation of a reason.
If it can be derived that

```
Valid(rule(r, conditions, concl))
Obtains(facts, case(c))
```

and it cannot be derived that

```
Excluded(rule(r), case(c))
```

then it is possible to derive that

```
Applicable(rule(r), case(c), facts, iconcl)
```

In plain language: if it can be derived that a rule is valid, and that its conditions are satisfied in a case, and if it *cannot be derived* that the rule is excluded in that case, then it can be derived that the rule is applicable to that case.

Let us look at an old example. If it can be derived that:

```
Valid(rule(theft1, thief(x), punishable(x)))
Obtains(thief(john), case(john's_theft))
```

and it cannot be derived that

```
Excluded(rule(theft1), case(john's_theft))
```

then it can be derived that

```
Applicable(rule(theft1), case(john's_theft),
  thief(john), punishable(john))
```

Notice the important difference between the first two conditions for the derivation, and the third one. It should be derivable that the rule is valid and that its conditions are satisfied to obtain the conclusion that the rule is applicable. It need not be derivable that the rule is not excluded to obtain that conclusion. It suffices that the exclusion of the rule cannot be derived.

The third condition for the derivation is weaker than the first two, in that it is satisfied in case nothing is known about the exclusion of the rule. At the same time, this reference to what is not derivable, rather than to what is given, makes the logic nonmonotonic. If a theory is extended, so that it becomes derivable that a rule is excluded, this means that the conclusion that this rule is applicable is invalidated.

24. From applicability to a reason for application

If a rule is applicable, this is a reason to apply this rule. If some facts are a reason for a particular conclusion, this is in RBL expressed by means of the `Reason` predicate. This predicate has three parameters. First comes the (possibly compound) fact that is the reason, then follows the conclusion of the reason, and finally it is indicated whether a reason for (pro) or against (con) the conclusion is concerned. For instance:

```
Reason(thief(john), punishable(john), pro)
Reason(juvenile(john), punishable(john), con)
```

(The fact that John is a thief is a reason why John is punishable, and the fact that John is a juvenile is a reason against John's being punishable.)

The statement that some fact is a reason for or against some conclusion cannot be a premise of an RBL-theory, but must be derived on the basis of a goal or the application of a rule.

The existence of reasons can be derived from the fact that a particular rule is applicable. From

```
Applicable(rule(theft1), case(john's_theft),
 thief(john), punishable(john))
```

it may be derived that

```
Reason(
 applicable(rule(theft1), case(john's_theft),
  thief(john), punishable(john)),
 applies(rule(theft1), case(john's_theft),
  thief(john), punishable(john)),
 pro)
```

(The fact that the rule `theft1` is applicable to John's case, with its condition-part instantiated to `thief(john)` and its conclusion-part instantiated to `punishable(john),` is a reason why this same rule, with the same instantiations, applies to John's case.)

From

```
Applicable(
 rule(juveniles_not_punishable), case(john's_theft),
 juvenile(john),   ~punishable(john))
```

it may be derived that

```
Reason(
 Applicable(rule(juveniles_not_punishable),
 case(john's_theft), juvenile(john), ~punishable(john))
 Applies(rule(juveniles_not_punishable), case(john's_theft),
  juvenile(john), ~punishable(john)),
 pro)
```

(The fact that the rule `juveniles_not_punishable` is applicable to John's case, with its condition-part instantiated to `juvenile(john)` and its conclusion-part instantiated to `~punishable(john),` is a reason why this same rule, with the same instantiations, applies to John's case.)

The first and the second parameter of the reason-statements are almost identical to each other, with as the only difference that the state of affairs that the rule is applicable in the first parameter is replaced by the state of affairs that the rule applies in the second parameter.

25. From the application of a rule to a reason for or against its conclusion

If a rule applies to a case, the instantiated rule conditions become a reason for the rule's conclusion. So from

```
Applies(rule(theft1), case(john's_theft),
 thief(john), punishable(john))
```

it may be derived that

```
Reason(thief(john), punishable(john), pro)
```

(The fact that John is a thief is a reason why John is punishable.)

And from

```
Applies(rule(juveniles_not_punishable), case(john's_theft),
  juvenile(john), ~punishable(john))
```

it may be derived that

```
Reason(juvenile(john), punishable(john), con)
```

(The fact that John is a juvenile is a reason against his being punishable.)
Notice that the rule-conclusions sometimes contain a negation, while the conclusions of reasons never contain a negation. The negation of the rule-conclusion is replaced by a reason against the conclusion. The rule `theft1` has a positive conclusion, and this is reflected in the fact that being a thief is a reason that pleads *for* (pro) the conclusion that John is punishable. The rule `juveniles_not_punishable` has a negative conclusion, and therefore the fact that John is a juvenile is a reason *against* (con) John's being punishable.

26. From reasons to a conclusion

The step from the presence of a reason to apply a rule to the conclusion that the rule applies is not a special one. It is like any other step from the presence of a reason for a conclusion to this conclusion. Not only this single reason has to be considered, but rather the sets of all reasons for and against the conclusion.

26.1 The general case

To facilitate the discussion about the reasons for and against a conclusion, we need terms that refer to the sets of these reasons. Here we encounter a new example of structured individuals, because RBL treats sets of reasons as logical individuals.

`reasons_pro(`*`state_of_affairs`*`)`

and

`reasons_con(`*`state_of_affairs`*`)`

denote respectively the set of *derivable* reasons that plead for *`state_of_affairs`*, respectively against *`state_of_affairs`*.
In case of the theory

```
Valid(rule(theft1, thief(x), punishable(x)))
Valid(rule(juveniles_not_punishable, juvenile(x), ~punishable(x)))
Thief(john)
Juvenile(john)
```

it holds that

`reasons_pro(punishable(john)) = {thief(john)}`

and

`reasons_con(punishable(john)) = {juvenile(john)}`

In order to decide on the basis of reasons which conclusion is to be drawn, explicit *weighing-knowledge* is needed. This weighing-knowledge indicates which set of reasons outweighs which other set with respect to a particular conclusion.
It is not impossible that the same two sets of reasons are relevant for different conclusions, but that their relative weight differs for these different conclusions. Therefore weighing-knowledge is always relativized to a conclusion.

In the present example, the relevant weighing-knowledge is laid down in the following RBL sentence:

```
Outweighs( {juvenile(john)}, {thief(john)}, punishable(john))
```

which means that concerning the issue whether John is punishable, (the set of reasons containing) the fact that John is a juvenile outweighs (the set of reasons containing) the fact that John is a thief.

This sentence illustrates the general form of weighing-knowledge. Weighing-knowledge is described in a sentence formed by the `Outweighs` predicate. This predicate has three parameters. The first two parameters denote sets of reasons, and the third parameter denotes a state of affairs. The weighing-knowledge expresses that with respect to the state of affairs of the third parameter, the first set of reasons outweighs the second set.

Notice that the parameters do not have a fixed order in the sense that either the reasons pro, or the reasons con are mentioned first. On the contrary, by means of the order of the sets of reasons it is indicated whether the reasons pro outweigh the reasons con, or the other way round. If the first parameter refers to the reasons that plead for the state of affairs mentioned in the third parameter, the reasons pro outweigh the reasons con. In this case the sentence that expresses the state of affairs of the third parameter follows as the conclusion.

If, however, the first parameter refers to the reasons that plead against the state of affairs mentioned in the third parameter, the reasons con outweigh the reasons pro. In that case, the negation of the sentence expressing the state of affairs of the third parameter follows as the conclusion.

So, if

```
reasons_pro(punishable(john)) = {thief(john)}
```

and

```
reasons_con(punishable(john)) = {juvenile(john)}
```

and if we also have that

```
Outweighs({thief(john)}, {juvenile(john)}, punishable(john))
```

then it may be derived that

```
Punishable(john)
```

However, if we have

```
Outweighs({juvenile(john)}, {thief(john)}, punishable(john))
```

then it may be derived that

```
~Punishable(john)
```

Notice that `reasons_pro(`*`state_of_affairs`*`)` and `reasons_con(`*`state_of_affairs`*`)` denote all the reasons for, respectively against *`state_of_affairs`* that are *derivable* from the theory. If a theory is extended, this can make it possible to derive more reasons both for and against a conclusion. If the extension of the theory leads to more rule exclusions or to more cases where goals are excluded, this can also lead to less reasons.

26.2 Empty sets of reasons

To derive a conclusion on the basis of reasons, it is always necessary to have weighing-knowledge that indicates which set of reasons outweighs the other set. Normally, this weighing-knowledge is not provided by RBL, but should be incorporated in the theory from which the derivations are made. For instance, it is not the task of RBL to indicate whether the fact that John is a thief outweighs the fact that John is a juvenile as a reason concerning the punishability of John. This information should be specified by the domain theory from which the derivations are to be made.

There is, however, one important exception. RBL provides the weighing-knowledge that is necessary to weigh a non-empty set of reasons against an empty set. If there are only reasons for a conclusion, or only reasons against, it should be possible to derive the conclusion, respectively its negation, even if the theory does not say anything about the relative weight of the reasons. That is why in RBL any non-empty set of reasons outweighs an empty set.

So, if there is only a reason for a conclusion, and no reason against it, this conclusion can be derived. The other way round the same holds: if there is only a reason against a conclusion, and no reason for it, the negation of this conclusion can be derived. In practice these situations, where there is one reason that pleads for a conclusion, and no reason that pleads against it, are very important. In fact, it is the type of situation that is classically dealt with by arguments of the form Modus Ponens.

For instance, if we have the rule that thieves are punishable and the fact that John is a thief, the argument to the conclusion that John is punishable would traditionally be constructed as a case of Modus Ponens. In RBL it becomes an argument in which there is one reason why John can be punished and no reason why he cannot be punished. Since the one reason for punishability outweighs the empty set of reasons against punishability, it is possible to derive that John is punishable, even if the domain theory does not contain any relevant weighing-knowledge.

27. Deriving a reason by application of a rule

It is possible to derive from

```
Applicable(rule(r), case(c), iconds, iconcl)
```

that

```
Reason(
 applicable(rule(r), case(c),iconds, iconcl),
 applies(rule(r), case(c), iconds, iconcl),
 pro)
```

Normally there will be no other reasons derivable concerning the application of the rule *r*, and then it is also possible to derive

```
Applies(rule(r), case(c), iconds, iconcl)
```

RBL makes it possible to derive from this that

```
Reason(iconds, iconcl, pro)
```

or, in case it was derived that

```
Applies(rule(r), case(c), iconds, ~iconcl)
```

it can be derived that

`Reason(`*`iconds`*`, `*`iconcl`*`, con)`

This completes the derivation of a reason for the rule-conclusion from the statement that the rule is valid, and the facts that satisfy the rule conditions.

28. From the validity of a goal to reasons

The reasons involved in weighing relations need not stem from rules or principles; they can also be based on goals. The derivation of a reason on the basis of a goal is relatively simple. If some action or state of affairs contributes to a valid goal, and it cannot be derived that the goal is excluded in this case, there is a reason why the action ought to be performed, respectively the state ought to be the case.

28.1 Deontic predicates and operators

To express that an action ought to be performed, RBL has a dedicated predicate `Ought` that has actors and action types as its parameters.[32] `Ought` stands for 'Ought to do'. Action types are denoted by constants or function symbols. They can be preceded by the sign ¬ to denote the refraining from this type of action. If `stealing` stands for an action of the type stealing, ¬`stealing` stands for refraining from stealing.
The use of the `Ought` predicate is illustrated in the sentence

`(`*`x`*`)(In_distress(`*`person`*`) → Ought(`*`x`*`, help(`*`person`*`)))`

which means that everybody ought to help a person in distress.
Analogously, the sentence

`(`*`x`*`)(`*`y`*`)((`*`x`*` <> `*`y`*`) → Ought(`*`x`*`, ¬hurt(`*`y`*`)))`

means that nobody ought to hurt another person.
A related predicate `O` stands for 'It ought to be the case that'. This predicate has states of affairs as its parameter. For instance, the sentence

`O(withheld(john, using_drugs))`

means that the state of affairs that John is being withheld from using drugs ought to be the case.

28.2 Reasons for and against actions

The relation between accepted goals and the reasons for actions based on them is illustrated by the following. If it can be derived that

```
Obtains(
 contributes_to(withhold(john, using_drugs), youth_protected),
 case(john's_drug_addiction))
```

[32] The discussion of reasoning with goals necessitates the introduction of deontic elements in the language of RBL. It is, however, not the purpose of this text to propose a system of deontic logic. Therefore, the discussion of the deontic predicates is kept minimal.

```
Valid(goal(protection_of_juveniles, youth_protected))
```

and it cannot be derived that

```
Excluded(goal(protection_of_juveniles), case(john's_drug_addiction))
```

then it can be derived that

```
Reason(
 contributes_to(withhold(john, using_drugs), youth_protected),
 ought(x, do(withhold(john, using_drugs))),
 pro)
```

(From the sentences that the goal `protection_of_juveniles` is valid, that withholding John from using drugs contributes to this goal, it can be derived that one ought to withhold John from using drugs, if it cannot be derived that the goal `protection_of_juveniles` is excluded in the case of John's drug addiction.)

If a type of action detracts from a goal, this leads to a reason against this action. If it is derivable that

```
Obtains(
 contributes_to(smoke, ~lives_long(x)),
 case(alice's_smoking))
Valid(goal(long_smokeless_life, lives_long(x))
```

and it is not derivable that

```
Excluded(goal(long_smokeless_life), case(alice's_smoking))
```

then it can be derived that

```
Reason(
 contributes_to(smoke, ~lives_long(alice)),
 ought(alice, ¬smoke),
 pro)
```

Notice that the fact that the smoking of Alice detracts from her long living is a reason *for* the conclusion that Alice ought to *refrain from* smoking, and not *against* the conclusion that Alice ought to smoke.

28.3 Reasons for and against states of affairs

If a state of affairs contributes to a goal-state, this leads to a reason why this state of affairs ought to be the case. If it can be derived that

```
Valid(goal(protection_of_parties_in_good_faith,
 parties_in_good_faith_protected))
Obtains( contributes_to(owner(gerald, clock),
   parties_in_good_faith_protected),
 case(transfer_clock))
```

and it is not derivable that

```
Excluded(goal(protection_parties_in_good_faith),
 case(transfer_clock))
```

then it can be derived that

```
Reason(contributes_to(owner(gerald, clock),
  parties_in_good_faith_protected),
 o(owner(gerald, clock)), pro)
```

Analogously, if a state of affairs detracts from a goal-state, this is a reason why this state of affairs ought not to be the case. If it can be derived that

```
Valid(goal(crime_prevention, less_crime))
Contributes_to(abolished(crime_prevention_team), ~less_crime)
```

and it cannot be derived that

```
Excluded(goal(crime_prevention), case(prevention_team))
```

then it can be derived that

```
Reason(
 contributes_to(abolished(crime_prevention_team), ~less_crime)
 o(~abolished(crime_prevention_team)), pro)
```

29. Derivability in Reason-Based Logic

After the informal discussion of RBL in the previous sections, I will now propose a definition of the derivability relation of RBL. Let T be an RBL-theory, that is a set of well-formed RBL-sentences which are allowed as premises for an argument. It then holds that:

A sentence can validly be derived from an RBL-theory T, if and only if this sentence is an element of the intersection of all extensions of T.

29.1 Extensions of RBL-theories

Extensions of RBL-theories are defined as follows[33]:

Let T be an RBL-theory, let Th(T) be the deductive closure of T, let E be a set of sentences, and let T_0 ... T_∞ be a sequence of sets of sentences such that:

1. T_0 = Th(T ∪
{(*reasonset*)(*state*)
(*reasonset* <> ∅ → Outweighs(*reasonset*, ∅, *state*)),
(*state*)(∃*c*)(case(*c*) & Obtains(*state*, case(*c*)) → *state*) })

T_0 is intuitively the original theory with the additional information that every non-empty set of reasons outweighs the empty set, and that if a state of affairs obtains in some case, this state of affairs obtains, and everything which can be deduced from this.

2. T_{m+1} = Th(T_m ∪ Δ(T_m)), where Δ(T_m) is the smallest set of sentences such that:

[33] The following definition of extensions of an RBL-theory is an adaptation of the second way in which [Reiter 1980] defines extensions of default-theories.

Assume that the sentence `Facts` minimally entails[34] the sentence `Condition` under some substitution σ, and that the term `iconcl` is the instance of the term `conclusion` under σ. Then the following holds:

a. If `Valid(rule(`*`rid`*`, `*`condition`*`, `*`conclusion`*`))` ∈ Tm, and
`Obtains(`*`Facts`*`, case(`*`c`*`))` ∈ Tm, and
not `Excluded(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`iconcl`*`)` ∈ E, then
`Applicable(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`iconcl`*`)` ∈ Δ(Tm).

b. If `Applicable(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`atom`*`)` ∈ Tm, then
`Reason(applicable(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`atom`*`),`
`applies(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`atom`*`), pro)` ∈ Δ(Tm).

c. If `Applies(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, `*`atom`*`)` ∈ Tm, then
`Reason(`*`facts`*`, `*`atom`*`, pro)` ∈ Δ(Tm).

d. If `Applies(rule(`*`rid`*`), case(`*`c`*`), `*`facts`*`, ~`*`atom`*`)` ∈ Tm, then
`Reason(`*`facts`*`, `*`atom`*`, con)` ∈ Δ(Tm).

e. If `Valid(goal(`*`gid`*`, `*`state`*`))` ∈ Tm, and
`Obtains(contributes_to(`*`action`*`, `*`state`*`), case(`*`c`*`))` ∈ Tm, and not
`Excluded(goal(`*`gid`*`),case(`*`c`*`))` ∈ E then
`Reason(contributes_to(`*`action`*`, `*`state`*`), ought(`*`x`*`, action), pro)`
∈ Δ(Tm).

f. If `Valid(goal(`*`gid`*`, `*`state2`*`))` ∈ Tm,
`Obtains(contributes_to(`*`state1`*`, `*`state2`*`), case(`*`c`*`))` ∈ Tm, and not
`Excluded(goal(`*`gid`*`),case(`*`c`*`))` ∈ E then
`Reason(contributes_to(`*`state1`*`, `*`state2`*`), o(`*`state1`*`), pro)` ∈ Δ(Tm).

g. If `Valid(goal(`*`gid`*`, `*`state`*`))` ∈ Tm, and
`Obtains(contributes_to(`*`action`*`, ~`*`state`*`), case(`*`c`*`))` ∈ Tm, and not
`Excluded(goal(`*`gid`*`),case(`*`c`*`))` ∈ E then
`Reason(contributes_to(`*`action`*`, `*`state`*`), (`*`x`*`)ought(`*`x`*`,¬action), pro)`
∈ Δ(Tm).

h. If `Valid(goal(`*`gid`*`, `*`state2`*`))` ∈ Tm,
`Obtains(contributes_to(`*`state1`*`, ~`*`state2`*`), case(`*`c`*`))` ∈ Tm, and not
`Excluded(goal(`*`gid`*`),case(`*`c`*`))` ∈ E then
`Reason(contributes_to(`*`state1`*`, `*`state2`*`), o(~`*`state1`*`), pro)` ∈ Δ(Tm).

Let `reasons_pro(`*`atom`*`)` be {*`s`*`: Reason(`*`s`*`, `*`atom`*`, pro)` ∈ E}
and let `reasons_con(`*`atom`*`)` be {*`s`*`: Reason(`*`s`*`, `*`atom`*`, con)` ∈ E}

i. If `Outweighs(reasons_pro(`*`atom`*`), reasons_con(`*`atom`*`), `*`atom`*`)` ∈ Tm,
then *`Atom`* ∈ Δ(Tm).

---

[34] A sentence A minimally entails a sentence C, if and only if C is true in all models of A, and there is no sentence B such that B is true in all models of A, while it is not the case that A is true in all models of B, and C is true in all models of B.

j. If `Outweighs(reasons_con(`*`atom`*`), reasons_pro(`*`atom`*`), `*`atom`*`)` $\in$ Tm, then *~Atom* $\in \Delta$(Tm).

3. E is an extension of T iff $E = \bigcup_{i=0}^{\infty} T_i$

Intuitively this definition of an extension of the theory T comes down to the following. Starting with a set of sentences (T_0), that consists of the original theory T and two axioms of RBL, a series of ever increasing sets of sentences (T_0 ... T_∞) is defined. Every set contains the sentences of the previous set plus what 'immediately follows' from that set (Δ(Tm)). Normally, at some moment there are no new conclusions possible anymore ($T_{m+1} = T_m$) The set E which is achieved at that moment is an extension of T .

A complication in this definition is that the sentences that can be added to a set to obtain the new set (Δ(Tm)) do not only depend on the first set (T_m), but also on the final set E that will be achieved. That this is the case follows from the conditions sub a, e, f, g, h, i, and j, which refer to the exclusions or the reasons which are present in the final extension of the theory.

This complication has at least two consequences. The first consequence is that it is not possible to generate the series of sets of sentences just by starting with the original theory. In addition to the original theory, an 'hypothesis' concerning the final extension is needed to generate the series of sets. An hypothesis concerning the extension of a theory is correct if the series of sets based on the original theory and the hypothetical extension ends with the hypothetical extension. Cf. figure 7.

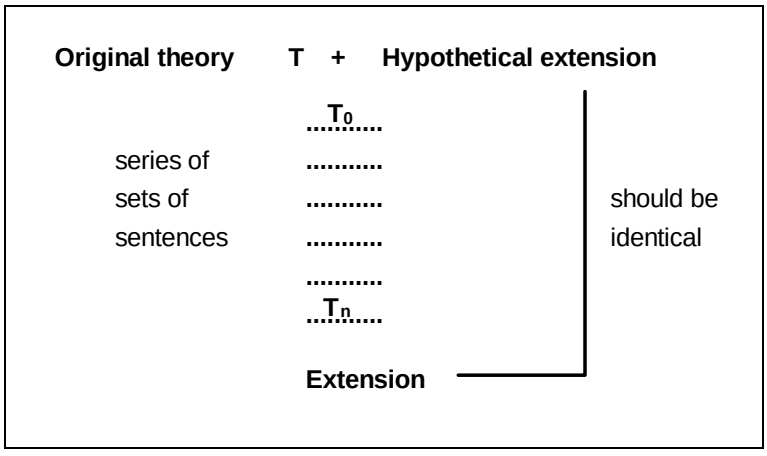

Figure 7: From the original theory to its extension

29.2 The number of extensions of a theory

The second consequence is that a theory can at least in theory have more than one extension. That is because there may be more than one combination of the original theory and a hypothetical extension that fulfils the demand that the resulting series of sets of sentences ends up with the hypothetical extension on which the series was based. Cf. figure 8.

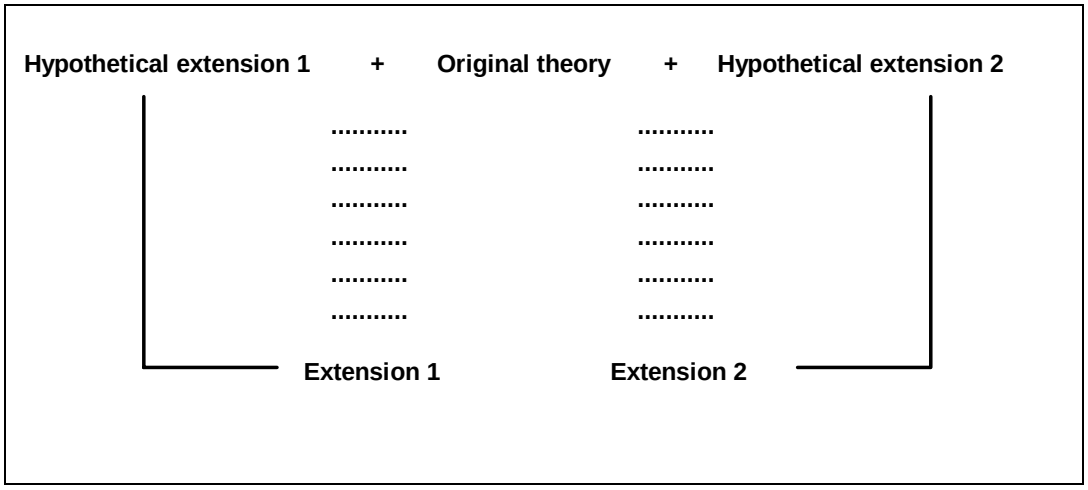

Figure 8: Multiple extensions

Moreover, it is also possible that a theory has no extension at all. Let me give an example of a theory which has no extension:

```
Obtains(a, case(1))
Valid(rule(1, a, excluded(rule(1), case(1)))
```

This theory has no extension, because if an extension contains the sentence

```
Excluded(rule(1), case(1))
```

it should for this very reason not contain that sentence, and the other way round. (If rule 1 is not excluded in case 1, it can be derived that it is excluded in that case. If, however, it is excluded, the very reason why it should be excluded, that is derivation on the basis of rule 1, does not obtain anymore.)

The following theory has two extensions:

```
Obtains(a, case(1))
Obtains(b, case(1))
Valid(rule(1, a, excluded(rule(2), case(1)))
Valid(rule(2, b, excluded(rule(1), case(1)))
```

The one extension contains, amongst others, the sentences

```
Applicable(rule(1), case(1), a, excluded(rule(2), case(1)))
Excluded(rule(2), case(1))
```

and not the sentences

```
Applicable(rule(2), case(1), b, excluded(rule(1), case(1)))
Excluded(rule(1), case(1))
```

The other extension contains, amongst others, the sentences

```
Applicable(rule(2), case(1), b, excluded(rule(1), case(1)))
Excluded(rule(1), case(1))
```

and not the sentences

```
Applicable(rule(1), case(1), a, excluded(rule(2), case(1)))
Excluded(rule(2), case(1))
```

Otherwise than in other logics, in RBL multiple extensions cannot be the result from inconsistencies, because RBL does not strive after consistency maintenance. It appears from the examples above that the occurrence of multiple extensions or the lack of any extension is the consequence of circularities in a theory.

As may be seen in the first example, if a rule directly refers to itself, this may block the existence of extensions. If the self-reference is indirect (one rule refers to another rule and vice versa), the result is that there are multiple extensions (cf. the second example).

If this hypothesis is correct, the intuitively correct conclusions from a theory can be obtained by not applying the rules which are involved in the self-reference. In the case of multiple extensions, the present definition of valid conclusions obtains this result by using the intersection of all extensions to define derivable sentences.[35] In this way, any rule that cannot be applied in some extension does not lead to conclusions in that extension. Its conclusion will therefore not be derivable from the theory.

30. Simple rule application

In the following sections I will illustrate how RBL can be put to use for legal applications. This involves a series of examples of how legal arguments can be cast in the form of reasons for and against particular conclusions. As a consequence, this part of the paper is particularly relevant for those who read this text with an interest in how RBL can be used for legal knowledge representation.

First I will show how the standard cases such as simple rule application and the weighing of reasons for and against a particular classification are handled. Then follow three extended examples that show how RBL deals with complex legal reasoning. These examples concern arguments about the relative weights of sets of reasons, arguments about the validity of legal rules, and an argument about the relative importance of conflicting case law where the interpretation of a statutory provision is concerned.

Most often, when we apply a rule, there are no complications at all. The conditions of the rule are satisfied and the rule conclusion follows just like that. A theory about reasoning with rules should be able to handle this situation without problems. This section deals with simple rule application and illustrates how RBL functions if there are no special circumstances that must be dealt with.

Suppose we have the rule that thieves ought to be punished, and the fact that John is a thief. From this information we want to derive that John ought to be punished.

This information can be represented by the following RBL-sentences:

```
Valid(rule(theft2, thief(x), o(punished(x))))
Obtains(thief(john), case(john's_theft))
∴ O(punished(john))
```

[35] Notice that the use of the intersection of extensions in RBL is inspired by the need to avoid the complications of indirectly self-referring theories, and not to obtain a skeptical inference relation. RBL is because it demands explicit weighing knowledge in the case of colliding reasons.

By following the definition of an extension of an RBL-theory in section 29, it can be seen that the deductive closure of the following sentences is an extension of the theory given above:

```
Valid(rule(theft2, thief(x), o(punished(x))))
Obtains(thief(john), case(john's_theft))
(reasonset)(state)
 (reasonset <> ∅ → Outweighs(reasonset, ∅, state))
Thief(john)
Applicable(rule(theft), case(john's_theft),
 thief(john), o(punished(john)))
Reason(
 applicable(rule(theft), case(john's_theft),
  thief(john), o(punished(john))),
 applies(rule(theft), case(john's_theft),
  thief(john), o(punished(john))),
 pro)
Applies(rule(theft), case(john's_theft),
 thief(john), o(punished(john)))
Reason(thief(john), o(punished(john)), pro)
O(punished(john)
```

Moreover, unless there are extensions which contain reasons not included in this extension, all sentences of this extension will be derivable. By inspecting the theory that consists of the two premises, it can easily be seen that there is no information which can underlie other reasons, because reasons must be based on valid rules or goals, and no other rules or goals than the rule `theft2` occur in the theory. Therefore, the sentence `O(punished(john))` will occur in the intersection of all extensions of the theory and is therefore derivable.

Notice that this argument makes use of meta-level reasoning. Presently there is no constructive proof theory for RBL.

31. The exclusion of rules

RBL uses the mechanism of exclusionary reasons to model exceptions that block the application of a rule. The law knows several situations in which such exceptions occur (cf. section 14). Sometimes the exception only makes it impossible to apply a rule. At other occasions the exception not only blocks the application of a rule, but also provides a reason for an incompatible conclusion. In the following two subsections I will discuss the prescription of punishable facts as an example of the first situation, and the rule that governs force majeure in contract law as an example of the second situation.

31.1 Undercutters

Reasons that only block the application of a rule, without having any other impact on the case at hand are sometimes called 'undercutters'.[36] The scope limitations of legal rules form the basis for

[36] The term stems from [Pollock 1987], and was also adopted in [Prakken 1993]. [Sartor 1991] distinguishes between exceptions to norms and exceptions to effects, where the former are undercutters, and the latter 'rebutters'.

several kinds of undercutters. One of these scope limitations concerns the scope of legal rules in time. In the case of criminal law, this means that a criminal law may not be applicable anymore, if the crime occurred sufficiently long ago. This situation is illustrated in our old example about John, the thief. Although John is a thief, he is not punishable, because his crime is prescribed.[37]

Let us formulate the relevant rules in an RBL-theory:

```
Valid(rule(theft1, thief(x), punishable(x)))
Valid(rule(prescription,
 criminal_law(rule(r)) &
  crime_type(rule(r), crime_type) &
  type(crime, crime_type) &
  prescribed(crime, case(c)),
 excluded(rule(r), case(c)) ))
```

The rule `prescription` boils down to the fact that if a particular rule is a rule of criminal law, if this rule is about a particular type of crime, and if a crime of this type is in a particular case prescribed, then the application of this rule is excluded in that case.

The facts

```
Criminal_law(rule(theft1))
Crime_type(rule(theft1), stealing)
Type(theft, stealing)
Prescribed(theft, case(john's_theft))
```

make the rule `prescription` apply to the theft of John's with as a consequence that it can be derived that

```
Excluded(rule(theft1), case(john's_theft))
```

and as a consequence it is not anymore possible to derive that

```
Applicable(rule(theft1), case(john's_theft),
 thief(john), punishable(john))
```

nor that

```
Reason(thief(john), punishable(john), pro)
```

As a further consequence, the sentence

```
Punishable(john)
```

cannot be derived anymore.

Notice that the facts that the rule `theft1` is a rule of criminal law, and that the crime was prescribed only block the application of the rule `theft1`, and have no other implication for the punishability of John whatsoever.

Notice also how the treatment of rules as logical individuals in RBL makes it possible to formulate general rules about when rules are excluded.

---

[37] In Dutch law, the prescription of crimes is actually regulated in a more complex way, namely through the non-admissibility of the public prosecutor. An adequate analysis of this phenomenon demands for a procedural approach. Cf. [Hage et al. 1994].

31.2 Rebuttals

Some exceptions to rules do not only block the application of the rule, but also provide the case with a legal consequence that is incompatible with the conclusion of the rule. Such exceptions are also called 'rebutters' [Pollock 1987].

In Dutch criminal law, theft is punishable with a fine of the third category, but burglary is punishable with a fine of the fourth category. Obviously, since burglary involves theft, application of the rule about burglary blocks application of the rule about theft. This can be formalised as follows[38]:

```
Valid(rule(theft3, thief(x), punishable(x, fine_type_3)))
Valid(rule(burglary1, burglar(x), punishable(x, fine_type_4)))
Valid(rule(burglary2,
 obtains(burglar(x), case(c)),
 excluded(rule(theft3), case(c)))
Thief(gerald) & Burglar(gerald)
```

Because Gerald is a burglar, both the rules `burglary1` and `burglary2` are applied. The first application leads to the conclusion that Gerald is punishable with a fine of the fourth category, while the second application leads to an exclusionary reason for rule `theft3`. As a consequence it cannot be derived that that Gerald is punishable with a fine of the third category.

Notice that the original rule that in the case of burglary, one is punishable with a fine of the fourth category has been translated into two RBL rules. The reason for this is that the application of this rule has two effects. First that the application of the rule theft is blocked, and second that there the burglar is punishable with a fine of the fourth category. RBL rules can (for historical reasons) only have one conclusion, and therefore two rules are needed to express the two effects of force majeure.

32. Elementary goal-based reasoning

An important legal goal in most Western countries is to protect the possibility for every citizen to express his or her opinion. Let us call this goal Freedom of Expression. Clearly, having a free press contributes to this goal. As a consequence of the goal Freedom of Expression, there should be a free press.

This simple argument, based on the application of a goal, would be expressed in RBL as follows:

```
Valid(goal(freedom_of_expression), freedom_to_express)
Obtains(
 contributes_to(free_press, freedom_to_express),
 case(press_freedom))
```

[38] Another formalisation which makes use of the Lex Specialis rule would be less ad hoc, but also a less clear example. Cf. [Hage and Verheij 1995] for this more principal approach.

The extension of this theory contains the following sentences:

```
Reason(contributes_to(free_press, freedom_to_express),
 o(free_press), pro)
Outweighs(
 {contributes_to(free_press, freedom_to_express)},
 ∅, o(free_press))
O(free_press)
```

and they would not contain the sentence:

```
Excluded(goal(freedom_of_expression), case(press_freedom))
```

Therefore, from this theory it can be derived that:

```
O(free_press)
```

Imagine, however, that a particular country is in war, and that there are serious reasons to assume that freedom of expression can be misused by the enemy for spreading unnecessary anxiety. Under these circumstances there is reason to restrict the goal to have freedom of expression. The goal is not completely abandoned; its effects are only temporarily suspended. Under these circumstances, the following sentence will be derivable:

```
Excluded(goal(freedom_of_expression), case(free_press))
```

This exclusion blocks the generation of the reason that a free press contributes to the freedom of expression, which would plead for the conclusion that there ought to be a free press. As a consequence, there is no reason left for the conclusion that there ought to be a free press, and this conclusion can therefore not be drawn anymore.

33. Weighing reasons

In legal reasoning, weighing reasons most often occurs when there are no legal rules to apply. Classification springs to mind as a field that is not always governed by rules of law and where reasons for and against a particular classification have to be weighed. Another situation where weighing reasons is important, is where legal decision makers have discretion. In the Netherlands, determining the measure of punishment (within certain limits) belongs to this field.

I will adapt an earlier example to show how weighing reasons in the field of criminal law can be represented in RBL. We have already met with the case of John, the thief. That he is a thief is a reason why John ought to be punished. Suppose, however, that John is a child of thirteen. This is a reason not to punish him. In some legal systems the age of John may even block the application of the normal rule about the punishment of thieves. In that case it should be represented by means of an exclusionary reason. If the age does not completely block the application of the rule `theft2`, and if the judge has discretion to determine whether the crime will be punished, the age of John will at least form a reason against punishing him.

The situation can then be represented as follows:

```
Valid(rule(theft2, thief(x), o(punished(x))))
Valid(rule(child, child(x), ~o(punished(x))))[39]
Thief(john) & Child(john)
```

On the basis of this theory we have:

```
reasons_pro(o(punished(john))) = {thief(john)}
reasons_con(o(punished(john))) = {child(john)}
```

However, to derive a conclusion regarding whether John ought to be punished, we need to know which set of reasons outweighs the other set. On the basis of the present theory, which contains no information concerning the relative weight of the reasons, it is not possible to derive whether John ought to be punished or not.

Suppose, however, that the information is added that being a child outweighs being a thief as a reason concerning punishment:

```
Outweighs( {child(x)}, {thief(x)}, o(punished(x)))
```

On the basis of this information, it is possible to derive that it is not the case that John ought to be punished:

```
~O(punished(john))
```

Suppose, moreover, that John is not only a thief, but also a mistreator, and that mistreators ought to be punished:

```
Valid(rule(mistreatment, mistreator(x), o(punished(x))))
Mistreator(john)
```

This information can be combined with the following additional weighing knowledge:

```
Outweighs( {child(x)}, {mistreator(x)}, o(punished(x)))
Outweighs( {mistreator(x) & thief(x)}, {child(x)}, o(punished(x)))
```

(Concerning the issue of punishment, being a child outweighs being a mistreator, but the combination of being a mistreator and a thief outweighs being a child.)

Here we encounter an important characteristic of RBL, that is the *accrual of reasons*. Two or more reasons that are individually outweighed by another reason, may together outweigh that other reason. So if John is both a thief, a mistreator and a child, the conclusion will be that John ought to be punished, because the facts that John is a mistreator and a thief together outweigh the fact that he is a child:

```
Outweighs(
 {mistreator(john) & thief(john)},
 {child(john)},
 o(punished(john)))
```

[39] Arguably, this rule should be replaced by a goal. For the purpose of the example, however, I have chosen for a rule (principle) that children ought not to be punished. For the same purpose, I have chosen to let the conclusion of the rule run that it is not the case that the child ought to be punished, rather than that it is the case that the child ought not to be punished.

34. Reasoning about weights

Whether one set of reasons outweighs the other set, or just the other way round, needs not solely depend on the reasons themselves and the conclusion for or against which they plead. It is also possible that there are other facts which, although they are not reasons for or against a particular conclusion themselves, influence the relative weight of the concerned sets of reasons.

34.1 An example

Such a situation is illustrated by the following real life example [Kantongerecht Rotterdam, June 12th, 1985]. A small supermarket had to dismiss one of its employees for financial reasons. For such a dismissal, the allowance of a judge is necessary. One of the employees, called Mary, has been longer in service, and this is a reason for the judge not to permit to dismiss her.[40]

```
Reason(longer_in_service(mary), permitted(x, dismiss(mary)), con)
```

The other employee, called Richard, on the other hand, has better papers for the job, and this is a reason not to dismiss him either.

```
Reason(better_papers(richard), permitted(x, dismiss(richard)), con)
```

Since it is clear that one of the employees has to be dismissed for financial reasons, a reason against permitting the dismissal of the one is also a reason for permitting the dismissal of the other[41]:

```
(reason)
 (Reason(reason, permitted(x, dismiss(richard)), con) →
  Reason(reason, permitted(x, dismiss(mary)), pro))
```

To determine what has to be decided, we will focus on the reasons that plead for and against permitting to dismiss Mary. We have the following two sets of reasons:

```
reasons_pro(permitted(x, dismiss(mary))) = {better_papers(richard)}
reasons_con(permitted(x, dismiss(mary))) = {longer_in_service(mary)}
```

The judge decided that, although Richard had better papers for the job, Mary was still sufficiently qualified, so that the better papers did not count for much. The fact that Mary had been longer in service should therefore tip the balance of reasons:

```
Reason(
 suitable_for_job(mary),
 outweighs(
  {longer_in_service(mary)}, {better_papers(richard)},
    permitted(x, dismiss(mary)),
 pro)
```

---

[40] In the formalisation of this example, I make use of the `Permitted`-predicate, which is the counterpart of the `Ought`-predicate.

[41] Strictly, the following sentences would be better represented as RBL-rules. To keep the example relatively simple, I have to chosen to use material implications.

Notice that the fact that Mary was suitable for the job was not considered as a reason no to dismiss her, but only as a reason why the seniority of Mary should outweigh the better papers of Richard.

34.2 Case-based reasoning as a way to reason about relative weights

One important function of case law is to indicate how sets of reasons are to be weighed against each other. A case that has already been decided can be seen as containing both a set of reasons that pleads for a particular decision, a set of reasons that pleads against this decision, and the decision about the issue, which is also a decision about which set of reasons outweighs the other set. (Of course, this only holds for cases in which the weighing of reasons was involved.) In this connection the following functions may be used:

```
reasons_for(decision, case(c)) =
 {f: obtains(f, case(c)) & reason(f, decision, pro)}
reasons_against(decision, case(c)) =
 {f: obtains(f, case(c)) & reason(f, decision, con)}
```
[42]

If we can find a case in which some decision was taken, this is a reason why the reasons in that case that pleaded for this decision, outweigh the reasons in that case that pleaded against that decision:

```
Valid(rule(case_based_reasoning,
 (∃c)(case(c) & decision(case(c)) = d  &
  reasons_for(d, case(c)) = pro_reasons &
  reasons_against(d, case(c)) = con_reasons),
 outweigh(pro_reasons, con_reasons, d)))
```

Intuitively, a decided case can provide a reason why the reasons in a new case have to be weighed in the same manner.

In this connection we can distinguish five situations. In all of them I assume that the issue at stake is the same for the case that was already decided, and the new case. I will use the following abbreviations:

`reasons_for(issue, case(old))` stands for the set of reasons in the decided case that pleaded for the thesis `issue`.

`reasons_against(issue, case(old))` stands for the set of reasons in the decided case that pleaded against the thesis `issue`.

`reasons_for(issue, case(new))` stands for the set of reasons in the new case that plead for the thesis `issue`.

`reasons_against(issue, case(new))` stands for the set of reasons in the new case that plead against the thesis `issue`.

`decision(issue, case(old))` stands for the decision on the issue that was taken in the decided case, that is either `issue`, or `~issue`.

---

[42] These functions should not be confused with the functions `reasons_pro` and `reasons_con`. Cf. section 26.1.

`decision(issue, case(new))` stands for the decision on the issue that should be taken in the new case, that is either `issue`, or `~issue`.

Given these abbreviations, the following situations and their consequences can be distinguished[43]:

Situation 1

```
reasons_pro(issue, old) = reasons_pro(issue, new)
reasons_con(issue, old) = reasons_con(issue, new)
```

Then there is a reason why `decision(issue, old) = decision(issue, new)`, because the situations are completely identical insofar as the reasons are concerned.

Situation 2

```
reasons_pro(issue, old) ⊂ reasons_pro(issue, new)
reasons_con(issue, old) ⊇ reasons_con(issue, new)
decision(issue, old) = issue
```

Then there is a reason why `decision(issue, new) = issue`, because in the new case there are even more reasons that plead for `issue`.

Situation 3

```
reasons_pro(issue, old) ⊆ reasons_pro(issue, new)
reasons_con(issue, old) ⊃ reasons_con(issue, new)
decision(issue, old) = issue
```

Then there is a reason why `decision(issue, new) = issue`, because in the new case there are only less reasons that plead against `issue`.

Situation 4

```
reasons_pro(issue, old) ⊃ reasons_pro(issue, new)
reasons_con(issue, old) ⊆ reasons_con(issue, new)
decision(issue, old) = ~issue
```

Then there is a reason why `decision(issue, new) = ~issue`, because in the new case there are only less reasons that plead for `issue`.

Situation 5

```
reasons_pro(issue, old) ⊇ reasons_pro(issue, new)
reasons_con(issue, old) ⊂ reasons_con(issue, new)
decision(issue, old) = ~issue
```

---

[43] The distinctions between the situations and the consequences attached to the distinctions were inspired by the HYPO-system for case-based reasoning. Cf. [Ashley 1991]. The parallel with the HYPO-system can even be increased if the presence to a higher degree of a reason for a conclusion is treated as an additional reason for the conclusion, if presence to a lesser degree of a reason for a conclusion counts as a reason less for the conclusion, and if reasons against a conclusion are treated analogously. The presence of facts in degrees is discussed in [Hage Reasoning].

Then there is a reason why `decision(issue, new) = ~issue`, because in the new case there are even more reasons that plead against `issue`.

34.3 General rules to reason about the weights of reasons

These different situations are based on an underlying style of reasoning. If we know that a set of reasons outweighs another set, and the stronger set is only strengthened by adding new reasons, the result will also outweigh its unchanged (or even weakened) rival. And analogously, if a set of reasons is outweighed by another set, and this set is weakened by deleting one or more reasons from it, the resulting set will also be outweighed by the other set that remained the same or was even strengthened.

Formally these rules look as follows:

```
Valid(rule(strengthening_reasons,
 outweighs(reasons1, reasons2, conclusion) &
  (reasons1 ⊆ reasons3),
 outweighs(reasons3, reasons2, conclusion)))
```

If a set of reasons outweighs another set, this is a reason why a superset of the first set will also outweigh the last set.

```
Valid(rule(weakening_reasons,
 outweighs(reasons1, reasons2, conclusion) &
  (reasons3 ⊆ reasons2),
 outweighs(reasons1, reasons3, conclusion) ))
```

If a set of reasons outweighs another set, this is a reason why a subset of the last set will also be outweighed by the first set.

These rules assume that reasons are independent of each other. Therefore, in a case where reasons are not independent, the application of these rules is excluded.

35. Reasoning about the validity of rules

Since, from an RBL-point of view, rules are individuals, and not sentences, they cannot stand on their own in a theory of RBL. Moreover, only a valid rule can be used in a derivation to create reasons. These two facts together are the reason why rules are introduced into arguments by means of the statement that they are valid. This has the additional advantage that it becomes possible to 'derive rules', because to derive a rule it is only necessary to derive the statement that it is valid.

RBL does not specify from which premises it can be derived that a rule is valid. In fact, it completely depends on the domain which criteria are used to determine the validity of rules (cf. [Toulmin 1958, pp. 36f. and 104] on the field dependence of backings). In the following subsections, I will discuss a number of examples.

35.1 Rules of recognition

According to Hart, a rule is a valid rule of law, if it is identified as such by a rule of recognition [Hart 1961, pp. 92f.]. An example of such a rule of recognition would be that a rule is a valid rule of law if it was made by the legislator. In RBL such a rule of recognition is expressed as follows:

```
Valid(rule(legislation,
 adopted_by_legislator(rule(r)),
 valid(rule(r))))
Adopted_by_legislator(rule(statutory_rule))
```

From these two sentences it is possible to derive (the validity of) the rule called `statutory_rule`:

```
Valid(rule(statutory_rule))
```

35.2 Interpretation

Another way to argue about the validity of a rule has to do with the interpretation of a legal source. Take for instance the rule that women and children under sixteen are not allowed to partake in the rodeo. This rule formulation is syntactically ambiguous, because at first sight it is not clear whether women over fifteen are allowed to partake. There are two possible interpretations of the rule formulation, only one of which expresses a valid rule (`P` is a predicate for permitted to be, analogous to the `O`-predicate):

```
rule(rodeo1,
 woman(x) v
  (child(x) & under_sixteen(x)),
 ~p(partakes_in_rodeo(x)))
rule(rodeo2,
 (woman(x) v child(x)) &
  under_sixteen(x),
 ~p(partakes_in_rodeo(x)))
```

Clearly the second interpretation (`rodeo2`) does not make much sense, because on this interpretation the condition that somebody is a woman is completely superfluous. This is a reason why the second interpretation does not expresses a valid rule.

```
Reason(
 superfluous(condition_woman_x), rule(rodeo2)),
 valid(rule(rodeo2)),
 con)
```

On the basis of this reason it can be derived that the rule `rodeo2` is not valid. Moreover if there are only two seemingly acceptable interpretations, the rejection of interpretation 2 forms a reason to accept interpretation 1:

```
Reason(
 ~valid(rule(rodeo2)),
 valid(rule(rodeo1)),
 pro)
```

This leads to the final conclusion that the rule expressed in the first interpretation (`rodeo1`) is a valid one:

```
Valid(rule(rodeo1))
```

35.3 Principles of interpretation

The law knows a number of traditional reasons that can be adduced in arguments about the interpretation of statutory rules [cf. MacCormick and Summers 1991]. These reasons are connected to 'maxims of interpretation', such as that the law should be interpreted according the normal meaning of its wordings, according to the historical intention of the legislator, according to the interests of society, or according to the system of the law. These maxims can in RBL be represented as rules that generate reasons why particular rules are valid ones. Take, for instance, the following rule that says that a rule is correctly interpreted if the interpretation is in accordance with the intentions of the legislator as they become clear from the preparatory works:

```
Valid(rule(genetic_interpretation,
 interpretation_of(text, rule(r)) &
  in_accordance_with(rule(r), intention_legislator),
 correct_interpretation_of(text, rule(r)) ))
```

Moreover, the correct interpretation of a legal source is a valid rule:

```
Valid(rule(source_validity,
 legal_source(text) &
  correct_interpretation_of(text, rule(r)),
 valid(rule(r)) ))
```

Historical interpretation and systematic interpretation can be dealt with by similar principles of interpretation.

35.4 Case-based interpretation

Case law is an important source of reasons for and against interpretations of legal texts. If in a case a particular interpretation of a statutory regulation has been adopted, this is a reason why this interpretation is the 'correct' one. Analogously, the rejection of an interpretation is a reason against this interpretation. So we have the principle of legal precedent:

```
Valid(rule(legal_precedent,
 (∃c)(case(c) &
  decision(case(c)) =
   correct_interpretation_of(source, rule(r))),
 correct_interpretation_of(source, rule(r)) ))
```

(If there is a case in which it is decided that the correct interpretation of a particular source is some rule, then this rule is the correct interpretation of that source.)

Remember that this is no more than a principle that generates reasons which may be outweighed by other reasons. If there is only one relevant case, it leads to only one reason concerning the interpretation of the legal source, and this one reason will probably be decisive in the sense that

the interpretation of the previous case is a sufficient reason to adopt this interpretation for future cases.[44]

However, it is possible that in different cases, different interpretations of a source have been adopted:

```
decision(case(1)) = correct_interpretation_of(source, rule(rule_1))
decision(case(2)) = correct_interpretation_of(source, rule(rule_2))
```

If we also have

```
Incompatible(
 correct_interpretation_of(source, rule(rule_1)),
 correct_interpretation_of(source, rule(rule_2)) )
```

we must make a decision which court decision concerning the interpretation of `source` is to be preferred.

To facilitate the reading of the rest of this example, I introduce the following abbreviations:

```
d1 = (decision(case(c1)) =
        correct_interpretation_of(source, rule(rule_1)))
d2 = (decision(case(c2)) =
        correct_interpretation_of(source, rule(rule_2)))
e1 = (court_of(decision(case(c1))) = court1)
e2 = (court_of(decision(case(c2))) = court2)
```

One good reason in this connection is that recent cases prevail over old ones:

```
Valid(rule(old_and_recent_cases,
 recent(decision(case(c1))) & old(decision(case(c2))),
 outweighs({d1}, {d2},
   correct_interpretation_of(source, rule(rule_1))) ))
```

Another good reason is that decisions of superior courts prevail over decisions of inferior courts:

```
Valid(rule(superiority_of_courts,
 e1 & e2 & superior_to(court2, court1),
 outweighs({d2}, {d1}),
    correct_interpretation_of(source, rule(rule_1))) ))
```

If only one of these (or other) reasons is present to prefer a particular court decision above another one, this decision will provide us with the reason to prefer a particular interpretation of the source. For instance, if one court is superior, while both decisions are approximately equally recent, the decision of the superior court will be preferred and the interpretation of this court will be chosen.

If, however, one court is superior, while the decision of the other court is more recent, we must weigh reasons about which court decision to adopt. Suppose that if there have been changes in society that are relevant for the issue at stake, the most recent decision is to be preferred, while if there were no such changes, the superiority is decisive. In that case, we will have the following reasons:

---

[44] Of course there may be reasons against a particular interpretation that have nothing to do with precedents.

```
Reason( relevant_changes,
 outweighs(
  {recent(decision(case(c1))) & old(decision(case(c2)))},
  {e1 & e2 & superior_to(court2, court1)},
   outweighs({d2}, {d1},
    correct_interpretation_of(source, rule(rule_1))) )
 pro)
```

(If there have been relevant changes in society, the fact that one decision is more recent than another outweighs the hierarchical relation between the courts as a reason concerning the correct interpretation of a source.)

```
Reason( ~relevant_changes,
 outweighs(
  {e1 & e2 & superior_to(court2, court1)},
  {recent(decision(case(c1))) & old(decision(case(c2)))},
   outweighs({d1}, {d2},
     correct_interpretation_of(source, rule(rule_1))) )
 pro)
```

(If, however, there have been no relevant changes in society, the hierarchical relation between the courts outweighs the fact that one decision is more recent than another as a reason concerning the correct interpretation of a source.)

Suppose that there have been relevant changes. In that case, the more recent case will be adopted as authoritative for the interpretation:

```
Outweighs( {d1}, {d2},
 correct_interpretation_of(source, rule((rule_1)))
```

This leads us to the conclusion that the interpretation of the more recent decision `(decision(case(c1)))` is the correct one:

```
Correct_interpretation_of(source, rule(rule_1))
```

This extended example not only shows how cases can serve as the justification for preferring one interpretation above another one, but also how it is possible to weigh reasons about the issue how reasons (for some other conclusion) ought to be weighed, and - in general - how RBL deals with what would in other logics be meta-level arguments.

C. Some logico-philosophical reflections on reason-based logic

The examples of the previous sections illustrated how RBL matches the model of legal reasoning that was exposed in parts A of this paper.

Since RBL is closely tied to First Order Predicate Logic, one may ask whether RBL is really a logic, rather than a formalisation of a model of legal reasoning in terms of FOPL. For instance, in one sense, RBL is no more than an extension of FOPL. All deductive arguments are valid in RBL. Moreover every well-formed sentence of FOPL is a legal sentence of RBL and the other way round. RBL 'only' adds a number of inferential possibilities to those of FOPL.

Nevertheless, the 'spirit' of RBL is completely different from that of FOPL. RBL is a nonmonotonic logic, because the part that is added to FOPL is nonmonotonic. Moreover, the

central notions in RBL-inferences are those of a reason and of the weighing of reasons, not that of a deductive argument.

The difference between the two logics is most prominent where the use of the material implication is concerned. Most real-life arguments will be represented in FOPL as arguments of the form Modus Ponens, where the relation between one sentence and another sentence that is derived from it is represented by a material implication. In RBL, most of these arguments will be represented as the application of a rule. As a consequence, almost all real life arguments - and not only legal ones or those which make explicitly use of rules - will be dealt with by the part of RBL that is built on top of FOPL. Where RBL is technically spoken only an extension of FOPL, it is for practical purposes an alternative.

In the line of the question after the difference between RBL and FOPL lie some other ones: Why does RBL deviate from other logics in that it assigns such a central place to the notion of a reason? And why does RBL not contain some more derivational possibilities which seem very natural? These and related questions are the topic of the final part of this paper. Of necessity, the discussion of these question leads the argument into philosophical waters, away from the more down to earth approach of legal knowledge engineers.

36. Why Reason-Based Logic is a logic

An argument is said to be deductively valid if it is not possible that the premises of the argument are true, while the conclusion of the argument is false [Copi 1978, p. 32]. I will leave the difficult question of what kind of possibility is involved here, only to notice that defeasible arguments based on the application of rules, goals and principles, are certainly not deductively valid. Of course it is possible to reconstruct them in a way that is deductively valid, but only at the cost of replacing the original argument by another one (cf. note 5). The defence of the use of nonmonotonic logics in the field Artificial Intelligence and Law shows that many authors are convinced that deductive logic is not the ideal means to evaluate reasoning with rules, etc.

Still there are arguments based on the application of rules, and within these arguments one can distinguish between good and bad ones. To make this distinction, we need a standard that is to be applied to arguments in order to classify them as either good or bad. Clearly, different arguments can be good or bad for different reasons, but it seems attractive to have a standard that is sufficiently abstract to encompass all arguments. The standard of deductive logic has turned out to be too specifically made for arguments which are not based on rules and has therefore to be rejected as the general standard for the evaluation of arguments. As an alternative I propose the following standard:

An argument from a set of premises to a conclusion is good, and the conclusion of such an argument is justified, if the reasons that can be based on the premises which plead for the conclusion outweigh the reasons that can be based on the premises and which plead against the conclusion.

Notice that this is a very abstract standard which leaves it to domain theories to specify what counts as a reason, when a reason can be based on a set of premises, and when one set of reasons outweighs another set. RBL can be seen as an elaboration of this standard by specifying how reasons relate to rules etc, and which reasons can be based on a set of premises. RBL does not specify under which circumstances one set of reasons outweighs another set, with the exception of when one set of reasons is empty and the other one is not.

Because RBL is so closely related with a general theory of what counts as a good argument and as a justified conclusion, I would say that RBL is a kind of logic, rather than a mere formalisation of the notions reason, rule, etc., and their mutual relations.

37. Reasons instead of arguments

A reason in RBL has some similarity with an argument in other logics.[45] Why replace the familiar notion of an argument with the relatively unfamiliar one of a reason? To answer this question we must take a closer look to the notion of a reason and the argument form of Modus Ponens to which it is related.

Suppose we have the rule that thieves are punishable and the fact that John is a thief. In FOPL the following argument leads to the conclusion that John is punishable:

$(x)(Thief(x) \rightarrow Punishable(x))$
$Thief(john)$
$\therefore Punishable(john)$

In this argument there are two premises which logically have the same function: they are both necessary for the application of the rule of derivation that allows Modus Ponens arguments (Detachment). That the premises have different logical structures is irrelevant for their identical logical roles.

In RBL, on the contrary, rules and facts have different logical roles. The rule is used to give the fact some additional value: the fact is made relevant for the conclusion. This is reflected in the two step-procedure of RBL, where the first step concerns the 'transformation' of brute facts into reasons, and the second step concerns the derivation of a conclusion on the basis of reasons. The argument (in an abbreviated form) would run:

$Valid(rule(r, thief(x), punishable(x)))$
$Thief(john)$
$\therefore Reason(thief(john), punishable(john), pro)$
$\therefore Punishable(john)$

In this way the difference between the two causes of defeasibility can be made clear. Pure undercutters operate only in the first step, where a fact is prevented from becoming a reason. Rebutters can operate in two ways. First they can solely consist of colliding reasons which have to be weighed. In that case their only function lies in the second step. Second, they can function in both steps, as is the case of conflicting rules. In that case the presence of a rebutter excludes the operation of the competing rule in the first step, and provides a reason for an incompatible conclusion which decides the matter in the second step.

---

[45] This becomes very clear if the notion of a reason in my definition of a good argument is replaced by the notion of an argument, and the outweigh-relation is replaced by the defeats-relation. Cf. [Prakken and Sartor 1995] and the discussion in section 40.

[Pollock 1987] also assigns a central place to the notion of a reason. He distinguishes between decisive reasons and defeasible ones. His defeasible reasons are rather similar to the reasons of RBL. Pollock pays relatively little attention to the relation between principles and rules on the one hand, and reasons on the other hand, however.

A second advantage is that the division of arguments in two steps makes it easier to recognise goals as the basis of reasons. We have seen that goals fulfil a role in legal reasoning which is very similar to that of principles. Yet, because goals do not have a structure that resembles that of the material implication, their function in legal reasoning has hardly been recognised in logical theories. Goals simply do not fit in the classical model of arguments.[46] As soon as the attempt to model legal reasoning as a kind of detachment is given up, there comes room for the logical role of goals. We have seen in section 32 how RBL can give goals a logical role which is similar to that of principles.

38. Argument-based logics

RBL has some important characteristics in common with argument-based logics. Argument-based logics are built on the idea that given a set of premises, there can be a number of arguments which possibly lead to conflicting conclusions. It is also possible that one argument has as its conclusion that a subargument of another argument does not hold. In such cases, these arguments are not compatible and they need to be compared in order to determine which argument 'survives'. In the end a number of arguments survive, and the conclusions of these arguments are justified on the basis of the original set of premises. Other arguments are defeated, and their conclusions are not justified. Finally some conflicts of arguments may end undecided, and the conclusions of these arguments so to speak 'hang in the air'.

A recent version of an argument-based logic is the logic proposed in [Prakken and Sartor 1995] (PS-logic for short). RBL and PS-logic can to a large extent be translated into each other. There is, however, one essential difference. RBL takes the notion of a reason as central, which almost automatically leads to the result that the comparison takes place between *sets* of reasons in the form of weighing reasons. Almost all argument-comparison logics compare arguments one by one[47], and as a consequence have difficulties if larger sets of reasons (arguments) have to be weighed.

In the following subsections I will compare RBL to PS-logic to show both how similar they are in some respects, and to argue why the reason-based approach is still to be preferred.[48]

38.1 Arguments

The central notion in PS-logic is that of an argument. An argument is a tree of rules, where the conclusion of the one rule is one of the conditions of the following rule.[49] Premises are rules without conditions.

---

[46] Of course, it is always possible to give them a place by adding the necessary premises, or by rewriting them as principles. Cf. note 5.

[47] By reconstructing an argument, so as to include two or more reasons into one argument, this characteristic of argument-based logics can be circumvented. Reconstruction of arguments to obtain more attractive logical results has some disadvantages, which are discussed extensively in [Hage Reasoning].

[48] I will omit some technical details which are not essential for the comparison between PS-logic and RBL. This means that the description of PS-logic is not sufficiently precise to evaluate PS-logic on its own. I expect, however, that it is sufficiently precise for the comparison with RBL.

[49] The terminology used here is not that of [Prakken and Sartor 1995], but that of the present paper. However, this should not make an important difference.

Reasons in RBL are comparable to the joint conditions of the last rule in a chain. This means that an argument in PS-logic corresponds to one reason in RBL, with the difference that an argument contains its conclusion, while a reason in RBL does not contain its conclusion yet, although it is known for which conclusion a reason pleads. Take for instance the argument $a \rightarrow b \rightarrow c$. In RBL-terminology, b would be a reason for c.

In PS-logic there are two kinds of rule conditions, strong and weak ones. Strong conditions are the ones which must be shown to be satisfied to apply the rule. Weak ones are assumed to be satisfied unless their contrary is shown; they are satisfied by default. The weak conditions which are part of an argument, are called the assumptions of this argument.

In the terminology of PS-logic, one might say that each RBL rule has two weak conditions, namely that the rule is not excluded, and that if there are reasons against the application of the rule, they do not outweigh the reason for application based on the rule's applicability.

38.2 Conflicts of arguments

In PS-logic arguments can conflict. There are two kinds of conflict. One argument is said to *rebut* another argument if the conclusions of the two arguments are complementary. One argument *undercuts* another argument if the conclusion of the former argument is the complement of one of the assumptions of the latter argument. An argument which undercuts or rebuts another argument is said to *attack* that other argument.

In RBL-terminology, two arguments rebut each other if they plead for and against the same conclusion. An argument undercuts another argument, if the conclusion of the former is that some rule which occurs in the latter cannot be applied.

If an argument undercuts another argument, the former is said to *defeat* the latter. In case of two arguments which rebut each other, the *stronger* argument of the two defeats the other one. (If both arguments are equally strong, they both defeat the other argument.) Which argument is the stronger is determined by another argument, which has as its conclusion that some argument is stronger than another argument. This argument with a conclusion concerning the relative strength of other arguments is comparable to weighing information in RBL.

38.3 Three categories of arguments

A conflict between two arguments cannot be solved by looking only at the arguments involved in that conflict. It is possible that some argument A1 is invalidated by an argument A2, and therefore cannot invalidate the argument A3 anymore. A3 'survives' because it is saved from defeat through the help of A2.

This is comparable with the situation in RBL where the reason R1 which pleads for a conclusion C does not have competition from a reason R2 which pleads against C, because the rule which would lead to R2 is excluded.

Because it can only be determined which arguments survive by considering all arguments at once, the definition of justified arguments, the ones which win and justify their conclusions, can only be defined in terms of sets of arguments.

Given a set of rules (including rules without conditions) it is possible to construct a number of arguments. These arguments are divided into three categories, the *justified* ones, the *overruled* ones, and the *defensible* ones.

The justified arguments are those arguments which are not attacked at all, or are only attacked by arguments which are defeated by a(nother) justified argument.

The overruled arguments are the ones which are defeated by a justified argument.

The defensible arguments are the arguments which are neither justified, nor overruled.

In RBL, overruled arguments would not lead to reasons at all.

Defensible arguments would in RBL fall in one of two subcategories. If the unsolved conflict, which makes that the arguments are not justified, belongs to a subargument which leads to a reason, the argument does not generate a reason. If the unsolved conflict concerns the reasons that plead directly for or against the conclusion, the arguments lead to reasons, but the necessary weighing information to decide the collision of reasons is lacking.

Take for instance the following arguments:

A1: $a \rightarrow b \rightarrow c \rightarrow d$

A2: $e \rightarrow \sim b \rightarrow f \rightarrow \sim d$

A1 is attacked by the argument $e \rightarrow \sim b$ and A2 is attacked by the argument $a \rightarrow b$. Assuming that there is no argument that determines whether the argument for or against b is stronger, the arguments A1 and A2 would merely be defensible in PS-logic. In RBL it would be impossible to derive either b or ~b, and consequently the reasons c and f, respectively for and against d, would not be derivable.

If, however, we take the arguments:

A3: $a \rightarrow b$

A4: $e \rightarrow \sim b$

the situation becomes different in RBL. We have two reasons, one which pleads for b, and one which pleads against b. Without additional weighing information it is impossible to decide which reason wins.

In both cases PS-logic and RBL lead to the same final conclusion, namely that the final conclusions of the arguments cannot be derived. The constructions used are somewhat different, however.

In the case of justified arguments, PS-logic and RBL are rather similar. A justified argument in PS-logic is one where the chain of reasons that leads to a conclusion is uninterrupted and where the final reason is stronger than any reason against the same conclusion. The main difference with RBL is that in PS-logic a justified argument justifies its conclusion on the basis of one single reason. This reason is stronger than any other reason which pleads in a different direction, but the conclusion is still based on one reason. In RBL, however, every conclusion is based on the weighing of the *sets of reasons* for and against that conclusion. Very often one set will consist of only one reason, and the other set of at most one reason, but this does not subtract from the principle that in RBL it is sets of reasons which determine the conclusions, and not single reasons. As a consequence, the weighing information in RBL deals with sets of reasons, rather than with pairs of reasons.

It would be relatively easy to modify PS-logic so that the relation of strength which now holds between individual arguments, would hold for sets of arguments. With that modification, PS-logic

would lead to the same results, and along essentially the same line as RBL.[50] The only remaining difference would be the way in which arguments are conceptualised. In RBL the notion of a reason takes the central place, and there is a two step-procedure that consists of collecting and weighing reasons. In PS-logic the notion of an argument plays the central role, and the two step-procedure consists of collecting and comparing all arguments. The main differences are first that in RBL reasons are weighed setwise, while in PS-logic arguments are compared individually, and second that the exclusion of rules plays in RBL its role in the first phase of collecting reasons, while in PS-logic it plays its role in the second phase of comparing arguments. The first difference leads to other differences, but can easily be remedied in PS-logic. The second difference makes in the end no difference at all.

38.4 Final remarks regarding the relation between PS-logic and RBL

PS-logic and RBL have a different conceptualisation of arguments, but despite this difference they are rather similar. There is one fundamental difference, namely that PS-logic compares arguments one by one, while RBL weighs reasons setwise. This difference is a natural consequence of the difference in conceptualisation. The notion of a reason almost automatically leads to the idea of weighing, while weighing is not confined to one item pro scale. The notion of an argument, on the contrary, is strongly related to that of a proof. Traditionally a proof guarantees the truth of its conclusion. In a nonmonotonic logic, this is not possibly anymore, because there can be arguments for incompatible conclusions. The natural suggestion in that case is that the strongest argument wins, which is precisely the approach adopted in PS-logic.

We have seen in the first part of this paper that reasoning with rules and goals can adequately be modelled by means of reasons that are weighed setwise. The comparison of arguments is less suitable because this approach cannot account for the accrual of reasons that plead for (or against) the same conclusion. That is why a reason-based conceptualisation of legal arguments seems more attractive than an argument-based conceptualisation. It is, however, possible to enhance the argument-based approach to allow setwise comparison of arguments. Then it also becomes possible in argument-based approaches to deal with the weighing of reasons, although the reason-based conceptualisation remains more attractive, precisely because the weighing of sets of reasons is a natural consequence of it.

Some other differences are that PS-logic has a proof-theory, which RBL presently has not, while on the other hand, RBL at this moment has more extensive facilities to reason about rules. I will not discuss these differences here, however.

39. Possible extensions of RBL

In this section I will discuss some possible extensions of RBL. One other extension has already been mentioned in section 34.3, where general rules for case-based reasoning were presented.

[50] This may not be fully correct for some deviant arguments in which rules are used that refer to themselves and to each other. Cf. section 29.2. For those cases, PS-logic might offer a slightly different solution than RBL.

39.1 Reasons for incompatible conclusions

Suppose that I receive a postcard from England which was signed by Tracey, and in which she writes me that she spends some weeks with her uncle there. This is a reason to assume that Tracey is in England. Suppose, moreover, that when I spoke Tracey a few weeks ago, she told me that she would leave for Australia and would stay there for a couple of months. This is a reason to assume that Tracey is in Australia.

Strictly spoken, the reasons why Tracey is in England and why she is in Australia have no logical relationship with each other. Still it seems that if I have reason to assume that Tracey is in Australia, this is also a reason to assume that she is not in England, and the other way round. This predicament can be solved by assuming that if two states of affairs are incompatible, a reason why the one state of affairs obtains is ipso facto a reason why the other state of affairs does not obtain:

```
(state_1)(state_2)(Incompatible(state_1, state_2) →
 (Reason(reason, state_1, pro) → Reason(reason, state_2, con)))
```

This cannot be turned around to the assumption assuming that if two states of affairs are incompatible, a reason why the one state of affairs does not obtain is a reason why the other state of affairs obtains, because this would commit the fallacy of the excluded middle.

It remains to be specified when two states of affairs are incompatible. The most obvious case is when they are each others opposite:

```
(state_1)(state_2)
 ( (State_1 ≡ ~State_2) → Incompatible(state_1, state_2))
```
[51]

For instance, the states of affairs that John is a thief and that he is not a thief are complementary and therefore also incompatible.

The lightest form of incompatibility of two states of affairs is when they happen not to obtain both:

```
(state_1)( state_1)
 ((State_1 ≡ ~State_2) →
 Incompatible(state_1, state_2))
```

A special form of incompatibility deals with the relative weights of sets of reasons:

```
(set_1)(set_2)
 Incompatible(
  outweighs(set_1, set_2,conclusion),
  outweighs(set_2, set_1,conclusion)
 ))
```

Sentences like these about incompatibility and its effects might be added a axioms to RBL.

---

[51] stands here for the traditional operator for logical necessity. Cf. [Chellas 1980, p. 4].

39.2 Epistemic rules

Some ways to derive rules seem so natural that it may seem strange that these ways to derive rules are not built into RBL. For instance if A is a reason for B, and if B logically entails C, then surely A must be a reason for C. Still, this cannot be derived in pure RBL. From the theory

```
Reason(a, b, pro)
 (b → c)
```

it cannot be derived that

```
Reason(a, c, pro)
```

Analogously, from

```
Valid(rule(r1, a, b))
 (b → c)
```

it is not possible to derive

```
Valid(rule(r2, a, c))
```

Why is this not possible? The answer is that the conclusion of this latter argument is a statement about the validity of a rule. It depends on the domain, under which circumstances a rule is said to be valid, and domain dependent knowledge is not part of RBL. Clearly, in a domain that contains the rule `r1`, there will normally be a reason to adopt the rule `r2`, but this need not be the case. That is why RBL does not allow the inference that leads to the conclusion that rule `r2` can be derived from the given theory.

However, this is not the whole story. If the rules `r1` and `r2` are epistemic rules, it is very attractive[52] to make it possible to derive

```
Valid(rule(r2, a, c))
```

from

```
Valid(rule(r1, a, b))
 (b → c)
```

To this purpose the following rule might be added to RBL theories, which makes it possible to derive a more specific form a more general rule:

```
Valid(rule(limited_transitivity_of_epistemic_rules,
 epistemic_rule(rule(r1)) &
  valid(rule(r1, cond, concl1)) &
   (concl1 → concl2),
 valid(rule(r2, cond, concl2)) ))
```

One can even imagine that the following inference relations between rules are allowed:

---

[52] The reasons why this is so attractive are discussed extensively in [Hage Reasoning].

```
Valid(rule(extended_transitivity_of_epistemic_rules,
 epistemic_rule(rule(r1)) &
  epistemic_rule(rule(r2)) &
  valid(rule(r1, condition, conclusion1)) &
  valid(rule(r2, conclusion1, conclusion2)),
valid(rule(r3, condition, conclusion2)) )).
```

40. Summary

In this paper I have argued for a particular way of conceptualising legal reasoning and for a logic that is based on this conceptualisation. The picture of legal reasoning defended here has two main characteristics. First it assigns a central place to the notion of a reason, rather than to the notion of an argument. In the line of this, rules, principles and goals are strictly distinguished from statements, which have a different logical role to play. And second, the picture of legal reasoning distinguishes two layers in the law, the interaction of which to a large extent determines the nature of legal arguments.

Not coincidentally, the conceptualisation of legal reasoning as reason-based, with its distinction between rules etc. and statements, makes it easy to deal with the interaction between the two layers of law. On the one hand, normally the applicability of a rule that replaces principles and goals excludes the application of these principles and goals. On the other hand, in exceptional cases, the application of a rule can be blocked by the violation of one of its underlying goals or principles.

Traditional logics, such as first order predicate logic, are not very suitable to deal with legal arguments as they are modelled in the picture of reasoning that is proposed here. That is why I proposed a special logic which assigns a central place to the notion of a reason: Reason-Based Logic. Although RBL is not a special logic for legal reasoning, it was nevertheless developed with the demands of legal reasoning in mind. By means of a number of simple and not so simple examples it was shown how RBL can be used to model familiar patterns of legal reasoning.

Finally some philosophical reflections on the nature of RBL were given, with the threefold purpose to argue why RBL is a logic and not merely a formal description of a model of legal reasoning, to compare RBL with its main competitor, the approach based on the comparison of arguments, and to explain why RBL is not extended with some natural-seeming additions.

ACKNOWLEDGEMENTS

The author wishes to thank Henry Prakken and Giovanni Sartor for their suggestions and criticisms which have improved this paper. Bart Verheij deserves credit for helping develop the formalism of RBL.

The author thankfully recognises the cooperation of the Maastricht Law Faculty in giving him exemption from part of his teaching tasks during 1993 and 1994.

The research for this paper was partly made possible by a grant of the Dutch Foundation for Knowledge Based Systems.

References

Alexy, R. (1978). *Theorie der juristischen Argumentation.* Suhrkamp Verlag, Frankfurt am Main.

Alexy, R. (1985). *Theorie der Grundrechten*, Nomos Verlagsgesellschaft, Baden-Baden.

Anscombe, G.E.M. (1957). *Intentions*, Blackwell, Oxford.

Ashley, K.D. (1991). Reasoning with cases and hypotheticals in HYPO. *International Journal of Man-Machine Studies*, vol. 34, p. 753-796.

Baker, G.P. (1977). Defeasibility and Meaning. P.M.S. Hacker and J. Raz (eds.) *Law, Morality and Society*. Clarendon Press, Oxford, pp. 26-57.

Chellas, B.F. (1980). *Modal logic; an introduction*. Cambridge University Press, Cambridge e.a.

Copi, I.M. (1978). *Introduction to Logic*. 5th ed. MacMillan Publishing Co., New York.

Dworkin, R. (1978). *Taking Rights Seriously*, 2nd ed., Duckworth, London.

Fuller, L.L. (1958). Positivism and Fidelity to law: A Reply to Professor Hart, *Harvard Law Review* 71, p. 630-672.

Gordon, Th. F. (1994): The Pleadings Game: An Exercise in Computational Dialectics, *Artificial Intelligence and Law* 2, pp. 239-292.

Hage, J.C. (1991). Monological reason based reasoning. (J.A. Breuker, R.V. de Mulder and J.C. Hage eds.) *Legal Knowledge Based Systems. Model-based reasoning*, Koninklijke Vermande BV, Lelystad, p. 77-91.

Hage, J.C. (1993). Monological Reason Based Logic. *Proceedings of the Fourth International Conference on Law and Artificial Intelligence*, ACM-press, Amsterdam, p. 30-39.

Hage, J.C. (1995). Teleological reasoning in reason-based logic. *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, ACM, New York, pp. 11-20.

Hage, J.C. (Reasoning), *Reasoning with rules*, to be published by Kluwer Academic Publishers in the Law and Philosophy series.

Hage, J.C., R. Leenes, and A. Lodder (1994). Hard cases; a procedural approach. *Artificial Intelligence and Law*, vol. 2, pp. 113-167.

Hage, J.C. and H.B. Verheij (1995). Reason-Based Logic: a logic for reasoning with rules and reasons. *Law, Computers and Artificial Intelligence*, vol. 3, nrs. 2/3, pp. 171-209.

Hare, R.M. (1952). *The Language of Morals*. Oxford University Press, Oxford.

Hare, R.M. (1963). *Freedom and Reason*, Oxford University Press, Oxford e.a.

Hart, H.L.A. (1961). *The Concept of Law*. Clarendon Press, Oxford.

MacCormick, D.N. (1978). *Legal Reasoning and Legal Theory*, Clarendon Press, Oxford.

MacCormick, D.N. and O. Weinberger (1987), *An Institutional Theory of Law*, Reidel, Dordrecht etc.

MacCormick, D.N. and Summers, R.S. (1991). *Interpreting statutes: a comparitive study*, Darthmouth Publishing Company Limited, Hants and Brookfield.

Peczenik, A. (1989). *On Law and Reason*, Kluwer Academic Publishers, Dordrecht e.a.

Pollock, J.L. (1987). Defeasible Reasoning. *Cognitive Science* 11, p. 481-518.

Prakken, H. (1993). *Logical tools for modelling legal argument*, Ph.-D.-thesis Amsterdam.

Prakken, H. and G. Sartor (1995). On the relation between legal language and legal argument: assumptions, applicable, *Proceedings of the Fifth International Conference on Artificial Intelligence and Law,* ACM, New York, pp. 1-10.

Raz, J. (1975). *Practical Reason and Norms,* Hutchinson, London.

Raz. J. (1979). *The Authority of Law; Essays on Law and Morality*, Clarendon Press, Oxford.

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

Sartor G. (1991). The Structure of Norm Conditions and Nonmonotonic Reasoning in Law, *Proc. of the Third International Conference on* Artificial Intelligence and Law*,* ACM Press, p. 155-164.

Searle, J. (1983). *Intentionality; An essay in the philosophy of mind*, Cambridge University Press.

Searle, J. (1995). *The construction of social reality*. The Free Press, New York.

Soeteman, A. (1989). *Logic in Law. Remarks on Logic and Rationality in Normative Reasoning, Especially in Law.* Kluwer Academic Publishers, Dordrecht e.a.

Soeteman, A.(1991). Hercules aan het werk; over de rol van rechtsbeginselen in het recht. Dolman e.a. (eds.) *Rechtsbeginselen,* Ars Aequi, Nijmegen, p. 41-56.

Strawson, P. (1971). *Logico Linguistic Papers*, Methuen, London.

Tammelo, I. and Schreiner, H. (1977). *Grundzüge und Grundverfahren der Rechtslogik*, Verlag Dokumentation Saur KG, München.

Toulmin, S.E. (1958). *The Uses of Argument,* Cambridge University Press, London, New York.

Valente, A. (1995). *Legal Knowledge Engineering. A modelling approach.* Thesis. IOS Press, Amsterdam.

Verheij, H.B. and J.C. Hage (1994). Reasoning by analogy; a formal reconstruction. H. Prakken, A.J. Muntjewerff and A. Soeteman (eds.), *Legal knowledge based systems; the relation with legal theory,* Koninklijke Vermande, Lelystad, pp. 65-78.