MODERATELY NATURALISTIC DEONTIC LOGIC

Jaap Hage
Maastricht University
Dept. of Metajuridica
P.O Box 616
6200 MD Maastricht
Netherlands
email: jaap.hage@metajur.unimaas.nl

Summary

Deontic naturalism is here taken to be the view that it is possible to give the truth conditions of deontic sentences by means of solely non-deontic sentences. This paper argues for a moderate version of deontic naturalism and describes a variant of deontic logic that takes this form of naturalism into account. The argument hinges amongst others on a strict distinction between deontic rules and deontic facts and on the treatment of rules, including deontic rules, as logical individuals. The logic is presented by model-theoretic means.

1 The 'naturalistic fallacy'

At the beginning of this century the Cambridge philosopher G.E. Moore introduced the 'naturalistic fallacy' in ethical theory (Moore 1903). Moore's own description of what this 'fallacy' amounted to was not very clear, but Frankena (1939) and Taylor (1961) improved Moore's work in this respect. The basic intuition behind the notion of the naturalistic fallacy is that purely factual sentences and purely evaluative or deontic sentences have no meaning components in common. As a result it is:

1. impossible to give the meaning of evaluative and deontic sentences in terms of purely factual sentences;
2. impossible to deduce evaluative and deontic sentences from purely factual ones; and
3. possible to agree about all purely factual sentences and still to disagree rationally about some evaluative or deontic sentence (Taylor 1961, p. 242f.).

Although non-naturalism, the theory that naturalism is wrong, has been very fashionable during the 20th century (Brecht 1959), Searle (e.g. 1969, p. 175f.) has argued that in a sense it is possible to deduce an ought from only factual and (almost) analytical sentences. Although I am not very happy with the way in which Searle presented his argument, I subscribe to the view that it is possible to infer (not: deduce) an ought from only factual sentences.[1]

This possibility is closely related to the possibility of giving the truth conditions of deontic sentences in terms of non-deontic sentences. In my opinion the semantics for deontic logic should reflect that the truth-values of deontic sentences in the end depend solely on the truth-values of non-deontic sentences. I do not adopt the view, however, that norms have truth-values. My view that deontic sentences have truth-values and that these truth-values in the end depend on the truth-values of non-deontic sentences presupposes a strict distinction between norms and deontic sentences. In contrast to what some other authors, e.g. Von Wright (1963 and 1998) and Alchourrón (1969 and 1993), seem to do, I do not regard deontic sentences as sentences about norms. I distinguish

[1] Actually, I already disagree with the terminology in which the problem is described. As I will argue in section 2, there is no distinction between factual and deontic sentences.

three categories: norms, sentences about norms, and deontic sentences. My naturalistic position regards deontic sentences.

This paper has the following structure. First, I present a theory of deontic facts that is moderately naturalistic. This theory also delineates the distinctions between norms, sentences about norms, and deontic sentences (sections 2 and 3). In section 4 I summarise these distinctions. The logical consequences of this moderate version of deontic naturalism are first discussed informally (section 5). Then I show how a model-theoretic characterisation can be given for a naturalistic version of deontic logic (sections 6-9). After a brief discussion of some examples (section 10), the findings of the paper are summarised and commented upon in the conclusion (section 11).

2 Deontic facts as a kind of reason-based facts

Many facts only obtain to the extent that other facts obtain. In this connection at least two cases can be distinguished. Sometimes one or more facts add up to some other, new fact. For instance, the facts that in chess the black king is threatened by a white piece, and that this threat cannot be taken away in one move, together add up to the fact that black is check-mated. This latter fact is in some sense nothing else than the combination of the former two, but in another sense it is a different fact, because if the rules of chess would have been relevantly different, the latter fact would not have obtained. Moreover, it is not only a matter of the meaning of the expression 'check mate' that procures the relation between on the one hand being threatened and not being capable to remove the threat in one move, and on the other hand being check-mated. The rules of chess concerning the issue of check mating might have been different, without a change in the meaning of the expression 'check mate'.[2]

Other examples of situations in which one or more facts add up to some new fact are that the fact that a soldier runs away at the approach of the enemy boils down to the soldier being a coward (or prudent), and that the composition and the use of colours in the picture make the picture into a beautiful one.

It is also possible that one or more facts by themselves do not add up to some new fact, but that in some sense they 'cause' this new fact to obtain. For instance, that I hit a winning service 'causes' it to be the case that I take advantage in the game of tennis we are playing. Or, that I contract to buy your house brings me under the obligation to pay you the price of the house.

In all of these cases, there is some substrate of facts that, thanks to some rules add up to, or cause, some other facts. These new facts cannot obtain without the basis provided by those other facts. I propose to call these new facts *reason-based facts*, because the facts on which they are based are the reasons why the new facts are present.

The phenomenon of reason-based facts has - under different names - often been described. Hare (1963) noticed that value sentences and ought sentences are 'supervenient' on other sentences. In connection with the law, MacCormick and Weinberger gave us an 'Institutional Theory of Law', in which legal facts are described as a kind of *institutional facts* (MacCormick and Weinberger 1986). This theory is elaborated in Ruiter's *Institutional Legal Facts* (Ruiter 1993). In (Searle 1995) a more general theory of institutional facts is presented.

2 For those who disagree with this example, I give a legal one. The rules how to obtain a property right may change, without a change in the meaning of the words 'owner' and 'property'. More thorough discussions of this phenomenon can be found in MacCormick 1974.

Reason-based facts can in their turn underlie new reason-based facts, as is illustrated by the chess-example above. The facts that the black king is threatened, and that the threat cannot be removed in one move, are both reason-based themselves.

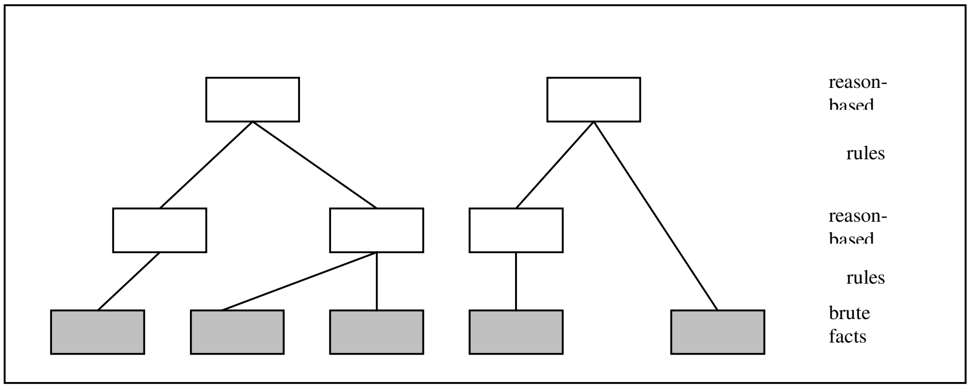

Figure 1: Reason-based facts built on top of brute facts

The notion of a reason-based fact is ambiguous. On the one hand it may mean a concrete fact that obtains because of its underlying reasons. On the other hand it may also mean a type of fact, that *can* only obtain on the basis of underlying reasons. The existence of a rule is an example of a concrete fact that may be reason-based in this first, broad sense, but does not belong to a fact-type that is reason-based in the second, narrow sense. A rule can exist because it was made in accordance with a procedure for making rules, but it may also have 'grown' in a certain society. The validity of a contract, on the contrary, is reason-based in the second, narrow sense. A contract cannot be valid if there are no reasons for it. In this paper, I am primarily interested in reason-based facts in the narrow sense.

Facts that are not reason-based in the broad sense are called *brute*. Reason-based facts are built on top of brute facts, and the way in which they are built on top of them is defined by rules (in a broad sense). For instance, the brute fact that I told you that I would pay you a thousand dollars is made into the reason-based fact that I promised to pay you a thousand dollars. The connection between this brute fact and the reason-based fact to which it amounts is created by the rule that saying that one promises counts, under suitable circumstances, as promising (Searle 1969, p. 57f.). The reason-based fact that I promised to pay thousand dollars in turn underlies the reason-based fact that I owe you a thousand dollars. This connection is made by the rule that one ought to do what one has promised. Rules are, in a sense, the cement of the universe by which facts are cemented together.[3]

The principal difference between brute facts and reason-based facts is that the former can obtain independently, while the latter always depend for their existence on reasons. Nevertheless, both brute and reason-based facts are part of reality, and the sentence that Mount Everest is the highest mountain is not more true than the sentence that the dollar is the monetary unit of the United States.

In my opinion, a special kind of reason-based facts in the narrow sense are the so-called *deontic facts*, such as that John ought to be punished, and that Gerald ought to pay Jane a thousand dollars. Deontic facts are to be distinguished from the deontic rules on which they are based. A deontic rule is for instance the rule that thieves ought to be pun-

---

3 See Mackie 1980, quoting Hume's Abstract of his A Treatise of Human Nature.

ished. I use the term *norm* for these deontic rules. Norms are to be distinguished strictly from deontic facts. Norms are also to be distinguished from deontic sentences, the sentences that express deontic facts.

Confusingly the sentence 'Thieves ought to be punished' is ambiguous between the formulation of a rule and the expression of the relation in the world between that somebody is a thief and that this person ought to be punished. Rule formulations denote, as I argue in section 3, logical individuals, while the expression of the relation in the world is a special sentence, that is true or false, depending on the world.

Deontic facts are just as much part of reality as are brute physical facts. Still it is clear that deontic facts are, in contrast to brute facts, based on norms[4], and in that sense they differ from the aforementioned brute facts. It seems that, as a consequence, deontic facts cannot fully be analysed in terms of brute facts, and that therefore non-naturalism is true.

This conclusion would be premature, however. The truth conditions of sentences expressing reason-based facts, including deontic facts, can be specified fully in terms of the existence of rules, the facts that satisfy their conditions, and the absence of exceptions. Although both the existence of rules, the facts that satisfy their conditions and exceptions may be reason-based themselves, these facts can in turn be analysed in terms of the existence of rules and facts that satisfy the rule's conditions. In the end the analysis in terms of truth conditions must end with brute facts, because otherwise the reason-based facts would, so to speak, 'hang in the air' (Searle 1995, p. 34/5).

Two points should be noted in this connection. First, the analysis of reason-based facts in terms of brute facts is not an analysis of what the reason-based facts *are*. Neither is the analysis an analysis of what sentences that express reason-based facts *mean*. The analysis only gives truth conditions. The reason-based facts themselves have an additional value that transcends their conditions of obtaining, and correspondingly, the sentences that express them have a surplus of meaning (Hage 1987). Consequently it is not possible to define deontic sentences in terms of non-deontic ones, if we want the definition to give the full meaning of the deontic sentence. But this impossibility does not prevent the possibility of a full specification of the *truth conditions* of deontic sentences in terms of brute facts.

Second, the existence of rules, including the existence of deontic rules, is assumed to be, in the end, a matter of brute fact. I discuss this assumption in the next section.

3 The existence of rules.

In this section I give a brief account of the modes of existence (or validity) of institutional and social rules.[5] I need this account to argue that the existence of rules, including deontic rules, is a matter of non-deontic fact. The reader needs not accept the full account given below, if only he accepts the conclusion that rules, including norms, exist as a matter of non-deontic fact.

Rules exist mainly in two ways, that is in the context of social institutions, and as social rules.[6] Rules exist in the context of institutions if they satisfy the criteria that hold in the context of that institution for the existence of rules. The obvious example of institutional rules are legal rules made by legislation. Because the legislator has within the in-

4 Actually I think that deontic facts can also be based on goals and principles. See Hage 1997. To make the point of this paper, these additional sources of deontic facts are not relevant.

5 The exposition in this section abbreviates an argument in Hage 1997, p. 49f. and 91f. A similar view is expressed in Lagerspetz 1995, ch. 7.

6 In Hage 1997 I distinguished also personal rules, but personal rules make up a category that is not hardly relevant in the present context. In section 5 I briefly mention them in the discussion of rational disagreement.

stitution of the law the power to make legal rules, rules made by the legislator according to the procedures for legislation are valid legal rules.[7]

If rules exist in the context of an institution, they cause certain connections between facts to hold within that institution. For instance, if the law has the institutional rule that thieves ought to be punished, the fact that somebody is (legally) a thief tends to go together with the fact that this person ought legally to be punished. Notice, however, that this connection between the types of fact that somebody is a thief and that this person ought to be punished is ontologically dependent on the existence of the rule. If there are no thieves and there is therefore no correlation between the fact that somebody is a thief and that he ought to be punished, the rule can still exist. (It would be very effective, one might even argue.) The connection between the fact that somebody is a thief and that he ought to be punished depends, on the contrary, on the existence of the rule.

I assume that social rules exist in the form of primarily mental behaviour and, in a derived sense, also the external behaviour of the members of a social group. A rule of the form 'If A then B' exists as a social rule in a group G, if and only if most members of G tend to accept the fact that B if they accept the fact that A, and think that other members of G should do the same. For instance, the rule that car drivers ought not to be drunk is a valid social rule in the community of Outopos if most members of Outopos tend to think that if somebody is a car driver (s)he ought not to be drunk. Moreover, they think that other members of Outopos should think the same.

From this assumption it follows that social rules also involve a connection between types of facts (in the example, being a car driver, and having the duty not to be drunk), but here there is no clear ontologically priority of the rules over the connection between the facts. On the macro level, the existence of the rule is the result of the connection that the group members make between the types of facts, while for individual members the existence of the rule is why they make the connection.

Institutional rules can only exist in the context of an institution, that can be governed by institutional rules, but must (to avoid an infinite regress) in the end be based on social rules (Hart 1961, p. 97f.). The existence of social rules depends in the end on the mental behaviour of the members of the social group in which the rule exists. Therefore I assume that the existence of both institutional and social rules ultimately rests in mental behaviour of humans, that is on brute facts.

The reader may have noticed that at the end of section 2 I gradually slipped into writing about the existence of a rule. If rules had the same logical standing as sentences, this would be a strange way of writing. We do not say that sentences exist in the context in which we say of rules that they exist. Nevertheless it is common parlance to say of rules that they exist, are binding, were issued or repealed, are effective or violated, etc. In short, in common parlance rules are often treated as if they were logical individuals, rather than sentences. Only if rules are used in rule-application do they seem to function as sentences, but even then the expression 'application' suggests the contrary, because sentences are not applied.

Besides common parlance there is a more fundamental reason why rules are not sentences. The sentences with which logic normally deals are true or false, depending on whether the states of affairs that they express obtain. In the terminology of Searle (1979, p. 3f.) descriptive sentences aim to fit the world. In the case of rules it is different. If the rule exists that thieves ought to be punished, this *makes* it true that thieves ought to be

---

7 The notion of validity as used here only means existence, and has no implications whatsoever concerning the issue whether the rule ought to be obeyed.

punished. The world fits the rule, rather than the other way round. In other words, rules are constitutive.

The same point is also illustrated by the ontological priority of, in particular, institutional rules above the connection between types of facts that they generate. The existence of institutional rules causes a connection between types of facts that can be described by descriptive sentences. The legal existence of the rule that thieves ought to be punished makes it true that legally thieves ought to be punished. Even if the rule and the descriptive sentence have the same formulation, they are different entities. The rule can exist, but it is not true or false. The descriptive sentence, on the contrary, is true or false, but its existence is quite a different matter. For these reasons I propose to treat rules as logical individuals, rather than as complete sentences. The sentence to the effect that a particular rule exists is on this view a logical atom.

4 Some distinctions

In the previous sections I have made a number of distinctions. It may be useful to summarise them and make some additions. The first distinction is between *sentences* and *the world* that is described by them. Sentences express *states of affairs*, and if these states of affairs obtain in the actual world, they are *facts* and the sentences that express them are *true*. States of affairs are either *concrete* or *generic*. Closed sentences express concrete states of affairs, while open sentences, such as the sentence that *x* is walking, express generic states of affairs.[8]

Some kinds of facts may obtain without the presence of human culture. The fact that Mount Everest is snow-covered would be a plausible candidate. Such facts are called *brute facts*. Other facts presuppose human culture, such as the fact that the dollar is the monetary unit of the United States. These facts, that presuppose other facts for their existence, are called *reason-based facts*. A special kind of reason-based facts are the so-called *deontic facts*, that is facts of the type that something ought (not) to be the case, or ought (not) to be done.

In the world it is possible to distinguish between facts and *individuals*. Facts are expressed by true sentences; individuals are denoted by *terms*. It is, however, possible to treat facts, or, in general, states of affairs, as individuals too, in which case they can also be denoted by terms.

Next to states of affairs there is another less usual kind of individuals, that is *rules*. There are many different kinds of sentences about rules, expressing for instance that a rule was created, that it exists, that it is modified, that it is effective, and that it was abolished. These sentences contain terms that denote the rule at stake. Rules with a deontic conclusion, such as the rule that thieves ought to be punished, and that it is forbidden to drive on the left hand side of the road, are called *norms*.

Norms can be given a linguistic expression, such as ‘It is forbidden to steal’. These expressions describe the content of the norm, but they are not sentences. They do not state that a norm has this or that content, but merely formulate that content. Neither do *norm-expressions* describe the world. The norm-expression ‘It is forbidden to steal’ does in particular not express the fact that it is forbidden to steal. There is, however, also a deontic sentence with the same formulation. This sentence expresses the fact that it is forbidden to steal.[9]

---

8 My notion of a generic state of affairs is similar to the notion of a generic state of affairs as used in Von Wright 1963, p. 26 and to the notion of a situation type as used in Barwise and Perry 1983, p. 53f.

9 This distinction is related to Alchourrón and Bulygin's (1981) distinction between the expressive and the hyletic conception of norms and, even more clearly, to Kelsen's (1960, p. 57f.) distinction, in a legal context, between 'Rechtsnorme' and 'Rechtssätze'.

From the foregoing distinctions it follows that one should distinguish between norms, norm-expressions, sentences about norms, deontic states of affairs, including deontic facts, and deontic sentences that express deontic states of affairs. In section 6 I distinguish between two branches of deontic logic, one that deals with the derivation of deontic sentences from sentences about norms, and one that deals with the logical relations between deontic sentences.

5 Logical consequences of moderate naturalism

The theory sketched above, that deontic facts are a kind of reason-based facts, implies a moderate form of naturalism. It is a naturalistic theory, because it holds that what ought to be done, or ought to be the case, ultimately depends on brute, non-deontic facts and on nothing else. This dependency goes so far that it is possible to specify the truth conditions of deontic sentences in terms of solely non-deontic sentences. Its naturalism is moderate, because all the implications of non-naturalism mentioned at the beginning of this paper still hold true. In my opinion it is impossible to give the meanings of deontic sentences in terms of purely factual sentences, to deduce deontic sentences from purely factual ones, and it is rationally possible to agree about all non-deontic sentences and still to disagree about some deontic sentence.

Let me be brief about the issue of definition. The naturalistic analysis given above only deals with the truth conditions of deontic sentences. It does not deal with their meanings. The meanings of deontic sentences do not coincide with their truth conditions (Hare 1963, chapter II). A definition of the full meaning of deontic sentences in terms of non-deontic facts is in my opinion impossible (Hage 1987, chapter IV).

The issue concerning rational disagreement is somewhat more complex. The basic idea is that it is possible to agree about all brute facts, and to disagree about the rules, and consequently to disagree about the conclusions of the rules. I have assumed that the existence of rules is in the end a matter of brute fact, so that agreement about the brute facts implies agreement about the rules. How, then, is it possible to agree about the facts, and not to agree about the rules?

In the end it is not possible to agree about the brute facts and not to agree about the rules. However, the facts at stake may be a matter of individual psychology. Sometimes different persons employ different rules, or standards. Suppose that you and I have different standards for social behaviour. You think that individuals ought to help people in need, while I think it is the task of the government. As a consequence, you judge that I ought to give the beggar on my way some money, while I judge that I need not do that (and that it is the task of the government to provide better social security). We disagree about whether I ought to give the beggar money, because we use different standards. Nevertheless we agree that according to your standards I ought to give money, and that according to my standards I need not. This is all that can rationally be said about the issue, because the question whose standards are 'objectively' valid makes no sense in the case of personal standards. Since we agree about all brute facts, we agree about the deontic facts, but we must admit that for me and for you there are different deontic facts. Our deontic facts are relative to our sets of standards. This situation is comparable with different legal systems that lead to different deontic sentences. Agreement on all brute facts leads to agreement on deontic sentences as relativised to a legal system.

If rules are treated as logical individuals, as they should in my opinion be treated, it is not possible to *deduce* the conclusion of a rule from the sentences that this rule exists and that its conditions are satisfied, For instance, it is not possible to deduce that John ought to be punished from the sentences that the norm exists that thieves ought to be punished, and that John is a thief. Nor does my version of naturalism involve another possibility of *deducing* deontic sentences from solely non-deontic ones. In the next section, however, I argue that the impossibility of deducing deontic sentences from solely non-deontic ones does not stand in the way of the possibility of deriving them non-deductively, but nevertheless validly in the sense that the truth of the premises guarantees the truth of the conclusion.[10]

Notice, by the way, that the sentence that a particular rule exists is non-deontic, because the mere existence of 'things' is not deontic, even if these 'things' are norms. Those who nevertheless think that norms have a special, deontic, mode of existence, should consider the mode of existence of non-deontic rules such as rules of competence. There is nothing deontic about the existence of the rule that the mayor of a town is competent to institute a fire brigade, and there is no reason why deontic rules have a different mode of existence than non-deontic ones.

6 A simple logic for the application of deontic rules

If rules are treated as logical individuals, reasoning with rules must be non-deductive. Since rules are often used to draw conclusions and these conclusions are considered to be justified on the basis of the rule and the facts that satisfy its conditions, a notion of logical justification must be involved that is broader than the traditional one. Some form of non-deductive logic for rule-application is necessary.

An advantage of treating rules as logical individuals is that it becomes possible to define the logical behaviour of rules from scratch. In my opinion the logic of rule-application is rather complex (Hage 1997), but for the purpose of this paper I simplify this logic considerably, to focus on the possibility of deriving deontic conclusions from non-deontic premises. Moreover, I confine my exposition to a deontic logic of the ought-to-be type. I do not think that an ought-to-be logic is adequate to deal with most deontic arguments. However, the issues concerning the derivation of deontic conclusions from non-deontic premises are in my opinion the same for ought-to-be logics as for ought-to-do logics, while ought-to-be logics can be introduced with less irrelevant complications.

The basic idea behind the logic I propose here is that instead of the rule formulation, the sentence that the rule exists is used to model rule-application. In this way rules are treated as logical individuals. Although rules by themselves do not function as premises, their internal structure is nevertheless used in making derivations. If the conditions of a rule are satisfied and there is no exception to the rule, the conclusion of the rule follows.

It may be objected that this is not a satisfactory account of rule application, since instead of the rule itself a sentence about a rule functions in the argument. This objection is to some extent correct. My approach is based on two objectives. First I want to able to have arguments in which both arguments about a rule and with that rule are possible. (See Hage 1997, p. 129) And second I want to stick to the convention that premises in arguments are sentences with truth values. In my approach both objectives are reached, at the cost of representing rule application in an unnatural way.[11]

Deontic facts are treated as reason-based facts, by demanding that a world in a model can only contain a reason-based fact if it also contains certain other facts, in particular the existence of the rule on which the reason-based fact is based, facts that satisfy this rule's conditions, and the absence of an exception to the rule. For instance, the sentence that John ought to be punished can only be true in a world that also makes true the sen-

[10] This presumes a distinction between deductively valid arguments and arguments that guarantee the truth of their conclusions on the basis of the truth of their premises. I will briefly return to this issue in section 11.

[11] Another compromise would be to allow norm formulations as premises in arguments under the condition that the sentences to the effect that these rules exist are true.

tences that John is a thief, that the rule exists that thieves ought to be punished (or some other rule with the same conclusion and facts that satisfy the rule conditions), and that there is no exception to this rule.

Exceptions to rules are also treated as reason-based facts, and sentences that express that there is an exception to a rule have similar truth conditions as deontic sentences.

In the following sections I describe Natuaralistic Deontic logic (NDL), a deontic (ought-to-be) logic, and I characterise this logic by model-theoretic means. There is more to deontic logic than the logic that deals with the derivation of deontic sentences from sentences about norms and about the facts that satisfy the conditions of these norms. In particular there is an important part of deontic logic that deals with the logical relations between deontic sentences. Therefore I include in NDL a simple logic for deontic sentences of the ought-to-be type, that is rather similar to the standard system of deontic logic. I call this logic Deontic Predicate Logic (DPL).

7 The language L_{NDL}

The language of NDL, L_{NDL}, is essentially the language of first order predicate logic, augmented with some conventions. I assume that all formulas of L_{NDL} have an uppercase initial, while all terms (including function expressions) have a lowercase initial. Variables of L_{NDL} are *italicised*. For the meta-language of L_{NDL} I use schematic sentences and terms, unless the contrary is clear from the context.

7.1 States of affairs

I adopt the convention that all formulas of L_{NDL} *express* states of affairs, and that true formulas express facts. If a formula contains free variables, it expresses a generic state of affairs.

For every state of affairs expressible in L_{NDL} there is a term that *denotes* it. In the case of atomic formulas this term is constructed by replacing the uppercase initial of the formula expressing the state of affairs by the corresponding lowercase initial and prefixing the resulting string by an asterisk (*). E.g. if `Thief(john)` expresses that John is a thief, `*thief(john)` denotes the state of affairs that John is a thief.[12] All terms that start with an asterisk are assumed to denote states of affairs. In this way the terms of L_{NDL} are subdivided into terms that denote states of affairs and terms that denote other individuals.

In the case of logically compound sentences, the replacement is executed for all atoms that are part of it. The asterisk is only added once. The term `*it's_raining ∨ shines(sun)` denotes the disjunctive state of affairs that either it's raining or the sun shines. This means that the symbols for the logical operators are also used for functors that map states of affairs on more complex states of affairs, such that the relation between logically compound sentences and the corresponding compound states of affairs is maintained.

7.2 Deontic predicates

L_{NDL} has a special one-place predicate `Ob`, that stands for 'it ought to be the case that ...'. This predicate ranges over states of affairs. So, if `Formula ∈` L_{NDL}, then `Ob(*formula) ∈` L_{NDL}. Notice that `Ob` is a predicate, not an operator, and that `*formula` denotes the state of affairs expressed by `Formula`..

---

12 Notice that this term does *not* denote the sentence 'Thief(john)'. The state of affairs expressed by a formula is not the reification of that formula. This means that formulas referring to states of affairs need not belong to a meta-language.

L_{NDL} has also the one-place predicates `Pb` and `Fb`, that range over states of affairs, and stand for respectively permitted-to-be and forbidden-to-be. They are, as usual, defined as follows:

`Pb(*state_of_affairs)` ≡def. `~Ob(*~state_of_affairs)`
`Fb(*state_of_affairs)` ≡def. `Ob(*~state_of_affairs)`

Because the deontic predicates can operate on any state of affairs, and a sentence that has a deontic predicate expresses a state of affairs just like all other sentences, a kind of iterated deontic modality is possible. The sentence

`Ob(*pb(holds_political_speech(francina)))`

means that it ought to be the case that it is permitted that Francina holds a political speech. Notice that the deontic 'predicate' within the scope of (another) deontic predicate has become a functor that maps a state of affairs on another state of affairs.

A formula of L_{NDL} is called a deontic formula if and only if either
a. it has either the form `Ob/Fb(*formula)` or `~Pb(*formula)`;
b. it is not a literal and one of its components is a deontic formula of L_{NDL}.

The states of affairs expressed by deontic formulas are called *deontic states of affairs*.

7.3 Rules

L_{NDL} has a functor ⇒, that has two, often generic, states of affairs as its parameters. Its value is a rule. For instance

`*thief(x) ⇒ *punishable(x)`

denotes the rule that thieves are punishable, while

`*true ⇒ *ob(~∃x(steals(x))`

may be interpreted as the expression of the norm that forbids stealing.[13] Rules with a deontic state of affairs as their conclusion parameter are called deontic rules, or *norms*. E.g.

`*thief(x) ⇒ *ob(punished(x))`

denotes the norm that thieves ought to be punished. This term is not a well-formed formula of L_{NDL}, let alone a deontic formula. In general, norms in the sense defined here are not deontic formulas, and neither do they denote deontic states of affairs.

To express that rules exist, L_{NDL} has the one-place predicate `Exists`. This predicate is defined as follows:

`Exists(entity) ≡def. ∃x(x = entity)`

An example of the use of this predicate is:

`Exists(*thief(x) ⇒ *ob(punished(x)))`

Notice that this sentence is an atom, and non-deontic, because the predicate `Exists` is not deontic. Notice also that sentences of the forms

`Ob(*r → q)`
`R → Ob(*q)`
`∀x(Ob(*r(x) → q(x)))`, and
`∀x(R(x) → Ob(*q(x)))`

are well-formed, but do *not* express rules.

---

13 Rules that have no condition part are constructed as rules that have the tautological state of affairs `*true` as their condition.

If the conditions of a rule are satisfied and there is no exception to the rule, the instantiated rule applies. This is expressed as follows:

```
Applies(*thief(john) ⇒ *ob(punished(john)))
```

Finally, the language L_{NDL} has a one-place predicate `Exception` that ranges over instantiated rules and expresses that there is an exception to the rule in question for the case to which the rule is instantiated. For instance:

```
Exception(*thief(john) ⇒ *ob(punished(john)))
```

8 The use of model-theory

The deontic logic NDL that I want to define by model-theoretic means, is an extension of predicate logic. The extension indicates how rules can be used in derivations and makes sure that deontic facts and exceptions to rules are reason-based. It also indicates how deontic states of affairs are mutually related.

In the characterisation of NDL, I will use model-theory in a somewhat exceptional way. Usually model-theoretic semantics is used to specify the meanings of logical operators and quantifiers by means of the truth values of the sentences in which they occur with respect to possible worlds in a model. The emphasis is on meaning, on truth values, and on linguistic entities such as operators and sentences.

In my approach the emphasis shifts to states of affairs and the worlds in which they obtain. There is a close connection between the truth of a sentence with respect to a world and whether a state of affairs obtains in that world. I exploit that connection by replacing the talk about the relations between the truth values of sentences with respect to a world by talk about the relations between states of affairs that obtain in a world.

Another difference in emphasis is that I use the definitions which states of affairs obtain given other states of affairs as constraints on logically possible worlds. For instance, a world only counts as logically possible if it satisfies the constraint that if a state of affairs of the form `*a & b` obtains, the states of affairs of the forms `*a` and `*b` must also obtain, and vice versa.

So, I am not taking the notion of a logically possible world as given and demand that if a state of affairs of the form `*a & b` obtains, the states of affairs of the forms `*a` and `*b` also obtain, and vice versa. Instead I characterise logics by means of the constraints that must be satisfied by worlds that embody this logic. A set of (not necessarily logically possible) worlds is a model of a logic, only if the worlds in that set satisfy the constraints that characterise this logic. In that case they are logically possible worlds. Subsets of these logically possible worlds can then in turn be used to characterise domain theories. The constraints that are used to characterise logics correspond to the properties of the valuation function that is used in traditional model-theoretic semantics.

I start out with the constraints for predicate logic. To these the constraints for DPL are added. The resulting set of possible worlds is then taken as the starting point for the identification of the worlds that are representative for NDL. Because a logical characterisation of reason-based facts has a lot in common with fixed point definitions used for nonmonotonic logics, I use a construction that is inspired by the way Reiter (1980) characterised default logic.

Because I do not focus on linguistic entities, and therefore also not on meaning, I prefer not to call my use of model-theory semantics. Nevertheless there is a close connection between my use of model theory and the way model-theory is put to semantic use.

9 Model-theory for NDL

The model-theoretic characterisation of NDL that I want to propose is the following:

L_{NDL} is the language of NDL. $L_{NDL} = \{S_1, S_2, ..., S_n\}$, where $S_1 ... S_n$ are all the well-formed sentences of L_{NDL}.

Let $*sa_i$ denote the state of affairs that is expressed by S_i. $*sa_i$ is then a state of affairs that is *possible relative to* L_{NDL}.[14]

Let the set SA be the set of all states of affairs that are possible relative to PL. And let W be the power set of SA. Intuitively W stands for the set of all worlds, the content of which is expressible in L_{NDL}. Every $w \in W$ is a subset of SA.

There are no other constraints on the states of affairs that are elements of the worlds in W. There are, for instance, worlds in W in which the states of affairs `*p` en `*~p` both obtain. Such worlds are possible relative to L_{NDL}, but they are not logically possible according to NDL. Worlds that are logically possible are subject to a number of additional constraints. The constraints that characterise worlds in W_{NDL} are given in three stages. First I give constraints that define the set of worlds W_{PL} that are possible with regard to predicate logic.

CONSTRAINTS ON WORLDS THAT ARE LOGICALLY POSSIBLE ACCORDING TO PREDICATE LOGIC

1. $*p \in w$ if and only $*{\sim}p \notin w$.
2. $*p \;\&\; q \in w$ if and only if both $*p \in w$ and $*q \in w$.
3. $*p \vee q \in w$ if and only if either $*p \in w$, or $*q \in w$, or both.
4. $*p \rightarrow q \in w$ if and only if either $*p \notin w$, or $*q \in w$, or both.
5. $*p \equiv q \in w$ if and only if either both $*p \in w$ and $*q \in w$, or both $*p \notin w$ and $*q \notin w$.

These constraints are the traditional constraints of propositional logic stated in terms of relations between states of affairs.

6. $*\exists x(r(x)) \in w$ if and only if there is an individual a, such that $*r(a) \in w$.
7. $*\forall x(r(x)) \in w$ if and only if there is no individual a, such that $*r(a) \notin w$.

These constraints give the traditional meaning of the quantifiers, again stated in terms of states of affairs.[15]

Finally there is a constraint to guarantee that terms that denote states of affairs expressed by logically equivalent sentences are co-referential:

8. If for all worlds $w \in W_{PL}$ it holds that $*p \equiv q \in w$, then $*p = *q$.

The worlds that satisfy these constraints are elements of W_{PL}. Since the constraints hold for worlds in W, $W_{PL} \subseteq W$.

---

[14] L_{NDL} may be thought of as the conceptual scheme by means of which the world is 'captured'. The worlds that are possible relative to L_{NDL} are then the worlds that can be captured by means of this conceptual scheme.

[15] To gain simplicity at the cost of precision, the formulations of the constraints 6 and 7 do not deal with compound formulas or quantifiers within the scope of the quantifiers.

CONSTRAINTS ON WORLDS THAT ARE LOGICALLY POSSIBLE ACCORDING TO DEONTIC PREDICATE LOGIC

Worlds that embody Deontic Predicate Logic, a simple logic for ought-to-be, must satisfy some additional constraints:

9. If `*ob(s)`$\in$ w then `*pb(s)` $\in$ w.
10. If `*pb(s)` $\in$ w then `*fb(s)` $\notin$ w.
11. If and only if `*ob(~s)` $\in$ w then `*fb(s)` $\in$ w.
12. If for all worlds w $\in$ W_{PL} it holds that `*p` $\rightarrow$ `q` $\in$ w,
    then it holds for all worlds w $\in$ W_{DPL} that `*ob(p)` $\rightarrow$ `ob(q)` $\in$ w.
13. If `*ob(p)` $\in$ w and `*ob(*q)` $\in$ w, then `*ob(p & q)` $\in$ w.

I will call the set of worlds that satisfy these constraints W_{DPL}. $W_{DPL} \subseteq W_{PL}$.

CONSTRAINTS ON WORLDS THAT ARE LOGICALLY POSSIBLE ACCORDING TO NATURALISTIC DEONTIC LOGIC

The set W_{DPL} is my starting point for the characterisation of the set of possible worlds W_{NDL}, which represents NDL. To specify this last set, I need some auxiliary constructs.

The first construct is a function that operates on possible worlds and maps them on the worlds that are the possible world-counterpart of the first world's deductive closure. I will call the function CL_{DPL} (closure under DPL), and it is defined as follows:
Let w $\in$ W_{DPL}, and s stand for a state of affairs. Then it holds that
CL_{DPL} (w) = {s | s $\in$ $\cap w_i$ for all $w_i \in W_{DPL}$, such that $w \subseteq w_i$ }
Informally: The closure under DPL of a possible world w is the intersection of all DPL-possible worlds of which w is a subset.

The second construct is a three-place relation R over worlds in W_{DPL}. This relation is defined as follows[16]:
$R(w_1, w_2, w_3)$ if and only if $w_2 = CL_{DPL}$ ($w_1 \cup$ additionset), where
additionset =

{ `*applies(i_conditions ⇒ i_conclusion) & i_conclusion` |
  `*exists(conditions ⇒ conclusion)` $\in$ w_1,
  `*i_conditions` $\in$ w_1, and
  `*exception(i_conditions ⇒ i_conclusion)` $\notin$ w_3
}

Intuitively this boils down to the following: For every rule that exists in world w_1, the conditions of which are satisfied in w_1, and to which there is no exception in w_3, the states of affairs that this rule applies and that its conclusion holds are added to the states of affairs of w_1. The result is w_2. w_2 is the closure under DPL of the result of the application of all applicable rules in w_1.

The third auxiliary construct is that of a finite sequences S_i of worlds. Such a sequence consists of one or more worlds in W_{DPL}, denoted as $S_i[1]$ ... $S_i[n]$, with n ≥ 1. The worlds in such a sequence satisfy the constraint that $R(S_i[n], S_i[n+1], w)$, where w is some element of W_{DPL}. This world w is called the *reference world* of the sequence. The world $S_i[1]$ is the *starting world* of sequence S_i.

[16] `*i_conditions` and `*i_conclusion` are the instantiations of respectively `*conditions` and `*conclusion` by means of some instantiation σ.

Given the definition of the relation R, it holds for every i and every n such that Si[n+1] exists, that $S_i[n+1] \supseteq S_i[n]$. Intuitively a sequence represents a series of possible worlds in which each world is the result of applying all applicable rules in its predecessor. The interesting case for my purposes is when a sequence ends with a subsequence of identical worlds. When that happens, the final worlds in the sequence are worlds in which all applicable rules have been applied, and the results do not give rise no new possibilities for rule application. Let us call such sequences *halting sequences*. The final world in a halting sequence is called the *fixed point* of that sequence. For a fixed point w_f it holds that $R(w_f, w_f, w)$.

A subset of the halting sequences consists of those sequences in which the fixed point is identical to their reference world. I will call these sequences *grounded sequences*. Those fixed points are then worlds in which no new rules can be applied, given the exceptions to rules that obtain in these worlds themselves. For these fixed points w_f it obviously holds that $R(w_f, w_f, w_f)$.

Such worlds w_f, ones in which the results of the application of all applicable rules are contained, are candidates for being elements of W_{NDL}. Yet, there is another requirement that they must meet. Worlds in W_{NDL} should not have 'free-floating' reason-based facts. Intuitively they should not contain reason-based facts that do not ultimately result from the application of existing rules to brute facts. This intuition is translated into the requirement that for every world in W_{NDL} there must be a grounded sequence, that starts with a world that does not contain any reason-based facts. The existence of the grounded sequence then shows that repeated application of the rules in the 'starting world' of the sequence results in the world that is to be in W_{NDL}. In other words, the fixed point can be the result of repeatedly applying the rules in a world that only contains brute facts.

To formalise this intuition we need a formal characterisation of reason-based facts. For the purpose of this paper I assume that there are only three kinds of reason-based facts, that is facts of the forms

`*ob(s)`[17],
`*applies(i_conditions, i_conclusion)`,
`*exception(i_conditions, i_conclusion)`.

Given this assumption we can formalise the following constraints on worlds w in W_{NDL}:

14. There must be a grounded sequence w_0 ... w, all elements of which are in W_{DPL}, such that
    a. there is no state of affairs s_1 of the form `*ob(s)` such that $s_1 \in w_0$, and
    b. there is no state of affairs s_2 of the form `*applies(i_conditions, i_conclusion)` such that $s_2 \in w_0$, and
    b. there is no state of affairs s_3 of the form `*exception(i_conditions, i_conclusion)`, such that $s_3 \in w_0$.

These constraints formalise the idea that an NDL-possible world should not contain free-floating reason-based facts.

15. If `*exception(i_conditions ⇒ i_conclusion)` $\in$ w, then
    `*exists(conditions ⇒ conclusion)` $\in$ w.

This constraints boils down to it that there can only be an exception to a rule if that rule exists.

[17] This includes facts of the form `*fb(s)` and `*~pb(s)`.

VALID INFERENCE IN NDL

To define valid inference for NDL I need an additional construct, that of the *minimal model* of a theory. Let TH be a theory, and let S be a sentence schema and let *s denote the state of affairs that is expressed by S. Then W(TH) = {*s | S ∈ TH}. The minimal *DPL-model* of a theory TH is then CL_{DPL}(W(TH)). It contains precisely those states of affairs that are expressed by the deductive closure under DPL of the theory TH.

A minimal *NDL-model* of a theory is a fixed point of a grounded sequence in which the minimal DPL-model of this theory occurs. Intuitively a minimal NDL-model is a possible world that satisfies the constraints of NDL, and that contains only those states of affairs that are required by the theory of which it is a model. A theory is NDL-consistent if and only if it has a minimal NDL-model. In the next section we will see that a theory can have more than one minimal NDL-model.

There are two ways in which valid NDL-inference can be defined. Either a sentence is a valid NDL-conclusion of a theory if and only if the state of affairs that it expresses occurs in *all* minimal NDL-models of this theory. This is a sceptical definition of valid NDL inference.

Or a sentence is a valid NDL-conclusion of a theory if and only if the state of affairs that it expresses occurs in *at least one* minimal NDL-model of this theory. This is a credulous notion of valid inference. For the present purposes it is not necessary to choose between these two variants.

10 Examples

I will illustrate NDL by means of four variations on a basic example. In all cases I assume that the world with respect to which the sentences hold true is an element of W_{NDL}, and that the sentences contain all the relevant information. There are, for instance, no unmentioned relevant exceptions.

VARIATION 1:

```
1.  Thief(john)
2.  Exists(*thief(x) ⇒ *ob(punished(x))
```

According to constraint 14, there is no exception to the rule that thieves are punishable, which makes that the following sentences are true:

```
Applies(*thief(john) ⇒ *ob(punished(john)))
Ob(*punished(john))
```

VARIATION 2:

```
1.  Thief(john)
2.  Minor(john)
3.  Exists(*thief(x) ⇒ *ob(punished(x))
4.  Exists(*minor(x) ⇒ *ob(~punished(x)))
5.  Exists(*minor(x) ⇒ *exception(thief(x) ⇒ ob(punished(x))))
```

If the premises 2, 4 and 5 are added to the first variation of the example, things change considerably. Now there is an exception to the rule that thieves ought to be punished. This rule does not apply anymore, and its conclusion is by default false (constraint 14). There is no exception to the rule that minors ought not to be punished (again constraint 14), this rule applies, and leads to the conclusion that John ought not to be punished.

VARIATION 3:

```
1.  Thief(john)
2.  Minor(john)
3.  Exists(*thief(x) ⇒ *ob(punished(x)))
4.  Exists(*minor(x) ⇒ *ob(~punished(x)))
```

If the fifth premise of variation 2 is left away, things become complicated. There is no exception to either one of the rules, so they should both apply and their conclusions should be true. However, the states of affairs `*ob(punished(john))` and `*ob(~punished(john))` cannot both obtain in a NDL-possible world. Variation 3 illustrates that a theory that is consistent in predicate logic can be inconsistent in NDL.

VARIATION 4:

```
1.  Thief(john)
2.  Minor(john)
3.  Exists(*thief(x) ⇒ *ob(punished(x))
4.  Exists(*minor(x) ⇒ *ob(~punished(x)))
5.  Exists(*thief(x)⇒ *exception(minor(x) ⇒ ob(~punished(x))))
6.  Exists(*minor(x) ⇒ *exception(thief(x) ⇒ ob(punished(x))))
```

This variation of the example illustrates that an NDL-theory can have multiple extensions. The following two sets of states of affairs can both be subsets of a model of this theory:

Set 1 = { `*ob(punished(john)),`
`*exception(minor(john) ⇒ ob(~punished(john)))` }
Set 2 = { `*ob(~punished(john)),`
`*exception(thief(john) ⇒ ob(punished(john)))` }

This possibility illustrates that in NDL-possible worlds, the brute facts need not determine the strictly reason-based facts. Whether this is an attractive characteristic of NDL is an ontological issue that falls outside the scope of this paper.[18]

Depending on whether a credulous or a sceptical notion of valid inference is used, the presence of the states of affairs in the sets 1 and set 2 are respectively derivable or not from the premises 1 to 6.

11 Concluding observations

Most systems of deontic logic do not distinguish between norms and deontic propositions. Von Wright (1963, p. 130f.) made the distinction and concluded that there could only be a logic of deontic propositions, because norms have no truth-value. Alchourrón (1963, 1993) also made the distinction and developed different logical systems for norms and for deontic propositions.

I agree with Von Wright that there cannot be a logic of norms that deals with the derivation of norms from other norms (and descriptive sentences). But there must be a logic of norm-application, because there clearly are logical standards that regard the application of norms in deriving deontic propositions. Moreover, the deontic propositions derived in that way stand in logical relations to other deontic propositions. That is why I think that there should be one deontic logic, that both deals with norm application and with the logical relations between deontic propositions. That logic should strictly distin-

[18] I find it interesting to notice that the multiple extension problem concerns not only the question what it is rational to believe of infer, but also what facts there are in the world. The distinction between on the one hand epistemology and logic, and on the other hand ontology seems to blur in case of reason-based facts.

guish between norms, propositions about norms, and deontic propositions. NDL is such a logic.

NDL is based on a number of ontological assumptions concerning the supervenience of deontic facts on non-deontic facts. If some fact supervenes on one or more other facts, this supervening fact is new in the sense that it was not included in the facts underlying it. A particularly important kind of supervenience is the supervenience of reason-based facts. Reason-based facts are facts whose existence presupposes the existence of other facts, their reasons. The supervenience of reason-based facts requires the existence of rules, which is in my opinion also a matter of (social) fact.

Rules and the facts that satisfy their conditions together constitute reason-based facts. Reason-based facts are new in the sense that they do not coincide with the existence of the rules and the facts that match their conditions. They are not new in the sense that they might have been absent where the rules and the facts that match their conditions are present.

The constitution of reason-based facts by their underlying rules and reasons can be *represented* as an argument. In that case, the argument is not deductively valid, because its conclusion contains new information in comparison to its premises. However, if we abstract from the defeasibility of rule-application, such an argument may guarantee the truth of its conclusion on the basis of the truth of its premises and is in that sense valid.

Deductive reasoning can be seen as reorganising the information contained in the premises of the argument (see Toulmin 1958, p. 123f.). Reasoning with rules is not a form of reorganising the information contained in the rules and the facts that satisfy their conditions. In fact, rules do not contain any information, but through their application they generate new facts and new information. Traditionally the two criteria for logical validity, that the information of the conclusion of the argument was already present in the premises, and that the conclusion must be true if the premises are true, go hand in hand. In the case of rule-applying arguments, these two criteria diverge, because the conclusions of these arguments contain new information, and yet they must be true if the premises are true.

There are many kinds of reason-based facts, facts that obtain thanks to their underlying reasons, but one important category are the so-called deontic facts. Like other reason-based facts, deontic facts supervene on other facts, thanks to the existence of rules. An argument that models the constitution of deontic facts typically contains the derivation of a deontic proposition from purely non-deontic propositions. There is nothing wrong with such an argument; it is not even a naturalistic fallacy.

It is possible to specify the truth-values of deontic propositions in terms of purely factual ones, and therefore it is also possible to derive deontic conclusions from purely factual ones. However, it remains impossible to define deontic propositions in terms of purely factual ones, because specifying the truth-value of a sentence falls short of defining its meaning.

These ontological assumptions are incorporated in a deontic logic by treating rules as logical individuals. This treatment was inspired by the different ontological role that rules play in comparison to facts (section 3), and by a similar treatment in common parlance. If rules are treated as logical individuals, several advantages result. First it becomes possible to reason about rules, without having to take recourse to the meta-language level. Second, reasoning with rules cannot be deductive, and must be specified from scratch. Since reasoning with rules is defeasible, this is an advantage, because all characteristics of this defeasibility can be built into the logic. And finally, if rules are logical individuals, the statement that a rule exists can express a brute fact. For a naturalistic theory of deontic propositions, the existence of at least some rules must be a matter of brute fact.

The logic presented above is extensional, and in that respect it resembles the logic of Lokhorst 1994. However, where Lokhorst motivates his approach with the relative ease of implementation, my motivation is philosophical instead. My main purpose is to give a logical account of a moderate form of deontic naturalism. If the resulting logic is relatively easy to implement, this is from my point of view an additional advantage.

Acknowledgement

The author wishes to thank Jörg Hansen, Paul McNamara, Henry Prakken, Giovanni Sartor, Bart Verheij and the anonymous referees for this volume for useful comments on earlier versions of this paper.

REFERENCES

Alchourrón, C.E. (1969). Logic of Norms and Logic of Normative Propositions. *Logique et Analyse* 12, p. 242-268.

Alchourrón, C.E. (1993). Philosophical Foundations of Deontic Logic and the Logic of Defeasible Conditionals. J.-J. Ch. Meyer and R.J. Wieringa (eds.), *Deontic Logic in Computer Science*. John Wiley & Sons, Chicester, pp. 43-84.

Alchourrón, C.E. and E. Bulygin (1981). The Expressive Conception of Norms. R. Hilpinen (ed.), *New Studies in Deontic Logic*, D. Reidel Publishing Company, Dordrecht e.a., pp. 95-124.

Barwise, J. and J. Perry (1986), *Situations and Attitudes*, MIT Press, Cambridge e.a.

Brecht, A. (1959). *Political Theory: The Foundations of Twentieth-Century Political Thought*, Princeton University Press, Princeton.

Frankena, W.K. (1939). The naturalistic fallacy, *Mind* 48. Also in Ph. Foot, *Theories of Ethics*, Oxford University Press, Oxford 1967, p. 50-63.

Hage, J.C. (1987). *Feiten en betekenis; een verhandeling over ontologie en praktische rede* (Facts and meaning: a treatise on ontology and practical reason), PhD-thesis, Leiden.

Hage, J.C. (1997). *Reasoning with rules*, Kluwer, Dordrecht.

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.

Kelsen, H. (1960). *Reine Rechtslehre*, 2nd ed., Franz Deuticke, Wien.

Lagerspetz, E. (1995), *The Opposite Mirrors*, Kluwer, Dordrecht.

Lokhorst, G.-J. C. (1994) Extensional Deontic Logics of Action and their Implementation, in A.J.I. Jones and M. Sergot (eds.), *Deon '94*, *Workshop Proceedings*. Tano, p. 145-164.

MacCormick, D.N. (1974). Law as institutional fact, *The Law Quarterly Review* 90, p. 102-129. Also in D.N. MacCormick and O. Weinberger, *An Institutional Theory of Law*, p. 49-76.

MacCormick, D.N. and O. Weinberger (1986), *An Institutional Theory of Law*, Reidel, Dordrecht e.a.

Mackie, J.L. (1980). *The Cement of the Universe*. Clarendon Press, Oxford.

Moore, G.E. (1903). *Principia Ethica.* Cambridge.

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

Ruiter, D.W.P. (1993). *Institutional Legal Facts. Legal Powers and their Effects*. Kluwer, Dordrecht e.a.

Searle, J. (1969). *Speech acts. An essay in the philosophy of language*. Cambridge University Press.

Searle, J. (1979). A Taxonomy of Illocutionary Acts, *Expression and Meaning*, Cambridge University Press, Cambridge e.a., p. 1-29.

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

Taylor, P.W. (1961). *Normative discourse*. Greenwood Press, Westport.

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

Wright, G.H. von (1963). *Norm and Action*, Routledge and Kegan Paul, London.

Wright, G.H. von (1998). Deontic Logic – as I See It. *Proceedings of Deon '98*, Bologna.