Applied legal epistemology

Building a knowledge-based ontology of the legal domain

Laurens Mommers

Applied legal epistemology

Building a knowledge-based ontology of the legal domain

Promotiecommissie

Promotor: Prof. dr. H.J. van den Herik (Universiteit Leiden,
Universiteit Maastricht)

Referent: Prof. mr. P.W. Brouwer (Universiteit van Amsterdam)

Overige leden: Prof. mr. P.B. Cliteur (Universiteit Leiden, Technische
Universiteit Delft)
Prof. mr. H. Franken (Universiteit Leiden)
Prof. mr. M.A. Loth (Erasmus Universiteit Rotterdam)
Dr. mr. A.H.J. Schmidt (Universiteit Leiden)

Applied legal epistemology

Building a knowledge-based ontology of the legal domain

For more information on the research presented in this dissertation, please visit http://research.mommers.net, or send an email to laurens@mommers.net.

Contents

Foreword

For the past five years, I have been intrigued by the question: what constitutes legal knowledge? In finding my way in the domains of positive law and jurisprudence – a rather difficult task given my non-legal background – my journey was considerably shortened by the many helpful comments and useful advice offered by various people. Although I now dare to claim that I have *some* sense of what the law is and of what a practitioner of law does, I do not wish to pretend that I know what constitutes legal knowledge. In the thesis I have introduced many possible distinctions for legal knowledge, and consequently I should know about knowing the law. Nevertheless, I would like to emphasise the fact that my answer to the question 'what constitutes legal knowledge?' is: I really, *really* do not know. Thus, a person who tries to find a single definition of legal knowledge in this thesis, will do so in vain.

   Writing a thesis is, by itself, a solitary activity – it is so, at least, in such disciplines as philosophy and law. However, I wrote mine in an environment that offered the necessary relief and the opportunity to put things into perspective: the department of law and computer science, part of the Faculty of Law in Leiden. After five years, I cannot imagine that I would have been able to finish a Ph.D. thesis in a place not offering the unique mixture of melancholy, interdisciplinary research and Friday-afternoon drinks. Within the practical and formal limits imposed on expressing one's gratitude in the foreword of a thesis, I would like to thank the following people: Jaap Hage, who made an effort in convincing me to study legal philosophy instead of only epistemology, Franke van der Klaauw, who provided urgent mental care in many cases, and introduced me to experts who proved to be important to my research project, and, of course, my parents, whose continuing support was essential during the research project.

Laurens Mommers

Leiden, April 23, 2002

1 Introduction

The current thesis is the result of a research project into legal knowledge representation. The research project started out as an investigation into the possibility of translating conceptual specifications of the legal domain into formal ones, using the language of situation semantics developed by Barwise and Perry (1983). It stands in the tradition of the development of general conceptual and formal specifications of the legal domain in the Leiden department of law and computer science (currently the Center for eLaw at Leiden). Starting with the challenging work by Van den Herik (1991), who, among other things, pointed out what factors complicate knowledge representation in the legal domain, the tradition was set with the work by Van Kralingen (1995) and Visser (1995), who built a conceptual model of the law and a formal model of the law respectively. Oskamp (1998) took a different, practically oriented viewpoint, modelling the elements relevant to determining a sentence in several domains within penal law.

1.1 The research lines

The research described in this book can be regarded as a follow-up to Van Kralingen's, Visser's and Oskamp's work, although it does not describe the relation between conceptual and formal models of law, as originally intended. Instead, it focuses on the relation between conceptual models of the law and a branch of research previously practised at the Leiden department by Hage (1987). He developed a view on the legal domain that can be characterised retrospectively as an ontology of law, based on legal-theoretical and general philosophical insights. Hage's thesis is mainly about philosophy, but it clearly provides handles for modelling activities in the legal domain.

   Thus, the current work arises from two different strands of research: first, the research into ontologies of law by the AI-and-law community, and second, the research into the structure of the legal domain by the legal-philosophical community. At the intersection of these two strands, it focuses on the concept of knowledge in the AI-and-law community. This concept is used in such phrases as 'legal knowledge representation', and its meaning is often taken for granted. In my view, in this context, the concept of knowledge deserves more attention for two reasons. First, the phrase 'legal knowledge representation' often refers to the representation of elements of the legal domain, and relations between these elements, rather than to knowledge about these entities and relations. Second,

scrutinising the concept of knowledge can yield valuable insights, that may help to build useful legal knowledge-based systems.

1.2 Three perspectives

For the reasons mentioned above, I elaborate on the role that the concept of knowledge plays in the legal domain. I do so from three perspectives: from general epistemology, from legal theory, and from AI and law. From the first perspective, I address the questions how knowledge is acquired, what it is about, and how it is justified. From the second perspective, I discuss the specific characteristics of legal knowledge. From the third perspective, I discuss the role the concept of knowledge plays in legal knowledge representation. By their nature, these subjects are strongly connected with the question what knowledge is about. With regard to the legal domain, this is the subject matter of the philosophical discipline called legal ontology.

Legal ontology scrutinises the existence of legal entities, such as rules, norms, and legal institutions, and the dependencies between these entities. Together, legal epistemology and legal ontology can provide an integrated view on the legal domain, thus facilitating the representation of knowledge. However, among those who practise legal epistemology and legal ontology, there is little consensus on what justifiable claims can be made within these disciplines. Moreover, prior to making such claims, one has to develop a general view on the law, and such attempts have been manifold (natural law theory, legal positivism *etc.*). A general view on the law inevitably comprises a view on the way in which the law can be known, and on the entities it consists of. Differences in the general views cause the lack of consensus on epistemological and ontological claims.

1.3 A model of law

In the thesis, I build a model of law that leaves open the possibility of expressing different views on legal epistemology and legal ontology, thus avoiding to take a stance in the legal-philosophical debate prior to building the model. In this sense, the model developed may be called a 'meta-ontology' of law − it allows for different views on what knowledge in the legal domain actually amounts to. To attain this, different ontological status layers and different epistemic roles are distinguished. The ontological status layers allow for different views on the existence of the law, and the epistemic roles allow for expressing different views on what knowledge amounts to in the legal domain.

Knowledge, I claim, may be regarded as the mark of a quality stamp. It is a mark of approval; it says that a belief or a skill conforms to a set of criteria, and that it deserves to be called 'knowledge' for that reason. The applicable set of criteria depends on the type of entity that we wish to qualify as knowledge, and the context in which we encounter that entity. For instance, if we wish to qualify a belief about the whereabouts of a suspect as knowledge, we may demand that this belief is true. However, if we wish to qualify a belief about the value of a piece of circumstantial evidence as knowledge, we demand that this belief is justified rather than true. Knowledge is a value predicate, a way to express the worthiness of an entity. Representing knowledge thus requires to make explicit the criteria by which the represented entities deserve their qualification as knowledge.

These criteria may apply to the acquisition, object and justification of knowledge. Thus, they do not only concern the content (object) of knowledge, but also the sources of knowledge (acquisition), and the reasons there are to believe its content (justification). Together, the criteria provide a framework for assessing whether to assign the quality mark. What is more, they provide valuable additional information on represented knowledge. For that reason, the concept of knowledge is useful, even if its meaning does not conform to the traditional view of having one set of criteria that determines all possible instances of the concept.

1.4 Problem definition and research questions

Having more insight into the nature of the concept of legal knowledge improves the quality of knowledge representation in the legal domain and provides more depth to the use of such phrases as 'legal knowledge representation'. Ideally, this phrase induces the distinction between the legal domain itself, knowledge about the domain, and the model in which the domain and the knowledge about the domain are represented (see figure 1.1).

```
(1) the model represents:     →      (2) knowledge about the
 – knowledge about the                    legal domain
    domain                                    ↑↓
 – the domain itself          →        (3) the legal domain
```

Figure 1.1. Distinction between model, knowledge and domain

In most existing models of law, there is no clear distinction between knowledge about the legal domain on the one hand, and the legal domain itself on the other hand, or there is focus on only one of the two elements. In this thesis, I attempt to explain the role knowledge plays with respect to the legal domain. The problem definition for this thesis is as follows:

> What role can a concept of legal knowledge, formulated from the perspective of epistemology, play in the representation of legal knowledge?

Knowledge about the legal domain may play two roles: as the object of a model of the legal domain (represented by the arrow between boxes 1 and 2 in figure 1.1), and as a potential part of the legal domain (because of the mutual dependence between knowledge about the legal domain and the legal domain itself, represented by the arrow between boxes 2 and 3 in figure 1.1).

The goal of the research is to develop an ontology of law that takes into account the concept of knowledge formulated as an answer to the problem definition. An ontology specifies what elements and relations we can find in the legal domain. It may form the basis for the representation of legal knowledge in computer systems. The desired effect of building an ontology is to reduce the representation effort that occurs each time when new knowledge is added to a system. Thus, an ontology is a framework in which a specific model (box 1 in figure 1.1) can be constructed. As a consequence of incorporating the concept of knowledge, the resulting ontology caters for the need to express relevant characteristics of knowledge about the legal domain.

The research questions that follow from the problem definition focus on four themes: general epistemology, legal epistemology, legal ontology, and knowledge representation. Four corresponding research questions are investigated:

> (1) What are the characteristics of knowledge about the legal domain, given the viewpoint of general epistemology?
> (2) What are the characteristics of knowledge about the legal domain, given the viewpoint of legal epistemology?
> (3) What characteristics of knowledge about the legal domain are useful as constituents for a model of the legal domain?
> (4) How can this ontology of law be used to represent knowledge about the domain of Dutch penal law?

The first question serves to reveal the characteristics that constitute knowledge about the legal domain. These characteristics are classified according to the distinction between acquisition, object, and justification of knowledge. The second

question serves to identify the limits of legal knowledge specification from the perspective of legal philosophy: what epistemic claims can be derived from different views on the law? The third question initiates a search for the elements that can be distinguished in the legal domain on the basis of relevant distinctions and characteristics applying to knowledge about the legal domain. A knowledge characteristic may shift to a constituent of the model because knowledge about the legal domain can reveal structural features of the domain itself. The fourth question is answered by the specification of a knowledge-based ontology of law. In this ontology, different existence claims and knowledge claims can be expressed, thus enabling to incorporate legal epistemology and legal ontology in the field of legal knowledge representation.

1.5 The structure of the thesis

The structure of the thesis follows the research questions outlined above. In chapter 2, to answer the first research question, I discuss the question what knowledge is from an epistemic point of view. I deal with some basic distinctions in the theory of knowledge. First, I shed light on three dimensions of knowledge: its acquisition, its object and its justification. Second, I discuss the distinction between doxastic and non-doxastic theories of knowledge, and between internalist and externalist theories of knowledge. Third, several types of beliefs and knowledge about the legal domain are discussed, based on their origination sources. Fourth, the phenomenon of an epistemic niche is used to explain the situation in which legal professionals acquire knowledge. By discussing these topics, I provide an answer to the first research question about the characteristics of legal knowledge from a general epistemic viewpoint.

In chapter 3, in accordance with the second research question, I attend the same subject. However, the viewpoint will be a legal-theoretical one. I start the chapter with a discussion of the intertwinement of the knowing subject (the person who has knowledge) and the known object (the thing he knows something about) in hermeneutical theories of law. Other views on knowing the law, derived from stances in legal philosophy, are discussed as well, among which natural law, legal positivism, and institutional theories of law. In addition, formal and material sources of law, and knowledge sources for the law are discussed. By discussing these topics, I provide an answer to the second research question about the characteristics of legal knowledge from a legal-epistemic viewpoint.

In chapter 4, I deal with the issue of legal ontology. As an addition to the discussion of epistemic claims from the point of view of different legal-theoretical

stances in chapter 3, I give an overview of ontological claims from these stances. Furthermore, I discuss in depth two objects of legal knowledge: systematisations and interpretations, as well as the elements that are part of the reasoning process that leads to establishing these objects. A detailed discussion of different modes of existence follows. Finally, two basic relations in the legal domain (namely counting as and causation) are dealt with. This chapter helps to provide an answer to the third research question, because a discussion of ontological approaches to the law enables me to explain how epistemology interferes with ontology in the current domain.

In chapter 5, the representation of legal knowledge is discussed. This chapter consists of three parts. The first part deals with the concept of representation. The second part is about meaning and reference, discussing the way in which classic views on meaning can be applied to legal concepts. The third part discusses three ontologies of law by listing prior work in the conceptualisation of the legal domain. Meaning can be an object of representation, as the meaning of an entity (for instance the meaning of a word) clarifies what role that entity plays. Making explicit meanings and representing them may thus enable a richer representation in information systems.

In chapter 6, a knowledge-based model of the law is presented, based on chapters 2 through 5. An answer to the fourth research question is given by describing a model that consists of entities, ontological status layers, epistemic roles, relations, acts and facts. The categories distinguished are clarified and, where necessary, further divided into subtypes. The model accommodates different concepts of knowledge, and different views on the ontology of law.

In chapter 7, I discuss central notions within Dutch penal law in order to have sufficient background information for the application of the knowledge-based model in chapter 8. For this purpose, some basic principles of criminal proceedings and penal law are discussed, as well as central notions of penal law, such as culpability and causality.

In chapter 8, I combine the findings of the chapters 6 and 7. In this chapter, I represent characteristics of Dutch penal law in terms of the knowledge-based model of the law. A representation language is chosen for this purpose, and it is explained how the elements of the knowledge-based model of law are represented in this language. Furthermore, inference rules regarding knowledge qualification and reasoning are proposed and discussed.

One final remark: wherever I use the pronoun 'he' ('his'), I use the pronoun as an abbreviation of 'he or she' ('his or her').

2 Legal knowledge from a general epistemic viewpoint

For centuries, knowledge has been subject of a lively debate among philosophers. What is knowledge?, How can we acquire knowledge?, and How can our knowledge grow? are three questions that are discussed in the philosophical discipline called epistemology. We can approach these questions in a normative manner. In this approach, we ask: what counts as knowledge?, what does it take to turn a mere belief into knowledge? The answer to these questions is given by an individual or a small group, arguing for certain choices on the criteria they deem necessary to qualify something as knowledge. The focus in this chapter is on the first research question, searching for the characteristics of knowledge about the legal domain from the viewpoint of general epistemology.

   In section 2.1, I offer an analysis of the different dimensions in the question what knowledge is. I distinguish three dimensions: acquisition, object and justification. Then, in section 2.2, the meaning of these dimensions is analysed with respect to legal knowledge. Subsequently, in section 2.3, I discuss two relevant distinctions in the theory of knowledge: doxastic and non-doxastic theories, and internalist and externalist theories. In section 2.4, different knowledge criteria are explained that can be attached to the dimensions that are discussed in sections 2.1 and 2.2. We may impose these criteria on mere belief in order to explain the difference between belief and knowledge. In section 2.5, I outline a typology of belief and knowledge. Knowledge qualification depends on the type of belief under scrutiny. In order to identify a particular belief or knowledge type, we have to know in what ways we can classify beliefs and knowledge. The classifications are listed in this section. Finally, in section 2.6, epistemic niches are discussed. These are (partially) controlled environments in which knowledge is acquired and processed. In this chapter, the above characteristics are used to explain in what manners we can scrutinise the concept of knowledge about the legal domain from a general epistemic viewpoint. In the next chapter, after I have discussed the characteristics of knowledge from a legal epistemic viewpoint, the resulting two sets of characteristics are employed to develop a framework for the characterisation of knowledge about the legal domain.

2.1 Dimensions of knowledge

The transition from belief to knowledge plays a major role in explaining the three dimensions of knowledge. By using criteria based on these dimensions, it is assessed whether a belief qualifies as knowledge. Only when it complies with

all applicable criteria, a belief may be called knowledge. Therefore, the three dimensions of knowledge are discussed relative to *beliefs*. The first dimension is acquisition: how is the belief acquired? A person can acquire a belief by different routes, some of which deserve more trust than others (subsection 2.1.1). The second dimension is its object: what is the belief's object? A person's belief has – supposedly – some kind of object; the belief is *about* something (subsection 2.1.2). The third dimension is justification: how is the belief justified? A person can be justified in believing something. He can, for instance, have good reasons for his belief. Both acquisition and object belong to the so-called *context of discovery*. The justification of a belief constitutes the *context of justification* (subsection 2.1.3). The three dimensions are discussed below.

2.1.1 Acquisition

A belief can be acquired from different sources. On a sunny Sunday morning, John acquires the belief that the sun is shining by looking out of the window. He learns from the Saturday newspaper that it will start raining before 2 pm. His daughter tells him that she watched the morning weather forecast on the weather channel, and that she learned it will not start to rain until the evening. Perception (looking out of the window) and testimony (reading the newspaper and listening to your daughter) are called belief sources (Audi 1998). A belief source is the process or phenomenon that a belief is based upon. A knowledge source is similar to that, except that in this case a piece of knowledge arises from the process or phenomenon. Audi (1998) distinguishes five sources of belief and knowledge: perception, memory, consciousness, reason, and testimony. I briefly discuss them below.

   This overview gives an idea of how we acquire beliefs. Extensive literature is available on most of these belief sources. An introduction is given in chapters 1 through 5 in Audi (1998). First, through perception, beliefs are acquired about what we see, smell, taste, hear, and feel. In most (but not all) cases, our senses provide an accurate picture of the world we live in, and on the basis of that accurate information we form correct beliefs. Second, we may preserve this information for later use in our memory. When we recall beliefs stored earlier, or base new beliefs on whatever we stored earlier, we use memory as a source of belief. Third, sometimes, a person is conscious of what he is doing and of what he is thinking. This characteristic enables a person to *reflect on* what he is doing and thinking. These reflections are themselves a source of belief, called consciousness. Fourth, reason is a source of belief. Our intelligent skills enable us to produce beliefs that are not directly based on other sources, but whose truth

is somehow obvious to us, or can be proved, for example mathematical theses. Fifth, we can acquire beliefs on the basis of something other people tell us. We need not experience or prove everything ourselves, we may also base our beliefs on what other people, or products of other people, tell us.

2.1.2 Object

Beliefs are generally *about* something. They reflect some view on how things relate to each other in reality. For instance, beliefs are about the weather, a book, or a judgement. In that case, the weather, book, or judgement forms the *object* of the belief. Beliefs differ in how they relate to objects. They vary in their degree of abstractness (an abstract belief about an arbitrary book from a library *versus* a concrete belief about the copy of this book you are holding). A belief is abstract when it generalises over individual objects; instead of referring to individual occurrences (tokens), it refers to object categories (types). A belief is concrete if it is connected with objects, *i.e.*, if it refers to individual objects (tokens).

The distinction between *type* and *token* deserves particular attention in the discussion of the object of belief. A type is a category of things or symbols. A token is a concrete specimen of a type. For instance, the previous sentence contains one token of the type 'concrete', and you are probably holding a copy of this book, which is a token of the publication 'Applied legal epistemology'. The object of belief is investigated in the philosophical discipline called 'ontology'. Ontology makes claims regarding the existence of individual things (tokens) and categories of things (types), and the kind of existence they have.

Ontological claims thus regard the nature of the object of belief. For instance, I can claim that the type 'judge' exists, or I can claim that a specific judge exists. Ontological claims may also vary according to the *kind of existence* they claim. It is an easy thing to claim that some object exists, but it is rather difficult, if not impossible, to say what existence amounts to. The kind of hard-boiled, physical, tangible existence of a rock is something quite different from the non-tangible, societal existence of the institution 'court'.

The distinction made between objects of belief is connected to the ontological view adopted. If my ontological view is such that types exist, my beliefs can be about types. If universals do not exist, my beliefs can still be about them, but in that case these beliefs do not refer to universals. So I may use the word 'book' to refer to a type, without referring to an existing universal. Instead, we could assume that there are only the names of types (words like 'book'), a position called nominalism. For a brief introduction to ontology and its claims, please refer to section 4.1.

2.1.3 Justification

Justification amounts to those circumstances in which the content of some entity or behaviour is sufficiently defended. Such a defence can be given in an explicit way: in terms of reasons for the content of an entity, or a proof of the content of the entity. A defence can also be given in a rather implicit way, for instance by establishing a high chance that the belief is true. Justification thus consists of all those factors that make us *believe* something. Justification is found in several forms. There are different types of justification. The typology I give in this subsection is partly based on Audi (1998, p. 2-3), who distinguishes between justification as a state of a belief, as a state of a person, and as a process.

I distinguish three main types of justification: justification as a state, justification as a process, and justification as a status. The first main type, justification as a state, is further divided into four subtypes. These are belief justification, personal justification, propositional justification (a proposition is an assertive sentence), and situational justification. The main types and subtypes of justification are summarised in table 2.1 and explained thereafter.

state	belief justification: state of justification for a belief
	personal justification: state of justification for a person
	propositional justification: state of justification for a proposition
	situational justification: state of justification for a belief not held
process	procedural justification: process in which the justification is formed
status	justification status: justifying role of a certain entity with respect to another entity

Table 2.1. Justification types

Before listing the subtypes of justification as a state, I have to explain the difference between a belief and a proposition. A belief is a proposition within the reach of a propositional attitude. For instance, if the proposition is '2 + 2 = 4', a corresponding belief might be: "I believe that '2 + 2 = 4'", or "I hold that '2 + 2 = 4'", where the propositional attitudes are 'I believe' and 'I hold'. A proposition is a sentence (in a natural or formal language).

Within the first main type of justification, there are four subtypes. (1) Belief justification is attained if certain criteria are met with respect to a specific belief.

The belief is, in other words, in the state of being justified. (2) Personal justification occurs if a person is actually justified in having a belief, and he knows that he is in this state of justification. (3) Propositional justification occurs if there are sufficient reasons for justifying the proposition. If certain criteria are met with respect to a proposition, this proposition is in the state of being justified. (4) Situational justification occurs if a person has sufficient reasons to justify a certain belief, but nevertheless does not hold that belief. For instance, John has consulted the marriage register and read Mary's name. However, he has not realised that this means that Mary is married. Thus, John would be justified in believing that Mary is married, but in fact, he does not believe that Mary is married.

The second main type of justification is justification as a process. A state of justification can, but need not be, the result of a successful process of justification. Such a process may consist of exchanging reasons, or applying certain rules, or any series of acts that aims at accomplishing a state of justification. For instance, the different steps in a penal trial aim at (among other things) reaching a clear picture of the actual facts. The rules that govern this process let the different parties present and explain their stances, and by presenting the evidence and responding to each other, ideally relevant and true statements are made as a conclusion.

The third main type of justification is justification as the status of an entity. It refers to the justifying role an entity can play. For instance, a fact can be qualified as a reason, and then its justifying role is based on a status layer of the fact. In the example given above, the fact that Mary's name is in the marriage register can be qualified as a reason for believing that Mary is married. Because it has the status of a reason, it performs a justifying role with respect to the belief that Mary is married.

To attain a state of justification for a belief, we often need reasons. Reasons generally help us to support some belief. Some reasons, however, do the reverse: they attack a reason. Such reasons are called *defeaters*. Reasons and defeaters play an important role in the justification process, and in reaching a state of justification. Reasons do so by their justifying function towards conclusions (for instance a proposition or belief). Defeaters do so by attacking reasons and thus by decreasing the justification of a conclusion. The following discussion of reasons and defeaters is based on Pollock (*cf.* Pollock 1974 and 1999). Reasons and defeaters play an important role in the discussion of legal knowledge in chapter 3 and, consequently, in the ontology described in chapter 6.

In Pollock's (1999) discussion of reasons, a reason is a ground for a belief,

which can be either conclusive or non-conclusive. A conclusive reason entails its conclusion (the belief). For instance, the reason 'three men ate poisoned fish' logically entails the belief 'two men ate poisoned fish'. A non-conclusive reason does not entail its conclusion. For example, the reason 'Mark hates fish' does not logically entail the conclusion 'Mark did not eat the poisoned fish'. Still, it seems to support that conclusion.

A conclusive reason supports its conclusion because a conclusive reason entails that conclusion logically. Non-conclusive reasons are inductive reasons and other (non-logical) grounds for belief (Pollock 1974, p. 36-39). Pollock (*ibid*.) claims that most reasons that matter (and are important for justification, and thus for epistemology) are non-conclusive reasons. These reasons, called *prima facie* reasons, can be defeated. This means that new information may force us to reject such reasons. The concept of 'defeater' is defined as follows (*ibid*., p. 38):

> "If P is a reason for S to believe Q, R is a defeater for this reason if and only if R is logically consistent with P and (P&R) is not a reason for S to believe Q."

The reason P only yields a justified belief Q if there is no defeater R, that, in conjunction with P, would cancel the reason to believe Q. There are two types of defeaters: *rebutting defeaters* and *undercutting defeaters*. A rebutting defeater is defined as follows (*ibid*., p. 38):

> "If P is a prima facie reason for S to believe Q, R is a rebutting defeater for this reason if and only if R is a defeater (for P as a reason for S to believe Q) and R is a reason for S to believe ~Q."

In this case, R is just another reason. It has basically the same status as P, but its conclusion is opposite to the conclusion of P. For instance, if I believe that it is raining outside because I saw the weather forecast predicting rain for the next two hours, a rebutting defeater for that belief is that I do not see rain falling when I look out of the window.

An undercutting defeater attacks the connection between the reason and the belief held as a consequence of the presence of that reason. An undercutting defeater is defined as follows (*ibid*., p. 39):

> "If P is a prima facie reason for S to believe Q, R is an undercutting defeater for this reason if and only if R is a defeater (for P as a reason for S to believe Q) and R is a reason for S to deny that P would not be true unless Q were true."

An undercutting defeater need not attack the belief itself or the reason for it,

but rather the assumption that the reason is *a reason for the belief*. For instance, I assume that there is a reason for me to believe that I will be happy next week. The reason is that I have read a prediction of this in my horoscope. An opponent may attack the reason (I will be happy because the horoscope says so) instead of the conclusion (I will be happy). For a further discussion of reasons and defeaters, see subsections 4.5.2 and 4.5.3. For an account of how reasons and defeaters may be used to determine the degree of justification for a conclusion, *cf.* Pollock (2001).

Having explained the different types of justification, and the way in which a state of justification can be attained by employing reasons and defeaters, I should stress that reasons also play a role in the other main types of justification. With respect to justification as a process (procedural justification), reasons constitute the main entity type governed by procedure (*qua* content and place in the procedure). Regarding justification as a status, reasons themselves carry a justifying status. To untie justification in general and justification purely based on reasons, I separate the two in the discussion of legal knowledge criteria (section 2.4). We find justification in general in the joint forces of three knowledge criteria: proper justification, reliability, and coherence. Justification based on reasons is found in the proper justification criterion.

2.2 Legal knowledge

The differences between regular knowledge (knowledge about the observable world) and legal knowledge can be traced back to the three dimensions of knowledge distinguished in the previous section. The acquisition (subsection 2.2.1), object (subsection 2.2.2), and justification (subsection 2.2.3) of legal belief are different from those of regular belief. Regular knowledge, the domain of traditional epistemology, is often acquired through perception, is often about tangible objects, and has justification demands matching with its acquisition, whereas legal knowledge is often acquired through testimony and interpretation, is about intangible objects, for instance institutions and norms, and imposes justification demands that conform to the acquisition of knowledge about those intangible objects. Below, I provide an overview of these differences.

2.2.1 Acquisition
Although legal knowledge may arise from the same sources as regular knowledge, the focus is somewhat different. In subsection 2.5.2, I elaborate on sources relevant for acquiring legal belief. For the moment, I focus on specific sources of

legal belief and knowledge, *i.e.*, sources that are classified because of their content rather than by the acquisition method employed. These are the so-called knowledge sources for the law. In order to explain what knowledge sources for the law are, I start to elaborate on two different types of sources of law: formal and material sources of law. Subsequently, I explain what knowledge sources for the law are, and finally, I clarify to what extent sources of law (formal and material sources of law) can be qualified as knowledge sources for the law.

There are two types of *sources of law:* formal sources of law and material sources of law. Formal sources of law are, according to Algra and Van Duyvendijk (1989, p. 19), the sources of positive law itself. These are statute law, treaties, and legal precedents (*ibid*.). Customary law is often also considered a formal source of law. Material sources of law are the origination sources of law, *i.e.*, those factors that contributed to the drafting and interpretation of positive law. Material sources of law themselves cannot be reduced to legal rules or legal norms. They form, however, the grounds for those rules and norms. For instance, as soon as a judge has made a decision in a case, and he has based his decision partly on the consequences his decision will have for the social structure (socioeconomic developments constitute a material source of law), he establishes a verdict (a legal precedent is a formal source of law).

Algra and Van Duyvendijk (1989, p. 20) distinguish the following material sources of law: political powers, civil servants, pressure groups, religious beliefs, moral beliefs, socioeconomic developments, geographical circumstances, and technological developments. The importance of this enumeration is that material sources of law are indeed the external factors relevant to the *origination* and *interpretation* of the law. Material sources of law, such as moral beliefs and technological developments, have an impact on both the content of law and on the way in which the content of law is understood.

Knowledge sources for the law are the sources through which we acquire knowledge about the law. Acquiring (explicit) knowledge about the law requires us to know two properties of the law: its *content* and its *validity*. Knowledge about the two properties is acquired in different manners for different legal–philosophical stances. In a legal-positivist stance, formal sources of law largely coincide with valid law. Thus, if one acquires knowledge of the formal sources of law, one will acquire knowledge of both the *content* and the *validity* of law. In a natural-law stance, however, this is not necessarily the case; the validity of law is also determined by principles that are not part of the system of positive law, and thus are not part of the formal sources of law. In such a stance, knowledge about the content of positive law is derived from the formal sources

of law, but knowledge about the validity of positive law is partly derived from principles outside positive law.

In legal-positivist and natural-law stances, knowledge about the content and validity of the law may also be derived from *material* sources of law. Material sources of law are helpful in interpreting the content of formal sources of law. Thus, sources of law (both formal sources of law and material sources of law) can function as knowledge sources for the law. However, their precise role may differ, depending on the legal-philosophical view of the law taken.

2.2.2 Object

The main difference between the object of legal knowledge and the object of regular knowledge, is that the object of legal knowledge largely consists of intangible institutions and entities, which brings about the danger of a confusion of the object of knowledge and the knowledge itself. I discern two categories within the object of knowledge about the legal domain. The first object category is *legally-relevant*, the second object category is *legal*.

The objects within the first category are situations in the world that are relevant for the legal domain, *i.e.*, entities, facts, acts, and practices that have not (yet) got assigned a legal status. The objects within the second category are situations in the world that are part of the legal domain, *i.e.*, entities, facts, acts, and practices that have been assigned a legal status. For instance, the object category of knowledge about the fact that John hit a pedestrian with his car is not legal. However, the object category becomes legal whenever the fact has the assigned legal status of criminal negligence. The fact that an object of knowledge is legal need not mean that the knowledge itself is legal. Neither does the fact that an object of knowledge is legally-relevant imply that the piece of knowledge itself is legally-relevant. That depends on the content of the piece of knowledge itself. For instance, knowledge about hitting a pedestrian with a car as criminal negligence may count as a piece of legal knowledge if it concerns the legal consequences of that fact.

The confusion of knowledge with the object of knowledge starts where the object of knowledge is more or less intangible, and the result of reasoning, interpretation, or an artifact resulting from social conventions. Whereas it is easy to distinguish the situation that John hits a pedestrian from the belief 'John hits a pedestrian', it is somewhat harder to distinguish an interpretation from a belief about that interpretation. The reason for this is that interpretations are not objects in the same way as we can regard, for instance, toys as objects: we cannot hold, feel, and look at interpretations from different angles (*i.e.*, not literally),

whereas in the case of toys, we can. We construct interpretations ourselves, and by doing this we 'make' knowledge. At the same time we add something to the world: a new interpretation, a new object of our knowledge. A further elaboration on the object of knowledge about the legal domain is given in chapter 4, that deals with the subject matter of ontology.

2.2.3 Justification

In the current subsection, I discuss legal examples of the three justification types distinguished in subsection 2.1.3: justification as a state, as a process, and as a status. After that, I discuss sources from which we can derive the content and structure of justification. Subsequently, I explain the analogy between the justification of legal decisions and the justification of legal belief. Finally, on the basis of my findings on these matters, I explain how justification of legal belief can be typified.

All three justification types distinguished in subsection 2.1.3 are found in the legal domain. Some of the instances of the types are actually institutionalised in the law. An example of justification as a state is the legitimate character of evidence (which is attained by acquiring evidence in a lawful manner). An example of justification as a process is the application of parts of civil procedural law, which guide two parties in exchanging arguments. An example of justification as a status is the legitimising force a piece of evidence exerts towards a conclusion (*e.g.*, evidence for finding a suspect guilty). From these examples, we can derive the main characteristic of justification with respect to legal belief. Basically, justification of legal belief is based on sources of law. In many cases, this means that it is rule-governed, *i.e.*, procedural rules determine in what cases justification as a state, process or status occurs. The legitimate character of evidence arises from the lawful application of legal rules regarding the gathering of evidence. The legitimate character of a civil trial partly arises from following the applicable rules of procedure. The legitimising force a piece of evidence exerts towards a conclusion may also arise from legal rules regarding the role of evidence. Of course, this is only valid insofar as the legal system concerned is rule-based, such as the Dutch one is to a certain extent.

Justificatory material in the legal domain is raised from several sources, distinguished by their content. I list the typology of Aarnio (1987, p. 123-131) as an example of the classification of justification material. He lists six sources of justification: (1) the law text itself, (2) the process that led to the accomplishment of the law text ('traveaux préparatoires'), (3) systemic interpretation in accordance with the legal system (coherence criteria), (4) court decisions (case

law), (5) doctrinal opinion (legal literature), and (6) practical reasons (goal-directed reasoning). Their justificatory nature is given by a legal tradition that also determines the prevalence of one source over another. Distinct legal systems may emphasise different sources of justification. In continental legal systems, emphasis is on legal rules. In Anglo-Saxon systems, past decisions (case law) are stressed.

Justification in the sense of the legitimisation of a legal decision is analogous to justification as a knowledge criterion. The legitimisation of a legal decision may add to its classification as knowledge. If a legal decision is sufficiently legitimised, the reasons given in that legitimisation may serve as a way to justify the belief about the decision. The legitimisation of the decision serves *justice*, the justification of the belief about the decision serves *truth*. The presence of legitimising reasons constitutes the rationality of a decision, helping it to be acceptable for the parties involved, thus helping it to be just. The presence of the same reason may to a certain degree help us to establish the truth of a belief about the decision: having good reasons available for a certain decision makes it plausible that the decision was indeed made, and thus, that the belief about it is true. Of course, this is only valid if we deem the person who makes the decisions capable of providing good reasons.

I distinguish two general features of justification in the legal domain. The first feature relates to the content of justification, the second comprises the structure of justification. The content of justification concerns the substantial reasons given for a legal proposition or a belief. This part of justification is based on the actual meaning of justifying entities, such as reasons and defeaters. The fact that some reason is a justifying reason for a proposition or a belief depends on its content, not on its structural features alone. Whether such a reason has a justifying content is to be assessed through the context in which it is used, the relevant facts, the content of both the reason and the proposition or belief it supports.

The structure of justification concerns the way in which reasons are handled and judged in order to assess whether the state of justification is attained. The structure of justification may be procedural or non-procedural. If structural justification is procedural, the occurrence (or the non-occurrence) of a resulting state of justification depends on the outcome of a process of exchanging arguments. The process is subject to a set of rules, which should help in developing a satisfactory defence of some stance. Procedural law comprises many rules that regulate the way in which statements and decisions are legitimised, and thus, it functions as a means to regulate procedural features of justification. If the structure of justification is non-procedural, it is static. A static justification structure

is, for instance, an inference scheme (deduction or induction). Such a structure does not require an exchange of arguments, just the filling in of the variables in the inference scheme.

justification features	content of justification: the content of justifying entities (for instance, the content of reasons and defeaters)
	structure of justification: - the form of a justifying structure (non-procedural) - the form a justifying procedure (procedural)
justification content sources	material sources of law: political powers, civil servants, pressure groups etc. (*cf.* subsection 2.2.1)
	formal sources of law: legal rules, case law, treaties, and customary law (*cf.* subsection 2.2.1)
justification structure sources	material sources of law: inference schemes (for instance deduction or induction, legal reasoning methods)
	formal sources of law: legal rules, case law, treaties and custom regarding procedure (*cf.* subsection 2.2.1)

Table 2.2. Content and structure of justification in the legal domain

state	content features: the content of the reasons leading to the state of justification
	structural features: the entity with respect to which the state holds (for instance, a conclusion), and the conditions under which the state holds
process	content features: content of procedural rules and the content of concrete procedures (for instance, the content of reasons exchanged in a trial)
	structural features: the way in which procedural rules govern the process
status	content features: content of the entities involved (for instance, reasons and conclusions)
	structural features: the entities with respect to which the status holds (for instance, a reason has a justification status with respect to a conclusion)

Table 2.3. Content features and structural features of three types of justification

In table 2.2, I give an overview of the sources of the content and structure of justification. In table 2.3, I provide an explanation of the content features and

structural features of the three types of justification: justification as a state, as a process, and as a status.

2.3 Two distinctions in the domain of epistemology

Epistemology is one of the philosophical disciplines that seem to require many distinctions in order to prove their own right to exist. Still, some of the distinctions are useful in that they are the easiest way to show how a domain is structured. Next, we have to admit that, of course, those distinctions are really only vague at best, and at their worst they form a no man's land where they remain untouched, unproved, and undefeated by rational discourse. As for other disciplines, distinctions are useful in epistemology to indicate the most important characteristics of divergent epistemic theories. Because many debates in philosophy focus on details, such distinctions are convenient for establishing a bird's eye view on epistemology. According to Pollock (1999), knowledge theories are currently classified along two main distinctions: doxastic *versus* non-doxastic theories (subsection 2.3.1), and internalist *versus* externalist theories (subsection 2.3.2). Not only can these distinctions be employed to classify epistemic *theories* in this section, they will also serve to classify epistemic *criteria* in section 2.4.

2.3.1 Doxastic and non-doxastic theories

Doxastic theories tell us that beliefs are ultimately justified by other beliefs. In doxastic theories, beliefs are related to each other to form a chain (a series of subsequent beliefs) or net (a structure of mutually related beliefs) of justification. A justification chain for a certain belief can be grounded by a final (basic) belief. In a justification net, a certain belief is linked to other beliefs in a web-like form.

Non-doxastic theories tell us that some external element plays a part, so that the justification of a belief can ultimately be given by something outside the realm of beliefs. In a non-doxastic theory one of the ways in which a justification can be given is in terms of the reliability of the mechanism (*e.g.*, human cognition) that acquires the belief. Non-doxastic theories fall into three main groups represented by direct realism, probabilism and reliabilism. These will be discussed in the next subsection, because, on their turn, non-doxastic theories are subject to the distinction between internalism and externalism.

There are two main types of doxastic theories: coherentism and foundationalism (Pollock 1999, p. 22*ff.*). Coherentism says that a belief is justified by its place in a net of beliefs. The relations between beliefs determine whether one of them

is justified. All beliefs have the same justificatory status; there are no beliefs that deserve priority, *qua* justification power, over other beliefs. In foundationalism, however, the existence of such basic beliefs is presupposed. In such theories, they are claimed to be suitable for starting a justification chain. Foundationalism assumes that the chain of justification stops at one or more basic beliefs. These basic beliefs are beliefs founded on our perception of the external world. They should justify themselves, for if they do not, the chain of beliefs that is based upon them will not be justified. The status of basic beliefs differs from the status of other beliefs. Basic beliefs have priority over other beliefs, in that they can end a justification chain. To give some water-related metaphors again: coherentism is best represented by a fishing net: all nodes in the net are more or less equal, and the net derives its strength from those nodes. Foundationalism can be compared to a ship that has dropped its anchor, grounding it firmly in the sea floor.

2.3.2 Internalist and externalist theories

A different distinction of knowledge theories is made between internalist and externalist theories (Pollock 1999, p. 24-27). Internalist theories state that a justification is always given in terms of the internal states of a person. These internal states are present in our cognitive system, and they can be beliefs or signals from our senses. Not all signals we acquire through our cognitive system are conscious. Moreover, even if they are conscious, they need not yield explicit beliefs. Audi (1997, p. 12) adds to this that internalist theories claim accessibility of these internal states for introspection: the belief-justifying internal state is accessible for the person who has the belief.

Foundationalism and coherentism are internalist theories (they are thus internalist *and* doxastic). An example of an internalist non-doxastic theory is *direct realism*. Direct realism is the view that some judgements about the world are not beliefs about that world, but that they are perceptual states, directly caused by the outside world (Pollock 1999, p. 87-88). Thus, there is no mediation of perceptual judgements by so-called perceptual basic beliefs (which are presupposed by foundationalist theories).

Externalism encompasses all views where elements other than mere beliefs and internal states are relevant to the justification process. So, if justification depends on the general reliability of one's cognitive system, an externalist view is employed: the tendency of one's cognitive system to yield correct beliefs is determined from an external perspective. Externalist theories do not adhere to the doxastic assumption because, ultimately, justification is not given in terms

of beliefs. Externalist theories evaluate cognitive procedures from an external perspective, so that forecasts can be made about the circumstances under which these procedures will yield reliable beliefs. Examples of externalist theories are probabilism and reliabilism. Probabilism explains epistemic justification in terms of the probability of the occurrence of individual beliefs, whereas reliabilism explains epistemic justification in terms of the general reliability of a cognitive mechanism that yields beliefs (*ibid.*, p. 100-119).

2.4 Legal knowledge criteria

The distinctions I have discussed in the previous section served to classify different theories of knowledge. Legal knowledge criteria, the subject matter of the current section, serve to assess whether a belief can be classified as legal knowledge. Two distinctions discussed in the previous section (doxastic/non-doxastic and internalist/externalist) help us establishing those criteria, because in answering the question how a certain theory of knowledge can be classified in terms of these two distinctions, we find an anchor for giving a more detailed account of justification in that theory. The term 'justification' is used here as a denominator for everything that adds to the degree in which a belief is supported. This is a broad concept of justification, and it will prove useful that this concept is subdivided into several knowledge criteria.

Under what conditions may we call something legal knowledge? This question has been revived by Gettier (1963), who explained that the definition of knowledge as true justified belief needed revision. Knowledge criteria define conditions under which we can qualify an entity as knowledge. I discuss four criteria that may establish a qualification relation between a belief and a piece of knowledge. These criteria are truth (subsection 2.4.1), proper justification (subsection 2.4.2), reliability (subsection 2.4.3), and coherence (subsection 2.4.4). They regard the correspondence of a belief with a part of reality, the presence of satisfactory reasons for a belief, the reliability of the acquisition of a belief, and the conformity of a certain belief with other, related beliefs, respectively. Except for truth, the knowledge criteria add to the fulfilment of the broad concept of justification, *i.e.*, if the criteria proper justification, reliability and coherence are fulfilled, this provides support for the conclusion that the broad concept of justification is fulfilled as well. The use of knowledge criteria for the evaluation of legal beliefs was discussed in Mommers (1999) and Mommers and Van den Herik (2000). In the former publication, it is also shown how such an evaluation framework can be used for the assessment of legal information systems.

Note that fulfilling knowledge criteria for a belief in order to qualify it as knowledge, is something different from qualifying a belief as a valid legal conclusion. In the former case, the correctness of the belief relative to its object is at stake. In the latter case, the assessment of criteria such as justice and coherence with the legal system is at stake. The two *may* coincide. The knowledge qualification of a belief about the content of a legal decision can be constituted by the validity of the legal belief. The belief that the Supreme Court qualified electricity as a good can be backed by showing that such a decision is just and coherent, although it can only be proved to be correct by showing the correspondence between the belief and the judgement.

2.4.1 Truth

'Truth' and 'true', as expressions of our daily language, refer to different phenomena. For instance, if we say that we want to know the truth about the Bijlmer disaster, we mean that we want an accurate and complete description of the plane crash, of its causes, and of the events that followed it. If we assert that Van Thijn acted as a true leader after the disaster, we mean that he did what we think a leader should do under certain circumstances. In case we state that what Van Thijn said was true, we mean that he made a statement that corresponded to an actual event – provided that we regard truth as a semantic, non-epistemic criterion. *Semantic* means that truth is a relation between sentences and reality. *Non-epistemic* means that truth does not depend on knowledge; if we are absolutely convinced that some proposition is true, that does not mean it *is* true. Instead, truth depends on the actual agreement between a proposition and reality.

If, however, we regard truth as an epistemic criterion, truth becomes a function of the presence of some form of justification for a belief. The truth of a belief then depends, for instance, on the presence of good reasons for that belief. As a consequence of identifying truth with some epistemic criterion, it easily becomes empty. So, for instance, when we identify truth with providing sufficient reasons, we may as well drop the notion of truth. But if we make the truth of a belief dependent on the fulfilment of several different epistemic criteria, for instance proper justification and coherence, truth can be a useful predicate.

Truth, understood as a non-epistemic criterion, is independent of our knowledge. It belongs to the realm of semantics. The non-epistemic versions of semantics describe the meaning of language in terms of the relation between language and reality. This means that the truth of a belief does not depend on the belief's relation with other beliefs. Instead, it only depends on the belief's relation with reality. The main non-epistemic truth criterion is correspondence

truth. This criterion says that for a sentence to be true, it should be in accordance with a situation in reality. Thus, the truth of some sentence does not depend on our opinion about its truth. The truth of a sentence can be postulated, but it can only be backed by giving reasons or proof for it, and never be proved unconditionally. To specify the notion of correspondence truth, I give a slightly abbreviated definition from Devitt (1991, p. 29) (the original definition is made dependent on types of sentences):

> "Sentences […] are true or false in virtue of: (1) their structure; (2) the referential relations between their parts and reality; (3) the objective and mind-independent nature of that reality."

The first part of this definition refers to the syntactic structure of the sentence under consideration. The second part concerns the referential relation between sentence parts and reality. The third part concerns the nature of reality: reality exists independent of what we believe about it. Together these elements mean that a sentence like 'the cat is sitting on the mat' is true whenever this sentence has the meaning that the cat is sitting on a mat, and there is indeed a cat sitting on a mat. The definition given does not include the existence of things and categories of things that are dependent on the mental. To put this constraint on correspondence truth is to exclude the possibility of determining the truth of sentences containing mind-dependent facts, for instance, 'Bill thinks the judge made the wrong decision'. This is an unnecessary deficit. Why?

The truth of a statement can also be established if it is about a human-constructed part of reality. So, after a soccer game, won by the Dutch team by 1-0, the statement 'one goal was scored in this game by the Dutch team' is true. To *know* this, we still have to employ epistemic means; after all, we see the goal being scored, or we hear the radio commentary, and thereby we draw conclusions about what happened and how the events are qualified. If the referee disallows a goal, it does not exist. Even with sentences about constructed reality, we have to know that reality before we can make (supposedly true) statements about it.

I adhere to an idealised semantic notion of truth, not to an idealised epistemic notion. This means that the question what is true is to a large degree independent from the question what can be justified. There are statements that we regard as true, but that are actually false. There are statements that we regard as false, but are actually true. The fact that there is no proof for a statement need not withhold us from leaving open the possibility of assigning a truth value to it. Moreover, *if* there is proof for a statement, we should reckon with the possibility that the statement is false after all.

An idealised, semantic notion of truth, as presented here, has a major disadvantage. If we succeed in separating epistemic considerations from semantic ones, *i.e.*, if we separate the relation between knowledge and reality from the relation between language and reality, our epistemology should preferably be such that it enables us to form correct beliefs about reality. Otherwise we would never be able to apply the concept of truth in real situations, *i.e.*, we would never be able to say (with an acceptable degree of certainty) whether some statement is true or false. Therefore, we have to find some way of linking our epistemic evidence (in the form of reasons and a reliability measure) to a state-of-affairs as it is in the objective world. A realist epistemology enables us to do so. In a realist epistemology we may hold the assumption that we form true beliefs whenever they are sufficiently justified.

In the legal domain, both epistemic and non-epistemic notions of truth play a part. For an extensive discussion of different conceptions of truth in this domain, I refer to Patterson (1996). Truth as correspondence is not very often considered a suitable criterion for application in the legal domain. For instance, Aarnio (1981) thinks that the question what is a just interpretation of a certain norm is a problem for correspondence. There are many interpretations possible for any norm, and correspondence truth seems to allow only one: the one that corresponds to the norm. Correspondence truth needs, in other words, the doctrine of the one correct solution (*ibid.*, p. 37). This doctrine states that for each norm, and each case to which it is applied, there is a correct interpretation. I elaborate on this in subsection 3.1.2. Aarnio's proposal is to replace truth with acceptance. Acceptance is defined relative to some (legal) community: the whole group, or only some part of it, adheres to a norm or a value.

Niiniluoto (1981) criticises Aarnio's rejection of correspondence truth. He says that Aarnio's mistake is to regard acceptance and values as parts of the relation between norm propositions and legal order. Niiniluoto states that, instead, acceptance and values are part of social reality (*ibid.*, p. 74). Norm propositions *presuppose* the existence of acceptance and values. Norm propositions are propositions about norms: given a certain norm (the state's authority should be accepted), we can state something about that norm ('the state's authority should be accepted' is a valid norm in The Netherlands). Norms are subject to acceptance. It thus becomes possible to determine the truth of norm propositions (*i.e.*, statements about norms). To put it differently: acceptance and values are relative to an auditory, but the moment they are established, they are prone to correspondence truth.

The result of taking Niiniluoto's stance is that, whereas we cannot say that

some value (the state's authority should be accepted) is true or not, we can do so for a norm proposition ('the state's authority should be accepted' is a valid norm in The Netherlands). This does, however, not solve the problem put forward by Aarnio. If, as Aarnio states, there are multiple interpretations possible for a norm, then we can state multiple norm propositions for a single norm. If these norm propositions are inconsistent with each other, they cannot all be true at the same time, in so far as we adhere to a correspondence notion of truth.

Yet, even in parts of the legal domain, truth is non-epistemic. For the parts of this domain where we need an epistemic truth notion, such as acceptability, maybe we should not try to apply truth at all. After all, it would be best to model our conception of truth after the actual use of this word, preferably its use in the legal domain, while removing possible inconsistencies and clarifying the notion. Thus, for instance, the construal of truth as acceptability is only suitable if it fits in with the conception of truth in the legal domain.

The larger part of the legal domain is constructed by human beings. However, entities in this part of reality often exist objectively, which means that truth remains non-epistemic: the truth of certain statements about such entities does not depend on our knowledge of those entities. The following example illustrates this. Assume that one day humanity ceases to exist. Assume that there is still a book, called 'Truths about humanity'. It consists of a list of statements. One of these statements is: 'In 1999, Dutch penal law said that killing a person on purpose counts as manslaughter'. The sentence is true. It will still be true if there is no human being to state that it is true and to justify it. It will still be true if some aliens land on earth, find the book, learn the language, find evidence (in other books), and regard the statement as true. Moreover, it will still be true if these aliens land on earth, find the wrong evidence, and think it is false. Thus, there remains a place for a correspondence notion of truth in the legal domain, even though parts of it are constructed by human beings.

Therefore, I adapt the simplified version of Devitt's (1991) notion of correspondence truth to accommodate human-constructed facts (for instance the existence of a legal rule 'killing a person on purpose counts as manslaughter') in the following way:

> A sentence is true or false in virtue of: (a) its structure; (b) the referential relations between its parts and reality; (c) the objective nature of that reality.

In this definition, only part (c) is altered. In the original definition Devitt (*ibid*.) refers to the "objective and mind-independent nature of that reality". The mind-independence demand is left out because it restricts the application

area of the correspondence truth criterion too much. To be able to apply the non-epistemic correspondence truth criterion we just need to guarantee that the parts of reality we are talking about are not true (or false) because we say they are true (or false), but because they have been established already when we state something about them. For instance, the establishment of such facts can be attained by institutional rules or by conventions.

To return to the discussion between Aarnio and Niiniluoto: does the doctrine of the one correct solution coincide with the correspondence truth criterion, or: need there be one exclusive interpretation of some norm to be able to employ correspondence truth? Let us first determine what correspondence truth exactly applies to. Correspondence truth only applies to sentences (in either a formal or a natural language). Take some sentences *s* and *t* that are interpretations of a legal regulation *l*. The question is whether *s* is a true interpretation of *l*, whether *t* is a true interpretation of *l*, and whether both can be true interpretations of *l*. The application of the notion of correspondence truth depends on the establishment of interpretations of *l*.

From such a viewpoint, determining truth in the legal domain involves the comparison of statements with facts, just like truth in other contexts. Legal truth follows upon the establishment of facts, and these facts are established by certain qualified legal professionals. Thus, truth itself is not 'human-made', but the reality it refers to is. One of the consequences of this is that the truth of two different interpretations of a norm is determined by their consistency with the actual norm, that is part of reality. If two interpretations can be consistent with reality, they can both be true, *if* truth is an applicable criterion.

This is precisely the problem. To apply the notion of truth to interpretative beliefs means that we have to establish new objects (interpretations). If we provide sufficient reasons for the content of these interpretations, we may establish the truth of propositions about those interpretations. For instance, if a qualified person, like a judge, rules in a case, and he justifies his verdict, he performs what Niiniluoto (1981) calls 'truth-constituting argumentation' (as opposed to 'truth-seeking argumentation'). In table 2.4, I give an overview of the situation in which there are two interpretations of a norm: a normative belief is formed about the norm. Subsequently, two interpretive beliefs are formed, either about the normative belief, or about the norm directly. The propositions that are based on these interpretative beliefs have a semantic relation with the norm, as well as with the interpretations that are established as objects on the basis of the interpretive beliefs. The former relation is interesting; if we claim the existence of such a relation, then interpretive beliefs can be true or false.

object → ↓ level	norm	interpretation 1	interpretation 2
ontological level	norm	interpretation 1	interpretation 2
epistemic level	normative belief reasons for belief	interpretive belief 1 reasons for belief 1	interpretive belief 2 reasons for belief 2
semantic level	norm proposition	interpretive proposition 1	interpretive proposition 2
truth	correspondence between norm proposition and norm	correspondence between interpretive proposition 1 and norm, and between interpretive proposition 1 and interpretation 1	correspondence between interpretive proposition 2 and norm, and between interpretive proposition 2 and interpretation 2

Table 2.4. Norm propositions and truth

Establishing an interpretative belief should be distinguished from establishing the truth of an interpretative belief. Moreover, the latter should be distinguished from the truth of a belief *about* that interpretative belief. Establishing an interpretative belief is a human activity, sometimes supported by an argumentation process. Even though the possibility of establishing the truth of an interpretative belief may be denied, we can still make true statements *about* those beliefs.

2.4.2 Proper justification

Beliefs are justified in different ways, depending on their content and on their type. A belief can be justified by other beliefs, such as beliefs based on sensory evidence, beliefs based on the statement of another person, or by facts. The fact that the streets are dry could form a *reason* for me to believe that is has not rained. In as far as the justification of a belief (or some other entity, such as an act or a proposition) is given in terms of explicit reasons, I call it *proper* justification. Proper justification refers both to the act of giving reasons and to the state of being justified. If I give some reason for my belief, I am in the process of justifying properly, and if I have sufficient reasons for my belief, the state of being justified properly is attained. A belief can be justified properly for one person while it is not for another. For instance, if I lack reasons for a belief, but a friend of mine has good reasons for a belief with the same content, he will be justified properly in his belief, whereas I will not be justified properly in my belief.

Proper justification is central to knowledge about the legal domain, as such knowledge is often about entities that derive their existence from reasoning.

Reasoning comprises the arrangement or reasons in such a form that a consistent argument structure results, and ideally, the proper justification criterion is fulfilled. In the description of proper justification in the legal domain, the focus is somewhat different from proper justification for regular knowledge. In traditional epistemology, proper justification is discussed mainly with respect to perceptual beliefs (*cf.* section 2.5). Proper justification in the legal domain also applies to other types of beliefs, such as interpretative beliefs (*idem*). For instance, if we need to interpret some law text, different sources of justification enter. Reasons in such a context may be based on material and formal sources of law (*cf.* subsections 2.2.1 and 3.3.1).

Defining proper justification in a realistic way, that is, employing a criterion that can actually be used in real life, can be attained. The following example definition of proper justification takes into account both reasons and defeaters (*cf.* subsection 2.1.3): there should be either a conclusive or a non-conclusive reason for a subject to believe a certain proposition *p*. Additionally, there should not be a defeater for this proposition, nor for the reason for the subject to believe that proposition.

In this definition, some arbitrary choices are made with respect to the depth of justification. For instance, in some cases, one non-conclusive reason is hardly sufficient justificatory material to justify a certain belief. And if there are reasons that are inconsistent with each other, how do we choose the right one? However, justifying a conclusion by giving a reason for it, and justifying these reasons by giving other reasons, and so on *ad infinitum*, does not fit in with our needs either. Therefore, a relatively arbitrary choice with respect to the depth of justification is necessary to apply the criterion in a sensible manner.

2.4.3 Reliability

The criterion of reliability mainly applies to perceptual beliefs (*cf.* section 2.5). The issue that induced the introduction of the reliability criterion is illustrated by the following example, drawn from Goldman (1976, p. 772-773), and paraphrased in Audi (1993, p. 188). A person named Henry enters a district where, along the roads, barns made out of papier-mâché appear. However, Henry cannot see the difference between real barns and fake ones. At the border of the district, he sees a barn, which is a *real* one. Now, Henry's belief that he sees a barn is true and it is justified, but is it knowledge? Goldman says we are inclined to say it is not. The criterion that should be complied with is the *reliability* of a belief, which is defined by Goldman as follows. Given a certain belief, there should be no (potential) event that could cause the same belief, while that belief,

which is justified and true, is not a piece of knowledge. Thus, in case of Henry, there should be no papier-mâché barn on the route. If there is one, Henry could form a belief about a real barn on the same route ('there is a barn right here'). He could be justified in believing it, and the belief would be true. But the belief would not be knowledge, because of the possible occurrence of an event that would prevent the belief from being transformed into knowledge. This criterion of reliability primarily concerns perceptual beliefs.

Audi lists a number of subcriteria that determine whether a perceptual belief complies with the reliability criterion (Audi 1993, p. 17):

"1 the acuteness of the senses relevant to forming, sustaining, and confirming the belief;
2 the normality of their operation at the time;
3 the appropriateness of the perceptual circumstances to the content of the belief;
4 the normality of the perceiver's responses to the sense(s);
5 the absence of a justified belief – or of justification *for* believing – that one or more of (1)-(4) fails to hold."

In brief, these criteria amount to the following. Criterion (1) says that we should have the power to discriminate a certain fact. So if I claim that I see a barn, I should be in close distance to it. In that case I can discriminate it clearly from its surroundings, and do not mistake it with some other object (a barn facsimile). Criterion (2) says that our senses should operate properly. Hallucinations and optical illusions can disturb what we perceive. Criterion (3) says that our perception should be appropriate to the kind of belief it is supposed to sustain. This means that when we perceive a colour, the lighting should enable us to distinguish red from green. Criterion (4) says that a disturbance between the act of perceiving and the forming of a belief should not occur. While we see a red thing, we may not belief it is red, just because we are confused in some way. Criterion (5) says that there should be no reason for us to believe that any of these problems occurs at the time (*ibid.*, p. 201-202).

If we want to define the reliability criterion relative to testimonial beliefs, we will have to find out whether the circumstances under which those beliefs are acquired reliably are different from those for perceptual beliefs. The reliable acquisition of perceptual beliefs depends mainly on the reliability of the cognitive apparatus that produces those beliefs. Apart from this, the circumstances under which the belief is formed, such as the lighting, determines whether the reliability of the cognitive system is sufficient; a reliable cognitive system should also be

able to discriminate whether the lighting is adequate, and if not, refuse to form a belief. In case a person forms testimonial beliefs (*cf.* section 2.5), the reliability of this person's cognitive system is relevant. But the reliability of the belief source is also important. While our environment cannot be deemed reliable or unreliable (that would be a category mistake), the sources of testimonial beliefs can. Thus, if I read a newspaper that generally has pretty accurate scoops on the extramarital life of the president, this adds to the reliable acquisition of my beliefs about the extramarital life of the president via that newspaper.

Reliability of legal beliefs should allow for the assessment of beliefs concerning legal affairs. But reliability has traditionally been applied to perceptual beliefs. These beliefs constitute only a small part of the beliefs relevant for the legal domain. Reliability is a measure for the integrity of the cognitive system. It is an externalist criterion, *i.e.*, it provides us with the chance that our cognitive system yields correct beliefs, and this measure is determined independently of our internal states. Audi's (1993, p. 17) five criteria for determining the reliability of a belief all hold for perceptual and testimonial beliefs in the legal domain as well. For the other types of belief, the reliability criterion has to be rephrased. In case of a memorial belief (*cf.* section 2.5) that is originally based upon perception, both the reliability of the perceptual apparatus and the reliability of the memory of a person are relevant. In case of a reasoned belief (*cf.* section 2.5), there should be an acceptable reasoning method, and in case of interpretative beliefs (*idem*), the way the belief is acquired should be acceptable as well. Therefore, I define reliability in the legal domain as follows:

> A legal belief is acquired in a reliable manner whenever the route by which it is acquired is acceptable by the standards in a given legal and social context.

Reliability is thus based upon the acceptability of the route by which a belief is acquired. In case of perceptual beliefs, reliability is measured in terms of the production of *truth*. In case of reasoned beliefs, the truth criterion may have to be dropped. Instead, the adherence to a procedure becomes important for the fulfilment of the reliability criterion. If there is no fixed procedure for the acquisition of a belief, such as with interpretative beliefs, the acceptability of its acquisition can be determined entirely by relating to the content of the interpretative belief. To establish the acceptability of the acquisition route of interpretative beliefs, we refer to the reasons given for these beliefs. Reliability is then reduced to the fulfilment of the proper justification criterion for a sufficient amount of interpretative beliefs produced by a certain person.

The difference between the reliability criterion and the proper justification criterion is that the former is located in the 'context of discovery', and the latter in the 'context of justification'. The former is about the production of beliefs, and the adherence to the procedures that perform this task, the latter is about providing good reasons for the result of the procedure. Distinguishing the two is difficult, because the adherence to a procedure also adds to the proper justification of a belief.

2.4.4 Coherence

In the current subsection, I start discussing general definitions of coherence. The main similarity between the two definitions I give (their internalist nature) is the outset of a brief discussion of the degree to which coherence constitutes truth. After that, I discuss legal philosophers' views on coherence, explaining the main components of their definitions.

There are different definitions of coherence, varying from the logical definition by Kirkham (1992) to the extensive definition that was derived from BonJour's (1985) *The Structure of Empirical Knowledge*. Kirkham's definition of coherence says that a coherent belief set should be consistent, and that there should be inductive or deductive implication relations among the beliefs (*cf.* Kirkham 1992, p. 104):

> "Each member of some set of statements or beliefs is consistent with any subset of the other statements or beliefs.
> Each statement or belief is inductively or deductively implied by the set of premises formed by all other statements or beliefs, or by each other statement or belief."

Bender (1989) gives a summary of the subcriteria of the coherence criterion he found in BonJour's (1985) *The Structure of Empirical Knowledge*. The following set of conditions must be fulfilled in order for a set of beliefs to be coherent (Bender 1989, p. 5):

> "(i) It is logically consistent,
> (ii) It has a high degree of probabilistic consistency,
> (iii) It has a significant number of relatively strong inferential connections among component beliefs,
> (iv) It is relatively unified, i.e., does not divide into relatively unconnected subsystems,
> (v) It contains few unexplained anomalies,
> (vi) It provides a relatively stable conception of the world and remains co-

herent (i.e. it satisfies (i)-(v) in the long run) and

(vii) It satisfies the Observation Requirement, i.e., it must contain laws attributing a high degree of reliability to a reasonable variety of cognitively spontaneous beliefs, including introspective beliefs."

When comparing Kirkham's (1992) and BonJour's (1985) coherence criterion, the striking difference is that Kirkham's criterion refers only to logical characteristics of the belief set, whereas BonJour's criterion also refers to the content of those beliefs, the degree to which the beliefs in the set are unified, the stability of the belief set, and the link between beliefs and the world (the Observation Requirement). However, in both definitions, coherence is still an internalist criterion, *i.e.*, it concerns the internal relations of some system. It does not concern the relations between the system and reality, unless there are beliefs in the system that concern reality (and then the Observation Requirement comes into play).

The internalist nature of the coherence criterion has a consequence for the relation between coherence and truth. The degree of coherence is often regarded as a direct measure for the degree of truth of the beliefs to which it applies. I reject this interpretation, because beliefs can be coherent without being about reality. But even if coherence is not used as a truth criterion itself, it can be used as a knowledge criterion in either of two ways. First, its purpose can be to give epistemic support for the truth of the beliefs involved (instead of being a direct measure for the truth of the beliefs, it gives support for the assumption that the beliefs involved are true by correspondence). This is attained by employing the assumption that beliefs cohere *because* they adequately represent a coherent reality. Second, its purpose can be to support the beliefs by the mere fact that they cohere. In this case, coherence directly supports the transition of belief to knowledge; it is an ideal in itself.

Coherence in law is defined in very different ways, but with very much the same intentions. MacCormick (1978, p. 152) says about coherence that "the multitudinous rules of a developed legal system should 'make sense' when taken together". The criterion usually consists of a demand of consistency plus a demand of the presence of reasonable content relations among beliefs. Thus, two beliefs cohere with each other when they do not logically exclude each other and we are able to say that they fit in with each other.

Peczenik (1989, p. 178-179) distinguishes between thirteen subcriteria constituting the notion of coherence. I summarise these subcriteria by three measures, *viz.* support, preference, and universality. The first measure, support, relates to the degree to which some statement is backed by reasons, and the degree to

which it is interconnected with other (universal) statements, and the degree to which a set of statements contains supported statements *etc*. The second measure, preference, indicates the degree to which principles relate to each other in a preferential order. The third measure, universality, denotes the degree to which concepts can be applied generally, resemble each other, are cross-connected *etc*. According to Peczenik, the list applies to statements, universal statements, concepts, theories, reasons, chains of reasons, and principles. But for the sake of simplicity, I assume that they apply to beliefs as well.

The criteria are used to sustain coherence as an ideal of mutual dependency of beliefs. The main assumption that should justify the choice of coherence as a knowledge criterion is as follows: a highly coherent theory of some part of reality reflects that part of reality, and as a corollary of that assumption, statements of a coherent theory are true statements (*cf*. Peczenik 1989, p. 184-186). There are two ways of attacking such a claim. First, a coherent theory does not necessarily provide an adequate description of the world that it should be about; the world may be completely different from the coherent picture we give of it. Second, starting from the world, how do we acquire a coherent theory about it? Surely not by just constructing a theory that is as coherent as possible. Instead, we should conform to empirical data about the world, and those data may be less coherent than we would wish.

But there is also an argument in support of employing coherence as a knowledge criterion. The legal domain is a part of reality constructed by human beings. If we try to theorise about that part of reality, it could be argued that coherence is a reasonable criterion to employ. Human beings usually try to apply as much consistency on the things they design as possible. The quality of some work (be it a law or a book) is often judged by coherence-type criteria. Therefore, coherence becomes a reasonable demand.

Peczenik (*ibid.*, p. 188) lists three limitations of the coherence criterion. First, coherence consists of multiple criteria, balancing those criteria does not always lead to a clear answer whether the coherence of a given system exceeds the coherence of a different system. Second, the coherence criterion cannot prevent that unjust beliefs remain present in a system of beliefs; coherence does not regard moral content. Still, Peczenik claims that a higher degree of coherence contributes to justice (*ibid.*). Third, all normative systems suffer from incompleteness; new norms and concepts are introduced or old norms and concepts are adjusted so as to fit in with new cases. Thus, even if there is a measure of coherence, its value varies over time.

A different set of three subcriteria of coherence is given by Alexy (1998, p. 41).

He regards the following elements as constitutive parts of coherence: consistency, comprehensiveness, and connection. The subcriterion of consistency says that a coherent set of propositions should not include a contradiction. The subcriterion of comprehensiveness says that a coherent set of propositions should contain "as many and as different propositions as possible" (*ibid.*, p. 42). Alexy claims that, ideally, a coherent theory covers as many parts of the world as possible. This means that it should contain propositions about all these different parts. The subcriterion of connection means that there should be as many relations as possible among propositions, where one proposition is a reason for another. If we compare these claims to the three subcriteria of Peczenik's, then we find that consistency and comprehensiveness lack in Peczenik's overview, whereas it seems that connection is similar to the support criterion.

Yet another approach to coherence is Brouwer's (1990). He lists four elements that constitute coherence (*ibid.*, p. 25-29). First, the part-whole relation between sources of law and legal norms features a degree of coherence, which is increased when that relation is less ambiguous. When two sources of law together constitute a certain norm, but they also give rise to a norm different from the first, this may decrease coherence. Second, the presence of logical relations (for instance, deducibility) or the absence of logical relations (for instance, inconsistency) may increase coherence. Third, if there are elements in a set that presuppose the presence of other ones, and those other elements exist, then coherence is increased. Brouwer (*ibid.*, p. 27) gives the following example (translated from Dutch): "If a person does A without permission from B, then he will be punished with C". This norm presupposes that B has the competence to issue a permission. If he does not, this will decrease overall coherence. Fourth, the mutual relations between legal norms and legal principles are a factor in determining coherence. Brouwer states that the presence of legal principles alone does not provide coherence. Legal principles can provide arguments pro or con some norm. They may increase coherence, if different legal principles can help us defend the same norm, but they also may decrease coherence, in case different legal principles provide reasons pro and con a norm at the same time. All four criteria can be characterised as support criteria, *i.e.*, they concern the relations among statements which sometimes belong to different categories. The other categories that I used to classify Peczenik's coherence criteria do not apply to Brouwer's criteria.

Summarising, coherence concerns the presence and nature of relations in a set of beliefs: its main ingredient is the support measure mentioned above. Except for the demand of consistency of the beliefs considered, content relation among

those beliefs play a central role in the concept of coherence. This takes me to the concluding remarks about the relation between the four knowledge criteria discussed in this section and the distinctions explained in the previous subsections. I have discussed four knowledge criteria: truth, proper justification, reliability, and coherence. I regard truth as a non-epistemic criterion, whose fulfilment can only be determined conditionally, through epistemic means. The connection between the other three criteria is that they are all epistemic criteria. If any of these criteria is fulfilled, this contributes to the fulfilment of the broad concept of justification, which covers all available epistemic means to support a belief.

In addition, there is a more specific relation between the criterion of proper justification and the coherence criterion. Proper justification is evaluated with respect to an individual belief: it concerns the reasons there are to support the content of individual beliefs. Coherence is evaluated with respect to larger sets of beliefs, and it concerns the nature of the relations among those beliefs. Thus, coherence evaluates relations among beliefs on a larger scale than proper justification, but essentially, they cover the same subject matter.

Starting from the three dimensions of (legal) knowledge in sections 2.1 and 2.2, we can indicate to which dimension each knowledge criterion belongs. First, truth belongs to the object dimension, because it evaluates the relation between the belief and its object. Second, proper justification obviously belongs to the justification dimension, as it concerns giving reasons in support of a belief. Third, reliability is found on the acquisition dimension, because it concerns an aspect of belief forming. Fourth, coherence is found on the justification dimension, because it concerns the support relations among beliefs, an increase of which may be regarded as an increase of justification of those beliefs.

The distinctions in section 2.3 provide further guidelines to specify the criteria. Along the doxastic/non-doxastic distinction, we find proper justification and coherence on the doxastic side, and truth and reliability on the non-doxastic side, because the former two criteria regard beliefs and relations between beliefs, while the latter two criteria do not regard beliefs directly, but rather the relation between beliefs and reality, and the way in which beliefs are acquired. Along the internalist/externalist distinction, proper justification and coherence are found on the internalist side, because they concern the relations between internal states (beliefs are a subset of internal states). Truth and reliability are found on the externalist side, because they review relations that cross the border between internal states and the external world.

2.5 Types of belief and knowledge in the legal domain

In the present section, I discuss distinctions that can be used to characterise beliefs and knowledge in the legal domain. Four distinctions are made to be able to do this: (1) between legally-relevant and legal, (2) between abstract and concrete, (3) between regular and moral, (4) between belief sources and (5) between practical and factual. The distinctions (1) through (4) apply to both belief and knowledge, while distinction (5) only applies to knowledge. I explain these distinctions below (subsection 2.5.1). Subsequently, I elaborate on the distinction between belief sources, which are called 'origination sources of legal belief' from there on (subsection 2.5.2).

2.5.1 Distinctions regarding legal belief and knowledge

As to the first distinction mentioned above, legally-relevant belief and knowledge, being regular belief and knowledge relevant to the legal domain, can be analysed within the paradigm of traditional epistemology. Legal belief and knowledge may be analysed within the same framework, but we should consider the special (normative) status of such knowledge. As explained in subsection 2.2.2, we should distinguish between belief and knowledge and the object of belief and knowledge. As an example, I stated that the object category of knowledge about the fact that John hit a pedestrian with his car is not legal. The object category becomes legal whenever the fact has the assigned legal status of criminal negligence. Whether belief or knowledge itself is legal, depends on its content. For instance, knowledge about hitting a pedestrian with a car as criminal negligence may count as a piece of legal knowledge if it concerns the legal consequences of that fact. All beliefs and knowledge that play a role in the legal domain, but whose content is not legal, are legally-relevant. All beliefs and knowledge whose content primarily concerns the legal status of some state-of-affairs are legal. For instance, a belief about an event in which a man is hit by a car is legally-relevant, and a belief about that event being culpable homicide is legal.

As to the second distinction, Hage (1999) distinguishes between abstract and concrete legal knowledge. This distinction can also be made for beliefs. Abstract belief and knowledge are about rules, principles, values and actions. Concrete belief and knowledge are about cases. These two types of belief and knowledge represent two sides of a range, between which the character of an actual belief or piece of knowledge may vary.

The third distinction is between regular knowledge and moral knowledge. Both types of knowledge are found in the legal domain. For instance, moral

knowledge is especially relevant to natural-law views, in which such knowledge is directly related to legal knowledge. Also, moral knowledge can be regarded as a material source of law, even in a legal-positivist view. More generally, there are analogies between the treatment of moral rules as indicated by Audi (1997) (see below) and the epistemology of legal rules. These analogies can be found in the justification criteria for legal rules (for instance the lack of susceptibility for truth of rules, and the alternative criteria that are available).

Although the distinctions between doxastic and non-doxastic, and between internalist and externalist theories of knowledge seem to be geared to regular (perceptual) knowledge, they also apply to moral knowledge. In this subsection, I explain how. Audi (1997, p. 13-17) explains the relevance of the internalism-externalism distinction to moral knowledge. The internalism-externalism distinction applies to both moral beliefs and actions. With respect to moral beliefs, justification can be acquired by reflection or introspection (internalism), or by an appeal to a general pattern (externalism) (*ibid.*).

To scrutinise the meaning of the distinctions for moral knowledge, I use one of Audi's (*ibid.*) examples of a moral belief. Assume that we have a belief that the promising rule is correct, *i.e.*, the rule that people should generally keep their promises is a morally valid rule. In a Kantian *internalist* view, such a belief can be justified by an appeal to the categorical imperative. The categorical imperative may be grasped by introspection, and its content may justify the belief in the promising rule. It boils down to the obligation that someone's own action (I promise Marie to pick up her child at the school bus stop) is translated to a rule (I should keep my promises). The rule should be applicable to all people in order to be valid (people should keep their promises). In a utilitarian *externalist* view, it is assessed to what extent a certain behaviour adds to happiness in general. If the promising rule adds to happiness, and people behave in accordance with that rule, this constitutes a reliable process, adding to the externalist justification of belief in the promising rule.

A similar approach can be taken towards the doxastic/non-doxastic distinction. Justification in a doxastic theory is always given in terms of beliefs. Justification in a non-doxastic theory is given – at least partly – in terms of other entities or criteria. Beliefs may contribute to the justification of some rule, for instance the promising rule. This is the case in the Kantian view. But other entities, such as acts, or internal states, may also add to the justification of a belief. For instance, the absence of an internal state could help justify the belief in the promising rule. The internal state may be the awareness of a counter-example of the promising rule (for instance, if keeping a promise would lead to serious

damage). Note that such awareness need not be explicit, in the form of a belief. If I have promised my neighbour that I will watch her car, and while I am waiting for her the garage is set on fire, so that I have to flee, my awareness of an exception to my promise need not be explicit. If such a defeater for my belief in the promising rule is not present, this may add to the justification of that belief.

The fourth distinction, between five belief sources, was discussed in subsections 2.1.1, and will be elaborated on with respect to the legal domain in subsection 2.5.2. From now on, I call these belief sources 'origination sources of belief', because they refer to the way in which the content of a belief is acquired. For now, it suffices to say that, although the origination sources of belief discussed for regular belief are relevant to legal belief as well, origination sources of legal belief may have a different character. For instance, the acquisition of legal belief may be regulated. If a judge has to disregard some part of evidence that he has already heard about, this is clearly a deviation from regular belief acquisition through testimony in a non-legal context, in which, normally, acquired belief does not have to be disregarded.

The fifth distinction, between practical and factual, is based on Ryle's (1949) distinction between practical and factual knowledge (or knowing how and knowing that). In a legal context, practical knowledge for instance concerns how to assess cases. Factual knowledge is explicit knowledge: what one regards as valid law can be expressed in language. Practical knowledge is often implicit; few lawyers can explain in detail how they assess cases. This kind of practical legal knowledge can still be made explicit to a certain extent, in the form of a reconstruction of the procedure followed. Even though this procedure may not be an exact reflection of what actually happens (it may even be completely different from that procedure), such a reconstructed procedure may prove useful in learning the skills.

2.5.2 Origination sources of legal belief

Knowledge sources for the law are to be distinguished from origination sources of belief: the former are about knowledge, and the latter about belief. Moreover, knowledge sources for the law are relevant sources for establishing legal norms. Origination sources of belief refer to the process of belief acquisition that is employed. So, for instance, knowledge of the Constitution can be acquired from the text of the Constitution (a knowledge source for the law, to be more precise, a formal source of law). The way this knowledge is acquired may vary. We can read the Constitution (perception), we may remember its content (memory), someone might tell us about its content (testimony), we may disagree with

some of its content (consciousness), or we may infer something about its content from other information (reason) (*cf.* Audi 1998).

To these origination sources of belief, I add *interpretation*. Interpretation is distinguished from reason by its apparent implicit nature. While reason is supposed to be explicit and thus verifiable, interpretation is implicit and only verifiable by its result and its justification. The result of an interpretative activity, the interpretation, can hardly be predicted. The reason why these belief types are distinguished, is that their origination may also have an impact on the way they should be turned into knowledge.

Traditional epistemology, in its examination of knowledge acquisition, is heavily biased towards perception as a source of belief, whereas legal belief, and to a lesser degree legally-relevant belief, are often based on testimony, reason and interpretation. The purpose of the present subsection is to find out in greater detail what origination sources of belief are involved in acquiring legal belief. In the discussion below, I have left out consciousness as an origination source of belief, because insofar as it plays a role in acquiring legal knowledge, it does not primarily concern its content.

Perceptual beliefs are beliefs that persons acquire directly through their senses. It would be naive to say that this category of beliefs does not require any form of interpretation, as perception is not a passive process. Let us, instead, say that to acquire a perceptual belief does not require any conscious act of reasoning, and that the state-of-affairs perceived is indeed the object of the belief. An example of a perceptual belief is the belief 'John is crossing the street', held by Peter, who sees John crossing the street. Another example is the belief 'the envelope is white', held by Susan, who holds a white envelope in her hands.

Memorial beliefs are formed on the basis of the memory of a person. For instance, a person who has once seen someone passing by on a yellow-painted bicycle, may remember this later. Accordingly, he forms a memorial belief that he once saw a person passing by on a yellow-painted bicycle. Acquiring beliefs from memory causes distortions due to the functioning of human memory.

Testimonial beliefs are derived from indirect evidence, *i.e.*, evidence that signifies that some state-of-affairs holds, but that does not directly refer to that state-of-affairs. To elaborate on the two examples given above: if I acquire a belief 'John crosses the street' when Peter tells me so, and I believe him, I acquire a testimonial belief. If I see a written statement saying that Susan thinks the envelope she holds is white, and I believe this statement, then I acquire a testimonial belief as well. These testimonial beliefs require an implicit or explicit inference step (*e.g.*, she tells me *p*, so *p*).

Reasoned beliefs arise from clearly defined reasoning schemes. Such beliefs, which are 'calculated' by clearly defined mathematical procedures (algorithms), for instance arise from deduction, induction, and argumentation procedures. The demand is that there should be a fixed, verifiable procedure with verifiable input. Together these will lead to a verifiable output.

Interpretative beliefs are the result of a process of imposing meaning on, for instance, acts, facts, or other beliefs. They provide a certain explanation or clarification of such objects, in different terms, like a translation. Generally, they are not derived by a simple inference from other beliefs or evidence. Instead, they depend on certain choices and valuations. The process leading to an interpretation is not completely verifiable, and sometimes it is not verifiable at all. Interpretative beliefs occur for instance when a judge sentences in a hard case (a case in which there appears to be no clear way to apply a valid rule of law to a specific case).

In the legal domain, the primary source of belief consists of texts, and reading a text can be both considered as acquiring knowledge via perception, and via testimony. In case just the content of the text is important, perception and interpretation are the origination source of belief. In case the object the text refers to is concerned, testimony is the origination source of belief. The two are closely related in the legal domain, where a law text expresses the content of legal norms. The law text is acquired by perception, its content by interpretation. Legal norms can be obtained by interpretation accordingly, from the content of the law text. Reason covers the different types of reasoning employed in the legal domain, for instance a–contrario and analogical reasoning. These reasoning types are generally accepted in the legal domain, even though, strictly speaking, they may yield logically invalid conclusions.

These three origination sources of belief (perception, reason, and testimony) are regulated to a certain extent by the law itself. When a judge has to disregard some part of evidence that he has already seen or heard about, this is clearly a deviation from regular knowledge acquisition through perception or testimony, where all reliable evidence should contribute to our knowledge. An example with respect to reason as an origination source of belief is the use of analogical reasoning. Such reasoning is not permitted in Dutch Penal Law, while it is permitted in Dutch Civil Law.

By far the most interesting types of legal belief – taken by this classification – are reasoned beliefs and interpretative beliefs. The former is the category of those beliefs whose origination is clear, *i.e.*, we know how they are derived from other beliefs. The latter is the category of those beliefs of which we do not

know exactly what factors played a role in their determination, or how these factors were valuated.

2.6 Epistemic niches

Some instances of knowledge acquisition in the legal domain can be characterised as knowledge acquisition in an *epistemic niche*, a (partially) controlled environment in which knowledge is acquired and processed. Van den Hoven (1995) gives examples of epistemic niches: control rooms of nuclear plants, surgery rooms, and war rooms. Such niches limit and structure the sources of knowledge available. A person in the control room of a nuclear plant has as his most important source of knowledge the instruments on the control panel. Although he can pick up a phone and check with employees elsewhere in the plant, the actions he performs might be based exclusively on information he reads from screens and meters. Such directions may arise from the content of applicable protocols. There may be very good reasons to use control rooms and applicable protocols, so that certain procedures can be followed accurately, but they limit the knowledge acquisition capabilities of the persons involved.

The same can be claimed of certain legal professionals, such as judges, who have to depend on knowledge supplied to them in a controlled environment. In this case, rather than walls and meters, rules and procedures limit epistemic opportunities for the persons involved. Instead of taking all available proof into account, a judge or jury must often decide on the basis of admitted evidence. Moreover, instead of being able to supply evidence in any stage of a trial, such actions are limited to certain stages.

The existence of epistemic niches has three major consequences. First, responsibility for belief acquisition tends to vaporise, as the epistemic starting point for the persons involved is predetermined. Second, procedural restrictions may change the reliability and justification of acquired beliefs: reliability may decrease or increase, and justification may be standardised (which may be qualified as either a positive or a negative consequence). Third, the content of the beliefs acquired will be dependent on the available access channels to knowledge sources.

In the legal domain, sometimes procedural restrictions are applied to knowledge acquisition. Therefore, epistemic dependence is created. As with the control room of a nuclear plant, that dependence is created intentionally. In the case of the plant, the goal is to create safety, while in the case of the legal domain, the goal might be to do justice in a rather efficient manner or to attain the goal of

due process. Below, I discuss two subjects. First, I explain what doxastic voluntarism is, and how this position relates to the different types of beliefs that can be distinguished based on the origination sources of beliefs (subsection 2.6.1). Second, I discuss the concept of epistemic dependence, explaining whether epistemic dependence threatens the knowledge qualification of beliefs (subsection 2.6.2).

2.6.1 Freedom of belief acquisition: doxastic voluntarism

The idea that the acquisition of beliefs is a process guided by rational considerations is called doxastic voluntarism (Van den Hoven 1995, p. 87*ff.*). The idea is that a person's beliefs are based on decisions to believe, rather than on spontaneous, automatically acquired information. If I see a dog cross the street, then I first see it, and then decide to believe that I see a dog cross the street. However, this is false in most cases of perceptual belief acquisition. We acquire beliefs without intermediary checks whether they are true or justified, or not. For most beliefs, it would be ridiculous to verify them additionally, as they are of minor importance, or they should be reacted upon immediately.

What is more important, is the fact that most non-perceptual beliefs, such as interpretative beliefs, are also acquired in a way that does not comply with doxastic voluntarism. Interpretative beliefs are often formed as a result of an unconscious process, unlikely to be governed by any voluntary act. Reasoned beliefs and testimonial beliefs are more likely to be based on voluntary acts. In case of reasoned beliefs, this is the act of reasoning. In case of testimonial beliefs, this is the act of believing what someone says or writes. The results of explicit reasoning steps are easy to reject or to accept. The same goes for testimony. However, most people would be inclined not to cast doubt on information they receive from certain sources, such as people and newspapers they trust.

Van den Hoven (*ibid.*) lists the arguments that defeat doxastic voluntarism, and he concludes that the adoption of beliefs is only under voluntary control insofar as the control concerns the environment in which the beliefs are acquired. This means that deciding to believe something is a matter of long-term planning of belief-acquisition conditions rather than an immediate act of will. I would like to challenge this observation, in that it is restricted to certain types of beliefs: (almost) all perceptual beliefs, and most non-perceptual beliefs, are acquired involuntarily. However, there are cases in which testimonial beliefs and interpretative beliefs are based on voluntary will. This leaves open the possibility of partial doxastic voluntarism, a desirable situation with respect to the legal domain, in which we want to maintain the autonomy of the judge.

2.6.2 Epistemic dependence

Epistemic niches tend to further epistemic dependence. This means that a person who operates in an epistemic niche has to rely on certain predetermined sources of knowledge. Epistemic dependence occurs if someone relies on other people's knowledge rather than on knowledge he acquired himself. In such cases, testimony is the knowledge source. Testimony is a suitable knowledge source if one has good reasons to believe the other person (Hardwig 1985, p. 336). To explain this principle, take the following example. If an expert witness testifies that a suspect's genetic fingerprint matches with the genetic fingerprint of a hair found on the victim's body, he should have good reasons to do so. These reasons might be his skills in doing this type of research, his knowledge of the equipment used, the reliable procedures employed *etc.* The judge may have good reasons for believing the expert witness, for instance the fact that he is known to be a reliable witness (his testimonies were never defeated), the explanation of his findings, and similar findings of another expert witness. So the judge has good reasons for believing that the expert witness has good reasons for believing that the genetic fingerprints match. Both reason sets may be entirely different, so that epistemic dependence arises. The judge epistemically depends on the expert witness. What, then, is the status of the beliefs he acquires through these channels? Following Van den Hoven (*ibid.*, p. 105), the question is: "But is it knowledge?"

Good reasons for believing someone else's beliefs may improve their truth-conduciveness. The epistemic dependence may, however, block a belief's qualification as knowledge. The question is whether it is sensible to distinguish between real knowledge (acquired directly by an individual) and testimonial knowledge (acquired through other sources), in a society where many beliefs are acquired indirectly through other sources of knowledge. If knowledge qualification is restricted to directly acquired perceptual knowledge, this is a denial of the common use of the predicate. The whole point of using the predicate 'knowledge' to qualify beliefs is to express the special nature of the belief under scrutiny.

My conclusion is that knowledge qualification cannot reasonably be denied to all beliefs arising in epistemically dependent situations. Such would be to deny the relevance of any kind of quality check mark for beliefs acquired in such situations.

2.7 The viewpoint of general epistemology

In this chapter, I discussed the characteristics of knowledge about the legal domain, given the viewpoint of general epistemology. I did so in order to provide an answer to the first research question. Characteristics of knowledge can be given by exploring three dimensions of knowledge: its acquisition, its object, and its justification. The same is valid for legal knowledge. The three dimensions help to explain the difference between regular knowledge and legal knowledge. Epistemological distinctions, knowledge criteria and beliefs can also be explained and typified in terms of these three dimensions. In the next chapter, they help to explore the different views that exist on knowledge from a legal epistemic viewpoint.

3 Legal knowledge from a legal epistemic viewpoint

This chapter is concerned with legal epistemology. Legal epistemology is the branch of general epistemology that is involved with developing an account of the nature of legal knowledge. I discuss this issue to find out the answer to the second research question, which amounts to finding the characteristics of knowledge about the legal domain from the viewpoint of legal epistemology. In order to give such an account of legal knowledge, I subsequently discuss the following subjects. First, in section 3.1, I provide an explanation of the hermeneutic tradition in legal philosophy. Then, in section 3.2, I outline the epistemic claims that can be inferred from some other important views in legal philosophy. Next, in section 3.3, I explain different sources of legal knowledge. Finally, in section 3.4, I make some concluding remarks on the issue. However, I first put some effort into explaining the present enterprise of discussing the concept of legal knowledge from a legal epistemic viewpoint.

The relevance of knowledge about the legal domain is twofold. First, there is the utility of knowledge. Compared to mere belief, knowledge can be used as a reliable ground for behaviour. It may also serve as a means of gaining authority relative to those only having belief. In the legal domain, knowledge provides grounds for authoritative decisions. Rather than basing one's decision on relatively unreliable beliefs, the basis for one's inferences should be knowledge.

Second, attaining knowledge is a goal that is worth aiming at as such, regardless of its utility. An argument with this content is put forward by Finnis (1980, p. 59-80). He claims that the pursuit of knowledge is a value, in the sense of a good: a goal that is worthwhile independent of any further utility in the achievement of survival, power, and popularity. The value of attaining knowledge is a principle of practical reasonableness, Finnis claims. It provides us with a direction in which we can lay out lines of argumentation. It can be used to generate new principles, and to direct the application of rules. In his discussion of the value of knowledge, he emphasises the importance of truth. Having knowledge presupposes truth, whereas beliefs can be true or false. Knowledge and truth are very close relatives, if we may regard the following quotation as representative of Finnis' opinion on the matter (*ibid.*, p. 61):

> "In explaining, to oneself and others, what one is up to, one finds oneself able and ready to refer to *finding out*, *knowledge*, *truth* as sufficient explanations of the point of one's activity, project, or commitment."

Finnis regards the value of knowledge as a self-evident principle. He asserts that self-evidence of some principle has little or nothing to do with our feelings of certitude about that principle. Rather, he claims, the self-evidence of a principle shows itself in its employment as a criterion for the assessment of feelings. A principle such as worthiness of knowledge cannot be proved. It can be adopted, though, on the assumption that its employment is fruitful, or rather, that, if it is not adopted, rational discourse becomes hard or impossible.

In sum, legal knowledge is a better starting point for making inferences (the utility argument), and it is worthwhile in its own (the value argument). I adopt the two arguments as backings for a further investigation into the concept of legal knowledge from a legal epistemic viewpoint in the sections below.

3.1 Object and subject intertwined: hermeneutics in the law

The dominant theory about judicial interpretation in The Netherlands in the second half of the 20ᵗʰ century has been Scholten's (1974) theory, which is inspired by hermeneutics. This section provides an elaboration on hermeneutic theories in the law. First, I provide some general characteristics of such theories (subsection 3.1.1), then I discuss an example of a hermeneutic view, namely Dworkin's theory (subsection 3.1.2). The purpose of the current section is to show that often, the seemingly clear-cut distinctions between categories of legal phenomena (for instance legal objects) cannot be maintained in every legal-theoretical stance. Actually, some of the central distinctions made in this thesis are challenged by hermeneutic theories of law, for instance the distinction between epistemology and ontology. This has a profound effect on building ontologies (conceptual models) of law; such ontologies presuppose basic ontological and epistemological choices.

3.1.1 Hermeneutic theories of judicial interpretation

Hermeneutic theories of law form the main stream in legal theory. Having said that, the differences among those theories are significant, and Dworkin's theory, discussed in the next subsection, is one with relatively few relativistic connotations, *i.e.*, the ontological and epistemological claims are stronger than one might expect in a hermeneutic theory. In the following discussion, ontological and epistemological claims in hermeneutic theories of law are discussed. According to Smith (1998, p. 224), in hermeneutic theories of the law, the law is defined as an interpretation of a practice. The interpretation should lead to an identification of what the law is. To clarify the nature of hermeneutics, I list

three distinctions that cannot be held in such theories: there is no clear distinction (1) between the descriptive and the normative, (2) between rules and facts, and (3) between the law and its interpretation.

The descriptive/normative distinction typically signifies the difference between what there is and what there ought to be. To describe something means to list its properties, while we assume that these properties are independent of their description. To prescribe something is to say that something ought to be the case, independently of the question whether it *is* the case: there is a norm, with which reality must comply. Hermeneutics says that such a distinction cannot be maintained, because we cannot describe reality independently of some idea about what reality should look like. In a reflection on the law, the descriptive/normative distinction means that there is no clear distinction between a rule of a law and a fact it applies to, but that there is an interaction between rule and fact, which yields a new characterisation of both the applicable rule and the fact. In the light of hermeneutic theories, brute facts (facts that need not be perceived or interpreted by a human agent) constitute an empty notion, because facts can only be characterised meaningfully in the light of a rule, and a rule can only be formulated meaningfully in the light of the fact.

The third distinction, namely between the law and the interpretation of the law, is related to the second distinction. However, in the third distinction, we regard the law as independent of its application. The distinction between the law and the interpretation of the law assumes that the law exists, and that its interpretations are about that independently existing law. Hermeneutics denies that this distinction can be made. In practice, the law is only 'found' or 'formed' in its interpretations. When we interpret the law without referring to individual cases, we systematise the law. Systematisation requires us to form an idea about the meaning of the individual rules of the law and their connections with each other, and to find out structural distinctions and patterns. Systematisation cannot be seen as an activity that leaves the original rules undisturbed. Instead, the systematisation activity is also an interpretation activity that establishes the content of the law.

Now that I have given a characterisation of the distinctions that can*not* be made in a hermeneutic theory of law, I give some properties of such a theory with respect to a concept of knowledge. Knowledge about the world, or about the law, is acquired through a form of *understanding* (Verstehen, *cf.* Smith 1998, p. 61). The subject/object distinction – the distinction between the person who knows something and the thing he knows something about – is denied. Before one can know anything, one has to have some form of knowledge about the

world; one has to be a part of that world. Any acquired knowledge is related to the position one has in the world; the knowledge that one has previously acquired. Thus, knowledge is neither completely subjective, nor completely objective. The same is valid for the meaning of language. The language one uses is related to the particular stance a person has in the world, and there is no objective relation between a sentence and some part of the world. Therefore, there is no such thing as objective truth either. The fact that all knowledge is the result of interpretation, and that this interpretation constitutes a part of reality, means that an objective distinction between a sentence and the fact that the sentence refers to cannot be established (*cf.* Smith 1998, p. 68).

3.1.2 Dworkin: constructive interpretation

In his book 'Law's Empire', Dworkin (1986, p. 45-86) gives an account of interpretation in the law. This account is a hermeneutic one, because it regards interpretation as a *constructive* activity. In order to explain the concept of interpretation, Dworkin distinguishes among four types of interpretation (*ibid.*, p. 50-51): conversational interpretation, scientific interpretation, artistic interpretation, and interpretation of social practice. Conversational interpretation occurs when we talk and listen to each other, and try to find out what other persons mean. Scientific interpretation is deemed not to be purposive, *i.e.*, the raw data should speak for themselves, and the interpretation of those data is led by their content. Artistic interpretation consists of attempts to discover the meaning of a work of art, which is projected in that work by the interpreter. The interpretation of social practice covers legal interpretation. It resembles artistic interpretation to the extent that it is about an object *created by* human beings, but *distinct from* those human beings. Thus, the types of interpretation are defined relative to their objects. People's words are the objects of conversational interpretation, nature's products are the objects of scientific interpretation, and people's products are the objects of artistic interpretation and interpretation of social practice. The agreement in object of artistic interpretation and interpretation of social practice leads Dworkin to accommodate them both under the denominator 'creative interpretation' (*ibid.*, p. 50).

The interpretation of social practice is an activity that scrutinises a certain case. It starts with a preinterpretive demarcation of what is the issue at hand: we have to find out what are the rules and the standards that determine the tentative content of the practice to be interpreted (*ibid.*, p. 65-66). Then the actual interpretation is carried out, by isolating those arguments that contribute to the justification of the main elements of the practice (*ibid.*, p. 66). Finally, in the

postinterpretive stage, an attempt is made to incorporate the deviant characteristics of the original practice into a new version of that practice that fits the justification established in the second stage (*ibid.*, p. 66). Dworkin expresses the view that the preinterpretive stage serves to isolate the practice, the interpretive stage serves to construct the interpretation (where the main direction of reasoning is from interpretation to practice), and the postinterpretive stage serves to reconstruct the practice (where the main direction of reasoning is from the practice to the interpretation, altering that interpretation if necessary).

Most notable in Dworkin's view on interpretation, is that he rejects scepticism. Interpretations can be "right or wrong", even "true or false" (*ibid.*, p. 78). Although he adheres to a constructive type of interpretation, where the interpretation does not necessarily reflect the intentions of (members of) a community, he claims that one interpretation can be better than another. According to Dworkin, there are two sceptical stances: internal scepticism says that there is no way to determine whether one interpretation is better than another, in fact, there is no way that an interpretation can earn priority over others, because no actual interpretation provides the unity required for such an enterprise (*ibid.*, p. 78). External scepticism, instead, makes an ontological claim, namely that there is no way to test the correctness of an interpretation, as there is nothing in reality it can be tested against (*ibid.*, p. 79-80). Dworkin emphasises that many sceptic attacks to claims on priority relations among interpretations (one interpretation is better than another) employ externally sceptic arguments, but make internally sceptic claims. However, the claim that there is no independent reality to which interpretations can be checked is quite a different claim from saying that no interpretation can be deemed better than another (*ibid.*, p. 83). Dworkin does not challenge the former claim, but he does dispute the latter one.

3.2 Knowing the law: epistemic claims in legal theory

The current section contains a discussion of epistemic claims that can be derived from five different theories about the law: natural law (subsection 3.2.1), legal positivism (subsection 3.2.2), institutional legal theories (subsection 3.2.3), a conventional-*cum*-institutional approach (subsection 3.2.4), and hermeneutic theories of law (subsection 3.2.5). These theories address the question what the law is: this ontological question should be answered before we can start making epistemic claims. Therefore, for each of the theories mentioned, I discuss the basic characteristics of its concept of law. Then, I explain under what conditions we can say that we know the law, based on that concept of law.

3.2.1 Natural law

Cliteur (1989, p. 363) provides six characteristics of the natural law view. I discuss them briefly below. First, there is law that has absolute validity. Second, this law can be derived from human nature or the nature of reality. Third, human reasoning capabilities are sufficient to know this law. Fourth, the validity of positive law depends on the test of content constituted by natural law. Fifth, positive law that does not pass this test is invalid. Sixth, the test consists of a set of speculative and metaphysical ideas. In a natural-law view, the epistemological claim is that abstract 'knowing that' of natural law, *i.e.*, abstract factual knowledge of legal norms, is possible. Concrete 'knowing that' can be established through knowledge of legal norms and knowledge of cases (knowing a case in the light of the legal norms that apply). A claim about abstract 'knowing that' of positive law is not made, but we may assume that such knowledge can be acquired.

Whether positive law is really valid law, can be known by testing positive law against natural law. This test consists of checking whether the content of positive law complies with the principles expressed in natural law. Thus, knowing the law from a natural-law perspective should be divided into knowing natural law and knowing positive law. Knowing natural law means knowing the abstract principles that natural law consists of. Knowing positive law means knowing the rules and norms that are part of the human-constructed legal system. Knowing the law then amounts to knowing those rules and norms of positive law that comply with the principles of natural law.

3.2.2 Legal positivism

The legal positivist view is summarised by Brouwer (1997) in five characteristics. First, authority is the source of the law. The validity of a norm as a legal norm is based on a human source. Second, the validity of a legal norm does not depend on the contents of that norm. Third, there is no necessary connection between the law and morality. The question what is law does not depend on what is morally justified. Fourth, legal norms claim to have the highest authority; norms that are not recognised in the legal system are legally irrelevant. Fifth, the meaning of a legal norm is a factual matter, not a moral one; it refers to the meaning the authorities provided that norm with.

In case of legal positivism, obtaining knowledge of valid law is possible through acquiring knowledge of the formal sources of law and knowledge of their being issued by the proper authorities. In the natural law view, the 'knowledge check' is performed by comparing positive law with moral principles, whereas in the legal-positivist view, the valid origination history of positive law

is the only criterion for the 'knowledge check'; knowledge of positive law can be obtained through knowledge about its being issued validly.

3.2.3 Institutional legal theories

Institutional theories of law give, according to Hage (1998), accounts of different themes in legal philosophy. However, they do have something in common, namely their primary point of departure. Institutional theories assume that legal reality supervenes the world that consists of material facts. Legal reality, however, unites the factual and the normative in its supervenience (*ibid.*, p. 141-142). An institutional theory of law consists of a framework describing the relation between brute facts and institutional legal facts. Brute facts are the object of traditional ontology, institutional legal facts are claimed to be the entities gaining existence through a system of rules. These legal facts are the object of legal ontology. Below I discuss two institutional theories of law.

First, the institutional theory of law developed by MacCormick and Weinberger (1986) claims that legal institutions, such as property and marriage, can yield institutional legal facts under certain conditions. These facts are brought about or terminated by way of institutive rules and terminative rules. Moreover, by means of consequential rules, their (legal) consequences are defined. Institutional legal facts exist in time between the point where they are established by an institutive rule, and the point where they are abolished by a terminative rule. In the meantime, the consequential rules hold for those facts. The existence of legal rules themselves can also be considered as institutional fact (*cf.* subsection 4.2.3).

Second, Ruiter (1993) proposes an institutional theory of law. In this theory, he emphasises the importance of speech act theories (such as in Searle 1969) for the establishment of institutional legal facts. Ruiter's theory covers the creation and abolishment of legal norms, based on the different types of speech acts distinguished by Searle. The idea is that a legal system can be analysed in terms of speech acts. These are acts that occur when people say certain things. For instance, when a civil servant says 'I hereby declare you husband and wife' under appropriate circumstances, these words have the result that the two people in front of him are married. Legal acts can be analysed in terms of such speech acts. In this example, the speech act implies a declarative legal act. A speech act can succeed or fail. For instance, when the person performing the speech act has no right to do so, the speech act fails. In that case, no legal act is performed (at least not the one leading to a marriage). Alternatively, the speech act is performed by the proper person, but the community in which the marriage has taken place,

does not recognise it. In that case, there is a legal act, but the necessary recognition of that legal act does not occur. The validity of legal norms thus depends on two characteristics; an 'ideal' characteristic and a 'factual' characteristic, where the former is achieved through the proper performance of legal acts, and the latter through social recognition of the norm (see Hage's discussion of Ruiter; Hage 1998, p. 137).

From the viewpoint of institutional legal theory, knowing the law requires knowing positive law, and knowing whether positive law complies with the specific demands that are set by institutional rules. In a legal-positivist approach, we have to know whether a rule was issued by a proper authority. In institutional legal theory, this demand is further specified by scrutinising the specific rules with which the validity of positive law is attained from the moment it is issued. In both cases, the validity of the knowledge obtained depends on its conformity with its object.

3.2.4 A conventionalist-*cum*-institutional approach

In Peczenik and Hage (2000), a conventionalist-*cum*-institutional (CI) theory of law is taken as a starting point for a discussion of the ontology of law. The CI theory fits in with the need for a more precisely defined object of legal knowledge, as it offers a detailed account of the different types of fact in the legal domain. The theory can be discussed by an explanation of three concepts on which the theory is based: supervenience, conventional fact, and institutional fact.

The first concept that is adopted in the CI theory is supervenience. Hage (1998, p. 128) defines supervenience as follows. A set of characteristics A is supervenient relative to another set of characteristics B if there cannot occur a change in set A without the occurrence of a change in set B, whereas there can occur a change in set B without the occurrence of a change in set A. Thus, while the supervenient characteristic (say, the colour red) can rest on different phenomena (a red sheet of paper versus a white sheet of paper illuminated by red light), the supervenient characteristic is always the same when the set of characteristics on which it rests remains the same (a red sheet of paper remains red).

The second concept adopted is conventional fact. The existence of a conventional fact rests on the conditions that a sufficiently large number of people within the relevant community believes that fact, and also believes that everyone believes that fact. In addition, the situation in which the relevant community believes that fact is partly reason for the people belonging to the community to perform actions. Those actions are meaningful because of the existence of the conventional fact (Peczenik and Hage 2000).

The third concept adopted is institutional fact. An institutional fact derives its existence from the following conditions. There should be a rule in the relevant population, of which the conditions are fulfilled, and there is no exception to that rule, and the conclusion of the rule is that the fact exists (*ibid.*).

Together, conventional facts and institutional facts can account for the structure of a legal system. The definition of an institutional fact can be applied in a recursive manner. The definition takes conventional facts, institutional facts and brute facts as its input, and produces conventional facts and institutional facts as its output. On their turn, the resulting facts may again be used as an input to the definition of an institutional fact. Rules, according to Peczenik and Hage (*ibid.*) may be considered as facts. Thus, rules can exist, either as conventional facts, or as institutional facts, or as both.

In the CI theory, knowledge of the law can be acquired by knowing positive law, and knowing whether positive law complies with the institutional and conventional demands imposed by the CI theory. The CI theory provides the means to describe the structural properties of a legal system, *i.e.*, to explain how facts and rules are interconnected. Suppose, for instance, that a legal rule is only valid if it is issued by a proper authority and if it is recognised by a population. In that case, knowledge of that rule can be obtained by knowing, in addition to its content, the fact that it is issued by a proper authority, and knowing the population's acceptance of the rule.

3.2.5 Hermeneutic theories of law

Hermeneutic theories have a different approach towards knowledge. As I explained in section 3.1, in these theories, there is no clear demarcation line between the knowing subject and the known object. As a consequence, the distinction between knowing that and knowing how is somewhat vague. In the hermeneutic view, knowledge is always the product of an interaction between the subject and the object of knowledge. Therefore, factual knowledge (knowing that), and a competence to do something (knowing how), are both partially based on knowledge that was already present. The epistemic claim that can be made about legal knowledge in general is that such knowledge is always dependent on the background knowledge one has. Legal knowledge cannot be acquired without some sense of what the legal domain amounts to. The process of acquiring such knowledge is a circle in which new knowledge is assessed from the perspective of the knowledge that is already present.

3.3 Sources of legal knowledge

The term 'legal knowledge' refers to all knowledge that is about the law. In the three subsections below, I clarify the sources of such knowledge. Such sources are content sources, *i.e.*, they provide the content of legal knowledge. I discuss the distinction between material and formal sources of law (subsection 3.3.1), and I deal with knowledge sources for the law (subsection 3.3.2). Finally, I explain in what sense reasoning can be regarded as a source of knowledge (subsection 3.3.3).

3.3.1 Formal and material sources of law

Sources of law are the starting point for the acquisition, *qua* content, of legal knowledge. Depending on the ontological stance taken, knowledge of the law may be derived directly or indirectly from sources of law. The concept 'source of law' is ambiguous. There is a distinction between formal and material sources of law. Formal sources of law are, as some authors state (*cf.* Algra and Van Duyvendijk 1989, p. 19), the knowledge sources for the law. These are statute law, treaties, legal precedents (*ibid.*), and customary law.

Material sources of law are the origination sources of law, *i.e.*, the factors on which the content of laws is based. Thus, material sources contribute to the drafting of new laws; they help determine their content. Material sources of law do not themselves amount to legal rules or legal norms. They form, however, the grounds for those rules and norms. Algra and Van Duyvendijk (1989, p. 20) distinguish the following material sources of law: political powers, civil servants, pressure groups, religious beliefs, moral beliefs, socioeconomic developments, geographical circumstances, and technological developments.

3.3.2 Knowledge sources for the law

Knowledge sources for the law are the sources through which we acquire knowledge about the law. There are two properties of having factual knowledge about the law: knowing its *content* and knowing its *validity*. As noted earlier, knowledge sources for the law are sometimes regarded as equivalent to the category of formal sources of law. However, I would prefer to make a clear distinction between formal sources of law and knowledge sources for the law.

In a legal–positivist stance, formal sources of law coincide with the sources of valid law. In such a stance, if one acquires knowledge of the formal sources of law, one will acquire knowledge of both the *content* and the *validity* of law. In a natural–law stance, this is not necessarily the case; the validity of law is determined by principles that are not part of the positive legal system. In such a

stance, knowledge about the content of law is derived from the formal sources of law and from principles outside law.

In legal–positivist and natural–law stances, knowledge about the content and validity of the law may also be derived from *material* sources of law. Material sources of law do not provide knowledge about valid law directly, but they are helpful in interpreting the content of formal sources of law. Both formal sources of law and material sources of law can thus function as knowledge sources for the law. Their precise role may differ, depending on the specific view on the law taken: in a legal–positivist stance, material sources of law may help explaining how certain formal sources of law acquired their specific content, and in a natural–law stance, material sources of law may provide part of the demarcation line between what is morally right and morally wrong, and thus, which part of law is valid and which is not.

3.3.3 Reasoning as a source of knowledge

Hage (1997) elaborates on the concept of reasoning in the legal domain, in which proper justification plays a central role. He states that there are two stages in legal reasoning. In the first stage, there is a set of principles. Principles have a more general nature than rules; they are less concrete. Principles are checked for their validity. If the conditions of a principle are satisfied, and the principle is not excluded, it can be used to generate reasons for or against the conclusion of an argument. These reasons for and against an argument are, on their turn, weighed against each other, and then lead to the conclusion of an argument. Approximately the same two-step process applies to reasoning with rules. The reasons resulting from the two two-step processes can be used in order to support the conclusion of a legal rule.

Hage's (1997) account of the nature of legal rules takes into account the aspect of defeasibility. In the case of legal rules, this means that a rule whose application conditions are satisfied, may still not be applicable. First, there are scope limitations for a legal rule, for instance that it applies only to facts that have taken place on Dutch territory. Second, exceptions exist which are formulated in the legal rule itself. Third, there are proper exceptions to legal rules, for instance the grounds for justification in criminal law. Instead of the traditional model of legal reasoning, in which a rule applies whenever all conditions of the rule are satisfied, Hage focuses on the justification of the conclusion of the rule. Legal reasoning is about giving reasons for a conclusion, and legal rules may (as a consequence of defeasibility at several levels) fail to generate their conclusion under certain circumstances.

Reasoning constitutes a source of legal knowledge, because it can generate new legal knowledge. In that respect, reasoning in the legal domain is a rational construction or a rational *re*construction (depending on one's ontological view) of the way in which, for instance, a legal decision is realised. Note that much of the actual work is done by an unconscious process which I call interpretation. A rational construction of the reasoning process underlying such a decision consists of a line of argument that reaches the actual decision, while a rational *re*construction should in some way reflect the actual steps leading to that decision; some kind of analogy should exist between the actual interpretation process and the rational reconstruction of that interpretation process. Clearly, the dividing line between the two is thin.

3.4 Combining the general and legal epistemic views

The previous chapter was based on the first research question. This question was: What are the characteristics of knowledge about the legal domain, given the viewpoint of general epistemology? The current chapter is based on the second research question. This question was: What are the characteristics of knowledge about the legal domain, given the viewpoint of legal epistemology? With respect to the first research question, I have explained the dimensions along which we can assess knowledge: its acquisition, object and justification. Moreover, I have outlined distinctions along which theories of knowledge can be classified, criteria used to assess whether something may be called knowledge, distinctions along which types of knowledge can be classified, and the assessment of the freedom with which knowledge is acquired. With respect to the second research question, I have discussed the intertwinement between the object and subject of knowledge, the way in which we can know the law from different legal-theoretical stances, and the sources of legal knowledge. The various distinctions made in chapters 2 and 3 are classified as related to one of the three main dimensions: acquisition, object and justification in table 3.1.

	acquisition	*object*	*justification*
distinctions in the theory of knowledge	doxastic - non-doxastic internalist - externalist		
distinctions regarding knowledge		legally-relevant - legal abstract - concrete regular - moral practical - factual	
knowledge criteria	reliability	truth	proper justification coherence
sources of belief and knowledge	belief sources	formal sources of law material sources of law	formal sources of law material sources of law

Table 3.1. Overview of distinctions regarding belief and knowledge

Given the main dimensions, we can immediately see that studying the object of knowledge is an important part of studying knowledge. However, it seems that knowledge of the law and the law itself are hard to distinguish. Hermeneutically-inspired theories are based on this assumption. If knowing and being are intertwined, then there is a bidirectional relation between the two, in which knowing constitutes being, and being constitutes knowing. The circle that seems to arise from this, is only a circle at first sight. If we refrain from taking a hermeneutic stance, we can find a clear difference between the way in which new legal entities are constructed, and the way in which knowledge about these entities is acquired. For instance, the realisation of a legal decision by a judge is an implicit process in which many types of knowledge are applied. The result is a decision, about which knowledge can be acquired through a different route. Rather than a circle, knowing the law constitutes two entirely different processes, pointing in opposite directions: the process of interpretation and reasoning leading to legal conclusions, thus establishing legal entities (from the epistemic to the ontological), and the process of obtaining knowledge of the law (from the ontological to the epistemic).

Anticipating on the next chapter, I would like to stress that there are two distinctions of crucial importance to disentangle knowledge of the law and the law itself. First, of course, there is the distinction between legal epistemology (discussed in the current chapter) and legal ontology (discussed in the next chapter). In the previous paragraph, I explained that the relations between these

two levels can be isolated as unidirectional processes that do not cause a circularity problem. Second, there is the distinction between the context of discovery (the acquisition and object dimensions of knowledge) and the context of justification (the justification dimension of knowledge). It seems that in the legal domain, the circularity is found here. Instead of discovery alone, acquiring legal knowledge involves justification as well. In the next chapter, these distinctions are further elaborated on in the discussion of different stances in legal ontology, and the clarification of the object of legal knowledge in the light of these stances.

4 Legal ontology

In this chapter, legal ontology is regarded as a branch of general ontology. Thus, following the subject matter of general ontology, legal ontology investigates the existence of elements in the legal domain. More specifically, it concerns the existence of such entities as laws, norms, and principles. But as it is about the legal domain, it should also regard non-legal entities that are relevant to this domain. Views on legal ontology are influenced by the philosophical insights into the law in general. Practising legal ontology means examining what objects there are in the legal domain, how they are brought into existence, how they relate to each other, and how they change through time. The results of this examination provide a part of the basis for the knowledge-based ontology in chapter 6.

In order to be able to express ontological claims from different theories of law, I first need to explain some basic concepts of ontology. I explain these concepts in section 4.1. Ontological claims (claims about the existence of individual objects or object types) depend on the theory of law adhered to. In section 4.2, I list the ontological claims of five theories of law, and I compare these claims. In legal ontology, we may discuss objects we encounter in the legal system, or we may discuss the legal system as a whole. Therefore, in section 4.3, I discuss two object types that play an important role in the legal domain (systematisations and interpretations), and in section 4.4, I explain the ways in which legal systems can exist. In some theories of law, argumentation constitutes the construction of legal entities. Hence, in section 4.5, I discuss elements of argumentation. In section 4.6, I give an overview of four ontological status layers that we can distinguish in the legal domain. Finally, in section 4.7, I explain how two relations (counting as and causality) can be used to represent different views on the existence of the law.

4.1 Introduction to ontology

Ontology is the discipline that investigates the nature of being. Two questions are central to ontology: first, what does it mean to exist?, and second, what things do exist? As a preliminary answer to the first question, I propose: the existence of a thing is its presence independent of the perception of human beings. The answer to the second question obviously depends on the way in which we define the nature of existence. The preliminary answer to the first question yields the following dilemma. Each attempt to make claims on reality depends on our ability to know that reality. In order to say something about the inde-

pendent existence of things, we need our senses and our brains. Thus, statements about reality are not independent of our ability to know.

4.1.1 Realism

Realism is a philosophical position that focuses on existence alone. A typical form of realism claims the existence of tokens of certain scientific or common-sense types (*e.g.*, atoms), independent of the mental (*cf.* Devitt 1991, p. 303). Realism is nothing else but that claim; it says nothing about whether tokens or types are knowable. For this reason, it also endangers itself: we cannot reasonably make any ontological claim without the knowledge needed to make that claim. Thus, a realist position demands an epistemology that enables us to know the world (at least partly). If we can establish such an epistemology, the realist doctrine is to be preferred, because, in combination with a realist epistemology, it enables us to say something about the world as it really is.

If we claim the independent existence of certain parts of reality, we make such claims, for instance, with regard to the object of our scientific or common-sense knowledge. If we claim that scientific entities (such as electrons) exist, a scientific realism emanates from that claim. If we claim the existence of common-sense objects (such as chairs), this yields a common-sense realism. To specify such types of realism, we have to include claims that regard the problem of universals. The problem of universals is the question what general terms in natural language (such as 'judge' and 'car') refer to. Concrete things are called 'tokens' or 'particulars', abstract things are called 'types' or 'universals' (in some positions, there is a difference between tokens and particulars, and between types and universals, but I use them interchangeably). Different views on the problem of universals lead to different forms of realism.

Using the type/token-distinction, we can make general ontological claims. We can either claim the existence of tokens, of types, or of both tokens and types. For instance, if we claim the existence of the type 'judge', we say that the term 'judge' refers to an objectively existing universal. If we claim the existence of tokens of the type 'judge', we say that specific judges exist objectively. And if we claim the existence of both the type 'judge' and tokens of this type, then we say that there are both an objectively existing universal and objectively existing judges.

4.1.2 Alternatives to realism

In philosophy, there are more theories about reality than just the hard-boiled kind of realism described before. The two most important alternative paradigms

for realism are constructivism and relativism. I regard both doctrines as anti-realist, in that they deny us the possibility of saying something about the objectively existing reality. Although constructivism claims that there is a reality, it also says that we are not able to investigate it in order to determine whether this reality contains certain structures. Reality is comparable to a Kantian 'Ding an sich'. People can create their own 'realities' by projecting structures into the 'Ding an sich'. Such 'realities' emanate from an epistemological source (*cf.* Putnam's internal realism for an example of constructivism; Putnam 1990). Relativism is a denial of the possibility to know anything at all about the real existence of things. Relativism does not make an ontological claim, but an epistemic one (we cannot know what really is). If we take a relativist stance, we will lose grip on reality as well as on our knowledge.

Realism, constructivism, and relativism cannot easily be placed on one scale, because they make claims in different philosophical disciplines: realism is primarily an ontological doctrine, while constructivism and relativism make epistemic claims. Ontology, epistemology, and semantics are, however, related to a high degree. Semantic theories depend on both ontological commitments and epistemology for the relations they describe between symbols and entities in reality, ontology cannot be practised without making epistemic commitments and *vice versa*. Ontological claims should be argued for by epistemic means, for humans can only gain knowledge about anything at all through epistemic means. Epistemic claims should be grounded in ontological commitments, for otherwise knowledge is not based upon reality. This is why a strict distinction between ontology and epistemology is difficult to maintain; but it is important for explanatory purposes.

4.1.3 Rules and norms

The difference between rules and norms has to be explained prior to discussing the ontological claims in different theories of law. I regard rules as entities that define a consequence of the occurrence of an object. Their general structure is: 'if there is an x, then there is a y', although, of course, their precise formulation may diverge. Rules may, for instance, qualify an occurrence of an object (an act, an entity, a fact) in a certain way, such that it gains a new status layer, or they may define the consequence for some kind of behaviour. Examples of rules are 'killing on purpose counts as manslaughter', and 'a person who commits manslaughter is punished with a prison sentence of at most fifteen years'.

I regard norms as entities that explicitly guide some kind of behaviour or action. The presence of a norm can be identified by the normative modalities it

contains, for instance ought-not, ought-to-do, or can. Norms can have a broad or narrow range of application. For instance, general norms such as 'one ought not kill other people' are widely applicable, whereas 'people working on the premises of the nuclear plant ought to carry a radiation detection card' is applicable to only a limited number of employees of the plant.

Whether rules and norms are legal often depends on the legal-theoretical stance taken with respect to the identification of law. For instance, if authoritative recognition is the main criterion, a rule is only a legal rule if it is issued by a proper authority. The difference between a legal rule and a legal norm is not always entirely clear. It is often more a matter of precise formulation than of the precise content. For instance, we may very well take a rule like 'a person who commits manslaughter is punished with a prison sentence of at most fifteen years' as implying the norm 'one ought not kill other people'.

Although rules and norms constitute two different categories, a rule may contain an embedded norm, or *vice versa*. A rule that says 'a person who commits manslaughter ought to be punished' contains an embedded norm 'ought to be punished'. A norm that says 'one ought to qualify killing on purpose as manslaughter' contains an embedded rule saying that 'killing on purpose counts as manslaughter'. This explanation of the difference between rules and norms may seem a bit artificial, but it helps to maintain a distinction between two different functions: the qualification task (ascribed to rules), and the behaviour-guiding task (ascribed to norms).

4.2 Ontological claims in different theories of law

Belief and knowledge are usually *about* something. This means that there is some state-of-affairs (or a property or relation) to which a belief is compared for the determination of its correctness. In other words: we may say that a belief is correct if it expresses something that is actually the case. But are there objects that the various entities we call 'the law' refer to? One of the issues in the natural law-legal positivism debate is precisely the ontological status of law: what are we talking about when we refer to rules and norms? In section 3.2, I discussed different theories of law, and the corresponding epistemic claims. In the present section, I explain the ontological claims that can be attached to natural law, legal positivism, institutional legal theories, the conventionalist-*cum*-institutional approach, and hermeneutic theories respectively (subsections 4.2.1 through 4.2.5). Subsequently, these claims are compared (subsection 4.2.6). At the end of that subsection, I provide a summary of the claims in table 4.1. The current sec-

tion serves to show that different views on the law lead to different ontological claims, and to analyse the modes of existence distinguished in these claims (*cf.* section 4.6).

4.2.1 Natural law

The general ontological claim that can be derived from natural–law theories is that some norms objectively exist. Therefore, their existence does not depend on subjective factors, like societal conventions or personal preferences. These norms provide a framework for the assessment of positive law. The existence of positive law may also be claimed, but its existence status is less important than the existence of natural law. The validity of positive law depends on the content of natural law. There is a conceptually necessary between morality and the validity of the law (*cf.* Cliteur 1989, p. 363). If some law is morally objectionable, *i.e.*, if it does not comply with the moral principles embedded in natural law, then that law is not valid. The source of validity of positive law is thus found in natural law, in addition to conventions or institutional conditions.

The source of existence of positive law, however, may be different. If validity is not a necessary condition for the existence of positive law, it may exist in other ways, so that we may account for the existence of morally objectionable legal rules, for instance by their efficacy. To be more precise in stating the ontological claim, in a natural law view, norms that do not satisfy the criteria of natural law are not part of the system of positive law. Rather, positive law consists of norms and rules, whose (moral) validity is checked against the norms of natural law.

4.2.2 Legal positivism

An ontological claim compatible with the general characteristics of legal positivism is that positive law is valid whenever it is issued by an authority. The absence of a conceptually necessary connection between morality and positive law means that a morally objectionable law can still be a valid law (*cf.* Brouwer 1997). The only link between morality and positive law is that positive law will often be based on moral principles, and that interpretations of positive law will also be influenced by those principles. However, there is not a conceptually necessary link.

Both norms and rules are part of positive law – here a major difference with natural law is found. The validity of positive law is derived from the issuing authority. The rules of positive law can be translated into norms. These norms say what one may, ought or ought not do. They derive their normative power from the validity of the rules of positive law.

4.2.3 Institutional theories of law

The ontological claim that can be attached to institutional legal theories in general is that law's existence is supervenient on brute facts. The law consists of institutional legal facts that combine the factual dimension of the law with its normative dimension (*cf.* Hage 1998, p. 141-142). MacCormick and Weinberger (1986, p. 38) claim that criteria of existence depend on the type of object whose existence is at stake. This implies that to say a stone exists means something different from saying that a marriage exists. The assumption that the meaning of the predicate 'exists' varies resembles the assumption that is made in this thesis about the varying meaning of the predicate 'knowledge'. Both assumptions are indebted to Wittgenstein (1953) for his introduction of the concept of 'family resemblance', which places the meaning of natural-language terms in its actual use.

MacCormick and Weinberger distinguish three modes of existence of norms (*ibid.*, p. 40). These are all related to their existence 'in the realm of human consciousness'. First, the members of a norm-issuing authority are conscious of the fact that something ought to be the case. Second, the members of a community in which norms are issued, are aware of the obligatoriness of norms. Third, a member of a community may have knowledge of some norm (of the fact that something ought to be the case) without actually behaving in accordance with the norm. Norms influence behaviour. The mere awareness of the norm then constitutes part of its existence. Norms motivate people to behave in certain ways. People behave in those ways for different reasons, among which the voluntary acceptance of their normative character, the imitative instinct of people, behaving according to patterns that are institutionalised in law, and by the threat of coercion (*ibid.*).

In addition, the authors distinguish three types of rules with which legal entities are manipulated: institutive, consequential, and terminative rules. Institutive rules bring into existence some kind of legal entity (*e.g.*, a contract). Consequential rules define consequences attached to there being in existence some legal entity (*e.g.*, one has to comply with the conditions in a contract, if there is a contract). Terminative rules put an end to the existence of some legal entity (*e.g.*, a divorce puts an end to a marriage) (*ibid.*, pp. 52-53). An instance (token) of a legal institution (*e.g.*, a particular marriage) exists in time, as consequence of it being created by an institutive rule, maintained by a consequential rule, and ended by a terminative rule. The institutions (as types) do not need that kind of existence (*ibid.*, p. 53). Institutions exist prior to any of their instances. A legal system must contain an institutive rule in order to initiate the existence of an

institution *as a type* (*ibid.*, p. 55). MacCormick and Weinberger (*ibid.*, p. 9) do not articulate the kind of existence they assume for institutions as types. Although MacCormick compares such existence with the existence of Platonic ideas (*ibid.*, p. 55), he immediately puts this remark into perspective as an attempt at jocularity (*ibid.*, note 7, p. 75).

The existence of rules and norms within institutional theories of law can be explained in a way similar to the existence of instances of legal institutions. Formally created rules and norms will generally start to exist as a consequence of an institutive rule, establishing the existence (validity) of a specific rule or norm (*ibid.*, p. 14). In addition to this characteristic, MacCormick and Weinberger (*ibid.*, p. 15) write:

> "[...] [N]orms so understood as pure thought objects, are merely possible norms, not entities which become real merely in virtue of being thought about. They become real only by becoming operative as part of an action-guiding system for some person or group."

The point they make is that norms become only real norms because of their efficacy in an action-guiding system of rules applicable to a community or to certain persons or institutions. I discussed three modes of existence MacCormick and Weinberger (1986) distinguish for norms. They have in common that they amount to *mental* existence: their residence is 'in the head', not just as normative content, but also as action-guiding entities. In the institutional framework provided by MacCormick and Weinberger, norms have this extra layer of existence as an addition to the validity they derive from the underlying rules.

4.2.4 A conventionalist-*cum*-institutional approach

In a conventionalist-*cum*-institutional approach, the additional ontological claim is that existence of legal facts is attained through recognition by a sufficiently large group of individuals, or by institutional rules, or by both. Thus, it introduces the conventional layer as constituting existence directly. In an institutional theory, existence is constituted by rules whose conditions are fulfilled and whose consequences establish or maintain facts. In a conventionalist theory, existence is constituted by agreement in belief among a relevant part of a community. Existence of a token of the consequence type of some legal rule (*e.g.*, a specific marriage) can be constituted by both the institutional dimension (its existence is triggered by a rule) and the conventional dimension (its existence is generally agreed upon by a community) (Peczenik and Hage 2000). In the example, the existence of a specific marriage can be both constituted by the legal

rule of which the application leads to the marriage, and by the acceptance of the marriage by a community.

As to the existence of rules and norms, both types of entities can acquire their existence from either conventional or institutional establishment. In most cases, rules derive their existence from institutional establishment, whereas norms generally derive their existence from conventional establishment.

4.2.5 Hermeneutic theories

As ontological and epistemological claims cannot be completely separated in hermeneutic theories, it is not possible to make purely ontological claims about the existence of legal entities. Still, the result of an interaction between the person who knows something and the object of his knowledge yields, for instance, an interpretation. This interpretation may be claimed to exist, even if it does not itself correspond to some object whose existence was established before the interpretation activity took place. Hermeneutic theories do not assume that the result of the interpretation activity is completely arbitrary. On the contrary, Dworkin's (1986) hermeneutic theory assumes that one interpretation can be better than another interpretation – better, in fact, than *all* other interpretations (*cf.* subsection 3.2.5). Such a claim is made from the internal point of view, though. This means that such claims primarily have meaning for the community within which the interpretation was made.

The dimension constituting existence of interpreted legal entities (norms *etc.*) within hermeneutic theories is the activity of *Verstehen*, carried out by an authority. The activity establishes such interpretations. In hermeneutic theories of law the activity of *Verstehen* is regarded a necessary step in articulating literally stated legal norms and rules. As there is no strict distinction between the descriptive and the normative, between rules and facts, and between the law and its interpretation, it is hard to leave room for clear existence claims. Still, it seems to me that the activity of *Verstehen* often induces individual or authoritative recognition as a mode of existence for an articulated norm or rule.

4.2.6 Comparing the claims

What is the difference between the ontological claims of the theories discussed? As regards the existence of 'the law', the claims apply to different objects, such as legal rules and legal norms. Because of this, in some of the views, a general ontological claim on the elements constituting the legal domain can hardly be made. Rather, we have to distinguish the different types of entities we regard as part of 'the law', and assess their ontological status. Moreover, the existence

claims have different natures. Most notable, however, is the fact that claims regarding the nature of the existence of objects influence claims regarding the existence of objects, and *vice versa*. If validity is taken to be the mode of existence of legal norms, then only valid legal norms exist. With respect to the nature of the existence claims, we have to distinguish between the different layers constituting existence (see section 4.6).

Ontological claims also depend on the perspective taken by the observers of entities; there is a distinction between the internal perspective and the external perspective. Hart (1961, 97*ff*.) explains the difference between the existence and the validity of a norm as becoming clear when we take a certain perspective towards the law. Validity is only a meaningful term from an internal perspective; it refers to the situation in which a legal rule complies with the rule of recognition (*cf.* subsection 4.4.2). Existence is a meaningful term from an external perspective, where the actual rule-following behaviour of a social group may be a measure for the existence of a rule. Its efficacy determines its existence.

In a traditional natural-law view, the existence status claim applies to the law as an entity existing independently from human beings. In that case, the validity of positive law depends on its conformity with natural law. The ontological status layer 'existence' applies to natural law itself. Moral principles are interwoven with natural law, and thus there is a link between the validity of law and the presence of moral principles. In legal positivism, the law is identified with positive law, and the validity of positive law depends on authorities promulgating positive law. In Hart's view, however, in addition to the validity of positive law from the internal perspective, existence can be a meaningful dimension when the law is regarded from an external perspective.

Institutional theories of law make explicit ontological claims about legal entities: the law is regarded as supervenient on brute facts. This supervenience can take the form of validity or constitution. The existence of normative entities, like norms, is thus given appropriate attention. However, by addressing the normative dimension of legal norms, one has not yet answered the question whether there is a link between law and morality. In their discussion of morality, institutional theories of law seem to conform with legal positivism. Apart from the institutional accounts that it shares with institutional theories of law, the conventionalist-*cum*-institutional theory of law also involves a different form of existence, namely the form of existence invoked by recognition. Recognition may count as a prerequisite for the existence of conventions: if most of the individuals belonging to a community believe that something is the case, and recognise this fact, the fact exists. Hermeneutic theories, finally, involve individual

and authoritative recognition rather than such claims as existence and validity.

In table 4.1, I provide an overview of the ontological claims in different legal–philosophical positions with respect to norms and legal rules, and the phenomena from which these ontological claims are derived.

view	ontological status of legal norms and legal rules	derived from
natural law	validity	moral principles
legal positivism	validity	proper issuing authority
institutional theories of law	validity/mental existence/ efficacy	institutional establishment/ moral awareness *etc.*
conventionalist/ institutional approach	efficacy/validity	conventions/institutional establishment
hermeneutic theories of law	recognition	Verstehen

Table 4.1. Ontological claims with respect to norms and legal rules

4.3 Two objects of legal knowledge

In this section, I discuss two objects of legal knowledge in more detail, *viz.* systematisations and interpretations. I set out with an explanation of the concept of systematisation (subsection 4.3.1). A systematisation is a categorisation of legal rules in terms of general criteria. Then I explain the concept of interpretation (subsection 4.3.2). An interpretation is the result of someone's explanation of some fact, phenomenon or symbol.

4.3.1 Systematisations

There exist different views in legal theory about the activity of systematisation, and, consequently, about the ontological status of the result of this activity (also called systematisation). I subsequently discuss the views of Aarnio (1987), Peczenik and Hage (2000), and Brouwer (1999).

Aarnio (1987, p. 136) regards systematisation of the law as a theoretical, scientific activity. The object of a systematisation, and its product, he claims, are both (part of) the legal system (*ibid.*, p. 136-137). The systematisation activity classifies legal norms into different classes, according to some conceptual schema that has been prepared for that purpose. If some systematisation remains unchallenged, this does not yield any problem. But in many cases different, competing and

mutually inconsistent (or incoherent) systematisations exist for the same legal system and legal rules respectively. This characteristic is essential to the legal domain, but it may defeat the consistency of legal reality.

In subsection 4.2.4, I discussed a conventionalist-*cum*-institutional theory of law. Peczenik and Hage (2000) embed this theory in a coherence theory of law, in which systematisations are accommodated. Legal dogmatics also connects the legal system "with its background in the form of morality and (political) philosophy" (*ibid.*, p. 6). Thus, not only internal coherence (coherence in the system of legal rules), but also external coherence (coherence among the legal system and its background) is aimed at. But what does such a theory say about the existence of systematisations? Can there be a correspondence between some belief about an entity and the entity itself (such as a systematisation)? Or is a systematisation only established when it is introduced and – maybe – believed by a certain number of people?

As an answer to these questions, Peczenik and Hage (*ibid.*) claim that a procedural theory of law can account both for knowledge of the law and existence of the law. Procedures can be used to acquire knowledge of the law by using those procedures for exchanging reasons for a certain opinion. Thereby we attain (a form of) justification for that opinion. Thus, contrary to the situation in which we have a justification for believing some brute fact, we have a justification for a procedurally constructed fact. In case we have a belief about a brute fact, the justification itself does not influence that fact (as the fact is independent of our beliefs about it, and independent of our justification for those beliefs). In case we have a belief about a procedurally constructed fact, the justification (the reasons we provide for the fact) constitutes – at least partially – the existence of that fact (*cf. ibid.*, p. 9-10 and Hage 1997, p. 30-31).

Brouwer (1999, p. 222-224) distinguishes between equivalent and non-equivalent systematisations. Equivalent systematisations consist of legal material that is merely ordered according to, for instance, date of issue. An equivalent systematisation does not add to the content of the basic material it is constituted by (*e.g.*, codified legal rules). In a non-equivalent systematisation, an interpretative step is taken. Such a systematisation is the result of interpreting legal material. In a non-equivalent systematisation, the criterion of coherence (or some other criterion) may lead us to reject some part of the basic material. Therefore, the systematisation is not equivalent to the basic material it is based on (*i.e.*, it does not consist only of that material). Brouwer's approach towards systematisation is a non-equivalent one.

Brouwer (*ibid.*) claims that the systematisation of positive law is an attempt at

rationality. The rationality consists in the demands that an individual legal decision should be sufficiently justified, and that the justifying reasons for individual legal decisions should cohere (*ibid*., p. 219-220). Justifying reasons derive their status from legal norms, and thus, according to Brouwer, one could also say that legal norms should cohere. The activity of constructing a systematisation starts from certain assumptions. On the basis of these assumptions, systematisation criteria are developed. The resulting systematisation depends on these criteria. Three elements play a role in determining such criteria: (1) the view on the elements present in the domain, that is subject to the systematisation, (2) the relations among these elements, and (3) the demands that are deemed necessary to attain a well-ordered system (*ibid*.).

The systematisation activity thus partly depends on the view one has on the ontology of the domain: what are the elements in that domain, and what are the relations among them? These are ontological questions that have to be answered *prior to* establishing a systematisation. In addition, we can make claims about the ontological status of systematisations themselves. From Brouwer's discussion, we may cautiously conclude that, insofar as we can attach any ontological status to a systematisation, it depends on ontological decisions made earlier, on a level below the systematisation. However, the systematisation's ontological status need not be the same as the ontological status of the components it is constructed from.

4.3.2 Interpretations

In this subsection, I subsequently discuss the views on the ontological status of interpretation put forward by Aarnio (1987) and Brouwer (1997).

Aarnio (1987, p. 136) regards interpretation as a practical activity, essential to legal practice. He claims that the object of interpretations, and the product of the interpretation activity are both part of the legal system (*ibid*., p. 136-137). In the interpretation activity, the interpretation of a specific norm is chosen from the realm of all possible interpretations. Both activities make up legal reality. If some interpretation remains unchallenged, this is not even problematic. But different, competing and mutually inconsistent (or incoherent) interpretations may exist for the same legal rules.

On the nature of interpretation, from a legal-positivist viewpoint, Brouwer (1997, p. 141) remarks that the meaning of a legal norm is given by the authorities responsible for establishing that norm and/or modifying it. The meaning of a norm is determined by the rules of language and the rules stating conventions for some of these terms (the meaning of legal terms may differ from the 'natural'

meaning of the same term, which is the case with terms like 'murder'). In most cases, the meaning of legal norms is clear, and sufficient to determine whether it applies to a certain case. There are also cases in which the meanings of legal norms are not clear. In such a case, a judge determines the meaning of a norm by filling it in for that specific case, and he thereby performs judicial construction (on the assumption that judicial construction starts where judicial application ends).

What matters in this subsection is the ontological status of the result of the interpretation activity (as opposed to the interpretation activity itself). But because of the nature of the interpretation activity, different demands may be imposed on the way the existence of the interpretation activity result is secured. A judge's recognition of an interpretation may outweigh a community's recognition of that interpretation. Because of the specific function and capabilities of a judge, either his recognition of his own interpretation, or the efficacy of his interpretation, may be deemed necessary to call the interpretation existent. For instance, a judge's decision to sentence a suspect to four years of imprisonment only exists when it is actually enforced (if efficacy is the main criterion for existence).

4.4 Two views on the existence modes of legal systems

As to the existence of legal systems, and of the different entities belonging to it, different claims were made in the history of legal philosophy. What is more, the leading philosophical positions center on existence claims. These were briefly discussed in section 4.2. Legal philosophy in the twentieth century is dominated by two different versions of legal positivism. Kelsen (1960) and Hart (1961) are the main representatives of this view in legal philosophy. Kelsen's positivism is more extreme than Hart's. Whereas Kelsen takes a validity-conferring basic norm as a starting point of a legal system, Hart employs a rule of recognition, which is embedded in an empirical view on the law. Recognition is an apparent social phenomenon: the legal system exists by virtue of its acceptance by a significant part of the community to which it applies. In the following subsections, I discuss Kelsen's Grundnorm (subsection 4.4.1), Hart's rule of recognition (subsection 4.4.2), and the difference between validity and existence (subsection 4.4.3).

4.4.1 Kelsen's Grundnorm

Kelsen's positivism arises from the assumption that the existence of a so-called basic norm (Grundnorm) is supposed, that provides the ground for the validity

of the norms in a legal system (Kelsen 1979, p. 203*ff*.). Kelsen states that the basic norm is a fictitious norm because it not only fails to reflect a factual situation (there is no act of will constituting such a norm), it is also inherently contradictory, because such a norm both poses the authority of its own content, and it requires that its authority is confirmed externally (*ibid.*, p. 206-207): the basic norm cannot attain the authority of its own content without external confirmation of that content.

The basic norm confers validity to the top of a pyramid. This pyramid reflects the hierarchy among norms. The validity of the norms at a lower level in the pyramid is ultimately derived from the basic norm. At the top of the pyramid we find the constitution. This is the constitution in a material sense, *i.e.*, the norms that enable authorities to create and abolish legal norms (Kelsen 1960, p. 228). Below the constitution, we find the general norms that were created by legislation and custom, as regulated by the constitution (*ibid.*, p. 230-235). These norms regulate what institutions and procedures are to be observed, and what the content of the individual norms is that have to be created by these institutions and procedures (*ibid.*). Below these general norms, we find the individual norms: the norms that apply to specific cases, derived from the nature of the general norm and the characteristics of the case (*ibid.*).

4.4.2 Hart's rule of recognition

In 'The Concept of Law', Hart (1961) sketches a moderate form of legal positivism, in which the so-called rule of recognition broadens the basis of the legal system from mere validity to existence. He distinguishes between primary rules and secondary rules. The former are typical legal rules such as obligations. The latter are rules that identify the former: a rule of recognition comprises of the conditions that need to be fulfilled in order to make some primary rule part of the legal system. Except for rules of recognition, there are rules of change, enabling, among other things, the enactment of rules, and rules of adjudication, conferring the power to make an authoritative decision on whether some primary rule has been broken to a certain person or institution and determining the procedure to be followed (*ibid.*, p. 92-94).

The rule of recognition is a validity-conferring source. Hart (*ibid.*, p. 97*ff.*) opposes validity of rules to existence of rules. He states that the term validity in most cases refers to a situation in which statements are made on legal rules from the internal perspective. Existence claims are made from the external perspective. Validity is only a meaningful term from an internal perspective; it refers to the situation in which a legal rule complies with the rule of recognition.

Existence is a meaningful term from an external perspective, where the actual rule-following behaviour of a social group may be a measure for the existence of a rule. Its efficacy determines its existence. However, there is *no* necessary connection between the validity of a rule and its efficacy. The validity of a rule is determined entirely by its compliance with the criteria defined in the rule of recognition. The rule of recognition arises from legal practice, and, in contrast with the validity it confers to legal rules, it is not itself valid. Its existence is not a matter of validity, but a matter of fact (*ibid.*, p. 107). In table 4.2, the ontological status of rules is summarised.

rules	internal perspective	external perspective
ontological claim	validity	existence
source of ontological status	rule of recognition	efficacy

Table 4.2. Ontological status of rules in Hart's view

4.4.3 The difference between validity and existence

Ruiter (1993, p. 17*ff.*) explains what he sees as the differences and similarities between Hart's and Kelsen's views with respect to the status and structure of legal systems. He claims that Hart's and Kelsen's views are very similar, and that they only have a terminological controversy, namely on the question whether to use the term 'validity' (Kelsen) or 'existence' (Hart) for characterising the status of a legal system. According to Ruiter, the accounts of validity and existence essentially amount to a single phenomenon: the status of a legal system is that of "a presumed source of legal validity, provided that it is accepted as such in practice" (*ibid.*, p. 19).

   In my opinion, such an attempt to regard two views as equal is not desirable. Validity is only one of the phenomena contributing to existence. Hart explicitly distinguishes between validity and existence, and between the perspectives from which these status layers can be perceived. Kelsen focuses on validity and efficacy (*cf.* Kelsen 1960, p. 10). Ruiter locates the element of recognition in the relation between a community and a legal system, and derives validity from the recognised legal system: the legal system is a source of validity for legal rules.

4.5 Weaving cloth: elements of argumentation

In distinguishing the entities present in the legal domain, we could be distracted from the relations that exist between those entities, and the ways in which both the entities and the relations between them are brought into existence. A phenomenon of considerable importance to the legal domain is argumentation. Argumentation is the process that leads to – or justifies – decisions, systematisations, interpretations *etc.* This process can be executed in different forms and configurations: an individual may develop an argument, or two individuals may do so in a dispute that is subject to a number of rules. For a discussion of dialogical reasoning in the legal domain, in which two parties exchange arguments, I refer to Gordon (1995), Leenes (1998) and Lodder (1998). The structure and content of an argumentation process can be regulated to different degrees. Argumentation in the legal domain is bound to certain conditions laid down in law. These conditions regard legal procedure and applicable material law. But argumentation is not restricted by these conditions alone. Rather, it is subject to rhetoric traditions that may be present in the societies in which legal systems subsist.

It is important to note that legal argumentation often does not seem congruous with logic. Logic, the discipline that investigates valid argumentation, is often too rigid in its demands on validity to be able to accommodate rhetoric elements in its formalisations. Attempts have been made to incorporate some of the typical aspects of legal argumentation. Such attempts resulted, *e.g.*, in the development and application of defeasible logics. These logics are not monotonous, which means that a valid conclusion may no longer be valid if new information is added to the body of information on which the conclusion was based. For accounts of defeasible reasoning in law, I refer to Prakken (1993, 1997), Verheij (1996) and Hage (1997).

Prakken clarifies the layers into which the application of logic to legal argumentation can be divided. First, in the *logic layer*, individual arguments are constructed, *i.e.*, pieces of information are arranged in such a way that they support a proposition. Second, in the *dialectical layer* conflicting arguments are compared. In this layer arguments may be defeated, and it may be determined whether an argument is justified. Third, in the *procedural layer*, disputes are regulated. Rules about the introduction and withdrawal of premises *etc.* are implemented on this level. Fourth, the *strategic layer* is the layer where it is tried to conduct disputes in a rational manner, *i.e.*, to settle a dispute following effective tactics (Prakken 1999, p. 86).

Logic does not play a substantial role in those areas where the content of

argumentation is most important, *i.e.*, where the main concern is to represent meaning. Most logical languages (for instance predicate logic) are hardly capable of representing the many subtleties in the meaning of natural language. It may function, however, as a means for structuring the argumentation process. The diverse structuring functions of logic are expressed in Prakken's (1999) typology of argumentation layers. In the following subsections, I discuss Toulmin's (1958) argumentation scheme (subsection 4.5.1), the notion of a reason (subsection 4.5.2), and the notion of a defeater (subsection 4.5.3). The goal of the current section is to show the role of different elements in the construction of an argument. These elements are a part of the knowledge-based model of the law in chapter 6.

4.5.1 Toulmin's argumentation scheme

Toulmin's (1958) seminal work on argumentation contains different notions; together they constitute a so-called argumentation scheme. These notions are *datum*, *warrant*, *backing*, *claim*, *rebuttal*, and *modal qualifier*. Data are facts that are appealed to in support of a conclusion (*ibid.*, p. 97). A datum can, for instance, be 'Harry was born in Bermuda' (*ibid.*, p. 98). Warrants are general, hypothesis-style statements. They relate two phenomena to each other, for instance in the form 'all x are y', or 'an x will be a y'. Toulmin (*ibid.*, p. 99) provides the example 'A man born in Bermuda will be a British subject'. Warrants can be supported by a so-called backing. A backing often has a factual nature. In case of the example above, the warrant takes the following form: "[a] man born in Bermuda will generally be a British subject" (*ibid.*, p. 105), and the backing consists of legal statutes on British nationality (*ibid.*). A claim or conclusion is the statement that should be justified through the other items ('Harry is a British subject', *ibid.*, p. 99). A rebuttal is a condition under which the claim does not hold (for instance the fact that Harry has become a naturalised American, *ibid.*, p. 105). The notion of a modal qualifier is used to indicate the expression with which the claim or conclusion is restricted. For instance, the phrases 'almost certainly' and 'presumably' may serve as modal qualifiers (*ibid.*, p. 105, 109).

Although many amendments have been made on Toulmin's (1958) scheme, its basic categories have played a major role in current argumentation research. It forms the basis for the notion of defeasible reasoning, not only by making exceptions an explicit part of arguments, but also by focusing on the nature of argument instead of on the possibilities of logic. Rather than using good-old hard-boiled deduction, Toulmin rationally reconstructed the nature of reasoning and concluded that different entities play a role: data (facts), warrants (rules

and principles), backings (facts counting as reasons), rebuttals (defeaters) *etc.* (*ibid.*). The terms between brackets can all be found in the knowledge-based model of the law in chapter 6.

Toulmin's scheme is a static one; it does not provide a procedural framework in which the building or exchange of arguments is analysed. Still, by distinguishing the different elements, he provides an onset to a procedural account of argumentation: he stresses the importance of the natural order of argumentation for his research into the subject (*ibid.*, p. 17).

4.5.2 Reasons

Reasons play a crucial role in argumentation. They can help us to justify a certain belief, decision or interpretation. For instance, a lawyer may be convinced of the guilt of his client by reasons as diverse as the sweat he sees on his client's forehead and the fact that his client has admitted a crime. Hage (1997, p. 11) claims that reasons form the most important and basic element of a theory of reasoning. The content of a reason determines its relevance and nature with respect to the entity it applies to. In most cases, a reason for a particular belief is not a reason for another belief. A reason for a belief may have a nature different from a reason for an action. Furthermore, a reason may either support the *belief* in a fact, or *constitute* that fact (Hage 1997, p. 60). Hage (*ibid.*) calls the former category epistemic reasons (*reasons to believe*), and the latter category constitutive reasons (*reasons for existence*).

Reasons are often explicit (*I* cannot *have* a reason for a belief without knowing what the content of that reason is, although there can *be* a reason for my belief without me knowing it), and because of this, reasons can be similar to beliefs. Apart from facts that qualify as reasons, beliefs can acquire a specific status: beliefs that are *qualified as* reasons. This may seem to contradict with the assumption that there are such things as constitutive reasons. After all, constitutive reasons operate at an ontological rather than an epistemological level. The presence of some rule may be an immediate reason for the establishment of, for instance, some duty, without someone knowing that duty. Raz (1975, p. 17) writes:

> "Beliefs are sometimes reasons, but it would be wrong to regard all reasons as beliefs. It should be remembered that reasons are used to guide behaviour, and people are to be guided by what is the case, not by what they believe to be the case. To be sure, in order to be guided by what is the case a person must come to believe that it is the case. Nevertheless it is the fact and not his belief in it which should guide him and which is a reason. If p

is the case, then the fact that I do not believe that p does not establish that p is not a reason for me to perform some action."

Raz' (1975) observations about the difference between beliefs and reasons are a ground for me to distinguish between personal and impersonal reasons. Personal reasons are beliefs or facts qualified as a reason by a particular individual (the fact that it rains is a reason for me to stay at work). Impersonal reasons derive their status as a reason from their particular meaning in a system of rules or from their place in a causal chain (the fact that someone has was born as the child of two British citizens is a reason for him to have the British nationality). The notion of an impersonal reason is based on a reconstruction of some process, in which certain elements clearly *function* as reasons. Impersonal reasons arise from systemic recognition: the recognition of something (a proposition, a fact) as a reason is based on the presence of rules that produce such reasons.

In order to provide a more fine-grained typification of reasons, I describe seven distinctions. The first distinction is made according to the type of subject who recognises the reason. The resulting reason types are personal reasons, impersonal reasons, and authoritative reasons. Personal reasons are personal beliefs qualified as reasons by individual recognition. A personal reason applies to a belief, *i.e.*, it is a reason for a certain belief. Impersonal reasons are things (propositions, facts) qualified as reasons by systemic recognition. Their content and validity depend on the content of rules in the system in which they are issued. Authoritative reasons are personal beliefs qualified as reasons by authoritative recognition. For instance, a decision of a judge can be backed (partially) by authoritative reasons: reasons the judge cannot derive from applicable law, but that he can propose and use due to his function. The content of authoritative reasons is not determined by the system in which they are used. However, their 'validity' (the fact that they are authoritative) is a consequence of the system; the system contains rules that leave discretionary power to individuals. The difference between impersonal reasons and authoritative reasons is analogous to the difference between deductive justification and second-order justification made by MacCormick (1978).

The second distinction is between explanatory and guiding reasons (Raz 1975, p. 18-19, *cf.* also Hage 1997, p. 35). They regard the nature of the relation between the reason and the object it applies to. Explanatory reasons are causes for behaviour or events. An explanatory reason for someone's behaviour really holds as a reason for that person's behaviour in a specific case. Guiding reasons are reasons why we should act in a certain way, or should refrain from that act;

they determine what ought to be the case (Raz 1975, p. 18). Raz (*ibid.*) claims that explanatory reasons presuppose the existence of guiding reasons; to explain one's behaviour in terms of explanatory reasons implies that one knows what guiding reasons there are in a specific case for one's actions. Reasons for physical actions, however, are not my main concern. Rather, reasons for beliefs, and reasons for decisions are at stake.

This takes us to a third distinction of reasons: the distinction which regards the object of the reason; the thing, the event, or the belief it is a reason for. I distinguish between reasons for beliefs, reasons for actions, reasons for decisions, reasons for classifications, reasons for interpretations, and last, but not least, reasons for reasons (for the former three reason object types, *cf.* Hage 1997, p. 60). This is not an exhaustive enumeration of objects, but it suffices for the legal domain.

The fourth and fifth distinctions are between the sources of reasons. We have to distinguish between content sources and origination sources. In a legal context, the content of reasons arises from formal sources of law, material sources of law, or mere deliberation by some person. The origination sources of reasons are the same as the sources of legal belief (*cf.* subsection 2.5.2): perception, memory, testimony, consciousness, reason, and interpretation. For instance, a reason for the fact that Napster furthers copyright infringement may be that it does not prevent its users from making available copyright-protected materials. The content sources of this reason are two observations: (1) the way in which Napster works (its users make available music), and (2) the fact that Napster refrains from acting upon these infringements. The origination sources of this reason can be perception and testimony.

The sixth distinction is between external modalities: what characteristic qualifies the relation between the reason and its object? Hage (1997, p. 60) distinguishes between epistemic, deontic and anankastic modalities, related to the corresponding types of reasons. Epistemic reasons are reasons to assume something, or reasons that make something probable (*ibid.*, p. 73-74). Deontic reasons are reasons that contribute to the establishing of obligations, permissions, or prohibitions (*ibid.*, p. 65-66). Anankastic reasons regard the necessity or possibility of facts (*ibid.*, p. 71-72). For instance, if a competent judge decides that Napster infringes copyright, this makes it probable that a damage claim from a record company will be rewarded. Thus, the external modality of the reason (the establishment of copyright infringement by Napster) is probability. The probability regards the likelihood that a damage claim will be rewarded.

The seventh distinction also regards modalities. In this case, however, the *inter-*

nal modalities of reasons are concerned. Reasons may contain modalities such as 'it is likely that...', 'this induces an obligation to...', and 'it is permitted to...'. The modalities that are used to specify the object that a reason *refers to*, may also be used to specify the reason itself. For example, in case a damage claim on Napster will probably be rewarded, and this constitutes a reason to believe that such claims can also be made on Napster's users, the internal modality of the reason is probability.

I have thus specified the characteristics of a reason: a reason has a content, a subject, an object, and a specific relation between subject and content, and between content and object. The content of a reason reflects (part of) its meaning relative to its subject and object; it reflects why it is a reason. The subject of a reason is a person, a group of persons or an authority, or there is no subject at all. The object of a reason is a belief, an action, a decision, a classification, an interpretation, or another reason. The nature of the relation between the subject and the reason depends on the nature of the subject; in case of a personal reason, the individual has a belief, which means that he has a 'believing attitude' towards the content of the reason. But he may also have a reason to act in a certain way without believing that reason. In case of an impersonal reason the relation should be causal or systemic (constitutive). In case of an authoritative reason the relation is systemic (constitutive); its systemic nature is derived from the discretion ascribed to the subject. The nature of the relation between the reason and the object is explanatory (description of causes) or guiding (normative).

recognising subject	reason type	object	content source
personal impersonal authoritative	explanatory guiding	belief action decision classification interpretation reason	formal sources material sources deliberation
origination source	external modality	internal modality	
perception memory testimony consciousness reason interpretation	epistemic deontic anankastic	epistemic deontic anankastic	

Table 4.3. Characteristics of reasons

In brief, the characteristics of a reason are: it has a content, a subject (in most cases), an object, and a specific relation between subject and the reason's content, and between the reason's content and object. Furthermore, the reason's content has two sources: one for its content and one for its origination, and there are modalities involved in two places: with respect to the content of the reason, and with respect to the object of the reason. In table 4.3, I give an overview of the characteristics of reasons.

4.5.3 Defeaters

A special class of reasons is formed by defeaters. Defeaters undermine reasons. Defeaters that directly attack the conclusions of an argument are called 'rebutting defeaters'. Defeaters that attack the relation between a reason and its conclusion are called 'undercutting defeaters'. With an undercutting defeater, the assumption is challenged that some statement or fact is indeed a reason for a conclusion (Pollock 1999, p. 196). Just as reasons, defeaters can be classified according to the distinctions made above. A defeater has a content, a subject, an object, and there is a specific relation between subject and content, and between content and object. Defeaters are not always regarded as reasons. Rather than considering defeaters as individual entities, Verheij (1996, p. 121) takes a holistic approach, in which a defeater is a relation between challenging and challenged arguments.

Defeaters can be both beliefs classified as defeaters and facts classified as defeaters. However, the object of a defeater can only be a reason or a belief. Defeaters do not apply to actions, decisions, classifications, and interpretations. A defeater D for a reason R is a reason why the *reason* R does not hold with respect to some object (*e.g.*, a belief). Focus is thus on R's status *as a reason*. A defeater for a belief R is a reason why the *belief* R does not hold. If some entity has multiple statuses (it is a belief and a reason at the same time), a defeater may apply to its status as a belief as well as to its status as a reason. Generally, if there is a defeater for the belief, this will influence the status of that belief as a reason as well. But if there is a defeater for the belief's status *qua* reason, this need not influence the status of the belief itself.

Defeaters can be classified according to the type of attack they are used for. With respect to a belief, defeaters can attack one of its status layers, its content, the assumption that it is held by a certain person, or the state-of-affairs it refers to. It is not a coincidence that these modes of attack are reminiscent of the knowledge criteria listed in section 2.3; reasons and defeaters (as a special class of reasons) are based on the same phenomena and relations as knowledge criteria.

Thus, the belief that Napster's users make copyright infringements may be defeated. The example below shows this.

> p: Napster's users make copyright infringements
> q: Napster's users make available copyright-protected material via the Internet
> r: The fact that Napster's users make available copyright-protected material via the Internet does not constitute a form of reproduction or publication, and hence, does not constitute copyright infringement
> s: Napster does not have any users any more

In this example, proposition q counts as a reason for proposition p; it inconclusively supports that proposition. If John believes that q, and John's believes that p, then proposition q is a reason for John to believe that p (that is to say: to hold the belief that p). Proposition r counts as an undercutting defeater for reason q. Thus, if John believes p on the basis of reason q, then r makes him not believe p on the basis that q no longer supports p. Proposition s counts as a rebutting defeater for reason q. Whereas q is a reason for p, q and s together cannot possibly support p.

Defeaters can be typified in terms of the characteristics of reasons listed in subsection 4.5.2. Thus, the nature of a defeater depends on, among other things, the type of subject, the nature of the relation between defeater and object, and its object. First, as regards the type of subject, the defeater is a personal reason, an impersonal reason, or an authoritative reason. The relevance of this distinction is caused by the fact that a defeater may only function as a defeater for an individual, but that it also may function as a defeater for a community, or for an authority. The recognition of a belief as a defeater by a legal authority may have major consequences for a judgement.

Second, the nature of the relation between a defeater and its object is relevant because, like a reason, a defeater may be either guiding or explanatory. Guiding defeaters say why something *ought not be* the case. Explanatory defeaters say why something *is not* the case. The difference boils down to a normative/descriptive distinction for defeaters.

Third, the different types of defeaters specified above (rebutting and undercutting) can be traced back to a defeater's object: if the object is the content of some proposition, then it is a rebutting defeater. If a defeater's object is the reason status of some proposition, it is an undercutting defeater. In the specification of a defeater, its sources and modalities also play a role.

The specification of a defeater in terms of these characteristics is comparable to the specification of a reason. The other three distinctions (sources of defeat-

ers, external modalities and internal modalities) are the same as those for reasons. However, external and internal modalities work negatively; because they are formulated negatively, their presence enables us to attack reasons instead of supporting them.

4.6 Ontological status layers

The discussion of ontological claims with respect to the law in general (section 4.2) and with respect to legal systems (section 4.4) is summarised in the current subsection. Rather than discussing ontological claims from different viewpoints in philosophy of law, I give an overview of ontological status layers that can be derived from those different viewpoints. I distinguish between the following status layers: validity and existence (subsection 4.6.1), recognition (subsection 4.6.2), constitution (subsection 4.6.3), and efficacy (subsection 4.6.4).

In brief, these status layers amount to the following. A legal entity (a legal rule, norm, or decision) is valid if a legal system implicitly or explicitly provides the criteria for the validity of that entity, and these criteria are fulfilled for that entity. A legal entity is legally recognised if it is accepted by an individual in his capacity as a legal professional, or by a legal institution. A legal entity is legally constituted if that entity is the result of applying a valid legal rule, or if the entity is the result of the use of a discretionary power backed by the legal system. A legal entity (a legal rule, norm or decision) is legally efficacious if it has a significant effect on the mental condition and the actions of a group of individuals, *i.e.*, if it affects their actions, and it does so because of its being part of a legal system.

Some status layers can only be applied to a limited number of legal entity types. Status layers can be combined. For instance, a legal rule may be valid, constituted, recognised and efficacious at the same time. It is legally valid because it was enacted through legal procedures, it is legally constituted because it is a rule that counts as a legal rule (because of its validity in a legal system), it is legally recognised if it is accepted by a relevant legal institution, and it is efficacious if it is actually followed by a community.

4.6.1 Validity and existence

To elaborate on the difference between validity and existence, we have to distinguish between the internal perspective and the external perspective on the law (for a discussion of this distinction, *cf.* Brouwer 1997, p. 109-117). The internal perspective is taken by a person who considers himself as a subject of a legal

system. For this person, the law consists of binding norms, that impose rights and obligations (*ibid.*, p. 110). Legal decisions and acts are being justified with an appeal to valid (binding) law. The external perspective is taken by someone who is not the subject of the legal system he observes, at least he does not consider the system in that capacity. Rather, he takes a descriptive and explanatory stance towards the law, or even a predictive one (*ibid.*, p. 111-112). The judgements such a person makes on the characteristics of a legal system are not legal judgements. The aim in employing an external perspective is to make true statements about the legal system instead of making justified legal decisions.

Validity and existence should be distinguished, as they do not necessarily co-incide. Validity arises from adequate validity-conferring sources; it is a status layer, imposed on certain objects. Validity may require the existence of the object it applies to. Existence is then employed as a term for a basic ontological claim with respect to objects (something *is*, but we do not know yet in what way). Their actual type of existence is constituted by the status layers that are imposed on them, such as recognition and validity. Still, validity and existence are intertwined, not only because validity is a potential characteristic of existence (and thus, it is one of the ontological status layers), but also because validity may invoke existence. Moreover, in Kelsen's legal-positivist view, validity is equal to existence: validity is the only way in which legal norms can exist. The validity of some object, for instance a legal decision, may contribute to its existence in the sense that the validity status is a necessary (and maybe even sufficient) condition for the existence of a legal decision. Without its validity, a legal decision may not be a proper legal decision, and thus not exist as such. So, rather than identifying existence with validity, I adopt a position in which existence in itself is an almost empty concept, to which other status layers grant meaning. This enables me to accommodate different ontological claims in the knowledge-based model of the law (*cf.* chapter 6).

4.6.2 Recognition

Recognition occurs if a person or institution acknowledges something as belonging to some category (*cf.* for instance Hart 1961, Hage 1998, Peczenik and Hage 2000). I distinguish four types of recognition, based on their subject (a person or institution whose recognition is relevant in the legal domain): individual recognition, collective recognition, systemic recognition, and authoritative recognition. Individual recognition occurs if an individual recognises the qualification of a situation (for instance, if I qualify a specific act of a person as manslaughter). Collective recognition occurs if a substantial part of some rel-

evant community recognises the qualification of a situation (for instance, if the community qualifies a specific act of a person as manslaughter). Systemic recognition occurs if the qualification of a situation is recognised through the application of some rule within a system, for instance a legal system (for instance, if I automatically get a fine for my speeding on a highway, in which case my behaviour is systemically recognised as an offence). Authoritative recognition occurs if some person or institution, that has acquired a specific (authoritative) role within a community, recognises the qualification of a situation (for instance, if a judge qualifies a specific act of a person as manslaughter). The difference between systemic recognition on the one hand, and the other three types of recognition on the other, is that systemic recognition does not require explicit human intervention; it is based on indirect human interference, for instance by the establishment of rules that take care of the indirect recognition.

The occurrence of systemic recognition is based on the assumption that there are rules that can be applied to cases without involving human judgement. Systemic recognition may be posited in cases in which we speak about legal acts, or offences, without needing a judge to qualify and recognise the facts as such. In such cases, the recognition is derived from a general act of recognition, for instance the establishment of a rule that enables the automated recognition of certain acts as offences (we could classify the automated transaction of speed limit exceeding offences as such). Systemic recognition may be combined with other types of recognition: for instance, the systemic recognition of a computer system's 'decisions' may be backed through authoritative recognition by the institution that has established the system.

In a legal positivist view, the validity of law depends on its being issued by a proper authority. This is a case of authoritative recognition, often combined with systemic recognition. The existence of judgements also depends on recognition. The recognition of judgement is derived from the recognition of law through systemic recognition, not only from the rules applied for the ruling, but also the rules determining a procedure. Even more importantly, from the fact that a judge has room to make his own ruling. Though it might be intuitively inattractive, my claim is that, since a judge has the authority of making his own decisions, he provides a part of the recognition of those decisions.

4.6.3 Constitution

New buildings and clothes are manufactured by putting bricks on one another and by cutting pieces of cloth and sewing them together. Other things, such as reasons, decisions and rules are not tangible. Things we cannot see, hear, or

smell, but that definitely play a role in our society. Without asking *how* they exist, we may claim *that* they exist, and we may attempt to find out how they were brought into existence. Entities such as reasons acquire their existence as a reason from their qualification or use as such. The qualification yields a constitution relation between the underlying entity and the new entity.

Constitution is the non-legal equivalent of validity. An entity is constituted whenever it is consequence of applying a rule or a certain practice. In order for an entity to attain constitution as a status layer, it should be initiated by the application of a rule or a practice. The constitution dimension helps to maintain a form of existence in cases where other ontological status layers do not apply. An institution, such as money, may not only derive its existence from efficacy and from legal validity, but also from constitution: qualifications attained from (non-legal) rules and practices applying to the nature and use of money.

4.6.4 Efficacy

The efficacy of a legal system can be used as a condition for its existence. The efficacy criterion can also be applied to individual legal rules. Raz' (1970, p. 5-26) account of Austin's theory of a legal system explains the use of this condition (*cf.* Austin 1954). Raz (*ibid.*, p. 16) summarises:

> "A system exists if its laws exist. One can, therefore, derive from what has hitherto been said on the existence of a law the following *criteria for the existence of a legal system:* A legal system exists if and only if (1) its supreme legislator is habitually obeyed, that is to say the laws of the system are by and large efficacious; (2) its supreme legislator does not habitually obey anyone; (3) its supreme legislator is superior to the subjects of every one of his laws relative to the sanction of that law. To these we should add a fourth condition, (4) that all the laws of the system were actually legislated and were legislated ultimately, by one person, or group."

In this efficacy-based account of the existence of legal systems, individual laws exist under the criteria that they are part of a legal system, and that this legal system is generally efficacious. This is claimed to ensure that individual laws that are not efficacious in themselves but are clearly part of a legal system, do indeed exist. Habitual obedience means that a substantive part of a community will generally conform to most laws that constitute a legal system. The other criteria (2 through 4) determine the hierarchy of the system, necessary to ensure its efficacy: there should be a legislating authority that is not accountable to any other authority, and that did indeed enact the laws that are part of the system.

Although this definition of a legal system is subject to criticism, for instance the comment that the existence of a legal system cannot be reasonably made dependent entirely on the presence of a single legislator, the importance of the concept of obedience (and, thereby, efficacy) should be stressed. The coercion that accompanies (or precedes, or follows) the establishment of a legal system may not be a necessary and sufficient condition for its existence, but it is still an important characteristic of such a system. Efficacy need not always be caused by coercion. Legal rules may work simply because they are reasonable conventions that are easy to follow. The consequence of not following such rules would have a negative effect on those breaking them, not through coercion, but by the natural and cultural consequences arising from the rule-abiding behaviour. Such consequences (being yelled at when passing a red traffic light) are not induced by the legal system itself.

Due to this observation, the following characteristics of a legal system should be distinguished: efficacy, coercion, and obedience. The relations among these three concepts are as follows. The coercion accompanying a legal system may cause its efficacy through enforcing obedience. Obedience may also be enforced through other means than mere coercion. Coercion does not necessarily enforce obedience. Thus, there are no necessary consequence relations between the three concepts; rather, the presence of coercion may function as an explanatory reason for obedience, and the presence of obedience as an explanatory reason for efficacy.

4.7 A meta-ontological conception of law

Rather than employing a specific view on the ontological status of the law, I provide a framework in which different ontological views can be accommodated. In this framework, there are three basic layers. The first layer consists of non-legal entities that are legally-relevant. The second layer consists of potentially legal entities whose existence status (*qua* legal entity) is not yet established. The third layer consists of legal entities. Transitions between these three layers are possible through varying sets of criteria. A potentially legal entity, such as a judge's decision, becomes a legal entity by checking whether the status layers required do indeed apply. Thus, for instance, an immoral decision is not classified as a legal entity in a natural-law view, whereas it is classified as such in a legal-positivist view. Entities on these three different levels are connected to each other through two different relations. These are the counts-as relation (subsection 4.7.1) and the causation relation (subsection 4.7.2), which are discussed below.

4.7.1 Counting as

The counts-as relation is based on Searle's (1995, p. 18) work about constitutive relations. It is used to link concrete objects or object types to each other. Searle uses it, among others, to point out the logical structure of institutional reality (*ibid.*, p. 79*ff.*). In his account of institutional reality, the formula 'x counts as y in c' represents a general status-assigning function. It says that under the circumstances referred to by c, a fact x is regarded as a fact y, transcending the mere fact x by the status assigned to it. This structure can be iterated and combined into complex systems of interrelated institutions (*ibid.*, p. 80). For instance, the relation can be used to provide insight into the structure of the judiciary, in which natural persons are assigned certain roles. According to Searle, there are different types of functions assigning status: symbolic powers create meaning, deontic powers create rights and obligations, honourific powers create status for its own sake, and procedural powers create conditions on the way to power and honour (*ibid.*, pp. 99-103). These different functions can all be put in the format of the counts-as relation. For instance, we may express symbolic power with the function 'the presence of a red light counts as a stop sign; it does so in the context of traffic and the presence of a traffic light'.

A theory about the structure of social reality, adjusted to the legal domain, was proposed by MacCormick and Weinberger (1986). They distinguish institutive, consequential, and terminative rules. Institutive rules bring into existence some kind of legal entity (e.g., a contract). Consequential rules define consequences attached to there being in existence some legal entity (*e.g.*, one has to comply with the conditions in a contract, if there is a contract). Terminative rules put an end to the existence of some legal entity (*e.g.*, a divorce ends a marriage) (*ibid.*, p. 52-53). An instance of a legal institution (*e.g.*, an occurrence of the concept of marriage) exists in time. Its existence starts when the concept instance is created as the consequence of an institutive rule. Its existence is maintained by a consequential rule, and it is ended by a terminative rule. The concepts and institution types themselves (for instance divorce and marriage) do not have that kind of existence (*ibid.*, p. 53). Searle's (1995) theory and MacCormick and Weinberger's (1986) theory have in common that instances of legal concepts are claimed to exist, and that the existence of these instances is supported by some kind of constitutive rules. Rules establish relations among concepts, and other rules produce instantiations of these concepts. The instantiations exist, and if we ascribe some kind of property to them, or if they are the argument of some relation, they constitute institutional facts. Searle provides no explicit solution for the question how such facts are brought into existence, whereas MacCormick

and Weinberger use their institutive rules for that purpose.

The constitutive relation functions in an intensional context: the fact that something counts as something else is often the result of a community *thinking* or *believing* that this is the case. This means that the terms that are filled in for x and y in 'x counts as y' often cannot be changed for extensionally identical expressions without changing the meaning of the expression. A different way to state the same is that it is not allowed to substitute co-referring terms in case they are under the reach of an intensional predicate (for instance, 'the community *thinks*'). Intensionality plays an important part in natural-language semantics. An example of intensionality is Frege's 'evening star - morning star' example (for a more extensive explanation, *cf.* Gamut 1991b): a sentence like 'the morning star is the evening star' is true, as both expressions refer to the same entity. But a sentence like 'John thinks the morning star is the evening star' need not be true, as 'to think' is an intensional verb. The extension of 'morning star' (the actual entity) is something different from the intension of 'morning star' (its meaning). Frege's distinction between extension and intension is a basis for conceptual definitions.

The intensionality of the context in which the counts-as relation is used, is important, because there may be changes in definitions of legal concepts. These changes are reflected in the application conditions of the concept, but they are not visible in the concept name. The counts-as relation enables us to express a situation in which the collective recognition of a constitutive relation between two concepts may depend on the concept names instead of their precise meaning. This enables the stable recognition of open-texture concepts, even through their meaning changes. Within the context of a constitutive relation that is recognised by a community, the terms that are related to each other cannot be interchanged with extensionally equivalent terms, as that could threaten the recognition of the constitutive relation.

In addition to the intensionality of the context in which it is used, the constitutive relation is non-transitive, asymmetrical, and irreflexive. These constraints are expressed in terms of relations between types; abstract categories of entities. The non-transitivity constraint to the counts-as relation means that, whenever there are two relations: s counts as s', and s' counts as s'', there is no relation s counts as s''. The asymmetry constraint means that whenever there is a relation s counts-as s', there is no relation s' counts-as s. The irreflexivity constraint means that if there is a relation s counts as s', then types s and s' are not equal.

An example illustrates the consequences of these characteristics. When we say that killing a person on purpose counts as manslaughter, we establish a re-

lation between conditions (someone kills someone else, and he does this on purpose), and the occurrence of 'manslaughter', which is supervenient on the fulfilment of the conditions. The conditions and 'manslaughter' are expressed as types, which are combined in a constraint, expressing a meaningful relationship between conditions and consequences arising from them. The antecedent type of the constraint expresses the conditions, while the consequence type expresses the 'concept name' (a legal type) itself.

The non-transitivity of the counts-as relation says that, if O counts as O', and O' counts as O'', then O need not count as O'' (there is no direct constitutive relation between the two types, even if all conditions are fulfilled). The asymmetry constraint says that if O counts as O', O' need not count as O. And the irreflexivity constraint says that O cannot count as O (if a reflexive constitutive relation existed, it would be meaningless). It is hard, if not impossible, to illuminate the first constraint with an example, because we would need differently named, intensionally identical legal concepts for that. The asymmetry constraint can be illustrated by inverting 'killing on purpose counts as manslaughter', as 'manslaughter counts as killing on purpose' is a useless expression. And irreflexivity is easily shown for the concept of manslaughter, as the phrase 'manslaughter counts as manslaughter' is also useless.

A constitutive relation between legal concepts is often based upon legal rules. A legal rule, for instance a rule that defines the conditions for an instance of manslaughter, contains the ingredients to define such a constitutive relation. But concrete legal cases are also constructed with the help of such relations. An instance of manslaughter is a result of qualifying certain facts as such; the abstract constitutive relation between two types becomes a concrete constitutive relation between two facts.

4.7.2 Causation

Rather than a matter of physics, legal causation is a matter of human construction. For instance: what can be reasonably attributed to a person? The circumstances under which a certain state-of-affairs or action is deemed to cause a different state-of-affairs or action in a legal manner, has been thoroughly discussed in legal theory. When can we rightly say that a person causes a certain fact or situation? The answer to this question is important in the determination of the liability of a person. In Remmelink (1996, p. 172*ff.*), four legal causality theories are discerned: conditio sine qua non, causa proxima, relevance theories, and adequacy theories. The latter three theories restrict the reach of the former. I discuss these theories of causality in order to find out what the legal causation

relation amounts to, which is necessary to give some kind of formal translation of the relation later on.

Von Buri's theory of causality neatly translates into the phrase '*conditio sine qua non*' (Remmelink 1996, p. 175-176). In this theory, each link in a causal chain leading to a certain fact is as good as any other link in that chain in its relation to the final consequence. The criterion with which we determine whether some phenomenon *a* causes another phenomenon *b*, is that if we removed *a*, *b* would not have taken place. Thus, *a* is a necessary condition (*conditio sine qua non*) for the occurrence of *b*. This theory has three major disadvantages. First, it does not directly limit the length of the causal chain. Thus, one may go back in time infinitely to find a cause for a phenomenon. This problem, however, can be overcome by taking into account the certainty with which a certain fact or state-of-affairs can be causally linked with the consequence. If there is a long causal chain, the link becomes less certain. Second, this theory attracts the epistemic problem that we can hardly know long causality chains. Third, should we not center on what can be reasonably attributed to a person, rather than regarding long causal chains?

Causality theories in which '*causa proxima*' plays a role, are similar to the '*conditio sine qua non*' theory in that they presuppose the possibility of knowing causal chains. However, they restrict these chains in length by employing the demand that the causal chain between the two phenomena among which causality is presupposed, should not be too long (*ibid.*, p. 176-177). Variants of such a '*causa proxima*' theory demand for instance that the most effective cause or the decisive cause should be selected.

The relevance theory (*ibid.*, p. 177-178) tries to isolate those causes that relate to intent or *culpa*. These causes often form the (implicit or explicit) core of the offence description, and were meant to be punishable by the legislator. If they have not made an explicit part of the offence description, they may be reconstructed as a part of the intention of the legislator. Remmelink uses the following example. Someone has been maltreated, and he is transferred to a hospital that burns to the ground. He dies in this fire. Can we qualify this as 'maltreatment inducing death'? We cannot, as the fire in the hospital is not a relevant cause for his death (it has not been proposed as a relevant cause by the legislator), even though he was in the hospital as a consequence of being maltreated.

Adequacy theories (*ibid.*, p. 178-181) are based on the idea that a phenomenon can be considered as a cause for a specific consequence, if this phenomenon would cause this consequence in general, *i.e.*, in most cases. Adequacy theories depend on our knowledge of laws or patterns: what would generally be the con-

sequence of a phenomenon? They also depend on our knowledge of what the cause in a specific case is. Subjective adequacy theories stress the knowledge of the perpetrator at the time of his act, while objective adequacy theories emphasise the actual situation. In objective adequacy theories the cause-consequence chain is reconstructed afterwards: what acts are suitable for bringing about the consequence that occurred? The knowledge and intention of the actor are not important here, just the general consequences of her act. So even if some consequence seemed unlikely to occur at some moment, its occurrence should be ascribed to the actor. Subjective theories state that the cause-consequence chain is evaluated through the foreseeability of the consequence by an individual.

The meaning of these different theories of causality can be clarified by an example: the 'slipper – egg skull' judgement ('pantoffel-eierschedel', Appeal Court Amsterdam, June 14, 1939, NJ 1940: 34). On February 5 1939 a man (W.G.) throws a slipper to the head of his 51-year-old wife. He hits her temple. She dies a few hours later. On the autopsy it appears that she has an abnormally thin skull. Her husband did not know this. When we consider this case through the different theories of causality, we notice the following. Seen from the perspective of '*conditio sine qua non*' we find that if the man had not thrown the slipper, his wife would not have died. Her death can thus be ascribed to him. From the viewpoint of the *causa proxima* theory, the result is the same: throwing the slipper does not only count as a *conditio sine qua non*, but also as a *causa proxima*. Seen from the perspective of the relevance theory, I assume that this theory does not solve the dilemma, because it is not clear whether the legislator wanted to punish the behaviour explained in the example. Seen from the viewpoint of 'adequate causation' we can take two stances. First, an objective one: only one person in two thousand has an egg skull. Still, in this case the woman has one, and the fact that she dies should be ascribed to her husband throwing a slipper to her head, because this will be fatal for almost anyone with an egg skull. Second, a subjective one: as the man did not know that his wife had an egg skull, he could not have foreseen the consequence of his deed. He cannot be punished for something that he could not foresee.

The goal of the discussion of causality theories is not to choose among those theories. Instead, I attempt to leave open the possibility of accommodating different causality theories in the knowledge-based model of law. From the overview of causality theories we may derive claims with respect to the ontological and epistemic characteristics of the causation relation. The causation relation can be regarded as having either an ontological or an epistemic basis. In the former case, causation reflects a relation in reality that we may – or may not

– know. In the latter case, causation is a relation that we project on reality. This projection need not be entirely arbitrary; it can be based on the conventions laid out in causality theories. The existence of such causation relations thus emanates from the epistemic: we may justifiedly believe that something causes something else in a legal manner, because we know the accepted causality theory within some jurisdiction. In subjective adequacy theories, for instance, what counts as an adequate cause depends on the current knowledge of a person. If we know that the adequacy theory is accepted within Dutch penal law, and we know that our suspect did not know his wife had an egg skull, then we have sufficient reasons not to project the causation relation between the act of the husband and the death of his wife.

In summary, the following factors play a role in the different causality theories: the presence of a *conditio sine qua non*, the proximity of the cause, the relevance of the cause, and the adequacy of the cause. The *conditio sine qua non* is found at the ontological level. The proximity, relevance, and adequacy of the cause, however, largely depend on human judgement. These judgements are not arbitrary; they are based on conventions on the matter. What we may learn from the discussion of different causality theories is that legal causation is only secondarily related to what we may denote as 'ontological causation'. Legal causation is not only based on the actual causal chain between the act and the consequence (as suggested by the *conditio sine qua non* theory), but also on that chain *as perceived by* some person. The alleged length of this causal chain (or the proximity of cause and consequence), the possibility of producing alternative readings of the same event, the alleged *communis opinio* about the extent to which a consequence can be ascribed to an act, the specific knowledge of the person involved in the act, her intention at the time of her act, the general intention ascribed to a person when she acts in a certain way, they all play a role in determining the causation relations.

Rather than regarding the legal causation relation as a purely physical, mechanical phenomenon, thus modelling it as an extensional relation, I choose to model it in accordance with the counts-as relation. If a phenomenon p is regarded to cause a different phenomenon q, then this is expressed as 'p counts as a cause for q'. We can model general patterns of legal causation in the same way: if phenomena of type P are generally regarded to cause phenomena of type Q, then we express this as 'phenomena of type P count as a cause for phenomena of type Q', or abbreviated: 'P counts as a cause for Q'.

4.8 Constituents for a model of the law

The third research question was: what characteristics of knowledge about the legal domain are useful as constituents for a model of the law? In this chapter, I explained different views on the existence of the law. These views have a far-reaching influence on what types of entities we can discern in the legal domain, and on the kind of existence they have. Additionally, I discussed some of these types of entities, and several ways in which legal systems and legal entities can be claimed to exist. From this discussion, I derived five ontological status layers and two relations that are central in the knowledge-based model of the law discussed in chapter 6.

5 Legal knowledge representation

After discussing the characteristics of legal knowledge, I now turn to the question how we can represent it, focusing on the ontological and epistemological characteristics of the legal domain. In section 5.1, I explain the concept of representation, the languages used for knowledge representation, and the characteristics of representing legal knowledge. In section 5.2, I discuss the concept of meaning, exploring the way in which the meaning of legal concepts can be pinpointed. In section 5.3, I outline three ontologies that were built specifically for the representation of the legal domain, or the representation of knowledge about the legal domain.

5.1 The concept of representation

The core of representation consists of symbolising something else; the object represented. The representation of an object (be it a physical entity or something else, such as knowledge) requires conventions or similarity. Conventions serve to indicate what the symbols used in the representation refer to, whereas similarity is based on 'implicit conventions' that are self-explanatory; *i.e.*, the resemblances between the representation and its object are such, that we do not have to know explicit conventions to be able to recognise the similarity between the two. There is no clear dividing line between convention and similarity. Similarity is often based on and strengthened by widely accepted conventions, such as perspective in art (for a discussion of representation in art, *cf.* Gombrich 1977). If the object of representation is not tangible, it becomes harder – if not impossible – to talk about similarity, because the object has no 'face' to represent. After all, similarity is often thought about as a depiction relation; and invisible objects are hard to depict. Invisible objects, such as institutions, cannot be represented through similarity. However, we are still able to use conventions to represent them.

Bench-Capon (1990, p. 11) provides the following definition of representation: "a set of syntactic and semantic conventions that makes it possible to describe things". The syntactic conventions determine how the language in which knowledge is represented is constructed from its smallest constituting parts, while the semantic conventions determine what the sentences constructed in that language mean. Bench-Capon's definition of representation illustrates one of the basic assumptions of artificial intelligence, namely that all representations are given in the form of some language, *i.e.*, a set of syntactic and semantic con-

ventions. Bench-Capon calls natural languages, such as English, the most powerful representation languages. It is, however, not at all clear whether it is possible to construct a set of syntactic and semantic rules to describe a natural language. A natural language is a suitable means for knowledge representation because we know how to use and understand it. However, we do not know the exact rules that exist for interpreting them, if there are any.

Instead of using representation languages with sets of syntactic and semantic rules, one could choose to employ other means to represent knowledge, such as neural nets. In such an exemplar-based approach, no set of syntactic and semantic rules is established in advance. Instead, neural nets are claimed to show an approach to learning similar to that of human beings. I do not pay attention to this form of representation. Rather, focus is on the ways to represent knowledge explicitly. For this purpose, I discuss representation languages (subsection 5.1.1), with which such explicit representation can be constructed. Subsequently, I explain the trade-off between the richness of representation and the practical use of a language in representing legal knowledge (subsection 5.1.2).

5.1.1 Representation languages

A representation language is a combination of syntactic and semantic rules with which expressions can be built or parsed and with which the meaning of the structures can be determined. There is a trade-off between two important features of a representation language: its expressiveness and its computational tractability. The more expressive a language is, *i.e.*, the more syntactic and semantic detail can be articulated through it, the harder it becomes to manipulate expressions in that language (*cf.* Bench-Capon 1990, p. 15). The need for computational tractability becomes acute when we try to infer new knowledge from the knowledge that has been represented in a computer system.

Proposition logic

Many representation languages are based on logic. Two of the most common logics are proposition logic and first-order predicate logic. These two types of logic form the basis for other types of logic, such as modal proposition logic and modal predicate logic. The language of proposition logic consists of symbols representing propositions: for instance p and q (propositions p and q). Furthermore, the language consists of symbols representing connectives: $\rightarrow$ (implication), $\wedge$ (and), $\vee$ (or), and $\neg$ (not) (*cf.* Gamut 1991a). If I want to express the sentence 'If John gets his diploma, he will be admitted to Berkeley', I use p for the first part of the sentence, q for the second part, and the connective $\rightarrow$ for the

conditional relation between them. The sentence would thus be represented as '$p \rightarrow q$'. The expressiveness of this representation is not very impressive; proposition logic enables us to represent structural relations among sentences or large sentence parts only. However, we can easily make inferences from this type of representation. Rather more interesting is the question if these are the inferences we wish to make. They are not. Proposition logic as a knowledge representation language is not sufficiently fine-grained to fulfil the representation needs of the legal domain – we have to able to express properties and relations. In first-order predicate logic, this is possible.

First-order predicate logic
First-order predicate logic adds some important features compared to proposition logic (*cf.* Gamut 1991a). The basic building blocks of first-order predicate logic (henceforth called 'predicate logic') suffice to speak of an ontology. The building blocks of the language of predicate logic are predicates, variables, object constants, quantifiers, and connectives. Properties and relations are represented by so-called *predicates*. A predicate says something about one or more objects. If it says something about one object, this 'something' is a property of that object. If it says something about the connection between two or more objects, it expresses a relation among those objects. Variables are place-holders for object constants. They thus enable quantification over multiple objects. Objects constants are symbols referring to objects or individuals.

Quantifiers enable us to state things about specific sets of objects, *i.e.*, to say that a property holds for at least one object, or that a property holds for all objects. Connectives roughly have the same function in predicate logic as in proposition logic, except that in predicate logic they can be used to combine the more detailed expressions containing predicates and the objects these predicates apply to, instead of combining sentence parts without any further structure (as in proposition logic). In first-order predicate logic, it is not possible to express properties of properties. We would need a higher-order predicate logic for that purpose.

Modal logic
To be able to give a detailed representation of specific predicates, modal logics can be employed (*cf.* Gamut 1991b). Modal logics enable us to represent particular modalities, *i.e.*, operators such as *possible*, *necessary*, *obliged*, *permitted etc.* With these modalities, we can represent sentences like 'It is necessarily the case that bachelors are not married' and 'It is not permitted to use vehicles in the

park'. Modal logics with operators like 'permitted' and 'obliged' are called *deontic* logics. They are used to reason about permissions and obligations. For instance, deontic logics enable us to derive the fact that something is permitted from the fact that something is obliged. Work on deontic logics in the area of AI and law has been done by, among others, Royakkers (1996) and Van der Torre (1997).

Natural language and situation semantics

Natural language is the language (for instance English or Dutch) that we employ to communicate with other human beings. It is the most widely applicable, flexible and most efficient means of communication we have at our disposal. But it also lacks the preciseness that we need to use it as a formal representation language. To overcome this problem, attempts have been made to represent the meaning of natural language, *i.e.*, to find methods to translate natural language into expressions in formally specified artificial languages. Whereas computers cannot make inferences directly on natural language yet, we can apply automatic inferencing techniques on represented natural language. Thus, we have to represent the meaning of natural language, which can be regarded as a representation language itself, in a formal language. I discuss one of these representation languages, *viz.* situation semantics, because the underlying ontological theory (situation theory) has contributed to the elements of the knowledge-based model of the law presented in this thesis.

Situation semantics was developed as a representation language for natural language in the general framework of situation theory, developed by Barwise and Perry (1983). Situation theory is a theory of information and meaning based on ecological realism. The central idea of ecological realism is that meaning is found in the world (in the objects that are part of it), and in the interaction between individuals and their particular environments (called 'ecological niches') (Gibson 1979). This is different from many other semantic theories, in which meaning is a relation between language and reality. Situation semantics has rather extensive ontological commitments. It recognises states-of-affairs, courses-of-events, individuals, relations, and space-time locations as elements of its ontology. States-of-affairs are situations as they are in the world. Courses-of-events are combinations of such situations, linked to different locations in space and time. Attempts to describe such states-of-affairs and courses-of-events force us to give their characteristics. In situation semantics, the resulting (incomplete) descriptions are called abstract states-of-affairs and courses-of-events.

An abstract state-of-affairs or course-of-events consists of one or more constituent sequences, each expressing a certain relation between different indi-

viduals or constants, and linked to a certain place and time. An infon can be compared to a predicate with arguments, with the added information of a link to a space-time location and a truth value. The elements of the ontology of situation semantics show its foundation in ecological realism: rather than abstract entities, real situations, individuals, and space-time locations are its central elements. In addition to abstract states-of-affairs and courses-of-events, situation semantics distinguishes situation types and event types. Whenever a certain part (a relation, an individual, or a space-time location) of an infon is not instantiated, a state-of-affairs becomes a situation type, and a course-of-events becomes an event type. For instance, 'John walks' constitutes an abstract state-of-affairs, whereas 'someone walks' constitutes a situation type.

A constraint is a relation between situation types. Barwise and Perry (*ibid.*, p. 97) introduce three types of constraints: necessary constraints, nomic constraints, and conventional constraints. Necessary constraints hold between relations. For instance, they express the necessary relationship between the situation type that Jane is guilty of theft and a situation type that someone is guilty of theft. Nomic constraints generalise over patterns caused by the laws of physics, *e.g.*, the relation between a type expressing that an object has a mass and the type expressing that this object exerts gravitational force. Conventional constraints indicate what relations exist between phenomena as a result of conventions in a society, e.g., the relation between a type expressing that a person has intentionally killed another one, and the type expressing that he has committed manslaughter. Conventional constraints are especially important in the legal domain, where conventions are an important factor in determining legal facts and legal consequences. Each constraint type is either conditional or unconditional. A conditional constraint is linked to a certain space-time location. An unconditional constraint is valid in all space-time locations. Constraints were already used in subsection 4.7.1 to express the characteristics of the counts-as relation. They are also suitable for expressing counts-as relations themselves.

Depending on the goal with which a representation language is built, its focus is on different parts of the structure and content of the object to be represented. In general, representation may have an ontological, an epistemic, or a semantic focus. If it has an ontological focus, it aims at representing the elements of a domain, and the relations among those elements. In brief, it attempts to represent facts. If it has an epistemic focus, it aims at representing knowledge. If it has a semantic focus, it is aimed at representing the meaning of language. Emphasis in situation semantics is on the ontology of situations, *i.e.*, on determining the objects that are referred to in language, and on classifying these situations.

A representation language is necessary to express knowledge and facts in terms of the knowledge-based model of the law. The ontology provides the basic conceptual framework, and the representation language is a means to express knowledge and facts in terms of that conceptual framework. However, representation languages often show ontological presuppositions as well. For instance, with each of the languages discussed, disregarding proposition logic, we are able to represent entities, properties, and relations. These basic commitments are already quite far-reaching (please note that predicate logic does not *require* such ontological commitments; this is a matter of choice). The language used to represent knowledge in the domain of penal law in chapter 8 is based on predicate logic. In choosing this representation language (reason-based logic), expressive power is taken as the primary criterion, instead of computational tractability.

5.1.2 Representing legal knowledge

'Knowledge representation' is generally regarded as either representing our knowledge about facts, or representing facts. The question is whether, in the legal domain, it suffices to represent objects and relations among them, and knowledge about those objects and relations. Other features are sometimes considered more appropriate for representation. The research into legal argumentation attempts, for instance, to establish the rules along which legal disputes are carried out. Such a procedural approach shows how legal decisions are accomplished, instead of merely positing that they are part of the domain to be represented. Thus, in addition to knowledge and entities, we may attempt to represent the way knowledge is acquired. Furthermore, we may try to add justifying reasons to represented facts. The elements that should be represented are discussed in other chapters, and an extensive overview is provided in chapter 6, which describes an ontology of law.

After determining the object of legal knowledge representation, the next issue is what language should be used for legal knowledge representation. Of course, the answer to this question depends, among other things, on the goal with which knowledge is represented (automated reasoning, building queries *etc.*), and the suitability of a representation language (its expressive powers, its tractability). In the research project underlying this thesis, several representation languages were used to represent characteristics of legal knowledge. Although making a choice for a representation language is not among the main goals of this thesis, evaluating representation languages helped me to determine the demands imposed on a conceptual model, the construction of which *is* one of the main goals of this thesis.

Following the goal of adequately describing a conceptual model of law, a number of demands applies to the choice of a representation language. First, the representation language may not limit expressive power in a manner that either distorts the representation of knowledge or leaves out essential characteristics. Second, the representation power of a language may not add wrong or distorted information because of superfluous expressive powers of the representation language. Third, the expressive power should be adjusted to the characteristics and complexity of the domain concerned, though the necessity to attain tractability for automatic inferencing can put limitations on this demand. Fourth, there should be some kind of isomorphy between the object represented and the representation, and thus between the domain to which the language applies, and the representation language. This isomorphy applies to either structure, content, or a combination of both. In section 8.1, I shall describe the representation language adopted in this thesis. The demands above were applied to the choice of this language, with the exception of the tractability demand.

5.2 The multiple meanings of 'meaning'

To be able to characterise an object properly, its representation often requires us to capture some of its characteristics. In a Fregean view (*cf.* Frege 1892), the expression of such characteristics in a representation constitutes the *meaning* of that expression. Meaning is related to knowledge in an important sense. To have knowledge of some object implies carrying beliefs in which representations of the object play a role. Knowing that a judge has sentenced Pete to two years of imprisonment often means that we have some idea of what a judge is, who Pete is, and what two years of imprisonment is, and even if we do not know Pete, that Pete refers to some person. In addition, these beliefs have to fulfil specific criteria in order to qualify as knowledge.

The meaning of 'meaning' has troubled generations of philosophers. I do not intend in any way to settle the dispute between proponents of different concepts of meaning. However, I do wish to establish understanding of the concept of meaning relative to the topics discussed in this thesis. A concept of meaning is needed to be able to speak about the representation of meaning in information systems or knowledge-based systems. For instance, a knowledge-based information system may need a concept of meaning based on necessary and sufficient conditions in order to make inferences as to whether a legal rule applies. An information retrieval system may not need any concept of meaning, if it only contains full-text search, in which case the meaning of the query and the mean-

ing of the documents stored are not relevant. Basically, if we wish to represent knowledge or information in a system, the concept of meaning will become relevant.

In brief, there are four views on meaning: meaning is defined in terms of (1) behaviour, (2) necessary and sufficient conditions, (3) objects referred to, and (4) use. A concept of meaning in terms of behaviour supposes that there is no such thing as 'meaning' separate from symbols and people's reactions on those symbols; the meaning of a term is found in the kind of reaction a person shows when he is confronted with that term. This concept of meaning is rejected, as it is obviously too simplistic to be adequate. A concept of meaning in terms of necessary and sufficient conditions says that the meaning of a term is defined with a set of conditions (the meaning of 'bachelor' is defined as an unmarried man). A concept of meaning in terms of objects referred to states that the meaning of a term is defined by the set of objects that it refers to (the meaning of 'bachelor' is defined as the set of men who are bachelors). A concept of meaning defined in terms of use, meaning is found in the actual use of a word or phrase (the meaning of 'bachelor' is determined by its use in everyday conversation).

However, meaning is hard to catch in a single phrase. The concept of meaning itself is an open-textured one, subject to the changes triggered by actual use of the concept. Attempts in legal knowledge-based systems to capture meaning often implicitly combine different meanings of 'meaning'. A knowledge-based system may combine an intensional definition (in terms of necessary and sufficient conditions) with an extensional definition (in terms of examples). A legal knowledge-based system may comply with the specific characteristics of the meaning of legal expressions. For example, such a system may reflect the changes in definitions that emanate from the open texture of legal expressions.

In the following subsections, I explain different views on meaning. Instead of leaving meaning out in the discussion of knowledge representation, I deem the role of meaning sufficiently important to integrate the concept in a discussion of the representation of legal knowledge, as the elaboration on a concept of meaning may reveal useful starting points for representation. Below, I discuss the concept of meaning in terms of necessary and sufficient conditions, and the related phenomenon called 'open texture' (subsection 5.2.1) and the concept of meaning in terms of use (subsection 5.2.2). Subsequently, I shift my attention to the relation of the concept of meaning with the type of concepts whose meaning we wish to establish (subsection 5.2.3). Finally, I explain how sense and reference play a major role in the representation of meaning (subsection 5.2.4).

5.2.1 Sense, reference and open texture

Frege (*cf.* Frege, 1892) developed a theory of meaning that distinguishes between two components of meaning: sense (Sinn) and reference (Bedeutung). The sense of an expression consists of the conditions under which it obtains: if we call an unmarried man a bachelor, then the concept 'bachelor' has as conditions: being a man and being unmarried. The presence of each of these conditions is necessary, and the presence of both of these conditions is sufficient for a person to be a bachelor. The reference of an expression consists of the set of all objects that fulfil the conditions that are part of its sense: 'bachelor' refers to all unmarried men. The sense of an expression determines its reference, and two expressions with the same sense have the same reference. But two expressions with different senses may also have the same reference. The paradigm example of this, as mentioned before, is the following: (1) the morning star, and (2) the evening star. We can describe the sense of 'morning star' as 'the brightest heavenly body in the eastern skies at dawn', and the sense of 'evening star' as 'the brightest heavenly body in the western skies at sunset' (Gamut 1991, p. 9). It is now well-known that the 'evening star' and the 'morning star' are identical; it is the planet Venus. So 'morning star' and 'evening star' have different senses, but identical references.

Meaning may change, *i.e.*, for instance, the conditions that determine the sense of a concept may change through time. This is sometimes referred to as the 'open texture' of a concept. Franken *et al.* (1995, p. 83) define open texture as the phenomenon that the definition of concepts can be revised on the basis of new insights. They give an example of an open-texture concept in the Lindenbaum-Cohen judgement. In this judgement, the definition of 'wrongful act' was changed. Before the judgement, a wrongful act was defined as an act that infringes another one's rights or an act that is in contrast with one's statutory duty. The new definition was that a wrongful act is an act or neglect that infringes another one's rights or is a violation of one's legal duty, or is a violation of the duty of care that should be regarded in social intercourse towards some person or property (*ibid.*, p. 84).

The open-texture nature of legal concepts can thus be more precisely defined as the possibility that elements of the definition of some concept may change, may be left out, or may be added at some point in time, from which point in time the new set of elements will constitute the definition of that concept. In other words: the sense of an open texture concept changes through time. The concept of open texture is not a purely legal phenomenon. It was introduced by Waismann (1952, p. 120), who distinguishes between open texture and vagueness

of empirical concepts. If we consider vagueness to be the unclear (or missing) demarcation lines of application of a concept, then open texture can be defined as the possibility of vagueness. Vagueness may apply to both the intension (sense) of a concept, in which case there is no (clear) set of necessary and sufficient application conditions for the concept, and to the extension (reference) of a concept, in which case we cannot (completely) determine the set of objects the concept refers to. Waismann writes that most or all empirical concepts have an open texture (*ibid.*, p. 121):

> "[…I]t is this texture which prevents us from verifying conclusively most of our empirical statements. Take any material object statements. The terms which occur in it are non-exhaustive; that means that we cannot foresee completely all possible conditions in which they are to be used; there will always remain a possibility, however faint, that we have not taken into account something or other that may be relevant to their usage; and that means that we cannot foresee completely all the possible circumstances in which the statement is true or in which it is false. There will always remain a margin of uncertainty. Thus the absence of a conclusive verification is directly due to the open texture of the terms concerned."

In the legal domain, we do not only regard empirical statements, but also statements containing normative concepts. The open texture of normative concepts is not a consequence of a non-exhaustive description of reality, but of the non-exhaustive description of the norm content. What is more important, is that the non-exhaustiveness is not a consequence of the inability to describe reality completely, but of an inability to determine all the different ways in which a certain legal concept can be applied. Hart introduced the concept of open texture in the legal domain; he discusses it in his 'The Concept of Law', in which we find the following fragment (Hart 1961, p. 124-125):

> "Whichever device, precedent or legislation, is chosen for the communication of standards of behaviour, these, however smoothly they work over the great mass of ordinary cases, will, at some point where their application is in question, prove indeterminate; they will have what has been termed an *open texture*. So far we have presented this, in the case of legislation, as a general feature of human language; uncertainty at the borderline is the price to be paid for the use of general classifying terms in any form of communication concerning matters of fact."

From this fragment, we may conclude that Hart explains vagueness rather than

open texture, *i.e.*, his description of open texture does not seem to distinguish between the vague borders of concept application and the radical changes of application conditions. Bix (1991, p. 66) claims that Hart's concept of open texture differs from Waismann's primarily in that Hart applies this concept on normative entities like rules and norms, instead of just to empirical concepts. Waismann (1952, p. 123) states that statements about an event have an open-texture character because something new and unforeseen may occur. In such cases we do not know whether we have applied all possible tests to check whether a certain event has occurred. One of Waismann's examples involves a cat (*ibid.*, p. 119):

> "Suppose I have to verify a statement such as 'There is a cat next door'; suppose I go over to the next room, open the door, look into it and actually see a cat. Is this enough to prove my statement? Or must I, in addition to it, touch the cat, pat him and induce him to purr? [...] What, for instance, should I say when that creature later on grew to a gigantic size?"

Waismann's statement concerns empirical statements. But similar obstacles bear on normative entities.

In my view, the following characteristics give an adequate impression of what open texture amounts to in the legal domain. Open texture is a characteristic of legal concepts. It amounts to the potential changeability of the conditions with which we determine whether a certain case, fact, or part of reality is referred to by a legal concept. The open texture character of legal concepts has a normative cause, *i.e.*, to check whether a certain instance is referred to by a legal concept is to make a deliberation not only about the question whether that instance is referred to by the concept, but also whether it should be referred to by that concept. A normative deliberation is made whenever the application of the standard conditions leads to an undesirable consequence. From an internal perspective, the change of open-textured concepts is a normative phenomenon. However, from an external perspective, it is an empirical phenomenon (*cf.* subsections 4.2.6 and 4.6.1).

5.2.2 Meaning as use

The view of meaning as use, introduced by Wittgenstein (*cf.* Wittgenstein 1953), always is subject to the danger of becoming a slogan rather than a serious idea. It has to be more strictly defined (or rather, explained by examples, as Wittgenstein did) to make sense, because otherwise, it raises questions such as: whose use constitutes meaning? and: what kinds of use constitute meaning? Putnam (1975,

p. 145) has an approach that may clarify the 'use' aspect of meaning:

> "[E]veryone to whom gold is important for any reason has to acquire the word 'gold'; but he does not have to acquire the method of recognizing if something is or is not gold. He can rely on a special subclass of speakers. The features that are generally thought to be present in connection with a general name – necessary and sufficient conditions for membership in the extension, ways of recognizing if something is in the extension ('criteria'), etc. – are all present in the linguistic community considered as a collective body; but that collective body divides the 'labor' of knowing and employing these various parts of the 'meaning' of 'gold.'"

The question whose use constitutes meaning becomes acute when only a small part of the community is able to determine whether some substance is really gold. Is only their use of the term 'gold' relevant? The majority of people, who do not distinguish gold from many other substances that look like gold, use the term 'gold' in a way that makes it impossible to determine its reference. Should we then exclude reference from our understanding of meaning? Any definition of meaning that is related to use yields such problems.

In the determination of the meaning of 'meaning' we can employ the approaches to the concept itself. For an intensional approach, this would mean that we can define the meaning of 'meaning' in terms of necessary and sufficient conditions. As 'meaning' itself is an open-textured concept, we need to consider the possibility of changes in this set of necessary and sufficient conditions. Meaning regarded as use also enables us to integrate open-texturedness into the concept of meaning. Because the use of the concept of meaning varies through time and through communities, we should incorporate use in its definition. The meaning of 'meaning' thus becomes dependent on the actual occurrence of the concept in natural language use, or, more specific, in the legal domain, or even in the domain of the employment of legal information systems.

5.2.3 Meaning with respect to types

The discussion of different meanings of 'meaning' can be further elaborated on by regarding the meaning of 'meaning' with respect to different types of concepts. For instance, there may be a difference between the meaning of natural kind terms (terms referring to things occurring in nature, such as water, gold, trees) and other terms, *e.g.*, social or legal type terms (please note that 'kind' refers to both countable and uncountable things, for instance sheep *vs.* water, whereas 'type' refers only to countable things). Institutional reality (of which the

legal domain forms a part) transcends individual perception and thought. The existence of a legal system is based on ontological status layers, most of which are intersubjective, such as recognition and efficacy. This observation influences the meaning of 'meaning' for different kinds. Meaning is related to the ontological and epistemic properties of a type: for instance, in a legal–positivist view, the meaning of legal types depends on criteria established in the law by certain authorities.

This means that the existence of an instance of a legal type depends on the presence of an authority that establishes its existence, and that the truth of statements about instances of such kinds depends on whether these statements comply with the actual presence of instances of legal kinds. This observation influences the way in which we can establish or discover the meaning of a specific kind. The observation that the meaning of a legal concept is established mainly by convention, and that this convention is formed and polished by its daily use, or its use by a group of legal professionals, may direct us towards different concepts of meaning, though. Given a legal–positivist view on the law, for instance, we may be tempted to introduce some kind of 'meaning is use' concept, and consider only its use within a limited group of legal professionals and politicians.

5.2.4 Sense and reference revisited

I already dedicated some words to the meaning of legal concepts, as opposed to the meaning of, for instance, natural kinds. By taking natural language as a starting point for legal expressions, we can hardly do without the content that it creates in the minds of people. The meaning of legal concepts does not seem to reside in the physical world containing mountains and stones, but rather in the heads of the members of a community, and in particular in the heads of legal professionals. At least, it seems that the intensions of legal concepts can be found there, and in law books. We are far from being able to interpret natural language mechanically, which means that the potential legal language has for creating content is hardly usable in an artificial context.

Sense and reference of legal concepts, contrary to sense and reference of natural kind terms, are for the larger part determined *post hoc*, *i.e.*, the meaning of legal concepts is developed in their use, partly by normative decisions. Statute law provides parts of the senses of such concepts, whereas case law may constitute additions to or changes in those senses, and, moreover, constitutes the cases or facts that concepts refer to. The extensions are determined by application of the senses that underlie the concepts, and by additional considerations that are

constituted by the current case. Those considerations may also find their way towards the senses of the concepts. In practice, this means that the change of the sense of an open-textured legal concept like 'good' is triggered by the lack of a fit between the 'old' sense and, for instance, the necessity to classify electricity as a good in order to protect a certain interest. To attain this protection, the classification of electricity as a good has to be put into correspondence with the sense and reference of 'good' in the current legal context. Thus, the sense of this concept is changed accordingly, and 'electricity' is added to its reference (in the context of this case).

Loth (1988, p. 55) rejects the Fregean view of meaning. He claims that the application conditions, that are part of the sense of an expression, and that constitute the meaning of that expression, are grasped by an interpretation act applied to the expression. However, these application conditions are linguistic expressions as well. To define the meaning of these application conditions means that we have to establish application conditions for these application conditions, and so on *ad infinitum*: we fail in establishing the reference of an expression. According to Loth (*ibid.*), this should be done by participating in a language game; in actually using a language together with other people. My opinion on this matter is that we cannot possibly demand that Fregean senses help us to establish their references all by themselves. We need the help our syntactic and semantic capabilities. In the interaction of our syntactic and semantic capabilities with a Fregean sense, the kind of content is created that refers to the world: *intentional* content.

This discussion of meaning with respect to legal types is thus still based on the sense and reference concept of meaning. Meaning as use, be it in the sense Wittgenstein gave to it, or in Quine's (Quine 1960, 1990) behaviourism-inspired sense, tempts us to revise our meaning concept into a concept that does not consider concepts alone; instead, we need to employ a more holistic sense of meaning, *i.e.*, a concept that applies to sentences and combinations of sentences. Words are not employed outside a context of other words, and (at least for spoken language) a context of use. Thus, we have to determine whether we can distinguish the meaning of separate words, and if we can, whether this is useful. In a legal context, it is indeed useful to consider the meaning of individual concepts, because they refer to the qualifications and institutions that carry so much import in this domain. Even though meaning as use is certainly not an empty slogan in the legal domain, use is restricted to a number of legal professionals, and much of the meaning of legal concept is actually codified, *i.e.*, at least the codified conditions can be constructed as a Fregean sense.

5.3 Conceptualisations of the legal domain

Ontologies are conceptual models of a specific domain. This use of the term 'ontology' in the context of computer science and artificial intelligence differs from the use of the term in a philosophical context (*cf.* Mommers, Schmidt and Oskamp 1997). Gruber (1993) defines an ontology as a specification of a conceptualisation, and, more specifically, as a description of concepts and relations that exist for an individual or a community of individuals. A conceptualisation is a representation of the world that is both simplified and abstract (*ibid.*). Ontologies are meant to provide a basic framework for knowledge representation: the entities and relations distinguished in an ontology provide a user with the means to represent knowledge in the domain that the ontology covers. Below, I discuss three ontologies of law: Valente's (1995) functional ontology (subsection 5.3.1), Van Kralingen's (1995) frame-based ontology of law (subsection 5.3.2), and Verheij and Hage's (1997) abstract model of the law (subsection 5.3.3). Next, I compare these three models (subsection 5.3.4).

5.3.1 Valente's functional ontology

Valente models an ontology of law from the perspectives of both information science, following the definition of 'ontology' given by Gruber, and of legal theory, following such legal theorists as Kelsen and Hart. His so-called functional ontology of law is introduced in Valente (1995). The functional ontology distinguishes between six different basic types of knowledge (*cf.* Breuker, Valente and Winkels 1997). These basic types are: normative knowledge, meta-legal knowledge, world knowledge, responsibility knowledge, reactive knowledge, and creative knowledge (Breuker *et al.* 1997, p. 27-30).

Normative knowledge is regarded the most obvious kind of legal knowledge. It consists of the elements of the legal domain prescribing the behaviour of the people constituting a society, and a description of the way social reality should look like (*ibid.*, p. 27). Meta-legal knowledge consists of the entities not directly regulating behaviour. These entities are empowering and derogating norms, and norms regulating the mutual relations among primary norms (*ibid.*, p. 27). World knowledge consists of elements telling us what the world looks like. Responsibility knowledge consists of the knowledge that links normative knowledge to reactive knowledge: in case some agent causes a certain event, responsibility knowledge helps to establish the extent to which the agent can be held responsible for what he did (*ibid.*, p. 29-30). Reactive knowledge consists of the sanctions imposed if an agent breaches a norm and is held responsible for that (*ibid.*, p. 30). Creative knowledge, finally, consists of information about newly created

institutions and other entities that arise from the application of the law (*ibid.*, p. 30).

5.3.2 Van Kralingen's frame-based ontology of law

Van Kralingen (1995) distinguishes between three kinds of so-called frames: concept frames, norm frames and act frames. I briefly explain the principles behind these frames and their contents. An act frame consists of fourteen elements: an act identifier, promulgation, scope, agent, act type, means, manner, temporal aspects, spatial aspects, circumstances, cause, aim, intentionality, and final state. To distinguish different acts from each other, each act has a unique identifier. This identifier constitutes the first slot. The promulgation slot of an act frame contains the source of the act description. This information serves a purpose as a means of reference. The scope slot contains the range under which the act description is applicable. The agent slot contains the individual or legal body that performs the act. The act type slot is the main predicate in an act frame. It is the act itself, regardless of intentions and circumstances.

The means slot specifies the material objects used in the act described, or gives more details about the act. In the manner slot, it is indicated in what way a specific act is performed. In the temporal aspects slot, the information given about time and time intervals is stored. In the spatial aspects slot, information is given about locations. Relevant aspects of the situation in which the act takes place, are collected in the circumstances slot. Of course, this enumeration of circumstances can never be complete. The cause slot specifies what reasons there are for the act to take place. In the aim slot, we find the goal the agent aims at in performing the specified action. The aim of an action is related to the action by an intention of the agent; this relation is teleological. The intentionality slot is used to represent this. In the final state slot, we find the results and consequences of the specified action.

The slots of a norm frame are described in Van Kralingen (1995, p. 57). They are respectively the norm identifier, promulgation, scope, conditions of application, subject, legal modality, and act identifier. The norm identifier slot contains a unique identifier for the norm. The norm type slot contains the norm type: a norm of conduct or a norm of competence. In this slot it is indicated whether the subject of the norm has a competence to do something, or the duty to perform the act referred to. A norm of competence confers powers, a norm of conduct determines what the norm subject ought to do or ought to refrain from. The promulgation slot contains the source of the norm. The scope slot contains the range of application of the norm. The conditions-of-application

slot contains the circumstances under which the norm is applied. Conditions of application limit the circumstances under which the norm can be applied to the norm subject. The conditions–of–applications slot contains predicates, complex conditions (conjunctions or disjunctions of predicates), or references to act descriptions. The subject slot contains the person or group of persons to whom the norm is addressed. The legal modality slot contains the legal modality (ought, ought not, may, can). Finally, the act identifier slot contains a reference to an act description.

The third type of frame, concept frames, consists of a concept, concept type, priority, promulgation, scope, conditions, and instances (Van Kralingen 1995, p. 106). The concept slot contains the name of the concept that is referred to. The concept type slot indicates what type the concept belongs to. Concept types are definitions, deeming provisions, factors or meta-concepts. The promulgation slot indicates where the original text of the concept description can be found. The scope slot contains the area of application for the concept. The conditions slot contains the application conditions for the concept (its intension). The instances slot is a set of instances of the concept (a part of its extension).

5.3.3 Verheij and Hage's abstract model of the law

Verheij and Hage (1997) develop a model consisting of three main elements: states–of–affairs, events, and rules. States–of–affairs are (possible or real) situations that can be described by descriptive sentences. Events bring about changes in the current states–of–affairs. Rules express direct relations among states–of–affairs. Thus, there are two different relations among states–of–affairs: either the transition from one state–of–affairs into another is *caused* by an event, or the transition from one state–of–affairs into another is the result of the application of a rule, which makes the former *constitute* the latter.

States-of-affairs have four characteristics. First, they have different *temporal aspects*. They can be either temporary (Wim Kok is the Prime Minister of the Netherlands, but this state-of-affairs has a beginning and an end in time) or durable (Wim Kok will always be the Prime Minister who was in function in the year 2000). Temporary states-of-affairs can be either normal (such as the example given) or momentary. Momentary states-of-affairs obtain only for a very brief period of time. This kind of states-of-affairs is expressed in English by, *e.g.*, the simple past: 'the bullet hit JFK's head' can never be expressed as 'the bullet was hitting JFK's head'. Second, states-of-affairs often *supervene* on other states-of-affairs. This means that a state-of-affairs exists only due to the presence of some other state-of-affairs. Third, states-of-affairs can have different modali-

ties, of which three are discerned: *anankastic*, *deontic*, and *probabilistic* states-of-affairs. Anankastic states-of-affairs concern the necessary, the possible and the impossible. Deontic states-of-affairs concern what is obligated, forbidden, or permitted. Probabilistic states-of-affairs concern what is probable, certain, or uncertain. Fourth, states-of-affairs differ according to the point of view taken. Points of view can differ across disciplines (*e.g.*, in case of a biological as opposed to a legal point of view), or within disciplines (*e.g.*, in case of a criminal law view as opposed to a civil law view) (*ibid.*, p. 5-7). Events, as noted above, bring about changes in states-of-affairs. Occurrences of events are momentary states-of-affairs. Events can cause different states-of-affairs.

5.3.4 A comparison of the three models

Although all three models aim at providing a conceptual framework for the representation of the law, they are very different. In Valente's (1995) model, primary focus is on different types of knowledge. Van Kralingen's (1995) model provides a detailed analysis of three main entity types in the legal domain. Verheij and Hage (1997) also provide three main entity types in the legal domain, but these focus on the dynamic properties of a legal system. The normative element, clearly present in both Valente's and Van Kralingen's models, respectively as normative knowledge and as norm frames, lacks in Verheij and Hage's model, although the latter model contains a clear account of rules.

In the models discussed, I have observed the absence of the following three elements, that are valuable in the framework of this thesis.

- In the models, there is no room for accommodating different legal-philosophical stances on the nature of law. Such stances influence what entities are to be distinguished in the domain, and the way in which knowledge about these entities is acquired.
- The difference between entities and knowledge about entities is not made within the ontologies discussed; there is no explicit distinction between an ontological level and an epistemological level.
- There is no way of distinguishing the normative difference between belief and knowledge, partly because there is no way to express attitudes (such as believing), and partly because justifying entities, such as reasons, are not regarded as a separate category – although they would fit in nicely with Hage and Verheij's (1997) ontology.

In the next chapter, I give an overview of a knowledge-based ontology of the legal domain, which can accommodate (1) legal-philosophical stances on the

nature of law, (2) the distinction between the ontological level and the episte-
mological level, and (3) the difference between belief and knowledge.

6 A knowledge-based model of the law

This chapter describes a knowledge-based model of the law. It provides a framework for the representation of different ontological and epistemological claims bearing on the representation of knowledge about the legal domain. The chapter builds on the findings of the previous chapters. The chapters 2 and 3 showed what legal knowledge amounts to. Chapter 4 provided different views on the ontology of law. Chapter 5 indicated what difficulties we run into when we try to represent meaning. The ideas acquired in these chapters are combined into a ontology of law (in the AI sense of this term, *i.e.*, a conceptual model). The ontology can accommodate different views on the existence of law. It can be used for the construction of a more specific ontology of law for a certain legal system with an accompanying view on the existence of legal entities. It should be noted, however, that the applicability of the ontology depends on the ontological view involved, and on the degree to which we can make that view explicit.

Although the ontological status layers adopted are based on widely acknowledged views in legal philosophy, this does not guarantee the availability of *all* status layers and object categories needed for specific ontologies. Rather, the idea is that building an ontology of law is a theory-laden activity. Hence, an ontology should provide sufficient opportunities for accommodating different views on the existence of the law. The ontology described in this chapter consists of the following six basic categories: entities, ontological status layers, epistemic roles, relations, acts and facts. Each of these categories has legally-relevant and legal representatives. In each of the sections 6.1 through 6.6, the legally-relevant and legal representatives of the basic categories of the knowledge-based model of law are listed and explained briefly. In section 6.7, an overview of the model is provided in the form of a number of tables. Finally, in section 6.8, I discuss the question what it takes to know the law.

6.1 Entities

The entities distinguished in the knowledge-based ontology do not have additional ontological status layers or epistemic roles. They are considered as the basic entities in the ontology, and can be assigned ontological status layers or epistemic roles. For instance, a belief (a basic entity) can be qualified as a reason (an epistemic role, *cf.* section 6.3).

6.1.1 Legally-relevant entities

Below, seven legally-relevant entities are listed and briefly explained: sentences, statements, propositions, beliefs, rules, concepts and artefacts.

Sentences – Sentences are meaningful sequences of words, put in a natural language. Sentences can be uttered by persons. Some types of sentences, uttered by persons, are statements.

Statements – An assertive sentence ('It is nine o'clock') becomes a statement ("It is nine o'clock") if it is uttered by some individual. An assertive sentence is a sentence that makes a claim.

Propositions – Propositions amount to what is expressed by assertive sentences: in case of the example, the assertive sentence ('it is nine o'clock') expresses that it is nine o'clock.

Beliefs – Beliefs are propositions that are within the reach of a propositional attitude (believe). Propositional attitudes express mental states or mental attitudes towards a proposition. In case of the example, the proposition (it is nine o'clock) falls within reach of the propositional attitude (I believe), thus resulting in the belief 'I believe that it is nine o'clock'. Characteristics of beliefs are listed in table 6.2.

Rules – Rules are linguistic expressions that relate two types of phenomena to each other (such as 'guiding dogs are not regarded as pets').

Concepts – A concept is an abstract idea, containing the characteristics of concrete instances of the concept. In most cases, a concept has a name. There are concepts without concrete instances to refer to. Concepts that refer to institutional entities do not always have tangible referents. The meaning of a concept (the abstract idea) is sometimes equated with its *intension*, *i.e.*, the necessary and sufficient conditions that – if fulfilled – constitute an instance of the concept.

Artefacts – Artefacts are tangible objects made by human beings. These should be distinguished from both natural objects (which are tangible but not made by human beings) and institutions (which are made by human beings but are not tangible).

6.1.2 Legal entities

In this subsection, eight legal entities are discussed: legal principles, legal norms, legal decisions, legal systematisations, judicial interpretations, judicial classifications, legal rules, and legal concepts.

Legal rules – Legal rules, just as rules, are expressions that relate two types of phenomena to each other. They are *legal* in the sense that they are valid and that they are part of a legal system.

Legal principles – Legal principles are general values that apply to certain areas within the law. The meaning of a legal principle is not fixed; it is often further specified when it is called upon in a specific discussion or case. An example of a principle is 'freedom of contract' within the law of obligations.

Legal norms – Legal norms are entities that explicitly guide some kind of behaviour or action. The presence of a norm can be identified by the normative modalities it contains. Legal norms are based on legal sources, and they govern behaviour. Some legal norms are explicitly expressed in legal sources (they contain normative modalities), other norms can be derived from legal rules (the rules can be formulated in a different way, using normative modalities). Rather than legal rules, which contain conditions for some kind of qualification, legal norms express duties, permissions, and competences (*cf.* subsection 5.3.2).

Legal decisions – Legal decisions are decisions in some legal case, made by a qualified legal authority under suitable circumstances (as prescribed in procedural law).

Legal systematisations – A legal systematisation is the result of an attempt to impose a more general structure on a relatively unstructured part of a legal system, or the implicit presence of such a structure in a part of a legal system. To attain or uncover such a structure, one has to employ or discover ordering principles or ordering rules. An example of an ordering principle is 'lex superior derogat legi inferiori'.

Judicial interpretations – A judicial interpretation is the result of an assessment that a judge (or a different qualified authority) makes of some complex of (legal) rules and (legal) norms. It is attained through a choice between conflicting norms and rules, or a specification or clarification of the meaning of norms and rules. The process by which a judicial interpretation is attained, is called interpretation (see also section 6.5). The way the process of interpretation works is partly unclear; the best way to clarify a judicial interpretation is to provide proper justification for it.

Judicial classifications – Somewhat similar to judicial interpretations are judicial classifications. Classification also involves a kind of interpretation. To classify some event or state-of-affairs as a legal act or a legal institution requires us to find and apply the conditions under which this classification can take place. It is not always clear what exactly these conditions are. Even if the conditions are known, they cannot be unambiguously applied to events or states-of-affairs. Although in this respect, classification and interpretation are similar phenomena, the object of the process of classification differs from that of interpretation. Classifications apply to facts, whereas interpretations apply to rules and norms.

Classifications abstract from reality, while interpretations make rules and norms more concrete.

Legal concepts – Legal concepts are concepts whose definition is given (at least partly) in formal sources of law. Additionally, they are used in a legal context. Many legal concepts have a meaning that differs from their meaning in a non-legal context. But even if the meaning of a concept does not differ from its natural-language meaning, its explicit introduction into a legal context provides it with the status of a legal concept (*cf.* subsection 5.3.2).

6.2 Ontological status layers

Ontological status layers are the ontological characteristics of legally-relevant and legal entities, acts, and facts (the latter two are described in sections 6.5 and 6.6. Sometimes more than one status layer applies to an object. For instance, a legal decision can be both valid and efficacious.

6.2.1 Legally-relevant ontological status layers

This subsection provides a brief overview of legally-relevant ontological status layers: existence, constitution, and recognition.

Existence – Existence amounts to 'real being'; the existence of some object. At this point, I ignore the philosophical problems attached to this notion. For an elaboration on the issue of existence, I refer to section 4.1. The issue is not important here, because existence fits in rather smoothly with a common sense notion of existence, in which both physical existence and the existence of ideas can be accommodated.

Constitution – When an object is constituted, it means that it is emergent on some other object, and that there is a relation that has established the constitution between the two objects, which has not been ended yet. An actual constitution relation can, *e.g.*, exist between a fact and a social fact; slapping a person in the face (a fact) constitutes an insult (a social fact), after the fact is qualified as such.

Recognition – An object is recognised if its existence is accepted by an individual, or a group of individuals.

6.2.2 Legal ontological status layers

In addition to the legally-relevant ontological status layers discussed in the previous subsection, there are four legal ontological status layers: legal efficacy, legal validity, and legal recognition.

Legal efficacy – A legal entity (a rule or a norm) is efficacious if it has a noticeable effect on the mental states and/or behaviour of a significant part of a community, and the effect amounts to a higher degree of conformity with the rule or norm (*cf.* subsection 4.6.4).

Legal validity – A legal entity is valid if it has been issued by a proper authority, or if it derives its validity from an entity that has been issued by a proper authority through (systemic) recognition (for instance, the automatic issuing of a speeding ticket can be deemed a valid legal decision) (*cf.* subsection 4.6.1).

Legal recognition – A legal entity is recognised if it is accepted by an individual in his capacity as a legal professional, or by a legal institution. The two types of legal recognition are authoritative recognition and systemic recognition. The status layers arising from these types are authoritatively recognised and systemically recognised (*cf.* subsection 4.6.3).

6.3 Epistemic roles

Just as existence layers, epistemic roles are claims regarding objects. However, these roles signify the function these objects have in acquiring or justifying knowledge. An epistemic role is applied to an object by employing the counts-as relation: for instance, a statement may *count as* a reason. In that case, the statement acquires an additional status layer or characteristic, which enables us to express the specific (epistemic) role that an entity plays.

6.3.1 Legally-relevant epistemic roles

Legally-relevant epistemic roles, discussed in this subsection, are reasons, defeaters, factual knowledge and practical knowledge.

Reasons – Reasons are statements, propositions or facts that are employed for the explanation or justification of some other statement, proposition or fact. An extensive typology of reasons is given in subsection 4.5.2.

Defeaters – A subclass of reasons is formed by defeaters. Defeaters are negative reasons, *i.e.*, they attack some belief or reason in such a way that it is no longer correct.

Factual knowledge – If an object complies with certain so-called knowledge criteria, it counts as factual knowledge.

Practical knowledge – If an ability complies with certain knowledge criteria, it counts as practical knowledge.

6.3.2 Legal epistemic roles

Apart from the legally-relevant epistemic roles, there are two legal epistemic roles: factual legal knowledge and practical legal knowledge.

Factual legal knowledge – If a legal object (a belief, a proposition) complies with certain so-called knowledge criteria, it counts as factual legal knowledge.

Practical legal knowledge – If a legal ability complies with certain knowledge criteria, it counts as practical legal knowledge.

6.4 Relations

Relations express interdependencies among phenomena. A relation may state the consequences of some event, or impose new roles on existing objects. For instance, the counts-as relation may impose epistemic roles on entities.

6.4.1 Legally-relevant relations

Three legally-relevant relations are briefly discussed in this subsection: causation, counting as, and recognition.

Causation – Causation is a two-place relation between phenomena: a certain event causes a different event or state-of-affairs if the occurrence of the former event is necessary and sufficient for the occurrence of the latter. Physical causation is described (not prescribed) by the laws of physics. Intentional causation amounts to the patterns of perceived causation in human behaviour (*cf.* Searle 1999). For instance, we say that some event (an attack) caused someone to behave in a certain way (strike back). This event, then, can be regarded as a necessary and sufficient reason for that person's behaviour.

Counting as – Counting as is a two-place relation that associates two types of phenomena to each other. The relation causes the phenomena of the first type to be regarded as phenomena of a different type, sometimes under additional contextual conditions. There are two types of legal counting as: establishment and maintenance. The first relation establishes the constituting relation, while the latter maintains it through time. The first relation triggers an event, while the latter constitutes a state-of-affairs.

Recognition – Recognition is a two-place relation that says what individual, what group of individuals, or what institution recognises a certain phenomenon. The object of recognition may be a fact (an event or a state-of-affairs), or an instantiated counts-as relation. Three types of legally-relevant recognition are distinguished: individual, collective and systemic. Individual recognition is recognition of some object by one person. Collective recognition is the recog-

nition of some object by the majority of some relevant community. Systemic recognition is the recognition of an object within a rule system.

6.4.2 Legal relations
Legal relations are legal causation, legal counting as, and legal recognition. They are briefly discussed below.

Legal causation – Legal causation is actually a subtype of the legal counts-as relation. Legal causes are events or state-of-affairs qualified as causes for certain phenomena within a legal context. What counts as a cause in a legal sense depends on the specific demands that are imposed in a legal system on, for instance, the attribution of a consequential event to the behaviour of a person.

Legal counting as – Legal counting as is similar to the counts-as relation described above, except that at least one legal type should be part of the relation. There are two types of legal counting as, just as there are for the normal counts-as relation: establishment and maintenance. The first relation establishes the legal constitutive relation, while the latter maintains it through time. Thus, the first relation triggers a legal event, and the latter constitutes a legal state-of-affairs.

Legal recognition – Legal recognition is similar to ordinary recognition, with the additional demand that the recognition needs to take place by a legal institution or by a legal professional in his capacity as such. There are two types of legal recognition: authoritative recognition and systemic recognition. Authoritative recognition occurs if a legal professional or legal institution has the discretionary power to recognise some phenomenon, and systemic recognition occurs if the recognition takes place by the application of a legal rule.

6.5 Acts

Acts play a role in, *e.g.*, the interactions between the parties in a trial. They lead to a decision through a process of asking questions, assessing statements, interpretation *etc.* Most of these are legally-relevant acts, and some of these are legal acts, in the sense that they have a legal status (they are qualified as legal acts). Apart from these, there are, *e.g.*, acts leading to breaches of norms. In this section, only a small selection of legally-relevant acts and legal acts is provided.

6.5.1 Legally-relevant acts
Legally-relevant acts briefly discussed in this subsection are: applying rules, making decisions, making systematisations, making interpretations, and making classifications.

Applying rules – The act of applying rules amounts to deciding whether the conditions in the antecedent of the rule are fulfilled, and if so, recognising the occurrence of the consequent of the rule.

Making decisions – Making a decision means to choose among alternatives, and, if necessary, to make up those alternatives. The actual decision is sometimes accompanied by a justification, which may add to its acceptability.

Making systematisations – Systematisations are made by applying ordering criteria. They serve to increase the perceived coherence in a system of rules. The nature of systematisations depends only partly on the system they are about. Rather, they depend on the choice made for applicable criteria. Making a systematisation means to establish ordering criteria and apply them.

Making interpretations – Making interpretations is the act that leads to the establishment of an interpretation of some kind.

Making classifications – Making classifications is the act that leads to the establishment of a classification.

6.5.2 Legal acts

In addition to the legally-relevant acts explained above, I discuss legal acts in this subsection: applying legal rules, making legal decisions, making legal systematisations, making judicial interpretations, and making judicial classifications.

Applying legal rules – The application of legal rules is only effective if it is performed by a qualified authority under certain conditions. The application of legal rules may result in classifications and decisions.

Making legal decisions – Making legal decisions is, of course, also bound to legal conditions. These conditions say what authorities may make such decisions under what conditions.

Making legal systematisations – Legal systematisations are made by a process of applying ordering criteria. Such systematisations have a legal nature insofar as they are based on criteria that emerge from the legal system itself, or if they are generally recognised as correct by legal professionals.

Making judicial interpretations – Judicial interpretations are made under certain conditions established in the legal system. The freedom of interpretation (such as by analogical reasoning) is, for instance, more restricted in the domain of penal law than in civil law.

Making judicial classifications – The act of making a judicial classification serves to provide a legal status to a fact. The process of making such classifications is thus bound to criteria concerning the person or institution carrying out the classification.

6.6 Facts

Facts are events and states-of-affairs as they occur in reality. An event brings about a change (for instance: 'John signed a contract on March 23rd 2001'), whereas a state-of-affairs refers to a static situation (for instance: 'The contract is valid according to Dutch law of obligations'). Facts involve objects, the characteristics of those objects, characteristics of characteristics, and relations between objects and between characteristics. An object alone is never a fact. Something has to be predicated for that object. For instance, the existence of an object is a fact, as well as the object's lying on the table.

6.6.1 Legally-relevant facts

The category of legally-relevant facts is represented by four types of facts: brute facts, recognised facts, conventional facts, and institutional facts.

Brute facts – Brute facts are facts that are tangible in some respects, *i.e.*, they involve real objects. Events are momentary incidents that may exert some form of influence on states-of-affairs or other events. States-of-affairs are relatively stable situations. Events and states-of-affairs both involve a certain characteristic or an act of some object or agent, or a relation between objects or agents (*cf.* subsection 5.3.3).

Recognised facts – Recognised facts are events and states-of-affairs that are recognised as such by an individual or a group of individuals. They are not tangible (as far as their existence as recognised fact is concerned). The fact that someone has an actual personal belief, for instance, is an individually recognised fact (if someone believes that the earth is flat, then the fact that he believes that the earth is flat is individually recognised). An actual collective belief is a collectively recognised fact.

Conventional facts – Conventional facts are events and states-of-affairs that are based on agreement within a community, or at least among a considerable part of that community.

Institutional facts – Institutional facts are events and states-of-affairs that are (as far as their institutional existence layer is concerned) not tangible. They arise from the application of institutionalised qualification rules (*i.e.*, qualification rules that are valid, seen from some system of rules).

6.6.2 Legal facts

In this subsection, I briefly discuss three types of legal facts: recognised legal facts, conventional legal facts, and institutional legal facts.

Recognised legal facts – Recognised legal facts are events and states-of-affairs

that have acquired their legal status by some form of recognition.

Conventional legal facts – Conventional legal facts are events and states-of-af-
fairs that are based on agreement within a community, or at least among a con-
siderable part of that community. The agreement should be supported by a part
of the legal system, or in turn, should constitute a part of the legal system.

Institutional legal facts – Institutional legal facts are events and states-of-affairs
that are (as far as their institutional existence layer is concerned) not tangible.
They arise from the application of legal qualification rules.

6.7 Overview of a knowledge-based ontology of law

In the following overview of the knowledge-based ontology of law, I name dif-
ferent types of relations, entities, ontological status layers, epistemic roles, acts,
and facts that we find in the legal domain (subsection 6.7.1). After that, I give an
overview of the characteristics of beliefs (subsection 6.7.2), factual knowledge
(subsection 6.7.3), and practical knowledge (subsection 6.7.4).

6.7.1 An overview of the basic categories
Table 6.1 provides an overview of the six basic categories of the model: entities,
ontological status layers, epistemic roles, relations, entities, acts, and facts.

type	legally-relevant	legal
entities	sentences statements propositions beliefs artefacts rules concepts	legal rules legal principles legal norms legal decisions legal systematisations judicial interpretations judicial classifications legal concepts
ontological status layers	existence constitution recognition	legal efficacy legal validity legal recognition
epistemic roles	reasons defeaters factual knowledge practical knowledge	factual legal knowledge practical legal knowledge

Table 6.1. Basic categories of the model

type	legally-relevant	legal
relations	causation counting as recognition	legal causation legal counting as legal recognition
acts	applying rules making decisions making systematisations making interpretations making classifications	applying legal rules making legal decisions making legal systematisations making judicial interpretations making judicial classifications
facts	brute facts recognised facts conventional facts institutional facts	recognised legal facts conventional legal facts institutional legal facts

Table 6.1. Basic categories of the model (continued)

6.7.2 Beliefs

Beliefs are further specified according to their content source, origination source, subject, and object. Table 6.2 lists possible values for these characteristics.

source (content)	source (origination)	subject	object
legal statutes legal principles case law	perception memory reason testimony interpretation	individual community authority	legal legally-relevant abstract concrete

Table 6.2. Beliefs

6.7.3 Factual knowledge

Table 6.3 provides a specification of factual knowledge. Factual knowledge can be specified according to the characteristics given in the table.

source (content)	source (origination)	subject	object	applicable criteria
legal statutes legal principles case law	perception memory reason testimony interpretation	individual community authority	legal legally- relevant abstract concrete	truth justification reliability coherence

Table 6.3. Factual knowledge

6.7.4 Practical knowledge

Practical knowledge can be specified according to four different characteristics: its source, subject, object and the applicable knowledge criterion. Table 6.4 provides an overview.

source	subject	object	applicable criterion
experience learning innate	individual community authority	legal legally-relevant	reliability

Table 6.4. Practical knowledge

6.8 Knowing the law

To know the law, one needs both practical and factual knowledge related to this object (*cf.* Mommers 1999). Practical knowledge includes the capacities to understand and interpret legal statutes and case law, and to classify legally-relevant acts and facts. Factual knowledge consists of statements, propositions and beliefs that comply with certain criteria. These beliefs are about sources of law and about legally-relevant acts and facts. Because of the different types of knowledge that we find in the legal domain, there is not a uniform set of criteria to determine whether some entity (a capacity, a statement, a proposition, or a belief) is knowledge.

Rather, we should regard 'knowledge' as a term that shows the characteristics of family resemblance (put forward by Wittgenstein, *cf.* Wittgenstein 1953). For such a term, there is not a fixed set of necessary and sufficient conditions that it has to comply with in order to be qualified as knowledge. Instead, the term 'knowledge' has a large number of uses that we should take into account by offering a detailed (but not exhaustive) typology of knowledge about the legal domain. After this, we can decide whether (and if so, under what conditions) each of the resulting types is qualified as knowledge.

This attempt at qualification of some entity as knowledge arises from the basic assumption that to have knowledge is more valuable than merely to have a capacity or a belief. Thus, the qualification of some entity as knowledge expresses a value judgement about that entity. The value judgement is the consequence of a two-step process. First, we classify an entity as belonging to a certain type (for instance abstract factual belief). Then, we apply the knowledge criteria applicable to that type of entity. If the entity fulfils the applicable criteria, it may be qualified as knowledge.

As I noted before, the three most important dimensions of knowledge are its acquisition, its object, and its justification. These three dimensions provide us both with the criteria to assess to which type of knowledge a certain entity potentially belongs, and the criteria to assess whether it actually is knowledge. These criteria were elaborated on in section 2.3 (namely truth, proper justification, reliability, and coherence). In the current section, I discuss the specific issue of what it takes to know the law. This question is closely related to the ontological claims made with respect to the law.

The framework presented in this chapter is suitable for representing the existence claims for the major positivist ontological positions in legal philosophy. The different ontological status layers distinguished in section 6.2 suffice for that purpose. To give an example, I explain what it takes to know the law in a conventionalist-*cum*-institutional approach. If we want to know the law, we will first need to find out what it is according to this approach. If 'the law' is understood as a set of legal rules whose existence is established by convention and authoritative recognition (*i.e.*, legal rules are conventionalist-*cum*-institutional facts), then we can know the law acquiring knowledge about the content of that set of legal rules, and acquiring knowledge about its actual establishment by convention and recognition.

Checking whether a rule is constituted by convention is an empirical matter; we have to ask a representative part of a group whether they agree with that rule. Checking whether a rule is constituted is a different matter. To do this, we have to find other rules of law validating the rule, or we have to establish the existence of a valid discretionary power on behalf of which the rule was issued. It seems absurd to demand that someone executes an empirical investigation in order to know a legal rule, but this follows naturally from taking a certain ontological stance with respect to legal rules. To avoid such absurdities, we could take a different ontological stance (a Kelsenian approach does not involve convention, and thus it eliminates the problem of establishing whether there are conventions), or we may (like Hart) acknowledge some ultimate rule of recognition, which confers the conventional element onto the whole system of rules, without the necessity to establish the conventional recognition of each individual rule. Thus, using the building blocks provided by the knowledge-based model of the law, it is possible to express the epistemological and ontological claims attached to different views on the law.

7 Dutch penal law: a domain description

In this chapter, I explain the characteristics of Dutch penal law in order to show how the framework presented in the previous chapter can be applied to an actual part of the law. Of course, I do not intend to represent the complete body of Dutch penal law. Rather, the discussion will serve to show the expressive power of the knowledge-based ontology. To attain this goal, in section 7.1, the questions are discussed on the basis of which the guilt and punishability of a suspect are determined in a trial. In section 7.2, I discuss some important principles that apply to penal law. In section 7.3, central concepts of penal law are dealt with. Finally, in subsection 7.4, the characteristics of a subdomain of Dutch penal law are described, namely violent crimes and crimes against life.

7.1 Guilt and punishability

In the following discussion of law of criminal proceedings, I restrict myself to the questions asked during the trial, starting with the formulation of the summons. The summons contains a charge (art. 261 sv) with the time and place of the fact that underlies the charge. The charge should contain the constituents of the offence description (Franken *et al.* 1995, p. 382). During the trial two groups of questions should be answered. The first group consists of four formal questions (*ibid.*, p. 382). The first question of this group is whether the summons is valid. If it is not valid, the summons is void. The second question is whether the judge is competent. If he is not, he is incompetent. The third question is whether the public prosecutor is amenable. If he is not, he is not amenable. The fourth question is whether there is a reason for an adjournment. If there is such a reason, the trial is adjourned.

Hereafter, the second group of four questions should be answered (*ibid.*, p. 382). The first question of the second group is whether it has been proved that the suspect has committed the fact as charged. If not, the suspect is acquitted. The second question is whether the proved fact is punishable. If it is not, the suspect is exempted of legal proceedings. There are different kinds of conditions that can be satisfied so as to yield punishable facts (Franken *et al.* 1995, p. 377). On the one hand there are conditions for punishability that are part of the offence description: the *constituents* of the punishable fact. On the other hand there are conditions for punishability that are not an explicit part of the offence description: *elements* of the punishable fact. Together they are called *conditions for punishability*. The third question is whether the suspect is punishable. If he is

not, he is exempted of legal proceedings. The fourth question is what punishment should be imposed. A denial of the first question thus leads to acquittal. A disaffirmative answer to the second and/or third question leads to exemption of legal proceedings (*ibid.*, p. 382).

I should clarify the notion of *evidence* that we need to determine whether the suspect has committed the fact as charged. The Dutch phrase 'wettig en overtuigend bewijs' can be translated into English as 'legitimate and convincing evidence'. The import of the phrase is that the judge should be convinced of the fact that the suspect has committed the act as charged. Also, there should be legitimate evidence for that fact. The judge determines to what extent he attaches significance to the statements given by, for instance, a suspect (*ibid.*, p. 383).

When the answers to the two groups of questions are given, we can determine the punishment that should be imposed on the suspect. The themes that are induced by the two groups of questions in this section are discussed in the subsequent sections, starting with a discussion of three groups of meta-principles.

7.2 Central principles of penal law

Besides the specific regulations that we find in statute law, we can also discern a number of principles. These precede the specific regulations; any application of a regulation that does not satisfy these principles may be disputed. I divide these principles into three groups: legitimacy and legality (subsection 7.2.1), subsidiarity and proportionality (subsection 7.2.2), and priority principles (subsection 7.2.3).

7.2.1 Legitimacy and legality

The first group contains the legitimacy and legality principles. The legitimacy criterion protects the suspect against an undesirable (*i.e.*, illegitimate) application of law. The legality criterion protects the suspect against an illegal application of law.

The legitimacy principle for penal law, on the one hand, demands that the way a judgement is argued for, is acceptable for us. The principle of legitimacy is basically what we call the *justification* criterion in philosophy. Without justification of a judgement we cannot say why a judgement is right. In this sense we can compare the legitimacy principle with the justification criterion for knowledge. Without justification we can have true knowledge, but the truth would be beyond our intellectual capacities to grasp if we did not have evidence for the

truth of a piece of knowledge. Without legitimacy a judgement may or may not in some sense be right, the rightness of the judgement only comes within our reach if we can legitimise it.

The legality principle for penal law, on the other hand, can be divided into five sub-criteria (Franken *et al.* 1995, p. 372). First, no fact is liable to penalty without applicable legislation. Second, no penalty can be imposed without a law permitting the imposition. Third, no penalty can be imposed without a *preceding* law, a law that has been effectuated before the act has been performed. Fourth, analogous reasoning is not allowed in the argumentation of the judge. Finally, laws should be sufficiently clear so that in principle anyone is able to understand them.

7.2.2 Subsidiarity and proportionality

The second group contains two criteria that can also be characterised as 'meta-principles'. These are subsidiarity and proportionality. Subsidiarity can be described as the rule that the least drastic means for attaining a certain goal is the most appropriate one. Proportionality says that the means should be in a reasonable relation with the intended goal (Nijboer 1995, p. 138). These principles are especially important when a person has committed a certain offence and the question arises whether some exemption ground from punishment may be applicable.

7.2.3 Priority principles

A third group of criteria helps us to determine what regulations we should apply in the assessment of a specific case. These criteria are the following. First, *lex specialis derogat legi generali;* special laws are preferred above general laws. Second, *lex posterior derogat legi priori;* laws that are promulgated at some point in time are preferred above laws about the same subject that were promulgated at some point earlier in time. Third, *lex superior derogat legi inferiori;* laws that are promulgated by some authority are preferred over laws on the same subject that are promulgated by a lower authority.

7.3 Central concepts of penal law

In the current section, I deal with concepts that play an important role in penal law. These are: facts (subsection 7.3.1), culpability (subsection 7.3.2), exemption grounds (subsection 7.3.3), causality (subsection 7.3.4), participation (subsection 7.3.5), and attempt (subsection 7.3.6).

7.3.1 Facts

With respect to the meaning of the concept 'fact' in Dutch penal law, I discuss three different possible situations that are described in artt. 55, 56, and 57 SR. Art. 55 par. 1 SR says that when a fact can be brought under more than one regulation, only one of these regulations can be applied, and this should be the one with the severest penalty. This article constitutes the situation in which two or more offences arise from the same act. Art. 56 par. 1 describes a so-called continuous act. A continuous act occurs when comparable acts are performed subsequently, and they all arise from the same intention. Art. 57 par. 1 SR says that when two or more criminal offences can be regarded as being independent of each other, they should each be punished, though the total penalty (imprisonment or detention) given should never exceed the severest penalty for an individual offence by more than one third. The situation in which one sentence is pronounced for various offences occurs when different acts have been carried out that cannot be classified under the same legal rule (Franken *et al.* 1995, p. 69).

As should be clear from these articles, it is important to know whether a fact can be seen as distinct from another fact. Until 1932, the Dutch Supreme Court considered one fact to be one "material behaviour", one "externally perceptible deed" (Jörg en Kelk 1994, p. 113). The old notion of 'fact' included a simplified view on facts, as a fact actually includes human intentions, interactions with the world, and the situation of the world itself (*ibid.*).

In a 1932 judgement (HR February 15, 1932, NJ 1932: 289) the Supreme Court formulated a different approach towards the notion of 'fact'. Two facts can be distinguished, even if they have taken place at the same place and time, if they can be seen as separate facts (their co-occurrence is not essential), if one fact is not part of the other, and if one fact cannot be seen as part of the circumstances under which the other took place (*ibid.*, p. 114). In practice, this means that a drunk driver whose car has defective headlights, can be prosecuted for two facts, a case in which one sentence is pronounced for two offences.

The notion of 'fact' in the above-mentioned articles is different from the notion of 'fact' in the *ne bis in idem* principle, expressed in art. 68 SR. Art. 68 par. 1 SR. The *ne bis in idem* principle says that a person can never be prosecuted twice for committing one offence (except in appeal cases). There are circumstances under which two facts in the sense of articles 55 and 57 yield one fact in the sense of article 68. For articles 55 and 57, we should take a look at the factual circumstances. For article 68, we have to consider applicable law. Whether a person can be prosecuted more than once given a complex of facts, depends on the question whether we can classify elements of this complex of facts under

different legal regulations. So, if someone kills two persons driving in his car, this will be seen as two facts in the sense of article 57, and as one fact in the sense of article 68. If someone drives too fast and ignores a red traffic light, this will be seen as an occurrence of two different facts in the sense of both article 57 and 68 (*ibid.*, p. 116-119). For the sake of clarity, I will call facts in the sense of articles 55, 56 and 57 *material facts* (not to be mistaken with 'fait materiel'), while I will call facts in the sense of article 68 *judicial facts*.

How can we characterise the difference between material facts and judicial facts? Material and judicial facts are both *construed* facts, in that they are not observer-independent. Construed facts are fact representations; they still refer to what happened in reality, but they are interpretations because their complexity partly depends on the legal rule they should comply with. In case of article 68 SR (*ne bis in idem*), the complexity of what is assumed to be 'one fact' in most cases depends on the question whether one regulation and one reproach of guilt are involved. In case of articles 55, 56, and 57 SR the complexity of single fact representations largely depends on the unity of that act in reality. Note, however, that the difference between material and judicial facts is only a matter of degree. In both cases the interpretation of what a fact consists of depends on the contents of legal rules.

What kind of consequence should be attached to this difference in these two types of facts in the knowledge-based ontology of law? We cannot give an ultimate representation of reality. However, the relations between objects and individuals we discern in representing knowledge can be considered as primitives. Combinations of these primitives constitute material and judicial facts. If we have different represented situations in our model, *e.g.*, a situation in which a certain person is drunk, and another situation in which this person is driving, both occur in a comparable moment in time, and at the same place, and there is a legal rule that says that drunk driving is a criminal offence, we may construe a new material fact by combining the two original situations. The new material fact consists of the two individual situations in which the person is drunk and he is driving. Given this new material fact, we may wonder whether in this case two or more offences arise from the same act. They do, because both art. 8 WVW (drunk driving) and art. 453 SR (public drunkenness) are applicable. Art. 453 SR applies to the agent being drunk at a particular location; he need not be driving. Because in this case two offences arise from the same act, the one with the severest penalty applies (in this case, according to art. 176 par. 3 WVW, three months of detention or a fine of the third category for drunk driving). An application of the *ne bis in idem* principle here prevents the agent from being prosecuted twice.

The following schema indicates under what circumstances material facts and judicial facts can diverge (based on Jörg and Kelk 1994, p. 119).

1 concurrence of one act
      one material fact, one judicial fact
2 concurrence of multiple acts
    2.1   similar regulations
        2.1.1   similar reproach of guilt
               multiple material facts, one judicial fact
        2.1.2   dissimilar reproach of guilt
               multiple material facts, multiple judicial facts
    2.2   dissimilar regulations
        multiple material facts, multiple judicial facts

With this schema, I conclude my discussion of the notion of fact in penal law.

7.3.2 Culpability

Jörg and Kelk (1994, p. 145-146) distinguish the following four forms of culpability. Culpability can be regarded as *having committed* some act, as *carelessness*, as *intent and carelessness*, or as *reproachability*. The following schema shows the mutual relations between these notions (*ibid.*, p. 146):

1 culpability
    1.1   legal notion
        1.1.1   having committed
        1.1.2   extensive notion (intent and carelessness)
               narrow notion (carelessness)
    1.2   theoretical notion (reproachability)

The legal notion of *having committed* can be more formally characterised as the condition that a suspect did something, or refrained from doing something, leading to a punishable fact. Thus, it is mainly a descriptive notion, while *carelessness* and *intent and carelessness* are more normative notions. Carelessness can be characterised as the qualification that a person can be held responsible for the consequences of his act, considering the knowledge he had at the moment of the act (*culpa*). Intent amounts to the degree to which a person aimed at the punishable fact. Reproachability amounts to the "absence of all circumstances that could excuse the suspect" (Remmelink 1996, p. 239).

In section 7.1, I discussed eight questions that should be answered in a trial. The first question of the second group was whether it had been proved that the

suspect had committed the offence he was charged with. The corresponding culpability notion is 'having committed'. The second question was whether the proved fact was punishable. Here the *fact* is considered, not the person, so a notion of culpability is not directly at stake. The third and fourth questions were whether the suspect was punishable and what punishment should be imposed. Here either of the two other legal notions of culpability is applicable: 'intent' or 'intent and carelessness'. Moreover, also the theoretical notion of punishability, reproachability, can be considered at this point.

7.3.3 Exemption grounds

There are two situations in which a suspect can be exempted from the application of criminal sanctions. These situations are the presence of a ground for justification, and the presence of an exemption ground from guilt (Franken *et al.* 1995, p. 383). Together, these constitute exemption grounds from punishment. There are general and specific exemption grounds from punishment. This distinction cuts through the distinction between grounds for justification and exemption grounds from guilt. We find specific exemption grounds in the applicable laws themselves (*ibid.*, p. 383). General exemption grounds from guilt and grounds for justification are force majeure as necessity, self-defence, the presence of a legal regulation, and the presence of an authorised official order.

General exemption grounds from guilt are being of unsound mind, psychological force majeure, unreasonable use of self-defence, and unauthorised official order that was considered to be given authorised in good faith (Jörg and Kelk 1994, p. 156). There are written exemption grounds and unwritten ones. The ones listed above are the written ones. There are two kinds of unwritten exemption grounds, corresponding to the two categories of general exemption grounds from punishability. The first one is absence of substantive illegality. The second one is absence of all guilt.

I give an overview of general grounds for justification (artt. 39-43 SR) and general exemption grounds from guilt respectively. I start with general grounds for justification. A person can call upon *force majeure* if he commits a fact he is forced to commit (art. 40 SR). We can discern between two main categories of force majeure: *absolute* and *relative* force majeure, where the latter one is an exemption ground from guilt, which will be discussed later. Absolute force majeure occurs if a person does not act himself, because he is used as a means, or when an *absolutely impeded omission* occurs. This is, *e.g.*, the case when a person should report a birth, but is physically barred from doing so because of a flood. In this case that person cannot do anything that would cause him to fulfil an ob-

ligation, which would lead to an omission offence under normal circumstances (Remmelink 1996, p. 297). In case of absolute force majeure the main criterion that needs to be fulfilled is that the person who appeals to this ground for justification should be in a physical position such that it is impossible for him to follow the law. *Force majeure as necessity* occurs if a person is forced to violate a legal rule in order to be able to attain a more important goal. This situation occurs, for instance, if someone trespasses by entering someone's property in order to extinguish a fire.

Self-defence occurs when a person commits an offence because he must defend his own or another's body, decency, or property against assault. The assault should be both illegal and immediate (art. 41 par. 1 SR). For a formalisation of self-defence we need the following criteria. There is an external force that is sufficient to say that a person has good reasons to commit a certain fact. The external force is exercised by a human agent. The external force is a physical one. The defence act should be directed against the assaulting person. Assault should be understood as follows (Remmelink 1996, p. 316): an already commenced factual attack, or a behaviour that is an immediate threat to such an attack. Immediateness exists in the relation between the assault and the reaction of the assaulted person. There should be a direct threat (be it the threat of a new assault or a threat of a continued assault) to the assaulted person, both in place and time. If there is no such threat, there can be no case of self-defence. The act that leads to self-defence should be illegal, *i.e.*, against objective law. Finally, the reaction should be proportional with the assault; a defence of one's property with physical violence while the assaulting person does not show any violent behaviour, is not justified (*ibid*., p. 313-325)

A person is justified in his behaviour if he acts according to a *legal regulation* when he commits a fact (art. 42 SR). A legal regulation is a rule that is issued by some power that has been assigned legislative authority by way of the constitution or legislation (*ibid*., p. 330). A legal regulation should *oblige* a person to a certain act, otherwise he cannot appeal to this ground for justification. However, the criteria of proportionality and subsidiarity still apply: if, for example, the use of violence is necessary to suppress riots, and there has been given an instruction to restore order, excessive (disproportional) violence is not allowed, although the use of violence itself is permitted. A related ground for justification occurs when a person executes an *authorised official order* (art. 43 par. 1 SR). This rule is complementary to the ground for justification in which compliance to a legal regulation is demanded. It extends this rule to those cases where an official orders a person to do something that is in itself illegal. If the official gives

the order in accordance with a legal regulation, we call it an authorised official order (*ibid.*, p. 337-340). If the execution of the order occurs in a proportional way, the person committing the fact as a result of the execution of the order is justified in his behaviour.

In addition to the four grounds for justification discussed above, there are four exemption grounds of guilt. First, a person is exempted from punishment if he is not accountable for his actions. This is the case if the committed fact cannot be ascribed to him because of either defective development or ailing disorder of his mental capabilities (art. 39 SR).

Second, a person is also exempted from criminal sanctions if he can call upon relative force majeure. Relative force majeure occurs mostly in case of psychological pressure. These pressures are not absolute, in the sense that they can be ignored. To be able to appeal to force majeure successfully in such a case, it is necessary that it is *reasonable* to give in to the psychological pressure. A life threat, for example, is a reasonable motive for a bank employer to hand money from the bank's safe to a robber, which would under other circumstances yield an offence.

Apart from psychological pressure, there are four other situations that may lead to a case of relative force majeure: moral pressure, affect, necessity, and scruple (Remmelink 1996, p. 299-309). These four situations have in common that they do not always yield a case of relative force majeure. First, moral pressure occurs when there is a non-physical pressure that appeals to the moral sense of a person. Moral pressure can lead to a case of relative force majeure if the element of psychological pressure, that is obviously apparent in all cases of moral pressure, is sufficiently strong. Second, an affect is a powerful emotion that may be considered as an overpowering psychological pressure from within the person himself (*ibid.*, p. 300). Third, in a case of emergency a person can sometimes deliberate on the consequences of his deed, and still decide to act in a way that makes him commit a punishable fact (*ibid.*, p. 300). This may then be justified by the occurrence of a situation where the intervention of the person concerned is of greater importance than the punishable fact that comes about as a consequence of the intervention. This case constitutes a ground for justification. Fourth, conscientious objections occur when a person has to choose between two options where both lead to an undesired situation, just as in a case of emergency. However, one of these options leads to a moral conflict for the person involved. This, however, is not taken as a valid excuse. An exception to this consists in explicit clauses in a law that allow an appeal to conscientious objections (*ibid.*, p. 306-307).

A third exemption ground of guilt is unreasonable use of self-defence. A person can call upon unreasonable use of self-defence if he acts in a case of self-defence, while additionally the boundaries of necessary defence are crossed. This should happen as an immediate consequence of an intense emotion, which should on its turn be the result of the assault.

A fourth exemption ground of guilt occurs if a person complies to an 'unauthorised official order that was considered to be given authorised in good faith', this can be regarded as a defence precluding the application of criminal sanctions (art. 43 par. 2 SR).

7.3.4 Causality

In subsection 4.7.2, theoretical notions of causality were discussed. Instead of trying to indicate what type of causality is generally accepted in Dutch penal case law, I only attempt to derive the formal characteristics of causality. The following characteristics can be discerned: the nature of a causal chain (length and interference), the nature of causal links (physical or intentional), the nature of intentionality (collective or individual), and the nature of interference (single cause or multiple causes). A causal chain has a certain length (the number of links between the first cause and the final consequence), and shows 'interference', *i.e.*, when some fact is not just caused by one preceding fact, but by multiple facts.

The nature of a causal links can be either physical or intentional. Physical causation, described by laws of physics, is (theoretically) predictable. Intentional causation occurs when a certain consequence arises through some mental act or mental state (Searle 1999, p. 104*ff.*). For instance, I decide to call the police, after which I actually call the police. Intentionality can be either individual or collective, the former when one person is involved, the latter when a group of individuals has the same intention. The interference can be expressed by single or multiple facts causing another fact. In case of multiple facts we could try to indicate what the individual relevance of the facts is. For instance, the intentional causation of the fact that there is a *communis opinio* about the necessity of a 'zero tolerance' policy can be caused by multiple incidents that are each discussed in national television channels and newspapers.

Note that any causal chain can be made complex in different degrees, and can be constructed in varying ways, in the sense that causal links can be added, deleted, or be arranged in a different manner. What causal links are mentioned depends on their relevance for the case. For example, in the 'slipper - egg skull' judgement ('pantoffel-eierschedel', Appeal Court Amsterdam, June 14, 1939, NJ

1940: 34, *cf.* subsection 4.7.2), we could represent the relation between the slipper hitting the woman's temple and the death of the woman in much more detail: what part of her head was hit, what kind of wound and/or fracture was caused, in what way did this lead to her death *etc.*

Above, intentional and physical causes were distinguished. This distinction should not be confused with Aristotle's distinction between *efficient* and *final* causes (Aristotle's *formal* and *material* causes are irrelevant for the subject under discussion) (Shaver 1985, p. 9). Efficient causes precede the result of their workings, while final causes are teleological, *i.e.*, they are the 'target' for preceding events. Both intentional and physical causes are efficient causes. 'Intentional' should not be understood as teleological here. It should be understood in the philosophical sense, namely that of 'aboutness'. An intentional cause is a *mental* cause; if someone is angry and wants to throw a slipper, his anger leads him to throw a slipper (physically). If we take a materialist stance, in which the mental is somehow built up from the material, intentional causes are ultimately based on physical causes. In the above-mentioned judgement, a final cause in Aristotle's sense would be the intention (in the normal sense) of the man to kill his wife to. To reach this goal, he throws a slipper to her.

7.3.5 Participation

In order to be able to characterise some fact as participation, it has to fulfil a number of criteria (Jörg and Kelk 1994, p. 170). A punishable fact has been committed, and there are at least two persons of which one has committed the fact and is a perpetrator, and another one can also be characterised as a perpetrator or as an accomplice. These persons have jointly (*i.e.*, interacting in some way) committed the fact. Each of the persons involved has played an essential role in the causal chain leading to the punishable fact.

Apart from the person who actually commits a criminal offence, there are several others that we can also call perpetrators (art. 47 par. 1 SR). These are persons that make another commit a fact, that co-commit a fact, or that provoke a fact. Apart from the persons we can call perpetrators, there is also the notion of 'accomplice'. An accomplice takes part in a criminal offence in a modest way, so that we may not call him a perpetrator. Together with the three categories above the accomplice forms the group of participators (*ibid.*, p. 172). I discuss each category below.

A person *commits* a fact if he fulfils all conditions of the description of the offence. A person *makes* another person *commit* a fact if the latter one is caused to commit that fact without being punishable for that matter. For example: a nurse

gives an injection to a patient, and she is unaware of the fact (and could not have known) that the syringe has been replaced by another one with a deadly poison. The patient dies, and the nurse has executed the deadly injection. She cannot be punished for that fact, but the person who made her do it, can. If two or more persons develop a plan to commit a punishable fact and execute that plan together, they *co-commit* that fact. We can distinguish three forms of co-committing a fact (Jörg and Kelk 1994, p. 177): (1) co-committing/co-committing occurs if two or more persons commit a fact together and are all held to be co-committing that fact, (2) co-committing/committing occurs if there is a main perpetrator and another perpetrator who co-commits the fact, and (3) committing/committing, where both are considered main perpetrators.

Provocation is described in art. 47 par. 1 sub 2 SR. The description contains the following elements. First, the provocation of the fact is done on purpose; second, the provocation is deliberately realised by gifts, promises, misuse of authority, violence, threat, deceit, or the providing of an opportunity, means or information; and third, the person who performs the fact is punishable – in contrast to a person who is made to commit a fact. The provocateur is punishable too, even though he has not directly participated in the actual committing of the fact.

Complicity can both occur during or before the actual committing of the fact. The former kind of complicity occurs if the accomplice co-operates in the committing of the criminal offence on purpose (art. 48 sub 1 SR). The latter kind occurs if the accomplice provides the opportunity, means, or information for another person to commit the criminal offence (art. 48 sub 2 SR). An accomplice never takes the initiative to perform the criminal offence; in that case he would *commit* the offence (Jörg and Kelk 1994, p. 180).

7.3.6 Attempt

If a person tries to commit a criminal offence, and he does not complete the offence due to some external cause, we speak about an *attempt*. Two types of attempt are discerned: punishable and non-punishable attempt. Additionally, there are two situations in which we cannot truly speak of an attempt: voluntary retiration and unsound attempt. Punishable attempt occurs if a criminal offence is committed, whereas non-punishable attempt occurs in case of an minor offence. Art. 45 SR par. 1 says: 'An attempt to a criminal offence is punishable, whenever the intent of the perpetrator has been revealed by a commencement of implementation'.

The interpretation of 'a commencement of implementation' is problematic. A well-known case in which this phrase had to be interpreted is the 'Eindhoven

arson' (HR March 19, 1934, NJ 1934: 450), in which the Supreme Court decided that there was no attempt to a criminal offence. In this case, a man called H. has done virtually everything to ensure that the house chosen will be burnt to the ground. He has put a gas stove slanting on a pillow, surrounded by cloths and toilet paper soaked with petrol. A gas pistol is attached to the gas stove, while a rope, dangling out of the window, is attached to the trigger. H. however, did not pull the rope, possibly because people had gathered around the house after they smell the petrol (Doomen 1990, p. 106).

H. is found guilty of attempted arson and sentenced to a prison sentence of four years. The Supreme Court, as remarked above, finds the man not guilty of attempted arson. Apparently the act of preparing the house for an arson in every detail is not the same as 'a commencement of implementation'. We can approach the content of this phrase in two ways: a subjective and an objective one. Subjective attempt takes into account mainly the *intention* of a person; what objective results his intention has, is not the important matter. Objective attempt occurs when we have traceable consequences of the attempt at the moment a person is prevented from accomplishing his attempt (Jörg and Kelk 1994, p. 159).

In the argumentation of the Supreme Court we can find a more detailed account of the criteria that should be fulfilled for a case of 'commencement of implementation' to occur. The first part is the original 'objective' criterion. There should be a deed that is necessary for the intended purpose (an arson in this case), that cannot be aimed at something else, and that stands in direct connection with the intended criminal act (*ibid.*, p. 162). The second criterion is added because the objective criterion is often valid for preparatory acts, not only for implementation acts. This criterion says that the deed should lead to the intended purpose – in accordance with the rules of ordinary experience – without further intervention of the perpetrator (*ibid.*, p. 163).

This second criterion seems too restricted. If some unforeseen event occurs, and the intended purpose is not attained, there will not be a 'commencement of implementation'. In the arson case, such events could be: the gas pistol does not function, the soaked cloths do not catch fire, the gas pistol does not ignite the fire, or someone prevents the hand of the perpetrator to pull down the rope (*ibid.*, p. 163). We can generalise these circumstances in the following clause: 'If the deed does not lead to the intended purpose, this should be due to some unforeseen fact or act that prevents the intended causal chain from taking effect.' It becomes clear now how the Supreme Court justified its decision: the second criterion does not apply: the intended purpose is not attained. The extra condi-

tion that applies to this criterion, and which could have yielded a commencement of implementation after all, does not apply either. If, on the other hand, H. would have commenced pulling the rope, but some of the bystanders would have prevented him from doing that, there would be a suitable commencement of implementation.

In case of *voluntary retiration*, an attempt does not occur. Voluntary retiration occurs if a person, originally intending to perform some criminal offence, ceases the execution of the act before it has been accomplished. Jörg and Kelk (1994, p. 167) formulate a criterion for the occurrence of voluntary retiration. This criterion reads as follows (translated from Dutch): 'Voluntary retiration occurs if the perpetrator ceases his attempt based upon a new deliberation of the same circumstances'. In a judgement of the Supreme Court (HR December 15, 1992, NJ 1993: 333) a man who refrains from raping a woman is considered to be retiring voluntarily. The woman kicked him and asked what point there was in raping her. We may wonder whether the above criterion applies to this case: there are new circumstances, namely that the woman offers resistance. Obviously, there is a new deliberation of these circumstances too, but we may wonder whether the man would have stopped his attempt if the woman had not offered resistance.

A case of *unsound attempt* occurs if the nature of an attempt is such that it can never lead to a completion of the criminal offence. On the one hand there can be an unsound *means* for the intended act, on the other hand the intended act can have an unsound *object*. Both means and object can be either *absolutely* or *relatively* unsound. Relative unsoundness occurs if the means or object is usually suitable for the intended goal, but does not suffice in the case at hand. Absolute unsoundness occurs if the means or object is never suitable for the intended goal (Jörg and Kelk 1994, p. 167).

The subjective interpretation of penal law says that we should consider the intentions of perpetrator. In this view, any attempt, sound or unsound, is punishable. The distinction between relatively and absolutely unsound attempts thus makes sense only in an objective interpretation of penal law. We may say that in an objective interpretation only an absolutely unsound attempt does not lead to punishment. An example of such an attempt is a woman that tried to poison her husband by giving him extracts of copper coins (Dutch Supreme Court, May 7 1906, W 8372). However, the amounts of copper were by far insufficient to cause her husband to die (*ibid.*, p. 168-169).

7.4 A domain choice: violent crimes and crimes against life

To limit the domain that should be formalised in chapter 8, I have chosen one specific group of offences: crimes against life and violent crimes. Titles 19 to 21 SR contain the regulations on these crimes. Title 19 contains regulations on crimes against life, title 19A contains regulations on abortion, title 20 contains regulations on maltreatment, and title 21 contains regulations on the culpable causation of death or physical injury. I limit myself to four types of crimes: manslaughter (subsection 7.4.1), qualified manslaughter (subsection 7.4.2), murder (subsection 7.4.3), and culpable homicide (subsection 7.4.4).

Cleiren and Nijboer (1994, p. 749) list the main characteristics of crimes against life. First, death is a constitutive element of the offence description. Thus, the death of a person is a necessary condition for the punishability of the suspect. Second, there should be a causal connection between the behaviour of the suspect and the death of the victim. Whether the death of the victim can be ascribed to the suspect is evaluated in accordance with the causality theory of reasonable attribution. Third, the time and place of the offence are often related to the time and place of the behaviour. However, sometimes the commencement of the consequence counts as the time and place of the offence. The time of the behaviour remains relevant for, among other things, the determination of the age of the perpetrator.

I list a number of features of culpable causation of death or physical injury (Cleiren and Nijboer 1994, p. 785). These offences have in common that they are caused by reproachable behaviour. The consequence of such behaviour is not intended, but the consequence is damaging to such a degree that the suspect should be punished. This form of guilt occurs if a person thinks or knows less or is less conscientious than man in general, and his behaviour leads to death or physical injury (*ibid.*, p. 786). Only those persons that have a certain profession should be judged in relation to the common standards in their profession (*ibid.*). Finally, if two or more persons are guilty of some fact, this does not reduce the degree of individual guilt (*ibid.*, p. 787). As regards the time and place of the offence, the remarks in the previous paragraph also apply in case of culpable causation of death or physical injury (*ibid.*, p. 788).

7.4.1 Manslaughter

The article on manslaughter (art. 287 SR) contains two constituents: a person A kills a different person B, and A does this intentionally. The act of killing can be performed in an unlimited number of ways (Cleiren and Nijboer 1994, p. 750). The causality theory employed is reasonable attribution. The consequence

(death) does not have to occur immediately after the act that can be seen as the cause of the death (*ibid.*, p. 751). The interest protected is human life. In practise, it can be hard to determine when life begins and when it ends (*ibid.*, p. 751). The norm expressed in the article on manslaughter is applicable to all people. The article not only forbids a person to kill; it is also applicable if a certain course-of-events leading to the death of a person can be prevented by an intervention. Finally, the kind of intent required ranges over all common types of intent, among which conditional intent. If the originally intended means for killing a person was not employed, but another means still leads to the death of that person, there is still intent (*dolus generalis*) (*ibid.*, p. 751-752). The constituents of manslaughter can be represented in the following schema.

manslaughter occurs if:

	1	a person (a) kills another person (b)
and	2	(a) does this intentionally

7.4.2 Qualified manslaughter

Qualified manslaughter differs from manslaughter in the following respects. As an addition to the article on manslaughter, qualified manslaughter (art. 288 sr) should be followed, accompanied or preceded by a punishable fact, committed with the intention to prepare or facilitate that fact. Or it should be committed to ensure that the person himself, or his accomplices, remains unpunished, or to ensure the possession of what has been illegally taken when he was caught in the act. The schema below shows the structure of the constituents in the article on qualified manslaughter.

qualified manslaughter occurs if:

		1	manslaughter (a) is
		1.1	followed by a punishable fact (c)
or		1.2	accompanied by a punishable fact (c)
or		1.3	preceded by a punishable fact (c)
and		2	manslaughter (a)
		2.1	committed by a person (b) with the intention to
		2.1.1	prepare a punishable fact (c)
or		2.1.2	facilitate a punishable fact (c)
or		2.2	when he (b) was caught in an act (f), (b) committed (a)
		1.2.1	to ensure the person himself (b) to remain

		unpunished
or	1.2.2	to ensure his accomplices (d) to remain unpunished
or	1.2.3	to ensure (b) the possession of what has been illegally taken (e)
or	1.2.3	to ensure (d) the possession of what has been illegally taken (e)

7.4.3 Murder

The third offence discussed is murder. Murder is a special case of manslaughter. If a suspect is charged with murder, and malice aforethought cannot be proved, a judge can declare manslaughter to be proved, even if manslaughter is a part of the charge. The central constituent of murder (malice aforethought) differs from intent. It means that a person kills someone after having thought this over in quiet and conscious deliberation (Cleiren and Nijboer 1994, p. 755). A few remarks should be made on this. First, malice aforethought can occur with all kinds of intent, including conditional intent (*ibid.*, p. 755-756). Second, it can occur in combination with the presence of a nervous exhaustion in a person. Third, means, person, and time and place of the offence are not a necessary subject of malice aforethought (*ibid.*, p. 756). This means that one may kill someone with malice aforethought without beforehand having determined the exact time and place of the offence, and without having established the means with which the offence will be performed, or even the person who will be the victim. We can represent the constituents of murder as follows.

murder occurs if:

| | 1 | a person (a) commits manslaughter (b) |
| *and* | 2 | this person (a) does so with malice aforethought |

7.4.4 Culpable homicide

Fourth, culpable homicide is described in art. 307 SR. This article says that a person is guilty of culpable homicide if another person's death can be ascribed to him. Central conditions in this offence type is the criterion that the death is caused in a culpable way (*ibid.*, p. 788). A considerable amount of time may elapse between the behaviour and the death of the person. The acts of persons with a special duty of care may be more eligible for a qualification as culpable homicide (*ibid.*, p. 789). Further characteristics of this offence are given in the general remarks on culpable causation of death or physical injury. The constitu-

ents of culpable homicide can be represented as follows.

culpable homicide occurs if:

	1	the death of a person (a) can be attributed to another person (b)
and	2	the death of that person (a) is caused by person (a) in a culpable manner

This concludes the brief description of crimes against life. These offences will be used as examples in chapter 8.

8 Dutch penal law: a domain representation

In this chapter, the ontology of law developed in the preceding chapters will be applied to the domain of Dutch penal law. First, in section 8.1, a logic is discussed that was originally developed for the representation of legal argumentation, and this logic is used for the representation of the elements distinguished in the ontology. Subsequently, in section 8.2, a number of inference rules is listed that can be used for reasoning with the logical representation of the domain. Then, in section 8.3, it is explained how knowledge qualification rules work. Finally, in section 8.4, the representation of elements, the inference rules, and the knowledge qualification rules are employed to model some aspects of the domain of Dutch penal law.

8.1 A logic for legal knowledge representation

In this section, I explain how a specific language can be used to represent a domain in terms of the ontology of law developed in this thesis. The logic employed for the representation of the domain is based on Hage's reason-based logic (RBL). RBL, and the reasons for its selection are discussed (subsection 8.1.1). Next, the way this logic can be used to represent each of the categories of the knowledge-based ontology of law is explained (subsections 8.1.2 - 8.1.7).

8.1.1 Reason-based logic

Reason-based logic is described in Hage (1996, 1997). It is an extension of first-order predicate logic (from now on called predicate logic). Furthermore, it is nonmonotonic, and its ontology is far more extensive than that of predicate logic. Monotonicity means that the conclusions drawn from a number of premises will never change when new premises are added. So, if we may conclude r from two premises $p \rightarrow r$ and p, we also may conclude r from the three premises $p \rightarrow r$, p, and q. In non-monotonous logics, adding premises may lead to different conclusions. So, while we may conclude r from a set of two premises $p \rightarrow r$ and p, we may not always conclude r from a set of three premises $p \rightarrow r$, p, and q. A non-monotonous logic, like RBL, enables us to represent only a limited part of a domain without barring the possibility of reasoning with that limited amount of represented knowledge in a manner more in line with human reasoning with incomplete information.

Some of the important ontological assumptions of RBL are that there are states-of-affairs. States-of-affairs are potential facts: if a state-of-affairs obtains, it

is a fact; if not, it is a non-fact (Hage 1997, p. 131). Reasons are facts that plead pro or con states-of-affairs (*ibid.*). These basic ontological assumptions deviate from the assumptions in this thesis. Events and states-of-affairs are concrete facts. If an event is not instantiated, it is an event *type*, and if a state-of-affairs is not instantiated, it is a state-of-affairs *type*. Events and states-of-affairs are both facts. An uninstantiated fact is a fact *type*.

Hage's model of legal reasoning consists of a more complex account of rule application than most other approaches. Instead of the deductive model, in which fulfilment of the conditions of a rule leads to the establishment of the consequence of that rule, Hage assumes that the following steps are necessary (*ibid.*, p. 152). First, the *applicability* of a rule is established, *i.e.*, it is checked whether, given the circumstances of a case, a rule can be applied. Second, if this is the case, the potential applicability of the rule is a reason to apply the rule. Third, it is determined whether the reasons in favour of the conclusion of the rules outweigh the reasons that plead against that conclusion. Fourth, if this is the case, and thus, if the rule ought to be applied, the conclusion of that rule can be derived (*ibid.*, p. 152-153, p. 161-164).

In brief, this four-step process amounts to: (1) establishing potential applicability, (2) transforming potential applicability into a reason for actual applicability, (3) weighing reasons pro and con the conclusion of the rule, and (4) deriving the conclusion of the rule from the fact that the rule ought to be applied. Rule application thus becomes a matter of generating and weighing reasons. This approach constitutes a way of expressing the defeasibility of legal reasoning: the applicability of a rule is not the only factor contributing to the establishment of the rule's consequence part. Instead, other reasons may block the establishment of the consequence, and thus the nonmonotonous nature of legal reasoning is expressed.

The ideas behind RBL will form the basis for the formalisation of the functions and entities that were introduced in chapter 6. First, just as rules are represented as terms in RBL, my extension of RBL accommodates functions such as the counts-as relation as terms. Second, a number of relations is added in order to provide the means for representing different types of beliefs and knowledge. These are $\Rightarrow_{ci}, \Rightarrow_{cp}, \Rightarrow_{ca}$ and $\Rightarrow_{c}$; respectively the relations *cause intentionally*, *cause physically*, *counts as*, and *constitutes*. The $\Rightarrow$ symbol (without an additional subscript) stands for the relation between condition and consequence in a rule, whereas $\supset$ stands for the relation between condition and consequence in a principle. Both symbols are borrowed from RBL.

My notational conventions are the following. Terms start with lower case

characters. Variables are printed in italic. Predicates start with upper case characters. Sets are typeset in bold. Predicate arguments are put between brackets. If necessary to solve ambiguity, brackets are also used in sentences. The difference between terms and sentences is illustrated as follows. The term walk(john) *denotes* the fact 'John walks', *i.e.*, it refers to the actual fact, whereas Walk(john) *expresses* the fact John walks; it describes what is the case, namely that John walks (Hage 1997, p. 131; Hage discusses states-of-affairs, while I generalise the discussion to facts). Thus, a term is the *name* of a fact, and a sentence *describes* a fact.

8.1.2 Entities

The entities that are part of the model were identified in section 6.1. The following legally-relevant entities should be represented in the formalism: sentences, statements, propositions, beliefs, artefacts, rules, and concepts. Additionally, there are the following legal entities: legal principles, legal norms, legal decisions, legal systematisations, judicial interpretations, judicial classifications, legal rules, and legal concepts. The (partial) formalisation in RBL of these sentences is discussed below.

Legally-relevant entities
Sentences – Sentences are meaningful sequences of words, put in a natural language. The characteristic of being a sentence is represented in the predicate Sentence. A name is provided for each sentence as an abbreviation, in addition to its content:

Sentence(*sentence,content*)

Statements – An assertive sentence becomes a statement if it is uttered by some individual. This is relevant for a formalisation and can be expressed by a predicate that has three arguments: the person who utters the statement, the name of the sentence that is uttered, and the name of the context in which the sentence is uttered:

Statement(*person,sentence,context*)

Propositions – Propositions amount to what is expressed by assertive sentences. The predicate Proposition has two arguments, expressing its name and its content:

Proposition(*proposition,content*)

Beliefs – Beliefs are propositions that are within the reach of a propositional at-

titude (the 'believe'-attitude). A belief is represented as a two-place predicate:

Belief(*belief*,*content*)

The content of a typical belief that John stole a car is represented as a structured logical individual:

belief: stole(john,car)

Each belief has several properties. These properties are expressed as predicates, relating a belief name to the specific characteristic the named belief has.

OriginationSource(*belief*,*origination_source*)
ContentSource(*belief*,*content_source*)
Subject(*belief*,*content_source*)
Object(*belief*,*content_source*)

For instance, for the belief above (named 'belief'), the following characteristics might hold:

OriginationSource(belief,testimony)
ContentSource(belief,case_dossier)
Subject(belief,anne)
Object(belief,legally_relevant)
Object(belief,concrete)

Artefacts – Artefacts are tangible objects made by human beings. They are represented by individual constants, and the fact that an individual constant refers to an artefact is expressed by the predicate Artefact:

Artefact(*artefact*)

Rules – Rules are linguistic expressions that relate two types of phenomena to each other. Rules are expressed as terms. Their conditions, the rule symbol and the consequence form a structured logical individual. The fact that this term is a rule can be further stressed by the predicate Rule. A rule can be abbreviated by replacing it with a rule name.

rule: condition ⇒ consequence
Rule(*rule*)

Concepts – A concept is an abstract idea, containing the characteristics of concrete instances of the concept. Concepts are expressed in terms of the counts-as relation between the conditions for their obtaining, and the resulting concept

type. The fact that some entity is a concept is expressed with the Concept predicate.

concept: *condition_expression* $\Rightarrow_{ca}$ *concept*
Concept(*concept*)

The conditions part can be a single condition or a complex of conditions, built with the syntactic conventions for structured logical individuals. The term *concept* is the name of the concept type that results from the application of the counts–as relation.

Legal entities
Legal rules – Legal rules, just as rules, are expressions that relate two types of phenomena to each other. In this sense, they are equivalent to rules. Additionally, there are conditions under which the rule is a legal rule. These depend on the ontological stance taken. For instance, the criterion can be legal validity.

legal_rule: *condition* $\Rightarrow$ *consequence*
LegalRule(*legal_rule*)

Legal principles – Legal principles are general values that apply to certain areas within the law. In line with Hage's RBL they are represented as structured logical individuals, containing a condition expression, a consequence expression, and the connective $\supset$. The principle gets a name in order to refer to it more easily. In case of a principle like 'freedom of contract', applicable conditions should be made explicit in order to be able to formulate it in terms of RBL.

principle: *condition_expression* $\supset$ *consequence*
Principle(*principle*)

Legal norms – Legal norms are entities that explicitly guide some kind of behaviour or action. They contain deontic modalities. Norms are formalised as follows. pd stands for permitted–to–do, fd stands for forbidden–to–do, od stands for ought–to–do, ob for ought–to–be (Hage 1997, p. 169).

norm: pd(*actor,action*)
norm: fd(*actor,action*)
norm: od(*actor,action*)
norm: ob(*goal*)
Norm(*norm*)

The actors, actions and goals playing a role in the legal norms are referred to by

the terms between brackets. To emphasize the fact that some structured logical individual is a norm, we can use the Norm predicate.

Legal decisions – Legal decisions are decisions in some legal case, made by a qualified legal authority under suitable circumstances. They result from a specific act by a specific person. The mere fact that a certain entity is a legal decision is expressed as follows. The term between brackets is the name of the legal decision:

LegalDecision(*legal_decision*)

Legal systematisations – A legal systematisation is the result of an attempt to impose a more general structure on a relatively unstructured part of a legal system, or the implicit presence of such a structure in a part of a legal system. Legal systematisations are referred to by terms, and they can be further identified by the following predicates:

LegalSystematisation(*judicial_systematisation*)

Judicial interpretations – A judicial interpretation is the result of an assessment that a judge (or a different qualified authority) makes of some complex of (legal) rules and (legal) norms. For judicial interpretations, the same is valid as for legal systematisations:

JudicialInterpretation(*judicial_interpretation*)

Judicial classifications – Judicial classifications abstract from reality, by classifying a fact or an act as belonging to some legal type. Judicial classifications are indicated by the predicate JudicialClassification:

JudicialClassification(*judicial_classification*)

Legal concepts – Legal concepts are concepts whose definition is given (at least partly) in formal sources of law. Legal concepts are formalised in a manner similar to that of ordinary concepts: a structured logical individual contains a name for the legal concept, the applicable conditions, and the name of the concept. The conditions and the concept are connected with the counts–as relation. In addition, the predicate LegalConcept indicates that the term indicated is the name of a legal concept.

legal_concept: *legal_condition_expression* $\Rightarrow_{ca}$ *concept*
LegalConcept(*legal_concept*)

8.1.3 Ontological status layers

All status layers are expressed as predicates. Thus, we can assign a status layer to a logical individual in the same way as we do with a normal property. This also means that one logical individual can have more than one ontological status layer.

Legally-relevant ontological status layers

> Exist(x)
> Constituted(x)
> Recognised(x)

Legally-relevant ontological status layers

> LegallyEfficacious(x)
> LegallyValid(x)
> LegallyConstituted(x)
> LegallyRecognised(x)

8.1.4 Epistemic roles

Epistemic roles are expressed as structured logical individuals or as predicates. Additionally, properties of the epistemic roles are expressed as predicates.

Legally-relevant epistemic roles
Reasons – Reasons are statements, propositions or facts that are employed for the explanation or justification of some other statement, proposition or fact. The typology of reasons in subsection 4.5.2 gives us a starting point for the formalisation of reasons. First, the reason is represented as a structured logical individual:

> reason: content

Second, its characteristics are given in the form of predicates that have as their arguments the name of the reason concerned, and the name of the subject, act, or other objects:

> ReasonFor(*reason,subject*)
> ReasonTo(*reason,act*)
> ReasonForTo(*reason,subject,act*)
> ObjectOrientedness(*reason,type*)
> ContentSource(*reason,source*)
> OriginationSource(*reason,source*)
> Subject(*reason,x*)

Object(*reason*,*type*)

InternalModality(*reason*,*type*)

ExternalModality(*reason*,*type*)

Third, we indicate what the origin of the reason is, *i.e.*, on which belief or fact it is based. For this purpose, we use the counts-as relation;

relation: belief $\Rightarrow_{ca}$ *reason*

or:

relation: fact $\Rightarrow_{ca}$ *reason*

Additionally, we may express the fact that a reason is recognised as such by an individual or a community:

Recognise(*x*,*relation*)

In order to be able to express the fact that the recognition of a reason concerns the recognition of that reason *as a reason to do something* requires a different format. Because ReasonTo is a predicate itself, it cannot function as an argument of the Recognition predicate. Instead, ReasonTo has to be translated into a term-like expression:

Recognise(*x*,reason_to(*reason*,*act*))

Defeaters – A subclass of reasons is formed by defeaters. Defeaters are negative reasons, *i.e.*, they attack some reason in such a way that it is no longer correct. They are expressed formally in a way analogous to reasons. Additionally, the object of the attack should be specified. A belief or a fact counts as a defeater, and this relation may be recognised. Also, it may be indicated whether a defeater is recognised as a defeater for a certain belief or act.

relation: *belief* $\Rightarrow_{ca}$ *defeater*

relation: *fact* $\Rightarrow_{ca}$ *defeater*

Recognise(*x*,*relation*)

Recognise(*x*,defeater_for(*defeater*,*object*))

Factual knowledge – If an object complies with certain so-called knowledge criteria, it counts as factual knowledge. Factual knowledge is an epistemic role a belief (or a different entity, such as a proposition) can play. The knowledge qualification relation says that a certain belief becomes factual knowledge.

factual_knowledge_qualification: *belief* $\Rightarrow_{ca}$ factual_knowledge(*belief*)

fkq_rule: true(*belief*) & acquired_reliably(*belief*) & justified(*belief*) $\Rightarrow$
applicable(factual_knowledge_qualification)

Practical knowledge – If an ability complies with certain knowledge criteria, it counts as practical knowledge. Reliability was chosen as a criterion for the assessment of practical knowledge. Translated to the knowledge qualification rules given above, the result is as follows.

practical_knowledge_qualification: *belief* $\Rightarrow_{ca}$ practical_knowledge(*belief*)
pkq_rule: acquired_reliably(*belief*) $\Rightarrow$
applicable(practical_knowledge_qualification)

Legal epistemic roles
Factual legal knowledge – If a legal object (a belief, a proposition) complies with certain so-called knowledge criteria, it counts as factual legal knowledge. The way this qualification is formalised is similar to the formalisation of legally-relevant knowledge, but the criteria employed may very well differ from the ones mentioned.

factual_legal_knowledge_qualification: *belief* $\Rightarrow_{ca}$ factual_legal_knowledge(*belief*)
flkq_rule: true(*belief*) & acquired_reliably(*belief*) & justified(*belief*) $\Rightarrow$
applicable(factual_legal_knowledge_qualification)

Practical legal knowledge – If a legal ability complies with certain knowledge criteria, it counts as practical legal knowledge. Practical knowledge is formalised in the same way as legally-relevant practical knowledge. However, the precise meaning of the reliability criterion may differ.

practical_legal_knowledge_qualification:
belief $\Rightarrow_{ca}$ practical_legal_knowledge(*belief*)
plkq_rule: acquired_reliably(*belief*) $\Rightarrow$
applicable(practical_legal_knowledge_qualification)

8.1.5 Relations

There are two different ways to express relations in RBL. The first is by using structured logical individuals, the second by using predicates. Both types are used in the discussion of legally-relevant and legal relations below.

Legally-relevant relations
Causation – Causation is a two-place relation between phenomena: a certain event causes a different event or state-of-affairs. Causation is represented by a

connective that combines events and their consequences (events or states-of-affairs) in structured logical individuals. This connective is $\Rightarrow_c$, and a causation relation is expressed as follows.

   causal_relation: *event* $\Rightarrow_c$ *effect*

Subtypes of causality can be represented as.

   event $\Rightarrow_{ci}$ *effect*
   event $\Rightarrow_{cp}$ *effect*

In these expressions, ci stands for 'causes intentionally', and cp for 'causes physically'.

Counting as – Counting as is a two-place relation that associates two types of phenomena to each other. The counts-as relation is represented in similar manner, where ca in the $\Rightarrow_{ca}$ symbol stands for 'counts as'.

   counts_as_relation: *x-term* $\Rightarrow_{ca}$ *y-term*

The counts-as relation can hold between different ontological types, for instance between beliefs and reasons.

Recognition – Recognition is a two-place relation that says what individual, what group of individuals, or what institution recognises a certain phenomenon. Recognition is expressed as a predicate. The recognition predicate takes as its arguments the subject of recognition (the person or institution who recognises), and the object of recognition.

   Recognise(*x,object*)

For instance, the fact that Judy recognises the fact that killing on purpose counts as murder is expressed as follows.

   Recognise(judy,killing_on_purpose(*x,y*) $\Rightarrow_{ca}$ murder(*x,y*))

This is a case of individual recognition: Judy is a single person accepting this counts-as relation. Collective recognition is expressed with the same predicate as individual recognition. In this case we quantify over the individuals belonging to a certain society (where a society is a collection of individuals, represented by the set society). Thus, we use the following expression.

   $\forall$*individual*∈**society**(Recognise(*individual,object*))

Legal relations
Legal causation – Legal causation differs from the other types of causation. There-

fore, it gets its own connective.

$$event \Rightarrow_{cl} effect$$

This expression can be translated into a different form:

$$(\text{Recognise}(authority, event \Rightarrow_{ca} \text{cause\_for}(effect)) \leftrightarrow \text{Valid}(event \Rightarrow_{cl} effect)$$

This expression says that *event* causes *effect* in a legal sense if and only if there is a legal *authority* who recognises the validity of the fact that *event* counts as a cause for *effect*. Note that this definition does not restrict in any way what is recognised as a cause for an effect. Specific restrictions emanating from legal causality theories can be used to narrow the scope of the formal definition of legal causation.

Legal counting as – The legal counts-as relation is expressed in the same way as the ordinary counts-as relation. The only difference is that the content of a legal counts-as relation is derived from some valid legal source. The formal expression is as follows.

$$counts\_as\_relation\colon x\text{-}term \Rightarrow_{ca} y\text{-}term$$

The fact that this counts-as relation is legal can be expressed by the predicate legal.

$$\text{Legal}(counts\_as\_relation)$$

Legal recognition – Legal recognition is expressed by a predicate too. The legal recognition predicate takes as its arguments the subject of recognition (the person or institution who recognises), and the object of recognition.

$$\text{LegallyRecognise}(x, object)$$

There are two types of legal recognition: systemic recognition and authoritative recognition. Systemic recognition is expressed in a separate predicate:

$$\text{SystemicallyRecognised}(object)$$

This predicate is connected to individual recognition as follows.

$$(\exists individual \in \textbf{legalsystem} \; \exists object1 \in \textbf{legalsystem} \; (\text{Recognise}(individual, object1)$$
$$\land \exists object2 \in \textbf{legalsystem}(\text{Valid}(object1) \rightarrow \text{Valid}(object2))$$
$$\leftrightarrow \text{SystemicallyRecognised}(object2))$$

The expression says that in order for an object to be recognised systemically, this object has to be a valid consequence of the validity of a different object that is

recognised individually. Thus, an object that is part of a legal system, and that is the valid consequence of a different object that has been individually recognised, is systemically recognised itself.

The second type of legal recognition is authoritative recognition. It is a special form of individual recognition. To be recognised authoritatively, the recognising subject should be a legal authority.

(Recognise(*subject*,object) ∧ LegalAuthority(*subject*)) ↔

AuthoritativelyRecognise(*subject,object*)

8.1.6 Acts

The difference between entities and acts should be reflected clearly in the formalism. Often, the result of performing an act is an entity; for instance, the result of making a decision is a decision, the result of making a classification is a classification, the result of applying a rule is the instantiation of the consequence part of that rule.

All acts are formalised in a similar way: by introducing a special predicate named after the act. This predicate has an identifier as its argument. Below, the general predicate and an example are given.

Act(*act*)

SignContract(act1)

The act identifier is used in two other predicates: the PerformAct and ResultFromAct predicates:.

PerformAct(*actor,act*)

ResultFromAct(*actor,act,result*)

In addition, the name of the actor and – in case of the ResultFromAct predicate – the result are involved as arguments.

Legal acts and legally-relevant acts are identified as such by applying the predicates Legal and LegallyRelevant to the act identifier.

Legal(act1)

LegallyRelevant(act1)

8.1.7 Facts

Facts are indicated with the Fact predicate. They are identified using structured logical individuals. Thus we can (1) indicate the characteristics of facts, and (2) construct complex facts from multiple parts, and still be able to refer to them

in first-order logic. The following expressions are used to specify the different types of legally-relevant facts and legal facts.

Fact(*fact*)

Legal(*fact*)

LegallyRelevant(*fact*)

Event(*fact*)

StateOfAffairs(*fact*)

Brute(*fact*)

Recognised(*fact*)

Conventional(*fact*)

Institutional(*fact*)

The difference between legally-relevant facts and legal facts is expressed by the predicates LegallyRelevant and Legal. These can be combined with other characteristics. If we want to express that *fact* is both legal and conventional, we use the following expressions.

Fact(*fact*)

Legal(*fact*)

Conventional(*fact*)

A complex fact may be built from other (simple) facts. We can build these expressions as structured logical individuals with the connectives used in RBL. Thus, a complex fact may look like the following.

Fact(*complex_fact*)

complex_fact: *simple_fact_1* & *simple_fact_2* & *simple_fact_3*

8.2 Inference rules

In this section, the upgrading of terms to sentences and downgrading of sentences to terms is discussed (subsection 8.2.1). Next, reasoning with different types of causation relations is explained (subsection 8.2.2), as well as reasoning with the counts-as relation (subsection 8.2.3), and with rules (subsection 8.2.4).

8.2.1 Upgrading and downgrading

The upgrading and downgrading inference rules are important, as they enable us to translate sentences into terms and *vice versa*. The translation occurs according to the following notational conventions. A sentence is translated into a term

by replacing all upper-case characters by lower-case characters, preceded by an underscore (except for the first character of the sentence). A term is translated into a sentence by replacing the first lower-case character by an upper-case character, and replacing all remaining lower-case characters, preceded by an underscore, by a corresponding upper-case character, leaving the underscore out. For instance, a sentence ExamplePredicate(*example*) translates into a term example_predicate(*example*).

Upgrading terms to sentences proceeds in the following way. If the following condition holds:

Obtains(*fact*)

then

Fact

Downgrading sentences to terms is done in the following manner. If the following condition holds:

Fact

then:

Obtains(*fact*)

8.2.2 Causality

The different forms of causation each have their own inference rules. These inference rules are listed for intentional causation, physical causation, and legal causation respectively.

Inferring effects from intentional causation relations – Inferring effects from intentional causation relations occurs in the same way as for physical causation. The difference is found in the means by which the derivation may be blocked: the type of defeating reasons involved differs. The conditions to be fulfilled in order to derive the obtaining of the effect are that both an event, and the intentional-causal relation between the event and an effect should obtain, and that the effect and the intentional-causal relation between event and effect should not be excluded.

If the following conditions hold:

Obtains(event) $\wedge$
Obtains(event $\Rightarrow_{ci}$ effect) $\wedge$
~Excluded(effect) $\wedge$

~Excluded(event $\Rightarrow_{ci}$ effect)

then

Obtains(effect)

Reasons for the exclusion of either the effect of the intentional cause or of the causal relation (in general or in a specific instance) are guiding reasons; reasons for an agent not to behave in a specific manner. For instance, if I want to call the police because I witnessed a crime, but I am afraid the perpetrator of the crime will threaten my life, this may be reason for me not to call the police.

Inferring effects from physical causation relations – Contrary to the nature of defeating reasons for intentional-causal relations, reasons for the exclusion of either the effect of the physical cause or of the causal relation (in general or in a specific instance) are reasons why that physical cause does not obtain or why the causal relation does not obtain. These are *explanatory* reasons. For instance, if I try call the police using my cell phone, but my cell phone does not work, I will not be able to contact the police. The derivation proceeds in the same way as for intentional causation, but the derivation leading to the conclusion that the effect or the physical-causation relation is excluded, differs in character, as the reasons employed are different. This derivation, however, is not further specified.

If the following conditions hold:

Obtains(event) $\wedge$
Obtains(event $\Rightarrow_{cp}$ effect) $\wedge$
~Excluded(effect)
~Excluded(event $\Rightarrow_{cp}$ effect)

then

Obtains(effect)

Inferring effects from legal causation relations – The legal causation relation is established by the recognition of a counts-as relation by a legal authority. As soon as the legal causation relation is projected *post hoc* on a certain course-of-events, and it holds between a specific cause and effect, the defeasibility of this relation is limited to the treatment of the case in an appeal. Additionally, the counts-as relation is defeasible. Defeat of this relation has as a consequence that the legal causation relation does not obtain. The derivation of the obtaining of a legal-

causation relation from the recognition of a counts–as relation and the obtaining of that relation proceeds as follows:

If

> *Recognise(authority,event ⇒$_{ca}$ cause_for(effect)) ∧*
> *Obtains(event ⇒$_{ca}$ cause_for(effect))*

then

> *Obtains(event ⇒$_{cl}$ effect)*

Subsequently, we can derive the obtaining a an effect from the obtaining of an event and the legal-causation relation between the event and the effect:

If

> *Obtains(event) ∧*
> *Obtains(event ⇒$_{cl}$ effect)*

Then

> *Obtains(effect)*

8.2.3 Inferring conclusions from counts–as relations

The derivation of the obtaining of the y-term of a counts–as relation depends on the obtaining of a counts–as relation and the x-term of that counts–as relation. If the y-term and the counts–as relation are not excluded, then the y-term will obtain. Thus, the following derivation holds.

If

> *Obtains(x-term) ∧*
> *Obtains(x-term ⇒$_{ca}$ y-term) ∧*
> *~Excluded(y-term) ∧*
> *~Excluded(x-term ⇒$_{ca}$ y-term)*

then

> *Obtains(y-term)*

8.2.4 Inferring conclusions from rules

The following account of the derivation of conclusions from rules is taken from Hage (1997, p. 141*ff.*). In the first step, the applicability of a rule is determined.

A rule is applicable if there is some rule whose validity is established, and whose conditions obtain.

Given a rule

> *rule: conditions $\Rightarrow$ conclusion*

If

> *Valid(conditions $\Rightarrow$ conclusion) $\wedge$*
> *Obtains(conditions)*

then

> *Applicable(conditions $\Rightarrow$ conclusion)*

In the second step, the applicability of the rule is translated into a reason that the rule ought to be applied.

If

> *Applicable(conditions $\Rightarrow$ conclusion)*

then

> *ReasonFor(applicable(conditions $\Rightarrow$ conclusion),*
> *ob(applies(conditions $\Rightarrow$ conclusion)))*

In the third step, it is determined whether the reasons why the rule ought to be applied outweigh the reasons why the rule ought not to be applied. If this is the case, the conclusion is derived that the rule ought to be applied.

If

> *Outweigh(reasons(ob(applies(conditions $\Rightarrow$ conclusion))),*
> *reasons(~ob(applies(conditions $\Rightarrow$ conclusion))),conclusion)*

then

> *OB(applies(conditions $\Rightarrow$ conclusion))*

In the fourth and final step, the conclusion is derived from the fact that the rule ought to be applied, and the fact that the conclusion is not excluded.

If

> *OB(applies(conditions ⇒ conclusion)) ∧*
> *~Excluded(conclusion)*

then

> *Conclusion*

8.3 Knowledge qualification rules

In the subsections below, the rules that constitute the transition between be-liefs and knowledge are discussed. First, it is explained how belief is qualified as knowledge in the formal framework introduced in this chapter (subsection 8.3.1). Then, I discuss under what circumstances different knowledge criteria are fulfilled (subsection 8.3.2). Third, it is explained how knowledge and belief can be transferred from one individual to another (subsection 8.3.3). Finally, I indicate the way in which we can use the knowledge qualification rules in reasoning with legal knowledge, while taking into account different stances in philosophy of law (subsection 8.3.4).

8.3.1 Inferring knowledge qualifications

Both propositions and beliefs can be qualified as knowledge. In the former case, we define knowledge relative to the characteristics of a proposition. In the latter case, the knowledge qualification is a characteristic of a belief held by a person. In order to define a knowledge qualification relation, we need to embed the counts-as relation in a rule. The conditions part of the rule contains the condi-tions that must be fulfilled in order to 'trigger' the consequence part of the rule; this consequence part consists of a counts-as relation between a proposition or a belief on the left-hand side of the relation, and the knowledge qualification with respect to that proposition or belief on its right-hand side.

As to the qualification of a belief as knowledge, two steps should be taken; first it is determined whether a belief is justified. This depends on the question whether a belief is acquired reliably, whether it is coherent with a relevant set of beliefs, and whether it is *properly* justified (*i.e.*, backed by reasons):.

> justification_qualification: *belief* ⇒$_{ca}$ justified(*belief*)
> jq_rule: acquired_reliably(*belief*) & coherent(*belief*) & justified_properly(*belief*) ⇒
> applicable(justification_qualification)

Second, the actual knowledge qualification depends on the truth of the belief and its status as being justified.

> knowledge_qualification: $belief \Rightarrow_{ca}$ knowledge($belief$)
>
> kq_rule: belief(x) & true(x) & justified(x) $\Rightarrow$ applicable(knowledge_qualification)

The derivation is conducted in several steps, following the inference schemata for rules and counts-as relations in subsections 8.2.3 and 8.2.4. The results of applying the inference rules (if they are all applicable) are:

> Obtains($x \Rightarrow_{ca}$ knowledge(x))
>
> Obtains(knowledge(x))

This framework also enables the application of different knowledge qualification rules for different types of knowledge. Both the justification qualification rule and the knowledge qualification rule can be adjusted. For instance, if reliability does not play a role in the justification of a certain type of knowledge, the alternative justification qualification rule would amount to the following.

> justification_qualification_1: $belief \Rightarrow_{ca}$ justified($belief$))
>
> jq_qualification_1: belief($belief$) & type($belief$,interpretation) & coherent($belief$) &
>
> justified_properly($belief$) $\Rightarrow$ applicable(justification_qualification_1)

The rule needs its own name, as it differs from the rule stated above. Additionally, the type of the belief should be an interpretation. This is indicated by the type($belief$,interpretation) part of the term. For some types of knowledge, the knowledge qualification rule itself would be different. For instance, for interpretative beliefs, the truth criterion may be left out.

> knowledge_qualification_1: $belief \Rightarrow_{ca}$ knowledge($belief$)
>
> kq_rule_1: belief($belief$) & type($belief$,interpretation) & justified($belief$) $\Rightarrow$
>
> applicable(knowledge_qualification_1)

Apart from inference rules for knowledge qualification in general, we have to propose rules for knowledge qualification for an individual in order to be able to say under which conditions an individual *knows* something. This is done by introducing a personal justification qualification rule, in which a proposition is linked to a person via the term believe(x,$belief$).

> personal_justification_qualification: $belief \Rightarrow_{ca}$ justified_for($belief$,x)
>
> pjq_rule: believe(x,$belief$) & knows(x,acquired_reliably(x,$belief$)) &
>
> knows(x,coherent($belief$)) & knows(x,justified_properly_for($belief$,x)) $\Rightarrow$
>
> applicable(personal_justification_qualification)

Thus, in order to justify a belief for a certain individual, this individual has to believe the content of the proposition, he has to know that the belief has been acquired reliably, that it is coherent, and that it is properly justified for him. In order for some person to know something, the proposition has to be true and it has to be justified for the individual, as was set out in the previous rule.

personal_knowledge_qualification: *belief* $\Rightarrow_{ca}$ knowledge_for(*belief*,*x*)

pkq_rule: knows(*x*,true(*belief*)) & justified_for(*belief*,*x*) $\Rightarrow$
applicable(personal_knowledge_qualification)

Most of the conditions in the qualification rules contain the term 'knows'. In order to prevent an infinite regress from occurring, we have to make sure that, if a person has to know that a proposition is true, or if he has to know that a proposition is justified for him, he can attain such a state without having to fulfil the same qualification conditions again. In that case, he would have to employ other criteria in order to know that a proposition is *true* than in order to *know* that proposition.

For instance, in order to know that it is raining, Peter has to know that it is true that it is raining, and he has to know that he is justified in believing that it is raining. If we would employ the same criteria in each instance of 'knowing', he would also have to know that he knows that it is raining *etc.* Instead, in order to attain a situation in which we can actually *use* the knowledge qualification in practical situations (such as in the assessment of data in computer systems), we have to end such infinite regresses. I will not go into this matter extensively, but infinite regresses may for instance be barred by introducing additional knowledge qualification rules for knowledge criteria, that impose less strict demands on propositions before they may be qualified as knowledge.

8.3.2 Fulfilling individual knowledge criteria

Knowledge criteria play a role in transforming mere belief into knowledge. Although the question whether a criterion is fulfilled often depends on human assessment, an attempt is made here to model knowledge criteria in RBL. First, the truth criterion below says that a sentence should correspond to a fact in order to be true. Of course, this does not solve the problem how to establish such correspondence. Rather, the assumption that truth depends on correspondence of a sentence with a fact is expressed:

truth: correspond_to(*sentence*,*fact*) $\Rightarrow$ true(*sentence*)

Second, the proper justification criterion lists the conditions determined for

that criterion in subsection 2.4.2. I choose to model defeat and challenge explicitly in the conditions of the rules. An alternative would be to leave defeat to the inference rules defined for RBL.

proper_justification: (reason_for(*reason*,*belief*) & conclusive_wrt(*reason*,*belief*)) ∨
(reason_for(*reason*,*belief*) & non_conclusive_wrt(*reason*,*belief*) &
~defeater_for(*defeater*,*belief*) & ~defeater_for(*defeater*,*reason*)) ⇒ justified(*belief*)

In the personal proper justification criterion, the conditions are made relative to an agent. In this case, defeat and challenge are also modelled explicitly.

proper_personal_justification: (reason_for(*reason*,*belief*) &
conclusive_wrt(*reason*,*belief*)) ∨ (reason_for_wrt(*reason*,*belief*,*agent*) &
non_conclusive_wrt(*reason*,*belief*) & ~defeater_for_wrt(*defeater*,*belief*,*agent*) &
~defeater_for_wrt(*defeater*,*reason*,*agent*)) ⇒ justified_for(*belief*,*agent*)

Third, reliability is defined recursively, relative to the chain of belief acquisition. The reliable acquisition rule demands that the current reliability is greater than 0.9. If this is the case, the belief has been acquired reliably.

reliable_acquisition: current_reliability(*belief*) > 0.9 ⇒ acquired_reliably(*belief*)

The version of the reliable acquisition rule that is made relative to an agent is called personal reliable acquisition. This rule demands that the reliability of the acquisition of some belief is greater than 0.9 for a certain agent.

personal_reliable_acquisition: current_reliability(*agent*,*belief*) > 0.9 ⇒
acquired_reliably(*agent*,*belief*)

The current reliability of a belief is the reliability with which it is acquired, multiplied by the reliability of acquisition in other steps in a chain of acquisition. Thus, if a belief is acquired through a chain of three steps, and in each step the reliability of acquisition is 0.95, the resulting reliability with respect to this belief at the end of the chain is $0.95^3 = 0.857375$, which is lower than 0.9. In this case, the belief is not acquired reliably relative to the final step.

Fourth, the coherence criterion may consist of the subcriteria consistency, support, preference, and universality. The former three of these subcriteria can be treated formally to a certain extent. Universality, in the way it was described in 2.4.4, can be hardly expressed in a formal manner. The consistency of a belief with a belief set can be translated in the demand that a belief set does not imply the denial of that belief. Support of a belief for a different belief exists when the former is a reason for the latter. A preference relation between two principles

exists if one entity generally outweighs another one. The expressions that represent these three subcriteria are given below.

consistency: ~implies(**entity_set**,~*entity*) ⇒ consistent_with(*entity*,**entity_set**)

support: reason_for(*entity_1*,*entity_2*) ⇒ supports(*entity_1*,*entity_2*)

preference: outweighs(*entity_1*,*entity_2*) ⇒ preference_relation(*entity_1*,*entity_2*)

The following definition determines whether a belief coheres with a set of beliefs that it is part of. A general notion of coherence needs a quantification notion which would enable us to express the presence of as many support and preference relations as possible. The definition demands that a belief is consistent with the belief set it belongs to, that there is a support relation between the belief and a different belief in the belief set, and that there is a preference relation between the belief and a different belief in the belief set.

coherence: consistent_with(*belief*,**belief_set**) &

(supports(*belief_1*,*belief*) ∨ supports(*belief*,*belief_1*)) &

(preference_relation(*belief*,*belief_2*) ∨ preference_relation(*belief_2*,*belief*))

⇒ coherent_with(*belief*,**belief_set**)

In this subsection, it was explained how knowledge criteria can be formalised to a certain extent. It should be noted, however, that for practical use of such criteria, a vast quantity of information about a domain has to be defined explicitly.

8.3.3 Transfer of belief and knowledge

If an individual wishes to communicate his beliefs, the content of these beliefs should be transferred to his audience. If an individual wishes to communicate knowledge, and make sure that it remains knowledge throughout this process, additional demands are imposed. To acquire belief solely by transfer, two conditions have to be met. First, the belief content *p* has to be transferred. Thus, an agent *a* transfers the belief content to agent *b*. Second, the agent *b* should acquire a belief with the same content. In order to transfer knowledge, the knowledge criteria applicable to the belief content have to be fulfilled. Each of these knowledge criteria can be fulfilled either directly or by transferring beliefs that constitute their fulfilment.

An account of knowledge transfer in which both beliefs and the content of applicable knowledge criteria are transferred, is subject to the criticism that beliefs alone cannot provide a knowledge qualification. What we need, is sufficient backing of the beliefs about the transferred content of fulfilled knowledge criteria. Without imposing restrictions on the length of the justification chain,

however, this would lead to an infinite regress. We have to assume that beliefs about the content of fulfilled knowledge criteria can be qualified as knowledge as well. This could be attained, for instance, by establishing the reliability of the source of the beliefs about the content of fulfilled knowledge criteria. The reliability of the source should ensure that those beliefs are true.

The transfer of a piece of knowledge thus amounts to transferring the belief underlying that piece of knowledge, say p, and the beliefs concerning the fulfilment of the knowledge criteria turning that belief p into knowledge: the belief that p is true, the belief that p is justified, the reasons for p, and the belief about the reliability of the acquisition route of p. The resulting belief will be turned into knowledge by these beliefs, if these beliefs are acquired reliably. This feature is expressed in the rule below by the demand that the belief content should be personal knowledge for the agent. The transfer_belief rule says that an agent b has a belief denoted by x if some other agent a has belief x, tells its content to b, and b understands it, and b believes it.

> transfer_belief: has_belief(a,x) & tells_to(a,b,x) & understands(b,x) & believes(b,x)
> $\Rightarrow$ has_belief(b,x)

The transfer_knowledge rule says that an agent b has a piece of knowledge (denoted by x) if some other agent a has that piece of knowledge, tells the content of that piece of knowledge to b, and it is personal knowledge for b. Under what conditions a piece of knowledge is personal knowledge for an agent was explained above.

> transfer_knowledge: knows(a,x) & tells_to(a,b,x) & understands(b,x) &
> personal_knowledge_for(b,x) $\Rightarrow$ knows(b,x)

8.3.4 Reasoning with legal knowledge

The knowledge predicate functions very much like a quality stamp. It basically says that a certain belief or statement is worthy. Once we have a piece of knowledge, we have a piece of information that is qualified for further reasoning. A piece of information can deserve its knowledge qualification from different sources and criteria. Once it has such qualification, it can be employed for further reasoning, regardless of its origination source. In this subsection, I focus on reasoning with legal knowledge from different legal-ontological perspectives.

As I remarked in subsection 3.2.1, knowing the law from a natural-law perspective should be divided into knowing the abstract principles of natural law and knowing the rules and norms of positive law. Knowing the law amounts to knowing those rules and norms of positive law that comply with the principles

of natural law. As an example, I take a legal rule and a derivable legal norm. The rule is: 'Killing a person on purpose counts as manslaughter'. The norm is: 'One may not kill another person intentionally'. These are a rule and a norm of positive law. A corresponding principle of natural law could be that one ought not kill human beings.

Thus, the conditions under which a belief about a rule of positive law can be turned into knowledge are as follows. First, we have to verify our belief about the rule of positive law. The content of the belief about the rule should correspond with the content of the rule. Second, we have to acquire knowledge of some principle of natural law that complies with the rule of positive law. Third, we establish the correspondence or coherence of the rule of positive law with the principle of natural law. The coherent_with expression could be used for establishing the relation between the rule and the principle.

> legal_knowledge_qualification_nl: *belief* $\Rightarrow_{ca}$ legal_knowledge_nl(*belief*)
>
> lkq_nl_rule: knowledge(*belief*) & corresponds_with(*belief,rule*) &
>
> knowledge(*belief_2*) & corresponds_with(*belief_2,principle*) &
>
> coherent_with(*belief,belief_2*) $\Rightarrow$ applicable(legal_knowledge_qualification_nl)

In case of legal positivism (*cf.* subsection 3.2.2), knowing the law is possible through knowledge of sources of law and their validity. Validity, in its turn, is determined by knowledge of the recognition of those sources of law by a proper authority. This can be expressed in an additional rule defining validity.

> legal_knowledge_qualification_lp: *belief* $\Rightarrow_{ca}$ legal_knowledge_lp(*belief*)
>
> lkq_lp_rule: knowledge(*belief*) & corresponds_with(*belief,rule*) & valid(*rule*) $\Rightarrow$
>
> applicable(legal_knowledge_qualification_lp)

The same goes for institutional theories of law (*cf.* subsection 3.2.2). In this case, the validity of a rule depends on its embedding in a system of institutional rules. This should be expressed in an alternative rule defining validity. The basic legal knowledge qualification rule for institutional theories of law is the same as for legal-positivist theories.

Hermeneutic theories have a different approach towards knowledge (*cf.* subsection 3.2.5). The distinction between knowing that and knowing how is vague in such theories, and accordingly, it is harder to define a knowledge qualification rule for individual beliefs. I will therefore not attempt to define such a rule regarding hermeneutic theories. It would require a thorough embedding of knowledge in its context, and therefore an account of the influence pieces of knowledge have on each other.

8.4 Representing knowledge about Dutch penal law

In this section, I apply the knowledge-based model of law to the domain of Dutch penal law. I do so by showing the role of knowledge in the representation of various components of penal law. A formal account of such components will generally lead to a loss of content compared to natural-language representations. The account given in the current section should show the central role of certain categories (entities, relations *etc.*) in the representation of legal knowledge. These categories are elements of the knowledge-based ontology of law described in chapter 6. In the current section, I discuss the representation of the questions leading to the establishment of guilt and punishability (subsection 8.4.1), of principles (subsection 8.4.2), central concepts of penal law (subsection 8.4.3) and violent crimes and crimes against life (subsection 8.4.4). Finally, I discuss how knowledge of specific statutes is established (subsection 8.4.5).

8.4.1 Guilt and punishability

At the start of a trial, the two main questions are: is a suspect guilty, and if so, is he punishable? The two groups of questions determining guilt and punishability of a suspect were discussed in section 7.1. Below, these questions will be expressed in the form of rules and norms. The first question is: Is the summons valid?

> norm_1: ob(valid(*summons*))
>
> rule_1: ~valid(*summons*) $\Rightarrow_{ca}$ void(*summons*)

The norm says that the summons ought to be valid. The rule says that if the summons is not valid, it counts as void. If there is an instance of the type *summons* that is not valid, then that instance counts as void. The next question is: Is the judge competent?

> norm_2: ob(competent(*judge*))

The norm expressed is that the judge ought to be competent. The third question is: Is the public prosecutor amenable?

> norm_3: ob(amenable(*public_prosecutor*))

The norm says that the public prosecutor should be amenable. The fourth question is: Is there a reason for adjournment of the trial? If there is such a reason, the trial ought to be adjourned.

> rule_4: reason_for(**reasons**,adjournment) $\neq \varnothing \Rightarrow$ ob(adjourned(*trial*))

In this case, the rule says that if there are reasons for the adjournment of the trial, then the trial ought to be adjourned.

After these four questions, a second group of four questions is introduced. The first question of this second group is whether it has been proved that the suspect has committed a fact as charged.

> norm_5a: ob(proved(commit(*suspect*,*fact*)))
>
> norm_5b: ob(acquitted(*suspect*))
>
> rule_5: ~proved(commit(*suspect*,*fact*)) ⇒ applicable(norm_5b)

The first norm (5a) says that it ought to be proved that the suspect committed the fact as charged. The rule says that if this is not the case, the second norm (5b) is applicable. This norm says that the suspect should be acquitted. The second question is whether the fact is punishable.

> norm_6a: ob(punishable(*fact*))
>
> norm_6b: ob(exempted_of_legal_proceedings(*suspect*))
>
> rule_6: ~punishable(*fact*) ⇒ applicable(*norm_6b*)

The first norm (6a) states that the fact ought to be punishable. The rule says that if the fact is not punishable, then the second norm (6b) is applicable. This norm says that the suspect ought to be exempted of legal proceedings. The third question is whether the suspect is punishable.

> norm_7a: ob(punishable(*suspect*))
>
> norm_7b: ob(exempted_of_legal_proceedings(*suspect*))
>
> rule_7: ~punishable(*suspect*) ⇒ applicable(norm_7b)

The first norm (7a) states that the suspect ought to be punishable. The rule says that if the suspect is not punishable, then the second norm (7b) is applicable. This norm says that the suspect ought to be exempted of legal proceedings. The fourth question is what punishment should be imposed.

> rule_8: (~void(*summons*) & ~incompetent(*judge*) &
>
> ~not_amenable(*public_prosecutor*) & ~ob(adjourned(*trial*)) &
>
> ~ob(acquitted(*suspect*)) & ~ob(exempted_of_legal_proceedings(*suspect*))) ⇒
>
> ob(applicable(norm_8))
>
> norm_8: ob(imposed(*punishment*,*suspect*))

The rule says that if all conditions (already discussed above) remain unfulfilled, *i.e.*, if the summons is not void, the judge is not incompetent *etc.*, then the norm is applicable. This norm states that a punishment should be imposed on the sus–

pect. Contrary to what the norm seems to imply, it is possible that a suspect will not be punished.

8.4.2 Principles

General principles of penal law are expressed using several entities from the knowledge-based model of law. These include rules and principles. The legitimacy principle says that any legal judgement should be properly argued for. The legality principle says that a judgement should conform to some basic demands that make the judgement *legal*, *i.e.*, for instance no fact is punishable without a prior law that makes that fact punishable. What is expressed as the term 'legal_basis_for' amounts to the conditions that no fact is liable to penalty without applicable legislation, no penalty can be imposed without a law permitting the imposition, no penalty can be imposed without a *preceding* law, analogous reasoning is not allowed in the argumentation of the judge, and laws should be sufficiently clear. The legitimacy and legality principles are expressed below.

> rule_1: properly_justified(*judgement*) $\Rightarrow_{ca}$ legitimate(*judgement*)
>
> legitimacy_principle: ob(legitimate(*judgement*))
>
> rule_2: legal_basis_for(*judgement*) $\Rightarrow_{ca}$ legal(*judgement*)
>
> legality_principle: ob(legal(*judgement*))

Each of the principles is expressed by a norm and a rule. The rules say under what conditions a judgement counts as legitimate and legal. The norms say that judgements ought to be legitimate and legal.

The subsidiarity and proportionality principles, expressing the necessity of employing the least drastic means available in a certain legal case, and the proportionality of the means to be employed with the goal to be reached, are expressed as follows.

> subsidiarity_principle: ob(minimal_wrt(***means,*** solve(*undesirable_situation*)))
>
> proportionality_principle:
> ob(proportional_wrt(***means,***solve(*undesirable_situation*)))

Both principles are translated into norms, and the criteria expressed by these principles are arguments of the norms. In both cases this is a criterion expressing a relation between the means employed to reach some goal, and the goal itself, which is solving an undesirable situation. In some cases, the principles should be formulated in a different way, stressing the fact that there ought to be employed minimal and proportional means in order to attain a *desirable* situation.

The following expressions are translations of the three principles establishing

the priority relations between conflicting rules: *lex specialis derogat legi generali*, *lex posterior derogat legi priori*, and *lex superior derogat legi inferiori*.

> specialis_principle: ob(applicable(more_specific_than(*rule_1*,*rule_2*)))
> posterior_principle: ob(applicable(issued_later_than(*rule_1*,*rule_2*)))
> superior_principle: ob(applicable(issued_by_higher_authority_than(*rule_1*,*rule_2*)))

These principles are translated as norms too. These norms say that there should be specific relations between two rules (expressed by the variables *rule_1* and *rule_2*).

8.4.3 Central concepts of penal law

In this subsection, I discuss the formalisation of six concepts that play an important role in penal law. These are: facts, culpability, exemption grounds, causality, participation, and attempt. Culpability of a suspect is determined in a trial following the two groups of four questions discussed under the denominator 'guilt and punishability'. As I discussed these questions extensively in subsections 7.3.2 and 8.4.1, I leave them out in the overview below.

Facts

In the legal domain, facts are often constructed entities. In section 6.6, I described how to build complex facts from simple facts.

> Fact(*complex_fact*)
> *complex_fact*: *simple_fact_1* & *simple_fact_2* & *simple_fact_3*

We can assign other characteristics to facts and complex facts too. The distinction between material facts and judicial facts made in subsection 7.3.1 to clarify the *ne bis in idem* principle, for instance, can be expressed as follows.

> MaterialFact(*fact_1*)
> JudicialFact(*fact_2*)

In an example case, simple facts may constitute material facts and judicial facts in different combinations.

> fact_1: drunk(john)
> fact_2: driving(john)
> fact_3: fact_1 & fact_2
> rule: drunk(*person*) & driving(*person*) $\Rightarrow_{ca}$ commits_criminal_offence(*person*)
> MaterialFact(fact_3)
> JudicialFact(fact_1)

JudicialFact(fact_3)

There are two judicial facts, because both the fact that John is drunk and the fact that John is drunk and driving are punishable (public drunkenness and drunk driving), and there is only one material fact, as the two simple facts cannot be regarded apart from each other. Application of the *ne bis in idem* principle would prevent the agent from being prosecuted twice, as there is only one material fact involved, whereas initially he could have been prosecuted for either of the two judicial facts.

Ontological status layers can be imposed on these facts. For instance, the act by which a judge recognises some fact, for instance the qualification of fact_3 as a judicial fact, is expressed by downgrading the predicate JudicialFact(fact_3) to a term judicial_fact(fact_3)), and taking it as an argument of the act LegallyRecognise.

LegallyRecognise(judge,judicial_fact(fact_3))

The result of this act is that the same fact gets the ontological status layer LegallyRecognised.

LegallyRecognised(judicial_fact(fact_3))

In similar ways, additional ontological status layers or specific legal characteristics can be attributed to a certain fact. The attribution of characteristics like JudicialFact and MaterialFact are useful for an obvious reason: as soon as we have qualified a fact in terms of such characteristics, we can also derive the applicability of the *ne bis in idem* principle.

Exemption grounds

Exemption grounds from punishment prevent the imposition of a sanction to a suspect. An overview of the distinctions relevant to exemption grounds was given in subsection 7.3.3. The questions in a trial in which exemption grounds play a major role were discussed in section 7.1, and the formalisation of the norms underlying these questions was the topic of subsection 8.4.1. I now turn to the formalised classification of facts as exemption grounds.

Rules 1 and 2 express the conditions under which a suspect is exempted from the application of sanctions.

rule_1: applicable(*exemption_ground*) $\Rightarrow$ exempted_from(*suspect,sanction*)

rule_2: *ground_for_justification* $\lor$ *exemption_ground_from_guilt* $\Rightarrow_{ca}$ *exemption_ground*

Rule 2 says that both a ground for justification and an exemption ground from guilt count as an exemption ground. Rule 1 says that if an exemption ground is applicable, then the suspect is exempted from the application of a sanction. The distinction between general and specific exemption grounds cuts through the distinction between grounds for justification and exemption grounds from guilt. The general or specific nature of an exemption ground is therefore added as an extra characteristic, expressed as a predicate.

GeneralDefence(*exemption_ground*)

SpecificDefence(*exemption_ground)*

As an example, I take two exemption grounds, namely being of unsound mind, and force majeure. These are expressed as follows:

rule_3: being_of_unsound_mind $\Rightarrow_{ca}$ exemption_ground_from_guilt

rule_4: absolute_force_majeure $\Rightarrow_{ca}$ ground_for_justification

Both are general exemption grounds, which is indicated by the predicate introduced above:

GeneralDefence(being_of_unsound_mind)

GeneralDefence(absolute_force_majeure)

Specific rules can be used to establish the actual presence of exemption grounds. In order to do so, a link should be made between the classification of a suspect's acts and applicable law, in that the specific behaviour of the suspect becomes the subject of a series of classifications, from 'brute' facts to a detailed characterisation in terms of the different types of exemption grounds.

Causality

Causality theories influence what acts are classified as causes for some phenomenon. Whether the man throwing a slipper to his wife's head could have foreseen her death (almost certainly not) and whether his act was a *conditio sine qua non* for her death (almost certainly) are two radically different interpretations of legal causality. The theory employed in a case determines what event is recognised as a legal cause for an effect. The recognition relation that underlies legal causality is expressed as follows (*cf.* subsections 4.7.2 and 6.4.2):

(Recognise(*authority,event* $\Rightarrow_{ca}$ cause_for(*effect*)) $\leftrightarrow$ Valid(*event* $\Rightarrow_{cl}$ *effect*)

Causality theories are hard to formalise. Probably the easiest theory to formalise

is the *conditio sine qua non* theory. This theory can be expressed in the following rule:

rule: ($fact \rightarrow effect$) $\Rightarrow$ ($fact \Rightarrow_{ca}$ cause_for($effect$))

This rule says that if the effect only occurs if the fact occurs, then the fact counts as a cause for effect.

Participation

The different forms of participation can all relatively easy be expressed in a formal way. First, there is the relation between the different subtypes of participation and participation itself. Rule 1 expresses the counts-as relation between the different subtypes of participation and the fact that a perpetrator participates in a fact.

rule_1: ($fact\_1 \vee fact\_2 \vee fact\_3 \vee fact\_4$) $\Rightarrow_{ca} fact\_5$
fact_5: participate_in($perpetrator,fact$)

Rule 2 expresses the condition under which a person commits a fact. It says that in order for a person to commit a fact, he should fulfil the constituents of the offence description.

rule_2: $fact\_6 \Rightarrow fact\_1$
fact_6: fulfil($perpetrator,offence\_description$)
fact_1: commit($perpetrator,fact$)

Rule 3 says that in order for a person to make another person commit a fact, he should cause that other person to commit a certain fact.

rule_3: $fact\_7 \Rightarrow fact\_2$
fact_7: causes_to_commit($perpetrator,person,fact$)
fact_2: make_commit($perpetrator,fact$)

In rule 4, the conditions are expressed under which a person co-commits a fact. In order to do so, there should be at least two different persons that can be considered to commit the fact.

rule_4: $fact\_8 \wedge fact\_9 \Rightarrow fact\_3$
fact_8: commit($perpetrator,fact$)
fact_9: commit($perpetrator\_2,fact$)
fact_3: co-commit($perpetrator,fact$)

Rule 5 expresses the condition under which a perpetrator is said to provoke

someone else to commit a fact.

> rule_5: *fact_10* ⇒ *fact_4*
>
> *fact_10:* provoke_to_commit(*perpetrator,person,fact*)
>
> *fact_4:* provoke(*perpetrator,fact*)

Complicity occurs if some person helps another person commit a fact. Rule 6 defines the conditions under which this phenomenon occurs.

> rule_6: *fact_11* ⇒ *fact_5*
>
> *fact_11:* help_commit(*accomplice,perpetrator,fact*)
>
> *fact_5:* accomplice_to(*accomplice,fact*)

Attempt

In establishing the occurrence of an attempt, we have to deal with the complex problem of revealing the intentions and acts underlying something that was never realised. After all, attempt does not show itself in a completed punishable fact – otherwise it would not be an attempt. The criterion for attempt formulated by the Supreme Court (discussed in subsection 7.3.6) can be formalised in the following manner.

> rule_1: *fact_1c* ⇒ *fact_1d*
>
> *fact_1a:* has(*perpetrator,intent*)
>
> *fact_1b:* commences(*implementation*)
>
> *fact_1c:* reveals(*fact_1a,fact_1b*)
>
> *fact_1d:* attempt ⇒$_{ca}$ punishable(*attempt*)

Rule 1 states that if the intent of a perpetrator reveals itself in a commencement of implementation, then the attempt that can be derived from this counts as a punishable attempt. Rule 2 defines voluntary retiration.

> rule_2: *fact_2b* ⇒ *fact_2a*
>
> *fact_2a:* voluntarily_retires(*person*)
>
> *act_2a:* ceases(*person,attempt*)
>
> *fact_2b:* performs_on_the_basis_of(*person,act_2a,*new_deliberation)

This rule says that in case a person ceases an attempt on the basis of a new deliberation of the circumstances given, he voluntarily retires. Rule 3 defines the concept of unsound attempt.

> rule_3: *fact_3b* ∨ *fact_3c* ⇒ *fact_3a*
>
> *fact_3a:* attempt ⇒$_{ca}$ unsound(*attempt*)
>
> *fact_3b:* unsound(*means*)

fact_3c: unsound(*object*)

fact_3d: ∀*case* ~suitable_for(*means,goal,case*)

fact_3e: ∀*case* ~suitable_for(*object,goal,case*)

rule_3a: *fact_3d* ∨ *fact_3e* ⇒ (unsound(*attempt*) ⇒$_{ca}$ absolutely(unsound(*attempt*)))

rule_3b: ~*fact_3d* ∨ ~*fact_3e* ⇒ (unsound(*attempt*) ⇒$_{ca}$ relatively(unsound(*attempt*)))

This rule states that if means or object are unsound, then the attempt counts as an unsound attempt. Rules 3a and 3b define the situations in which an attempt is objectively and subjectively unsound. An attempt is objectively unsound if there is no case in which a means or object could be sufficient to attain a certain goal. An attempt is subjectively unsound if the goal cannot be attained in the case at hand.

8.4.4 Violent crimes and crimes against life

Below, it is demonstrated in what way a statute can be translated into the ontology described in chapter 6. The first example is art. 287 SR:

> "Anyone who deliberately takes the life of someone else, is – being guilty of manslaughter – subject to a prison sentence that does not exceed 15 years, or a fine of the fifth category."

The analysis of this article is as follows. Two act types are defined. The first act type consists of a person *a* killing another person *b* intentionally. The first person is not equal to the second, hence *a* ≠ *b*. Then, there is a rule comprising of the counts–as relation between the first act type and the second act type (*act_1* and *act_2*), and a rule that says that if there is a case of manslaughter, then the suspect is punishable with a certain maximum sentence. The specification of this punishment is given as the fact type *fact_1*.

act_1: intentionally(kill(*a,b*)) ∧ *a* ≠ *b*

act_2: commit(*a*,manslaughter)

rule_1: *act_1* ⇒$_{ca}$ *act_2*

rule_2: *act_2* ⇒ *fact_1*

fact_1: punishable_with(*a*,{max_15_years(*prison_sentence*) ∨

max_5th_category(*fine*)})

The English translation of the article on qualified manslaughter (art. 288 SR) is:

> "Manslaughter, followed, accompanied or preceded by a punishable fact, and committed with the aim to prepare or to facilitate that fact, or – if caught in the act – to ensure oneself or other participators to that fact that

they remain unpunished or to ensure the possession of the illegally acquired good, is subject to a life sentence, a prison sentence that does not exceed 20 years, or a fine of the fifth category."

This article is translated in formal expressions as follows. Three act types refer to manslaughter, a punishable fact, and to qualified manslaughter. The fact types express the conditions under which the article applies, and the consequences of its application. Rules 1 and 2 define the stages in which the punishment of a person is determined; first the conditions are given under which manslaughter occurs, then the potential punishment is listed.

act_1: commit(*a*,manslaughter)

act_2: commit(*a*,punishable_fact)

act_3: commit(*a*,qualified_manslaughter)

fact_1: intention_of_is(*a*,*act_1*,(prepare(*a*,*act_2*) ∨ facilitate(*a*,*act_2*))

fact_2: caught_in_the_act(*a*,*act_2*) &

intention_of_is(*a*,*act_1*,ensure(*a*,remain_unpunished)) ∨

intention_of_is(*a*,*act_1*,ensure(accomplices,remain_unpunished))

fact_3: intention_of_is(*a*,*act_2*,ensure(possession_of_illegally_taken_material))

fact_4: followed_by(*act_1*,punishable_fact) ∨

accompanied_by(*act_1*,punishable_fact) ∨ preceded_by(*act_1*,punishable_fact)

rule_1: (*fact_1* ∨ *fact_2* ∨ *fact_3*) & *fact_4* $\Rightarrow_{ca}$ *act_3*

rule_2: *act_3* $\Rightarrow$ *fact_5*

fact_5: punishable_with(*a*,life(prison_sentence) ∨

max_20_years(prison_sentence) ∨ max_5th_category(fine)

The translation of the article on murder (art. 289 SR) is as follows:

"Anyone who intentionally and after deliberation takes the life of another person is guilty of murder and is subject to a life sentence, a prison sentence that does not exceed 20 years, or a fine of the fifth category."

This article is translated into the following formal expressions. The acts types 1 through 3 define the conditions for murder to occur, whereas act type 4 is the actual murder act. Fact type 1 describes the possible punishment. Rule 1 defines the conditions under which murder occurs. Rule 2 defines the resulting punishment.

act_1: kill(*a*,*b*) ∧ *a* ≠ *b*

act_2: commit(*a*,with_malice_aforethought(*act_1*))

act_3: commit(*a*,on_purpose(*act_1*))

act_4: commit(*a*,murder)

fact_1: punishable_with(*a*,max_20_years(prison_sentence) ∨ life(*prison_sentence*) ∨ max_5th_category(*fine*))

rule_1: *act_1* & *act_2* & *act_3* ⇒$_{ca}$ *act_4*

rule_2: *act_4* ⇒ *fact_1*

The translation of the article on culpable homicide (art. 307 SR) is:

"A person who can be blamed for the death of someone else, is punished with a prison sentence or detention of no more than nine months, or a fine of the fourth category."

This article is translated into the following formal expression:

act_1: commit(*b*,culpable_homicide)

fact_1: guilty_of(*b*,death(*a*)) ∧ *a* ≠ *b*

fact_2: punishable_with(*b*,max_9_months(prison_sentence) ∨ max_4th_category(fine))

rule_1: *fact_1* ⇒$_{ca}$ *act_1*

rule_2: *act_1* ⇒ *fact_2*

Rule 1 states that in case a person is guilty of the death of another person, this counts as a case of committing culpable homicide. Rule 2 says that an occurrence of culpable homicide is punishable with either a prison sentence or a fine.

8.4.5 Knowing statute law

In subsection 8.3.4, I explained how to define the knowledge qualification relation within a framework of legal positivism.

legal_knowledge_qualification_lp: *belief* ⇒$_{ca}$ legal_knowledge_lp(*belief*)

lkq_lp_rule: knowledge(*belief*) & corresponds_with(*belief*,*rule*) & valid(*rule*) ⇒ applicable(legal_knowledge_qualification_lp)

The manslaughter rules are repeated below.

act_1: intentionally(kill(*a*,*b*)) ∧ *a* ≠ *b*

act_2: commit(*a*,manslaughter)

rule_1: *act_1* ⇒$_{ca}$ *act_2*

rule_2: *act_2* ⇒ *fact_1*

norm: ought_not(*act_2*)

fact_1: punishable_with(*a*,{max_15_years(*prison_sentence*) ∨ max_5th_

category(*fine*)})

The first manslaughter rule (intentionally killing another person counts as man-slaughter) is used as an instantiation for the variables in the knowledge qualification rule:

legal_knowledge_qualification_lp: ($act\_1 \Rightarrow_{ca} act\_2$) $\Rightarrow_{ca}$

legal_knowledge_lp($act\_1 \Rightarrow_{ca} act\_2$)

lkq_lp_rule: knowledge($act\_1 \Rightarrow_{ca} act\_2$) & corresponds_with($act\_1 \Rightarrow_{ca} act\_2, act\_1$

$\Rightarrow_{ca} act\_2$) & valid($act\_1 \Rightarrow_{ca} act\_2$) $\Rightarrow$ applicable(legal_knowledge_qualification_lp)

Each instance of the variable *belief* is filled with the content of rule 1, assuming that the content of that belief is the given rule. The variable *rule* is filled with the content of the rule. The demands that the belief is knowledge, that the belief corresponds with the rule, and that the rule is valid are fulfilled. Thus, if there are no defeating reasons present, the consequence of the rule applies, which means that the rule 'legal_knowledge_qualification_lp' is applicable. The variable *belief* is filled with an instance of the rule, and thus we can derive that the rule counts as legal knowledge in a legal–positivist stance.

8.5 Representing knowledge about the legal domain

In this chapter, it has been shown how the knowledge-based model of the legal domain can be used as the basis for formal representations of legal knowledge. I explained the way in which the representatives of the model's basic categories can be modelled in terms of Hage's (1997) reason-based logic. Furthermore, I discussed the ways in which inference rules can be used to reason with representations, how knowledge qualification can be formalised, and how specific knowledge from the domain of Dutch penal law can be represented. Rather than showing that such representations can be meaningfully employed in knowledge-based systems, I tried to explain the way in which certain features of legal knowledge can be represented: *e.g.*, the specific transition from legal belief to legal knowledge. By making explicit a larger part of the structural features of knowledge about the legal domain, a more precise representation of such knowledge can be attained.

Still, having explained the lines along which we can discern major characteristics of knowledge about the legal domain, it may easily be seen that such characteristics can play a role in assessing the content of legal information systems. Moreover, they can play a role *in* such information systems. Meta-information

on the structure and content of knowledge about the legal domain, *e.g.*, on the justification of represented knowledge, constitutes an important asset for users of those systems. Such meta-information may not only prove useful in adding justificatory material to conclusions of information systems that use simple production rules for inferencing, it a may also add to the functionality and quality of legal databases and accompanying search engines.

There is, of course, a trade-off between the richness of expressions in formal languages (the degree to which they can provide detailed meaning), and their usefulness in automated inferencing (the degree to which they can be employed to derive conclusions). Therefore, the application of characteristics of knowledge, such as the ones expressed in the model in this thesis, should not beforehand be deemed restricted to formalised representations of the law. Providing information on, for instance, the precise justifying role of some part of a judgement, could very well be useful as meta-information on that judgement, without the involvement of any formalisation.

9 Conclusion

In this chapter, I draw conclusions on the basis of the content of the previous chapters. The conclusions are provided in accordance with the problem definition and research questions. The problem definition for this thesis was: What role can a concept of legal knowledge, formulated from the perspective of epistemology, play in the representation of legal knowledge? The research questions were:

(1) What are the characteristics of knowledge about the legal domain, given the viewpoint of general epistemology?

(2) What are the characteristics of knowledge about the legal domain, given the viewpoint of legal epistemology?

(3) What characteristics of knowledge about the legal domain are useful as constituents for a model of the legal domain?

(4) How can this ontology of law be used to represent knowledge about the domain of Dutch penal law?

Below, I provide a number of conclusions answering the research questions.

The first research question
The first research question regards the characteristics of knowledge about the legal domain given the viewpoint of general epistemology. With respect to this question, I list five conclusions:

1) The three dimensions of knowledge – acquisition, object and justification – are dimensions of knowledge about the legal domain too. In the legal domain, these dimensions have different meanings, given the specific nature of such knowledge. The acquisition of knowledge about the legal domain is partly driven by rules. These rules determine what sources may be employed, and in what way knowledge about these sources should be acquired. The object of knowledge about the legal domain is characterised by the facts that it is intangible (most of its objects cannot be touched), and that it is rule-based (the existence of many objects is based on some kind of explicit or implicit rule), or based on reasoning and interpretation (it is derived from other objects by a process of reasoning). The justification of such knowledge is also driven by rules. Legitimisation demands impose specific demands on the way in which legal knowledge is justified.

2) One of the main characteristics of knowledge about the legal domain is that acquisition, object and justification are intertwined to a high degree. Proper

acquisition of knowledge in the legal domain may lead to instant justification. Proper acquisition and justification can create new objects, such as interpretations and judicial decisions. Rather than a unidirectional relation between, for instance, perception and the object of perception, knowledge about the legal domain is often acquired in a bidirectional process between the subject and the object to be known. Knowledge of such objects can be acquired through knowledge of their acquisition and justification: if we know the origination history of an interpretation, we may claim to know that interpretation. This phenomenon causes the intertwinement of acquisition, object and justification.

3) Knowledge about the legal domain is prone to the distinction between externalist and internalist knowledge theories, and between doxastic and non-doxastic knowledge theories. First, legal knowledge, such as knowledge of an interpretation, fits in with an internalist and doxastic framework: beliefs about interpretations are justified through other beliefs. Second, legally-relevant knowledge, such as knowledge about the facts of a case, fits in with an externalist non-doxastic knowledge theory: in such a theory beliefs are justified by an appeal to external criteria like reliability. Handling all types of knowledge about the legal domain in the same manner, employing a single concept of knowledge, would mean not to recognise the diversity of knowledge types actually present.

4) Knowledge criteria are the barrier between mere belief and knowledge; they define the conditions under which a belief is qualified as knowledge. Truth, proper justification, reliability and coherence were discussed. The legal domain allows for the use of these criteria, though they have to be adapted to suit the specific characteristics of the different types of knowledge about the legal domain. Knowledge about the legal domain can be divided into legally-relevant knowledge and legal knowledge. Furthermore, there are distinctions between abstract and concrete knowledge, and between factual and practical knowledge. A piece of knowledge can be classified through these three distinctions. Accordingly, different sets of knowledge criteria are applied.

5) Some knowledge acquisition situations in the legal domain can be classified as epistemic dependency situations: the acquisition is done within an epistemic niche. Although epistemic dependence can threaten the truth-conduciveness of beliefs, this need not be the case if practical and procedural restrictions are carefully designed with the goal of truth-conduciveness in mind.

The second research question
With respect to the second question, the characteristics of knowledge about the legal domain should be established, given the viewpoint of legal epistemology. I

provide three conclusions:

1) Different legal-theoretical theories influence the way in which knowledge can be acquired. These theories include or imply views on the acquisition, object and justification of knowledge, and thereby tend to impose different conditions on knowledge qualification. A clear distinction of knowledge and the object of knowledge, for instance, may facilitate knowledge acquisition, while the entanglement of knowledge and its object may complicate matters.

2) Formal and material sources of law may be considered as knowledge sources for the law, if additional requirements are met. Such requirements depend on the legal-theoretical stance taken. Apart from formal and material sources of law, such knowledge acquisition processes as reasoning and interpretation prove to be very important sources of knowledge for the law, precisely because knowledge and object of knowledge are intertwined to a certain degree in each of the legal-theoretical views discussed.

3) To disentangle the intertwinement between knowledge and its object, the distinction between epistemology and ontology has to be maintained. The relations between the epistemic and ontological levels can be isolated as unidirectional constitution relations that do not cause a circularity problem: knowing may constitute being and *vice versa*. The ability to maintain this distinction depends on the legal-theoretical stance taken; a hermeneutical theory of law does not allow for it.

The third research question

The third research question regards the usefulness of the characteristics of knowledge about the legal domain as constituents for a model of law. With respect to this question, I provide three conclusions.

1) With respect to the object of knowledge, the following is observed. Due to the intertwinement of knowledge of the law and the law itself, knowing the law requires a view on the entities constituting the law. The search for characteristics of knowledge about the legal domain thus implies a search for the ontological constituents of the domain. From the discussion of ontological stances with respect to the law, we can derive five existence status layers: existence, validity, recognition, constitution, and efficacy. By distinguishing different ways in which objects can exist, it becomes possible to accommodate different legal-theoretical stances on the existence of objects in a model of the legal domain.

2) With respect to the justification of knowledge, the following can be concluded. Reasons (and defeaters) can function as basic elements for legal argumentation. This means that they can be used as basic building blocks to

represent larger argumentation structures. Moreover, guiding reasons may constitute existence status layers for certain objects. Thus, the intertwinement of the epistemic level and the ontological level becomes visible in more detail: justification in the form of reasons may constitute certain aspects of the existence an object (for instance, a decision).

3) In the relations between objects (both tangible and intangible ones) two relations play a major role: counting as and causation. The former serves to assign a new status to an object. The latter indicates the consequences of the occurrence of an event. The counting-as relation is a useful mechanism for the description and explanation of the existence of (intangible) structures of social and legal entities. Regarding the causation relation, a distinction is made between physical and intentional causation, catering for the need to explain the interaction between the mental and the physical world.

The fourth research question

The fourth question regards the actual use of the ontology of law for the representation of knowledge in the domain of Dutch penal law. With respect to this question, I list two conclusions.

1) An analysis of the legal domain leads to the conclusion that there are six basic categories: entities, ontological status layers, epistemic roles, relations, acts and facts. These categories, and subtypes belonging to these categories, were used to represent characteristics and statutes of Dutch penal law. In fact, these characteristics and statutes were translated into expressions of reason-based logic, employing the categories of the knowledge-based ontology of law.

2) From this application of the knowledge-based ontology of law, and the application of a slightly adapted version of reason-based logic, it may be concluded that it is possible to express an extensive ontology of law in a formal language, without losing to much expressive power. However, in this case, a higher expressiveness has a trade-off: a lower usefulness of these expressions for automated inferencing.

The problem definition

Above, I have drawn conclusions with respect to the research questions. Still, I have not yet given an explicit answer to the main question. The problem definition was: What role can a concept of legal knowledge, formulated from the perspective of epistemology, play in the representation of legal knowledge? The implicit answer to this question consisted of a discussion of legal epistemology, legal ontology, knowledge representation, penal law, and a knowledge-based

model of the legal domain. What began as a theoretical exercise into the nature of knowledge about the legal domain, was gradually made concrete in a listing of categories of entities, relations and ontological status layers and epistemic roles.

All these subjects were meant to provide insight into the content, structure, and borders of knowledge about the legal domain. They lead to the most important conclusion of this thesis: if we wish to use a concept of knowledge as the basis of an ontology of law, we have to accept the fact that we have to define content, structure, and borders of knowledge about the legal domain ourselves, based on a view on the nature of law. By explaining the relation between knowledge of the law and the law itself, following different views on the nature of law, it becomes possible to provide basic categories of a model. The model allows – to a certain extent – for the accommodation of the epistemic and ontological commitments of different legal-theoretical positions, thereby bringing legal theory to the domain of legal knowledge representation.

References

Aarnio, A. (1981), 'Truth and Acceptability of Interpretative Propositions', in: *Rechtstheorie, Beiheft* **2**, p. 33-51.

Aarnio, A. (1987), *The Rationable as Reasonable. A Treatise on Legal Justification*, Dordrecht: D. Reidel Publishing Company.

Alexy, R. (1998), 'Coherence and Argumentation or the Genuine Twin Criteri-aless Super Criterion', in: Aarnio, A., R. Alexy, A. Peczenik, W. Rabinowicz, and J. Wolenski (1998), *On Coherence Theory of Law*, Lund: Juristförlaget i Lund, p. 41-49.

Algra, N. and Duyvendijk, K. van (1989), *Rechtsaanvang. Enkele hoofdstukken over recht en rechtswetenschap voor het onderwijs in de Inleiding tot de Rechtswetenschap*, Alphen aan den Rijn: Samsom H.D. Tjeenk-Willink.

Audi, R. (1993), *The structure of justification*, Cambridge: Cambridge University Press.

Audi, R. (1997), *Moral Knowledge and Ethical Character*, New York/Oxford: Oxford University Press.

Audi, R. (1998), *Epistemology. A contemporary introduction to the theory of knowledge*, London and New York: Routledge.

Austin, J. (1954), *The Province of Jurisprudence Determined*, New York: The Noonday Press.

Barwise, J. and J. Perry (1983), *Situations and Attitudes*, Cambridge, Mass./London, England: The MIT Press.

Bench-Capon, T.J.M. (1990), *Knowledge representation. An approach to Artificial Intelligence*, London: Academic Press.

Bender, J.W. (1989) (ed.), *The Current State of the Coherence Theory. Critical Essays on the Epistemic Theories of Keith Lehrer and Laurence BonJour, with Replies*, Dordrecht, Boston, London: Kluwer Academic Publishers.

Bix, B. (1991), 'H.L.A. Hart and the Open Texture of Language', in: *Law and Philosophy* **10**, p. 51-72.

BonJour, L. (1985), *The Structure of Empirical Knowledge*, Cambridge, Massachusetts: Harvard University Press.

Breuker, J., A. Valente, and R. Winkels (1997), Legal Ontologies: A Functional View, in: Visser and Winkels (1997), p. 23-36.

Brouwer, P.W. (1990), *Samenhang in recht. Een analytische studie*, Groningen: Wolters-Noordhoff.

Brouwer, P.W. (1997), 'Rechtspositivisme', in: Cliteur, P.B., B.C. Labuschagne, and C.E. Smith (eds.) (1997), *Rechtsfilosofische stromingen van de twintigste eeuw*, Arnhem: Gouda Quint, p. 109-147.

Brouwer, P.W. (1999), 'Systematisering van recht', in: P.W. Brouwer, M.M. Henket, A.M. Hol, and H. Kloosterhuis (eds.) (1999), *Drie dimensies van recht. Rechtstheorie, rechtsgeleerdheid, rechtspraktijk*, p. 219-237.

Cleiren, C.P.M. en J.F. Nijboer (eds.) (1994), *Tekst en commentaar strafrecht*, Deventer: Kluwer.

Cliteur, P.B. (1989), *Conservatisme en cultuurrecht*. Leiden University: Ph.D. thesis.

Devitt, M. (1991), *Realism and Truth*. 2nd edition, Cambridge, Mass./Oxford: Blackwell.

Doomen, J. (1990), *De hollende kleurling. Het Nederlandse strafrecht in negen verhalen*, Amsterdam: Contact.

Dworkin, R. (1986), *Law's Empire*, London, UK: Fontana Press.

Finnis, J.M. (1980), *Natural law and natural rights*, Oxford: Clarendon Press.

Franken, A.A. (1995), *'Hetzelfde feit'. Over samenloop van strafbare feiten en het non bis in idem-beginsel*. Ars Aequi Cahiers Strafrecht deel 4, Nijmegen: Ars Aequi Libri.

Franken, H., P.W. Brouwer, P.B. Cliteur, M.C.E. van Gorkom-Meeuwsen, M.E.W. Neleman-Dobschütz, I.L.E. Sanders, C.E. Smith, and R.Ch. Verschuur (1995), *In*Leiden *tot de rechtswetenschap*, Arnhem: Gouda Quint.

Frege, G. (1892), 'Über Sinn und Bedeutung', in: *Zeitschrift für Philosophie und Philosophische Kritik* **100**: 25-50.

Gamut, L.T.F. (1991a), *Logic, Language, and Meaning. Vol 1 Introduction to Logic*, Chicago/London: The University of Chicago Press.

Gamut, L.T.F. (1991b), *Logic, Language, and Meaning. Vol 2 Intensional Logic and Logical Grammar*, Chicago/London: The University of Chicago Press.

Gettier, E. (1963), 'Is justified true belief knowledge?', in: *Analysis* **23**, p. 121-123.

Gibson, J.J. (1979), *The ecological approach to visual perception*, Boston: Houghton Mifflin.

Goldman, A.J. (1976), 'Discrimination and Perceptual Knowledge', in: *Journal of Philosophy* **73**, p. 771-791.

Gombrich, E.H. (1977), *Art & Illusion. A study in the psychology of pictorial representation*, London: Phaidon.

Gordon, T.F. (1995), *The Pleadings Game – An artificial intelligence model of procedural justice*, Dordrecht: Kluwer Academic Publishers.

Gruber, T.R. (1993), 'A Translation Approach to Portable Ontology Specifications', in: *Knowledge Acquisition* **5**, p. 199-220.

Hage, J.C. (1987), *Feiten en Betekenis. Een verhandeling over ontologie en praktische rede*, Leiden University: Ph.D. thesis.

Hage, J.C. (1996), 'A theory of legal reasoning and a logic to match', in: *Artificial Intelligence and Law* **4**: 199-273.

Hage, J.C. (1997), *Reasoning with Rules. An Essay on Legal Reasoning and Its Underlying Logic*, Dordrecht: Kluwer Academic Publishers.

Hage, J.C. (1998), 'Institutionele rechtstheorieën', in: *Rechtsfilosofie & Rechtstheorie*, **27**: 127-143.

Hage, J.C. (1999), *Commentaar op 'Legal information systems as a source of knowledge'*.

Hardwig, J. (1985), 'Epistemic Dependence', in: *Journal of Philosophy,* **82**: 335-349.

Hart, H.L.A. (1961), *The Concept of Law*, Oxford, UK: Oxford University Press (ninth impression 1978).

Herik, H.J. van den (1991), *Kunnen computers rechtspreken?*, Arnhem: Gouda Quint.

Hoven, M.J. van den (1995), *Information Technology and Moral Philosophy. Philosophical Explorations in Computer Ethics*, Erasmus University Rotterdam: Ph.D. thesis.

Jörg, N. and C. Kelk (1994), *Strafrecht met mate*, Arnhem: Gouda Quint.

Kelsen, H. (1960), *Reine Rechtslehre*, Vienna: Verlag Franz Deuticke, 1960 (2nd edition).

Kelsen, H. (1979), *Allgemeine Theorie der Normen*, Wien: Manzsche Verlags- und Universitätsbuchhandlung.

Kirkham, R.L. (1992), *Theories of Truth. A Critical Introduction*, Cambridge, Mass./ London, England: The MIT Press.

Kralingen, R.W. van (1995), *Frame-based Conceptual Models of Statute Law*, Leiden University: Ph.D. thesis.

Leenes, R.E. (1998), *Hercules of Karneades, Hard cases in recht en rechtsinformatica*, University of Twente: Ph.D. thesis.

Lodder, A.R. (1998), *DiaLaw. On legal justification and dialog games*, Maastricht University: Ph.D. thesis.

Loth, M.A. (1988), *Handeling en aansprakelijkheid in het recht*, Leiden University: Ph.D. thesis.

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

MacCormick, D.N. and O. Weinberger (1986), *An Institutional Theory of Law. New Approaches to Legal Positivism*, Dordrecht: Reidel.

Mommers, L, A.H.J. Schmidt, and E.W. Oskamp (1997), Controversies in the ontology and law debate, in: Visser and Winkels (1997), p. 1-5.

Mommers, L. (1999), *Knowing the law. Legal information systems as a source of knowledge, ITeR series, no. 25*, Deventer: Kluwer.

Mommers, L. and H.J. van den Herik (2000), 'Knowledge Criteria for the Evaluation of Legal Beliefs', in: *Information & Communications Technology Law* **9**: 139-148.

Niiniluoto, I. (1981), 'Truth and Argumentation in Legal Dogmatics', in: *Rechts-theorie, Beiheft* **2**: 53-76.

Nijboer, J.F. (1995), *Strafrechtelijk bewijsrecht*, Nijmegen: Ars Aequi Libri.

Oskamp, E.W. (1998), *Computerondersteuning bij Straftoemeting*, Leiden University: Ph.D. thesis.

Patterson, D. (1996), *Law and Truth*, New York and Oxford: Oxford University Press.

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

Peczenik, A. and J.C. Hage (2000), 'Legal knowledge about what?', in: *Ratio Juris* **13**: 325-345.

Pollock, J.L. (1974), *Knowledge and Justification*, Princeton, New Jersey: Princeton University Press.

Pollock, J.L. (1999), *Contemporary Theories of Knowledge*, London *etc.*: Hutchinson.

Pollock, J.L. (2001), Defeasible reasoning with variable degrees of justification, in: *Artificial Intelligence* **133**: 233-282.

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

Prakken, H. (1997), *Logical tools for modelling legal argument: a study of defeasible reasoning in law*, Dordrecht *etc.*: Kluwer Academic Publishers.

Prakken, H. (1999), On Formalizing Burden of Proof in Legal Argument, in: H.J. van den Herik, M.-F. Moens, J. Bing, B. van Buggenhout, J. Zeleznikow, C. Grütters (1999), *Legal Knowledge Based Systems. Jurix 1999: The Twelfth Conference*, Nijmegen: GNI, p. 85-97.

Putnam, H. (1975), 'The meaning of 'meaning'', in: K. Gunderson (ed.), *Language, Mind, and Knowledge, Minnesota Studies in the Philosophy of Science,* vol. VII. Minneapolis: University of Minnesota Press, pp. 131-193.

Putnam, H. (1990), Realism With a Human Face, Cambridge, Mass.: Harvard University Press.

Quine, W.V.O. (1960), *Word and Object*, Cambridge, Mass.: MIT Press.

Quine, W.V.O. (1990), *The Pursuit of Truth*, Cambridge, Mass./London, England: Harvard University Press.

Raz, J. (1970), *The Concept of a Legal System. An Introduction to the Theory of Legal System*, Oxford: Clarendon Press.

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

Remmelink, J. (1996), *Mr. D. Hazewinkel-Suringa's Inleiding tot de studie van het Nederlandse strafrecht*, Arnhem: Gouda Quint.

Royakkers, L.M.M. (1996), *Representing legal rules in deontic logic*, Tilburg University: Ph.D. thesis.

Ruiter, D.W.P. (1993), *Institutional Legal Facts. Legal Powers and their Effects*, Dordrecht: Kluwer Academic Publishers.

Ryle, G. (1949), *The Concept of Mind*, London: Hutchinson.

Scholten, P. (1974), *Mr. C. Asser's handleiding tot de beoefening van het Nederlands burgerlijk recht. Algemeen deel*, Zwolle: W.E.J. Tjeenk Willink.

Searle, J.R. (1969), *Speech Acts: an essay in the philosophy of language*, Cambridge: Cambridge University Press.

Searle, J.R. (1995), *The Construction of Social Reality*, London, UK: Penguin Books.

Searle, J.R. (1999), *Mind, Language and Society*, London: Weidenfeld & Nicolson.

Shaver, K.G. (1985), *The Attribution of Blame. Causality, Responsibility, and Blameworthiness*, New York: Springer Verlag.

Smith, C.E. (1998), *Feit en rechtsnorm: een methodologisch onderzoek naar de betekenis van de feiten voor de rechtsvinding en legitimatie van het rechtsoordeel*, Leiden University: Ph.D. thesis.

Torre, L.W.N. van der (1997), *Reasoning about Obligations. Defeassibility in Preference-based Deontic logic*, Erasmus University Rotterdam: Ph.D. thesis.

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

Valente, A. (1995), *Legal knowledge engineering: a modelling approach*, University of Amsterdam: Ph.D. thesis.

Verheij, H.B. (1996), *Rules, Reasons, Arguments. Formal studies of argumentation and defeat*, Maastricht University: Ph.D. thesis.

Verheij, H.B. and J.C. Hage (1997), States-of-affairs, Events, and Rules: an Abstract Model of the Law, in: Oskamp, A, R.V. de Mulder, C. van Noortwijk, CA.F.M. Grütters, K. Ashly, and T. Gordon (1997), *Legal Knowledge Based Systems. Jurix: The Tenth Conference*, Jurix, p. 3-19.

Visser, P.R.S. (1995), *Knowledge Specification for Multiple Legal Tasks. A Case Study of the Interaction Problem in the Legal Domain*, Leiden University: Ph.D. thesis.

Visser, P.R.S. and R.G.F. Winkels (eds.) (1997), *Proceedings of the First International Workshop on Legal Ontologies, LEGONT '97, July 4, 1997, University of Melbourne, Law School*.

Waismann (1952), 'Verifiability', in A. Flew (1952), *Logic and language (first series)*, Oxford: Basil Blackwell.

Wittgenstein, L. (1953), *Philosophical Investigations*, Oxford, UK: Blackwell, translated from *Philosophische Untersuchungen* by G.E.M. Anscombe.

Dutch translations of legal terms

abortion — afbreken van zwangerschap

absence of all guilt — afwezigheid van alle schuld

absence of substantive illegality — afwezigheid van materiële wederrechtelijkheid

absolutely impeded omission — absoluut verhinderd nalaten

accomplice — medeplichtige

acquittal — vrijspraak

adjournment — schorsing

affect — affect

amenable — ontvankelijk

attempt — poging

authorised official order — bevoegd gegeven ambtelijk bevel

being of unsound mind — ontoerekeningsvatbaar

carelessness — onvoorzichtigheid

charge — tenlastelegging

co-commit — medeplegen

commencement of implementation — begin van uitvoering

commit — plegen

competent — bevoegd

complicity — medeplichtigheid

conscientious objections — gewetensdrang

constituent — bestanddeel

continuous act — voortgezette handeling

crimes against life — levensdelicten

crimes against life and violent crimes — levens- en geweldsdelicten

culpability — schuld (*culpa*)

culpable causation of death or physical injury — veroorzaken van de dood of van lichamelijk letsel door schuld

culpable homicide — dood door schuld

death — dood (overlijden)

detention — hechtenis

Dutch Supreme Court — Hoge Raad

emergency — noodtoestand

evidence — bewijsmateriaal

exemption ground	uitsluitingsgrond
exemption ground from punishment	strafuitsluitingsgrond
exemption ground of guilt	schulduitsluitingsgrond
exemption of legal proceedings	ontslag van rechtsvervolging
fine	geldboete
force majeure	overmacht
force majeure as necessity	overmacht als noodtoestand
formal source of law	formele rechtsbron
freedom of contract	contractvrijheid
guilt	schuld
imprisonment	gevangenisstraf
intent	opzet
judgement	arrest
ground for justification	rechtvaardigingsgrond
knowledge source for the law	kenbron van het recht
law of obligations	verbintenissenrecht
legal regulation	wettelijk voorschrift
legality	legaliteit
legitimacy	legitimiteit
legitimate and convincing evidence	wettig en overtuigend bewijs
make commit	doen plegen
malice aforethought	voorbedachte raad
maltreatment	mishandeling
manslaughter	doodslag
material source of law	materiële rechtsbron
misuse of authority	misbruik van gezag
moral pressure	morele druk
murder	moord
offence description	delictomschrijving
participation	deelneming
proportionality	proportionaliteit
providing of an opportunity	verschaffen van gelegenheid
provocation	uitlokking
psychological force majeure	psychische overmacht
publication	openbaarmaking
punishability	strafbaarheid

qualified manslaughter	gekwalificeerde doodslag
reasonable attribution	redelijke toerekening
reproach of guilt	schuldverwijt
reproachability	verwijtbaarheid
reproduction	verveelvoudiging
self-defence	noodweer
source of law	rechtsbron
special duty of care	bijzondere zorgplicht
subsidiarity	subsidiariteit
summons	dagvaarding
two or more offences arising from the same act	eendaadse samenloop
unauthorised official order	onbevoegd gegeven ambtelijk bevel
unauthorised official order that was considered to be given authorised in good faith	onbevoegd gegeven ambtelijk bevel dat te goeder trouw bevoegd gegeven werd geacht
unreasonable use of self-defence	noodweerexces
unsound attempt	ondeugdelijke poging
various offences for which one sentence is pronounced	meerdaadse samenloop
void	nietig
voluntary retiration	vrijwillig terugtreden
wrongful act	onrechtmatige daad

Index

Summary

This thesis concerns the representation of knowledge about the legal domain. It is based on the assumption that philosophical theories of knowledge can contribute to the quality of knowledge representation. In chapter 1, I provide an introduction to the questions discussed in this book, focusing on the role of the concept of knowledge in representing the legal domain.

In chapter 2, I discuss the subject what knowledge is in accordance with an epistemic viewpoint. I show the differences between three dimensions of knowledge: its acquisition, its object, and its justification. Five acquisition sources are distinguished: perception, memory, consciousness, reason, and testimony. These account for the broad spectre of channels through which we can acquire knowledge. The object of knowledge is the domain of ontology. As properties within the object dimension of knowledge I name the type/token distinction and the kind of existence concerned. These properties cover the ways in which we can refer to and generalise over concrete objects. As to the justification dimension of knowledge, I carefully distinguish between the different meanings of justification: justification as a state, as a process, and as a status. Additionally, I explain the role of reasons and defeaters in justification. The same dimensions of knowledge are discussed with respect to a special class of knowledge: legal knowledge.

Furthermore, I discuss the distinction between different types of knowledge theories. They regard the ways in which the acquisition and justification of knowledge is attained. The qualification of beliefs as knowledge involves the establishment of knowledge criteria, and the distinction between different belief types. Various belief types may demand the imposition of different knowledge criteria. Finally, I use the phenomenon called *epistemic niche* to explain the situation in which legal professionals acquire knowledge.

In chapter 3, I discuss legal knowledge from a legal-theoretical viewpoint. I start the chapter with a discussion of the hermeneutical stance in the theory of law. As an example of such a stance I discuss Dworkin's view on knowing the law. In addition, I derive epistemic claims from several stances in legal philosophy: natural law, legal positivism, institutional theories of law, a conventionalist-*cum*-institutional approach, and hermeneutical theories. Finally, I explain what formal and material sources of law are, and how they can function as knowledge sources for the law.

In chapter 4, I deal with the constituting entities of the legal domain, *i.e.*, with the ontological characteristics of law. I give an overview of ontological claims

with respect to the law from different legal-theoretical stances. Furthermore, two objects of legal knowledge are discussed: systematisations and interpretations; moreover, the elements that are part of the reasoning process that leads to establishing these objects are explained. For that purpose, I discuss Toulmin's seminal work on argumentation, and I give an elaborate view on reasons and defeaters. Subsequently, a distinction is made between different ontological status layers. Such ontological status layers express the different modes of existence that we can find for objects in the legal domain. The status layers distinguished are: existence, validity, recognition, constitution, and efficacy. Finally, I introduce two basic relations that take care of the connections between different objects in the legal domain: counting as and causality. Counting as establishes constitution relations between objects and facts, and is thus the primary basis of institutional facts. Causality in the legal domain is very much subject to theoretical considerations of the question what ought to be considered as a cause for some phenomenon.

Chapter 5 is devoted to the representation of legal knowledge. In this chapter, I discuss three subjects. The first subject is the concept of representation. Several representation languages are discussed, as well as demands that can be imposed upon a representation language for the legal domain. The second subject is the question how classic views on meaning can be applied to legal concepts. The third subject consists of legal ontologies, *i.e.*, conceptual frameworks for the representation of legal knowledge.

In chapter 6, I introduce a knowledge-based model of the law. This model consists of entities, ontological status layers, epistemic roles, relations, acts and facts. Ontological status layers are the ontological states that an object can have. Epistemic roles are knowledge-related states an object may acquire. For all entities, relations *etc.*, subtypes are distinguished. The combination of entities, relations, acts and facts with ontological status layers and epistemic roles enables the expression of epistemic and ontological views with respect to legal entities.

In chapter 7, some central notions within Dutch penal law are explained. This chapter serves to have sufficient knowledge about the legal domain for an example of the concrete application of the knowledge-based model of law. The questions underlying guilt and punishability, other main principles and concepts of penal law, and some specific articles in the domain of crimes against life are explained.

In chapter 8, I use the ontology presented in chapter 6 to represent some of the characteristics of Dutch penal law. For this purpose, I choose the representation language Reason–Based Logic. This language is adapted to fit in with the

specific ontological and epistemic demands imposed on the basis of the knowledge-based model of law. I indicate how the elements of the knowledge-based model of law should be represented in this language. Furthermore, inference rules regarding knowledge qualification and reasoning are introduced.

In chapter 9, I draw conclusions, grouped as answers to the research questions and the problem definition. The problem definition of this thesis is: What role can a concept of legal knowledge, formulated from the perspective of epistemology, play in the representation of legal knowledge? The main conclusion is, that if we wish to use a concept of knowledge as the basis of an ontology of law, we have to define content, structure, and borders of knowledge about the legal domain, based on a view on the nature of law. By explaining the relation between knowledge of the law and the law itself, following different views on the nature of law, it becomes possible to provide basic categories of a model. The model provided allows – to a certain extent – for the accommodation of the epistemic and ontological commitments of different legal-theoretical positions, thereby bringing legal theory to the domain of legal knowledge representation.

Samenvatting

Dit proefschrift gaat over de representatie van kennis van het juridische domein, en is gebaseerd op de vooronderstelling dat filosofische theorieën over kennis kunnen bijdragen aan de kwaliteit van kennisrepresentatie. In hoofdstuk 1 leid ik de vragen in die in dit boek besproken worden, waarbij ik de nadruk leg op de rol van het concept van kennis bij de representatie van het juridische domein.

In hoofdstuk 2 bespreek ik, vanuit een epistemologisch gezichtspunt, de vraag wat kennis is. Daarbij ga ik in op de verschillen tussen drie dimensies: de verkrijging, het object en de rechtvaardiging van kennis. Er worden vijf bronnen voor de verkrijging van kennis onderscheiden: perceptie, geheugen, bewustzijn, rede en getuigenis. Het object van kennis wordt onderzocht in de ontologie. Belangrijke onderscheidingen met betrekking tot het object van kennis zijn het verschil tussen type en token, en tussen verschillende vormen van bestaan. Zij betreffen de manieren waarop we naar concrete objecten kunnen verwijzen, en hoe we er algemene uitspraken over kunnen doen. De rechtvaardiging van kennis valt uiteen in verschillende typen: rechtvaardiging als toestand, als proces en als status. Daarnaast bespreek ik de rol van redenen en weerleggingen in rechtvaardiging. Dezelfde dimensies van kennis worden bovendien besproken met betrekking tot een speciale soort kennis: juridische kennis.

In hoofdstuk 3 bespreek ik juridische kennis vanuit een rechtstheoretische invalshoek. Een toelichting op het hermeneutische standpunt in de rechtstheorie is aanleiding voor een nadere uitleg van Dworkin's gezichtspunt op kennis van het recht. Ik geef aan welke epistemologische claims kunnen worden afgeleid uit verschillende rechtstheoretische perspectieven, te weten natuurrecht, rechtspositivisme, institutionele rechtstheorie, conventioneel/institutionele rechtstheorie en hermeneutische theorieën. Daarnaast licht ik toe wat formele en materiële rechtsbronnen zijn, en hoe zij kunnen fungeren als kenbronnen voor het recht.

In hoofdstuk 4 ga ik in op de entiteiten die deel uitmaken van het juridisch domein. Daarbij geef ik een overzicht van ontologische claims met betrekking tot het recht vanuit verschillende rechtstheoretische gezichtspunten. Ik bespreek twee objecten van juridische kennis in meer detail: systematiseringen en interpretaties. Dergelijke objecten zijn het resultaat van een redeneerproces. In dat redeneerproces zelf kunnen we bepaalde entiteiten onderscheiden. Toulmin's standaardwerk over argumentatie vormt de aanzet tot een uitgebreide discussie van dergelijke entiteiten, zoals redenen en weerleggingen.

Vervolgens wordt het onderscheid geïntroduceerd tussen verschillende ontologische statuslagen, die uitdrukking geven aan de verschillende manieren waarop objecten in het juridische domein kunnen bestaan. De onderscheiden statuslagen zijn: bestaan, geldigheid, erkenning, constitutie en effectiviteit. In aanvulling op deze statuslagen zorgen twee belangrijke relaties voor de verbanden tussen verschillende objecten in het juridische domein: gelden als (counting as) en causaliteit. Gelden als vestigt constitutierelaties tussen objecten en feiten, en vormt daardoor de belangrijkste bestaansvoorwaarde voor institutionele feiten. Causaliteit in het juridische domein is sterk verbonden met de vraag wat als oorzaak moet worden aangemerkt voor een bepaald fenomeen.

In hoofdstuk 5 besteed ik aandacht aan de representatie van juridische kennis. Ten eerste bespreek ik het concept van representatie. Verschillende representatietalen komen aan de orde, naast eisen die aan een representatietaal voor het juridisch domein gesteld kunnen worden. Ten tweede komt de vraag naar voren hoe traditionele opvattingen over betekenis kunnen worden toegepast op juridische concepten. Ten derde worden drie juridische ontologieën besproken. Dit zijn conceptuele raamwerken voor de representatie van juridische kennis.

In hoofdstuk 6 bespreek ik het kennisgebaseerde model van het juridisch domein. Het model bestaat uit entiteiten, ontologische statuslagen, epistemologische rollen, relaties, handelingen en feiten. Ontologische statuslagen zijn de verschillende manieren waarop een object kan bestaan. Epistemologische rollen zijn rollen die een object kan hebben met betrekking tot kennis. Voor alle basiscategorieën van het model worden subtypen onderscheiden. De combinatie van entiteiten, relaties, handelingen en feiten met ontologische statuslagen en epistemologische rollen maakt het mogelijk om epistemologische en ontologische gezichtspunten uit te drukken met betrekking tot juridische entiteiten.

In hoofdstuk 7 volgt een overzicht van enkele centrale kenmerken van het Nederlandse strafrecht. Dit hoofdstuk dient om voldoende kennis over het juridische domein weer te geven om tot een concrete toepassing van het kennisgebaseerde model van het juridische domein te komen. Hierbij komen onder meer de vragen aan de orde waarmee de schuld en strafbaarheid van een verdachte worden bepaald, naast andere belangrijke principes en concepten van het strafrecht, en enkele specifieke artikelen met betrekking tot levensdelicten.

In hoofdstuk 8 gebruik ik de ontologie uit hoofdstuk 6 om enkele eigenschappen van het Nederlandse strafrecht te representeren. Bij de representatie koos ik voor de representatietaal Reason-Based Logic. Deze taal wordt aangepast om te voldoen aan de specifieke ontologische en epistemologische eisen die op basis van het kennisgebaseerde model worden gesteld. Ik geef aan hoe

de elementen van het model in deze taal kunnen worden gerepresenteerd. Daarnaast introduceer ik een aantal afleidingsregels voor kenniskwalificatie en redeneren in het algemeen.

In hoofdstuk 9 trek ik conclusies. Daarbij volg ik de onderzoeksvragen en de probleemdefinitie. De probleemdefinitie van dit proefschrift luidt: welke rol kan een concept van juridische kennis, geformuleerd vanuit het standpunt van de kenleer, spelen in de representatie van juridische kennis? De belangrijkste conclusie is als volgt. Als we een concept van kennis willen gebruiken als basis voor een ontologie van het recht, dan moeten we de inhoud, structuur en grenzen van kennis over het juridische domein definiëren, waarbij we ons baseren op een visie op de aard van het recht. Door de relatie tussen kennis van het recht en het recht zelf te bespreken, waarbij verschillende visies op de aard van het recht worden aangehouden, wordt het mogelijk om basiscategorieën van een model te bepalen. Het model dat in dit proefschrift wordt beschreven, laat tot op zekere hoogte toe dat de diverse epistemologische en ontologische claims van verschillende rechtstheoretische posities een plaats kunnen krijgen binnen één model, waardoor rechtstheorie binnen het bereik van juridische kennisrepresentatie wordt gebracht.

Curriculum vitae

Laurens Mommers was born in Doetinchem on the 23rd of August, 1972. From 1985 to 1991, he went to the Ludger College in Doetinchem. After he received his diploma, he went to the University of Nijmegen. He studied political science, philosophy (master's degree in 1996), and computational linguistics (master's degree cum laude in 1996). At the department of Law and Information Technology in Leiden, he started working on a Ph.D. thesis on legal knowledge representation. Apart from doing his Ph.D research, he also worked as a lecturer and researcher at the department.

SIKS dissertation series

1998-1 Johan van den Akker (CWI) – DEGAS – An Active, Temporal Database of Autonomous Objects

1998-2 Floris Wiesman (UM) – Information Retrieval by Graphically Browsing Meta-Information

1998-3 Ans Steuten (TUD) – A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective

1998-4 Dennis Breuker (UM) – Memory versus Search in Games

1998-5 Eduard Oskamp (RUL) – Computerondersteuning bij Straftoemeting

1999-1 Mark Sloof (VU) – Physiology of Quality Change Modelling; Automated modelling of Quality Change of Agricultural Products

1999-2 Rob Potharst (EUR) – Classification using decision trees and neural nets

1999-3 Don Beal (UM) – The Nature of Minimax Search

1999-4 Jacques Penders (UM) – The Practical Art of Moving Physical Objects

1999-5 Aldo de Moor (KUB) – Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems

1999-6 Niek J.E. Wijngaards (VU) – Re-design of compositional systems

1999-7 David Spelt (UT) Verification support for object database design

1999-8 Jacques H.J. Lenting (UM) – Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation

2000-1 Frank Niessink (VU) – Perspectives on Improving Software Maintenance

2000-2 Koen Holtman (TUE) – Prototyping of CMS Storage Management

2000-3 Carolien M.T. Metselaar (UVA) – Sociaal-organisatorische gevolgen van kennistechnologie; een procesbenadering en actorperspectief.

2000-4 Geert de Haan (VU) – ETAG, A Formal Model of Competence Knowledge for User Interface Design

2000-5 Ruud van der Pol (UM) – Knowledge-based Query Formulation in Information Retrieval

2000-6 Rogier van Eijk (UU) – Programming Languages for Agent Communication

2000-7 Niels Peek (UU) – Decision-theoretic Planning of Clinical Patient Management

2000-8 Veerle Coupé (EUR) – Sensitivity Analyis of Decision-Theoretic Networks

2000-9 Florian Waas (CWI) – Principles of Probabilistic Query Optimization

2000-10 Niels Nes (CWI) – Image Database Management System Design Con-

siderations, Algorithms and Architecture

2000-11 Jonas Karlsson (CWI) – Scalable Distributed Data Structures for Database Management

2001-1 Silja Renooij (UU) – Qualitative Approaches to Quantifying Probabilistic Networks

2001-2 Koen Hindriks (UU) – Agent Programming Languages: Programming with Mental Models

2001-3 Maarten van Someren (UVA) – Learning as problem solving

2001-4 Evgueni Smirnov (UM) – Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets

2001-5 Jacco van Ossenbruggen (VU) – Processing Structured Hypermedia: A Matter of Style

2001-6 Martijn van Welie (VU) – Task-based User Interface Design

2001-7 Bastiaan Schonhage (VU) – Diva: Architectural Perspectives on Information Visualization

2001-8 Pascal van Eck (VU) – A Compositional Semantic Structure for Multi-Agent Systems Dynamics

2001-9 Pieter Jan 't Hoen (RUL) – Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes

2001-10 Maarten Sierhuis (UVA) – Modeling and Simulating Work Practice Brahms: a multiagent modeling and simulation language for work practice analysis and design

2001-11 Tom M. van Engers (VU) – Knowledge Management: The Role of Mental Models in Business Systems Design

2002-01 Nico Lassing (VU) – Architecture-Level Modifiability Analysis

2002-02 Roelof van Zwol (UT) – Modelling and searching web-based document collections

2002-03 Henk Ernst Blok (UT) – Database Optimization Aspects for Information Retrieval

2002-04 Juan Roberto Castelo Valdueza (UU) – The Discrete Acyclic Digraph Markov Model in Data Mining

2002-05 Radu Serban (VU) – The Private Cyberspace: Modeling Electronic Environments inhabited by Privacy-concerned Agents

2002-06 Laurens Mommers (UL) – Applied legal epistemology: building a knowledge-based ontology of the legal domain

Applied legal epistemology examines several practical implications of theoretical insights into the nature of legal knowledge. In this book, a small part of Dutch penal law is modelled, using a conceptual framework based on theoretical investigations into legal knowledge. Knowledge can be regarded with respect to three main dimensions: its acquisition, its object, and its justification. In each dimension, the specific characteristics of legal knowledge become apparent.

Apart from a theoretical exercise into the nature of legal knowledge, this book provides opportunities for lifting the representation of legal knowledge to a higher level. Specific additional information on the status of represented knowledge can, for instance, be used to determine whether it is sufficiently justified. It may thus prove relevant for the legal-philosophically inclined as well as for those putting their effort into building legal knowledge-based systems.