



# The law as a dynamic interconnected system of states of affairs: a legal top ontology<sup>†</sup>

JAAP HAGE AND BART VERHEIJ

*Department of Metajuridica, Universiteit Maastricht, P.O. Box 616, 6200 MD Maastricht, The Netherlands. emails: jaap.hage@metajur.unimaas.nl, bart.verheij@metajur.unimaas.nl*

In this paper, an abstract model of the law is presented that has three primitives: states of affairs, events and rules. The starting point of the abstract model is that the law is a dynamic system of states of affairs which are connected by means of rules and events. The abstract model can be regarded as a top ontology of the law, that can be applied to legal knowledge representation. After an elaboration of the three primitives, the uses of the abstract model are illustrated by the analysis of central topics of law. Then we discuss heuristic guidelines for legal knowledge representation that are suggested by the abstract model. The paper concludes with a comparison with related work. The appendix contains a formalism for the abstract model.

© 1999 Academic Press

## 1. Modeling the law

In this paper, we present an abstract model of the law that is based on two crucial characteristics of the law. The first characteristic is that the law is a dynamic system of states of affairs. The law evolves over time: regulations change, contracts are signed, property rights are acquired, etc. The second characteristic is that the law is an interconnected system of states of affairs. The elements of the law are not independent of each other, but hang together in a rule-like way: stealing is punishable, the signing of a contract gives rise to obligations.

Our abstract model of the law can be regarded as a top ontology. Ontologies have recently attracted considerable interest of the field of knowledge representation in general<sup>‡</sup> and in the field of Law and Artificial Intelligence in particular<sup>§</sup>. Motivations for the development of ontologies, or explicit specifications of domain conceptualizations (Gruber, 1995), include knowledge sharing and knowledge reuse (cf. Bench-Capon & Visser, 1997). Our motivation to the development of an ontology is to provide an explicit view of the legal domain with the aim to find heuristic guidelines for legal knowledge representation.

<sup>†</sup>The present paper extends Verheij and Hage's (1997) paper at the tenth JURIX conference.

<sup>‡</sup>Compare the special issues on ontologies of the International Journal of Human-Computer Studies in Vol. 43, 1995 & Vol. 46, 1997 with papers by, e.g. Gruber (1995), Guarino (1995, 1997), Hobbs (1995), Sowa (1995), Van Heijst, Schreiber and Wielinga (1997*a, b*).

<sup>§</sup>Compare Bench-Capon (1989), McCarty (1989), Valente (1995), Van Kralingen (1995), Visser (1995, 1997), Den Haan (1996), Bench-Capon and Visser (1996, 1997), Visser, Bench-Capon and van den Herik (1997), and Visser and Bench-Capon (1998). Compare also the proceedings of the First International Workshop on Legal Ontologies (eds. Visser and Winkels, 1997).

The abstract model of the law as proposed in this paper can be summarized as follows:

- The law is a system of *states of affairs*.
- The law is *dynamic*: the obtaining states of affairs are subject to change due to the occurrence of *events*.
- The law is *interconnected*: there are (directed) connections between the obtaining states of affairs based on *rules*.

The model uses three primitives.

- *States of affairs*. A state of affairs can be characterized as a possible part of the world as expressed by a (descriptive) sentence. An example is the state of affairs that the contract has been signed as expressed by the sentence ‘The contract has been signed’.
- *Events*. An event causes a change of the obtaining states of affairs. An example is the event of signing some contract by which the state of affairs that the contract has been signed starts to obtain.†
- *Rules*. A rule is a directed connection between states of affairs. An example is the rule that, if the contract has been signed, obligations of the contractors towards each other emerge.

We start with a description of the abstract model in Section 2–5. The core of this paper consists of Sections 6–12 in which we illustrate the uses of the model by analysing some central legal topics. In Section 13, we reconsider the main elements of the abstract model after its elaboration in the examples. In Section 14, we discuss heuristic guidelines for legal knowledge representation as suggested by the abstract model. The model is put in perspective by the discussion of related research in Section 15. The paper is summarized in Section 16. In Appendix A, a formalism for the abstract model is provided.

## 2. Two types of connections between states of affairs

Our model distinguishes between two types of connections between states of affairs: *causation* and *constitution*. Causation involves the lapse of time, while constitution is timeless. An example about a sales contract illustrates the two types of connections.

Suppose that A sells his car to B by signing a sales contract. The signing of the contract is an event causing that a contractual bond between A and B comes about. The relation between the signing of the contract and the existence of the contractual bond between A and B is one of *causation*. The contractual bond brings with it that A is obligated to transfer the ownership of his car to B, and that B is obligated to pay A the price of the car. The relation between the existence of the contractual bond and the obligations of A and B towards each other is one of *constitution*.

In the case of causation, an event changes which states of affairs obtain. States of affairs appear or disappear.‡ Graphically, causation is depicted as a horizontal connection between states of affairs (Figure 1).

†The notions of states of affairs and events as we use them are related to, but not fully identical with those used by von Wright (1963, p. 25f.).

‡State transitions also played an important role in the model used by Gardner (1987).



FIGURE 1. Causation.




FIGURE 2. Constitution.

In the case of constitution, a state of affairs obtains thanks to another state of affairs that obtains. There is a rule that connects the states of affairs. Graphically, constitution is depicted as a vertical connection between states of affairs (Figure 2).<sup>†</sup>

In Section 5.1, we show that there are not only rules of constitution, but also rules of causation.

In the rest of this paper, we elaborate the abstract model of the law based on the distinction between constitution and causation, and show it to be beneficial for modeling the law.

### 3. States of affairs

It is convenient to view the law (and the world) as a system of states of affairs. A *state of affairs* can be characterized as a possible part of the world expressed by a (descriptive) sentence.<sup>‡</sup> We take the notion of a state of affairs rather broadly. Examples of states of affairs are the following.

1. It is raining.
2. George Washington was the first president of the USA.
3. The sun will rise tomorrow.
4. John has taken away Gerald's car.
5. John is a thief.
6. Meryl is under a contractual bond toward Jane to pay her \$100.
7. Meryl ought to pay Jane \$100.
8. A minor cannot make a valid will.
9. It is uncertain whether O.J. Simpson killed his wife.

<sup>†</sup>Visser (1995, pp. 92f., 155) makes analogous distinctions.

<sup>‡</sup>The close relation between states of affairs and sentences implies that the expressive power of the chosen language determines which states of affairs are possible.

10. From the point of view of civil law, O.J. Simpson killed his wife.
11. From the point of view of criminal law, O.J. did not kill his wife.

Obviously, some states of affairs obtain, while others do not obtain. For example, the states of affairs that  $3 + 4$  equals  $7$  obtains, while the state of affairs that Bill Clinton was the first president of the USA does not obtain. States of affairs that obtain are called *facts* and are expressed by true sentences. States of affairs that do not obtain are called *non-facts* and are expressed by false sentences.

As the examples show, states of affairs can be in different tenses (examples 1–3), can supervene on each other (examples 4/5, 6/7), can have different modalities (examples 7–10) and depend on a point of view (examples 10–11).

### 3.1. TEMPORARY AND DURABLE STATES OF AFFAIRS

The examples of states of affairs 1–3 above are in different tenses. We regard the law as a *dynamic* system of states of affairs: the obtaining states of affairs can change over time. For instance, the state of affairs that Bill Clinton is president of the USA obtains today (July 1998), but did not obtain in 1967. Some states of affairs can stop or start obtaining, others cannot. For instance, the state of affairs that George Washington was the first president of the USA obtains and will always obtain, since it is a state of affairs about the past.

States of affairs that can stop or start obtaining are said to be *temporary*, otherwise *durable*. An example of a temporary state of affairs is that it is raining; an example of a durable state of affairs is that the French Revolution took place in the 18th century. States of affairs that deal with the past are always durable, because the past does not change. For obvious reasons, tautological states of affairs are also durable. Temporary states of affairs that only obtain for a moment are *momentary*. A momentary state of affairs is for instance that John hits Gerald.

Temporary states of affairs which deal with the present, such as the state of affairs that it is raining, are called *states*. In Section 8, we show that different kinds of rights can be thought of as states.

### 3.2. SUPERVENIENCE

In the examples above, state of affairs 5 depends on state of affairs 4. The state of affairs that John is a thief obtains *due to* the state of affairs that John has taken away Gerald's car. It is said that the state of affairs that John is a thief *supervenes* on the state of affairs that he has taken away Gerald's car (Jones, 1995).

Supervenience of a state of affairs on another state of affairs is a rather common phenomenon. It can, amongst others, be based on definitions. For instance, something counts as a motor vehicle in the sense of the Dutch Traffic Law (*Wegenverkeerswet*) if and only if it satisfies a number of conditions.

In general, modal states of affairs, discussed in the next subsection, always supervene on other states of affairs. For instance, the state of affairs that Meryl ought to pay Jane \$100 (example 7 above) supervenes on the state of affairs that Meryl is under a contractual bond toward Jane to pay her \$100 (example 6 above).

### 3.3. MODALITIES

Examples 7–9 illustrate different modalities. We distinguish three categories of modal states of affairs: anankastic, deontic and probabilistic states of affairs. (Here we do not regard tense as a modality.)

*Anankastic* states of affairs (von Wright, 1963, p. 10) have to do with the necessary, the possible and the impossible. For instance, the state of affairs that the released stone *must* fall, is anankastic. Other examples are the states of affairs that hydrogen and oxygen *can* react, that the Democrats *cannot* win the elections and that the conclusion of a deductively valid argument with true premises is *necessarily* true.

A specific anankastic state of affairs in the law has to do with competence. To perform particular acts in the law, such as engage into a contract, to issue a governmental order or to legislate, the person who performs the act must have the competence to do so. If the competence is lacking, the particular juristic act cannot exist at all, or is void and has no legal consequences. In other words, competence has to do with what an actor can or cannot do.† We return to juristic acts in Section 10.

*Deontic* states of affairs have to do with the obligated, the forbidden and the permitted. Examples are that Meryl ought to pay Jane \$100, that smoking is prohibited in public buildings and that John is allowed to take a day off.

Two basic categories of deontic states of affairs are usually distinguished: deontic states of affairs of the *ought-to-do* type and of the *ought-to-be* type. Examples of the first category are that car drivers ought to drive on the right-hand side of the road, that public officers are prohibited to accept bribes and that John is permitted to walk in the park. Examples of the second category are that car drivers ought to be sober, that it is forbidden that high public officers are members of Parliament and that it is permitted that Jane walks in the park.‡

Deontic states of affairs should be distinguished from the non-modal states of affairs on which they supervene. An example is the state of affairs that there is a contractual bond between two parties, which underlies the state of affairs that one party has to pay the other.

*Probabilistic* states of affairs have to do with the probable, the certain and the uncertain. Examples of probabilistic states of affairs are that it will probably rain, that the train definitely will be late and that Jane might pay her bill.

Probabilistic states of affairs should be distinguished from anankastic states of affairs: the reasons why something is necessary are not those which make something probable or certain. The announcement that the train will be late makes it highly probable that the train will be late, but does not make it necessary.

†In the law, competence is sometimes assumed to be a state of affairs of the deontic modality. On that assumption, competence is considered to imply primarily the *permission* to perform an act in the law. However, it is better to consider the *capability* to perform the act as the primary modal state of affairs implied by competence.

‡The state of affairs that it is permitted *that* Jane walks in the park is of the *ought-to-be* type, while the state of affairs that John is permitted to walk in the park is of the *ought-to-do* type. The first concerns a (specific) state of affairs involving Jane, the second a (generic) act by John.

### 3.4. POINTS OF VIEW

The examples of states of affairs 10–11 depend on a point of view. Points of view include the logical, the physical, the biological, the social and the legal point of view. As examples 10 and 11 show, the legal point of view encompasses the points of view of civil and criminal law.

It should be noted that states of affairs can belong to more than one point of view. For instance, the state of affairs that John should be punished can belong to the social, the moral and the legal point of view. Moreover, states of affairs from different points of view can conflict. For instance, the states of affairs that O. J. Simpson killed his wife and that he did not kill his wife belong to the point of view of civil and of criminal law, respectively. Because these facts belong to different points of view, the conflict does not lead to an inconsistency.

## 4. Events

Events cause changes in the total set of obtaining states of affairs. For instance, if it starts to rain, the state of affairs that it is raining starts to obtain. Other examples of events are the following.

1. The starting of the European Economic and Monetary Union.
2. The apple's falling to the ground.
3. Jane's dying.
4. John taking away the car of Gerald.
5. The Supreme Court annulling the judgement of the Court of Justice.
6. An international treaty being ratified.
7. The transfer of the ownership of a house.

Notice that the *occurrence* of an event is itself a (momentary) state of affairs, for instance the state of affairs that John takes away Gerald's car.

A special kind of events are *acts*: events that consist of the intentional behaviour of an individual (examples 4–7). A special category of acts are the so-called *juristic acts* (examples 5–7). Juristic acts are discussed in Section 10.

### 4.1. THE EFFECTS OF AN EVENT

By an event, one or more states of affairs *State of affairs*<sub>1</sub> stop obtaining and other states of affairs *State of affairs*<sub>2</sub> start to obtain (Figure 3). For instance, if the event that it starts to rain occurs, the state of affairs that it is not raining stops obtaining, and the state of affairs that it is raining starts to obtain.

We will use rectangular boxes to denote states of affairs, and rounded boxes to represent events. Arrows indicate the directed connection between states of affairs. If the




FIGURE 3. By an event, states of affairs stop and start to obtain.

state of affairs that stops to obtain by an event is trivial or irrelevant, it is not shown (cf. Figure 4).

Since the occurrence of an event is itself a state of affairs, there is another way (Figure 5) to depict the event of Figure 3.

To indicate that the occurrence of an event is a special state of affairs related to an event, it is shown as a rectangular box containing a rounded box.

An event can have effects on more than one level. For instance, the event of signing a sales contract trivially results in the state of affairs that the sales contract has been signed. The same event also has the (derived) effect that the signing parties engaged into a contractual bond. Moreover, the contractual bond between the parties involves that the one party has an obligation toward the other party, which in turn involves that the party under the obligation has a duty to perform some action. The relations are depicted in Figure 6. The vertical arrows are examples of constitution.




FIGURE 4. The initial state of affairs is sometimes not shown.




FIGURE 5. The occurrence of an event as a state of affairs.

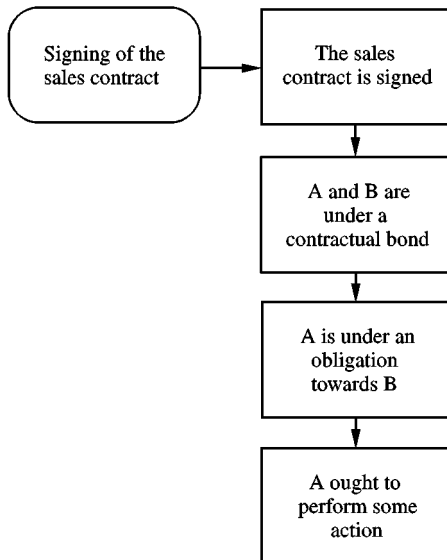


FIGURE 6. An event can have derived effects.

4.2. SUPERVENIENCE OF EVENTS

Events can supervene on other events, just as states of affairs can supervene on other states of affairs. This is illustrated by the example of the signing of a contract that indirectly leads to the existence of a contractual bond (cf. Figure 6). The event of signing of the sales contract implies the event of engaging into a contractual bond. We say that engaging into a contractual bond *supervenes* on the signing of the contract.

Each of the derived effects of the signing of the sales contract in Figure 6 can be regarded as the result of an event that supervenes on the signing of the contract, as shown in Figure 7.

In Figure 7, arrows seem to be used in a new way, namely between supervening events. However, if the alternative way of depicting events (as in Figure 5) is used, it turns out that the supervenience of events can be regarded as a special case of the supervenience of states of affairs; cf. Figure 8.

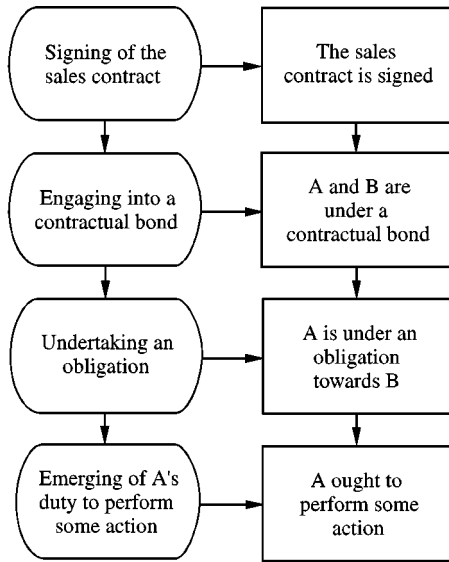


FIGURE 7. An event can supervene on another event.




FIGURE 8. Two ways of depicting the supervenience of events.



## 5. Rules

A directed connection between states of affairs is called a *rule*. It is, for instance, a rule that if a contract is signed, a contractual bond between the contracting parties has come into existence. The formulation of a rule should be distinguished from the state of affairs that this rule exists. It is possible to formulate all kinds of rules, but obviously not all of these possible rules exist. The existence of a rule is a particular state of affairs, which may obtain or not. Connections between states of affairs can only be based on rules which actually exist.

The reader should be aware of other philosophical and legal connotations of the term “rule” that might be confusing. Rules in our sense include many divergent phenomena, such as physical laws, rules of evidence, power conferring rules and legal norms.† For instance, Newton’s law of gravitation is in our terminology a rule, because it connects the states of affairs that two bodies have masses  $m_1$  and  $m_2$ , and the state of affairs that these bodies attract each other with a force equal to  $Gm_1m_2/r^2$  (where  $G$  is the gravitational constant and  $r$  is the distance between the gravitational centres of the bodies).

It might be a rule of evidence that if three independent witnesses saw someone commit the crime, this person counts as having committed the crime. This hypothetical rule connects the states of affairs that Peter, Paul and Mary saw Snoop kill Ice and that Snoop counts as having killed Ice.

It is a power conferring rule that if the legislator attributes some legal body with the competence to perform a particular juristic act, this body can perform that act. This rule connects, for instance, the states of affairs that the legislator gave the community council the power to make by-laws, and that the community council can make by-laws.

A rule consists of a condition part and a conclusion part. The *condition part* consist of one or more generic states of affairs (as expressed by a sentence with variables), while the *conclusion* consists of one single generic state of affairs.‡ In applying the rule, the generic states of affairs are instantiated. For instance, it might be a rule that thieves ought to be punished. The condition part of the rule is the generic state of affairs that someone is a thief; the conclusion part is the generic state of affairs that someone ought to be punished. If the rule is applied to the case of the thief John the condition part of the rule is instantiated to the state of affairs that John is a thief. The conclusion part is correspondingly instantiated to the state of affairs that John ought to be punished.

### 5.1. RULES OF CONSTITUTION AND RULES OF CAUSATION

In Section 2, we discussed two fundamental types of connections between states of affairs, that is constitution and causation. This distinction corresponds to a similar distinction between types of rules.

If one state of affairs constitutes another one, there is a *constitutive rule* underlying the connection.§ An example is the rule that someone is checkmated if the King is threatened and the threat cannot be taken away in one move. The state of affairs that the King is

†Rules in our sense are comparable to constraints in the sense of situation semantics. See Barwise and Perry (1983, p. 94f).

‡The state of affairs of the rule conclusion may be logically compound.

§Notice that our use of the term “constitutive rule”, which is opposed to a causal rule, deviates from Searle’s (1969) use which distinguishes between constitutive and regulative rules.




FIGURE 9. The occurrence of an event as a state of affairs.

threatened and the threat cannot be taken away in one move is the *reason* that someone is checkmated.

A state of affairs can be brought about by an *event*. Rules that govern the relation between an event and the effects that result from it are called *causal rules*. An example is the rule that heating an object (an event) makes that the heated object is warmer than before. The event does not have to be a purely physical event. For instance, signing a sales contract is the (legal) cause for the existence of a contractual bond.

Since the condition part of rules can only contain states of affairs, there is no place for events in the rule conditions. Therefore, causal rules must attach consequences to the *occurrence* of an event, which is a state of affairs, possibly in combination with other states of affairs. For instance, there might be a causal rule that if somebody has the competence to make regulations (a state of affairs in the rule's condition) and exercises this competence (a state of affairs, viz. the occurrence of an event, in the rule's condition), the regulation that was made is valid (a state of affairs in the rule's conclusion). This construction is depicted in Figure 9.

The causal rule connecting the states of affairs that *L* is competent and that *L* makes regulation XYZ to the state of affairs that regulation XYZ is valid is represented as a circle (cf. Figures 2 and 5).

## 5.2. DEFEASIBILITY

Although rules are formulated in the 'If..., then ...'-form, they do not guarantee their conclusion if their conditions are satisfied. A rule that guarantees that its conclusion obtains if its conditions obtain is called strong, otherwise weak. The application of weak rules is defeasible. The usefulness of the notion of a rule is considerably enhanced by this possibility of defeasible rule application.

Two main types of defeasibility of rule application have been distinguished. First, the connection between the conditions and conclusion of a rule may be blocked for some reason. For instance, the connection between condition and conclusion of the rule 'If the weather is good on Sunday, the highways are full' is blocked if there is a driving restriction because of an ozone alert. A legal example would be that application of the rule that thieves ought to be punished is blocked if the thief is a minor. Such reasons blocking the application of a rule are called undercutters (Pollock, 1987; Prakken, 1997,

p. 102), exclusionary reasons (Raz, 1975; Hage, 1997) or just exceptions to rules (Sartor, 1991; Verheij, 1996).

Second, rules can have incompatible conclusions, so that they cannot all lead to their conclusions. For instance, if the conditions of the rule ‘If the weather is good on Sunday, the highways are full’ and ‘If there is an international soccer match, the highways are empty’ obtain, the state of affairs that the highways are full can be undetermined. In the law, this type of defeasibility is related to priorities between legal rules (as for instance in cases of *Lex Superior*) and the weighing of opposing reasons resulting from legal principles.

In Section 7.1, we discuss an example of an exception to a rule from the point of view of our abstract model. Hage (1996, 1997) and Verheij (1996) (among others) discuss the topic of defeasibility more extensively.

### 5.3. RULES AND PRINCIPLES

Dworkin (1978) has argued that the intuitive differences between-reasoning with rules and principles in the law require a logical distinction. As an example of a typical legal rule, he mentions ‘A will is invalid unless signed by three witnesses’, while ‘No man may profit from his own wrong’ would be a typical legal principle.

There appear to be three differences between rules and principles. First, legal rules seem to lead directly to their conclusion, whereas legal principles merely seem to lead to a reason for their conclusion. Second, legal rules and principles seem to behave differently in cases of conflicts: whereas a conflict of legal rules leads to a contradiction, a conflict of legal principles leads to opposing reasons that can subsequently be weighed. Third, legal rules are independent of each other, while legal principles can interact, as in the case of weighing.

In our abstract model of the law, both legal rules and principles are instances of rules: they provide directed connections between states of affairs. The only difference is that in the case of legal rules the connection is apparently stronger than in the case of legal principles. Logically, there are several ways to make the distinctions between legal rules and principles explicit. For instance, Sartor (1994, p. 189) argues that the distinctions disappear in a defeasible context, Verheij, Hage and van den Herik (1998) give an integrated view on legal rules and principles, in which the intuitive differences appear at the extremes of a spectrum, while Hage (1997) treats legal rules and principles as logically distinct.

### 5.4. GOALS

Goals play an important role in the law: criminals are punished with the *goal* to protect society, but the punishment should not be too severe to *prevent* the social isolation by a long period of imprisonment.

Since goals give rise to connections between states of affairs, we discuss them under the general heading of rules. Goals underlie reasons for deontic states of affairs. Their functioning is related to that of principles (Alexy, 1985), in that they generate reasons which plead for or against a particular (deontic) conclusion. Goals are less determinate than principles, however, because they do not explicate which means ought to be chosen

to obtain the goal. For instance, the goal to protect society can underlie reasons why criminals ought to be imprisoned, but also reasons why poverty should be combated.†

We consider goals as *underdetermined* rules: whereas a rule has a condition and a conclusion, a goal can be better conceived as one-half of a rule, that only consists of a condition that can support a plethora of conclusions. Each way of achieving the goal leads to another completion of the rule with a conclusion.

Our use of the term rules is slightly ambiguous. Rules as opposed to principles and goals should be distinguished from rules as a primitive of the abstract model. If both are used in the same context, we speak of rules in the strict sense and rules in the broad sense, respectively.

### 6. Signing a sales contract

In Sections 6–12, we illustrate the uses of the abstract model of the law by analysing some central legal topics.

As a first example of the application of our abstract model, we elaborate the example of signing a sales contract, that was used throughout the discussions above. Figure 10 extends Figure 7.

We have eight states of affairs, four events and three rules. Four of the states of affairs form the initial state, when (1) the sales contract is not signed by A and B; (2) A and B are not under a contractual bond; (3) A is not under an obligation towards B; and (4) it is not the case that A ought to perform some action.

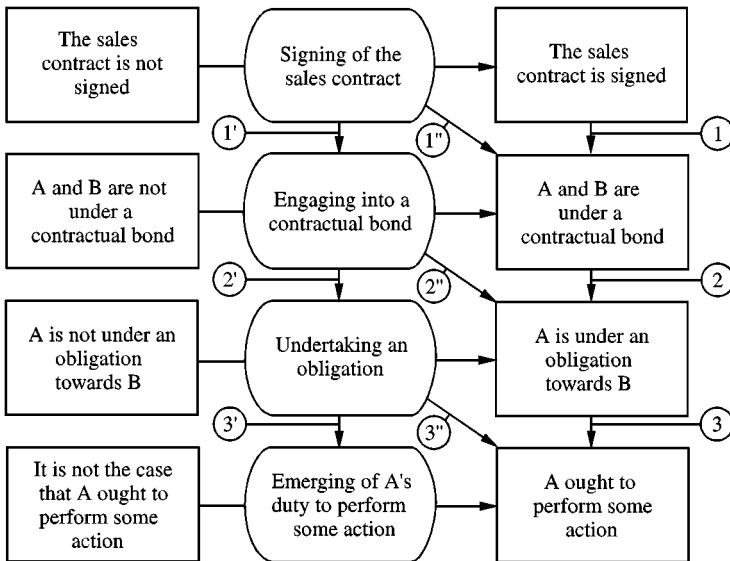


FIGURE 10. Signing a sales contract.

†Hage (1997, pp. 232, 233) discusses the logical behaviour of goals.

In this initial state, four events take place: (1) A and B's signing of the sales contract; (2) A and B's engaging into a contractual bond; (3) A's undertaking of the obligation towards B to pay him the sales price; and (4) the emerging of A's duty to pay B the sales price.

The events lead to the four states of affairs that form the final state: (1) The sales contract is signed by A and B; (2) A and B are under a contractual bond; (3) A is under an obligation towards B; and (4) A ought to perform some action.

The states of affairs in the final state supervene on each other: the state of affairs that A ought to perform some action supervenes on the state of affairs that A is under an obligation towards B, which in turn supervenes on the state of affairs that A and B are under a contractual bond, which supervenes on the state of affairs that the sales contract is signed by A and B.

The connections between these states of affairs result from three rules.

1. A signed sales contract leads to a contractual bond.
2. A contractual bond implies obligations of the contracting parties towards each other.
3. An obligation implies the duty to perform the contents of the obligation.

The events also supervene upon each other, just as the final states of affairs. The emerging of A's duty to pay B the sales price supervenes on A undertaking the obligation towards B to pay him the sales price. A's undertaking of this obligation supervenes on A and B's engaging into a contractual bond, which in turn supervenes on the signing of the sales contract.

The connections between these events result from three rules, closely related to the three rules above.

- 1'. Signing a sales contract is a form of engaging into a contractual bond.
- 2'. Engaging into a contractual bond implies the undertaking of obligations of the contracting parties towards each other.
- 3'. Undertaking an obligation implies the emerging of the duty to perform the contents of the obligation.

In the figure, three more rules are marked, that non-trivially connect the events and the final states of affairs.

- 1". Signing a sales contract leads to a contractual bond.
- 2". Engaging into a contractual bond implies obligations of the contracting parties towards each other.
- 3". Undertaking an obligation implies the duty to perform the contents of the obligation.

There are also the trivial connections between the events and the states of affairs that start to obtain by them, e.g. the event of signing the contract that leads to the state of affairs that the contract has been signed. Notice that the non-trivial effect of an event (as results from the rules 1", 2" and 3") is the trivial effect of its supervening event.

The rules in a triplet such as 1'/1" are closely related, and are in practice not distinguished.

### 7. Classification

An important topic in law is classification. To make a legal rule applicable, a factual situation must be classified, so that it falls under the rule’s conditions. It is important to note that in the law, classification is not just determining whether something falls under the meaning of a word, but also assignment of a particular status. The possible outcomes of classification encompass diverse states of affairs. Something or somebody may be classified as, for instance, a vehicle, tortuous, force majeure, the cause of particular damages, mens rea, competent to issue licenses and liable to be punished.

In our abstract model, classification is treated as a special case of constitution. In other words, classification is based on constitutive rules. As examples of classification, we discuss subsumption and imputation.

#### 7.1. SUBSUMPTION

One type of classification is subsumption of a concrete object under an abstract category. The determination of whether some object classifies as a vehicle is an example that has become traditional.

Assume that there is a rule that the use of vehicles in the park is prohibited, and also a rule that defined vehicles as objects on wheels which are meant for transportation. Can roller-skates be classified as vehicles in the sense of the first rule?†

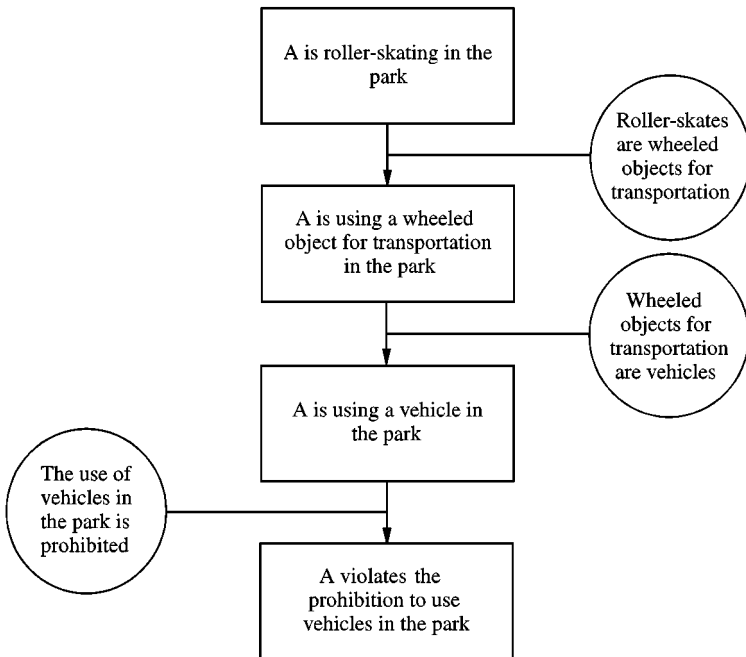


FIGURE 11. Classification as subsumption.

†In this example, we assume that there is no special rule that determines whether roller-skates are vehicles.

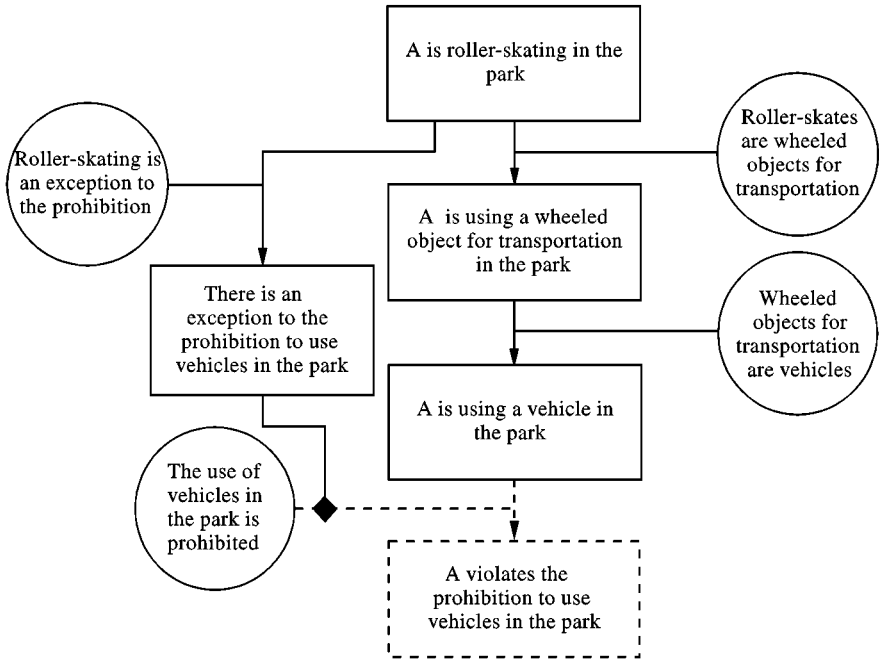


FIGURE 12. An exception to a rule.

Since roller-skates are objects on wheels, meant for transportation, and therefore vehicles, someone roller-skating in the park is violating the prohibition to use vehicles in the park. In our abstract model, we get Figure 11.

Let us now assume that there is a rule that roller-skating is an exception to the prohibition to use of vehicles in the park is forbidden. The resulting exception blocks the connection to the state of affairs that A violates the prohibition, as in Figure 12. (Note that the blocking of the connection is depicted as a vertical line ending in a diamond.)

The existence of the exception is just another state of affairs that supervenes on the state of affairs that A is roller-skating in the park.

7.2. IMPUTATION

As a second example of classification, we discuss the classification of a tort as the cause of damages. In the Netherlands, a tort is classified as the cause of damages if the tort was a necessary condition (*conditio sine qua non*) for the damages and the damages can reasonably be imputed to the tort. In our model, imputation is depicted as in Figure 13.

8. Rights

We discuss three kinds of rights in our abstract model: claims against some concrete person (*iura in personam*), property rights (*iura in re*) and human rights. It turns out that the three kinds of rights are *states*, i.e. momentary states of affairs (cf. Section 3.1).

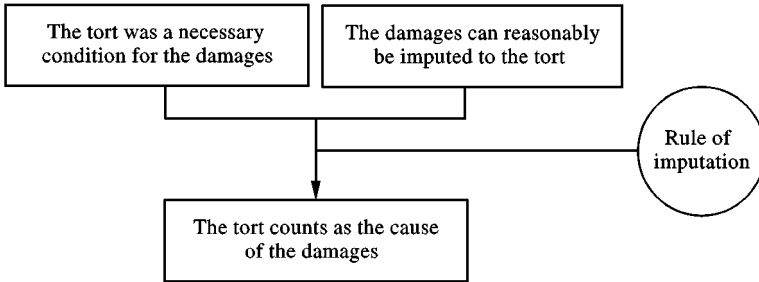


FIGURE 13. Classification as imputation.

### 8.1. CLAIMS

In his paper *Tû-tû*, Ross (1957) writes the following:†

“We find the following phrases, for example, in legal language as used in statutes and the administration of justice:

1. *If a loan is granted, there comes into being a claim;*
  2. *If a claim exists, then payment shall be made on the day it falls due;*
- which is only a roundabout way of saying:
3. *If a loan is granted, then payment shall be made on the day it falls due.*

That “claim” mentioned in (1) and (2), but not in (3), is obviously [ ... *omission added, JH & BV*] not a real thing; is nothing at all, merely a word, an empty word devoid of all semantic reference.”

Here Ross provides an account of phenomena like claims as mere intermediaries between facts: the intermediary is only a manner of speaking, and does not really exist. While rejecting this reductionist consequence, MacCormick and Weinberger (1986) adopt the idea that certain legal states of affairs function as an intermediary between other (legal) states of affairs. They describe a particular category of legal concepts, called *institutional legal facts*, in our terminology related to states of affairs supervening on other states of affairs (MacCormick & Weinberger, 1986, p. 52/3). Institutional legal facts have certain features in common.

For each of them, the law contains rules that lay down when, e.g. a contract, a corporation or an obligation of reparation, comes into existence. These rules are called *institutive rules*. The law also contains rules that attach further legal consequences in case these concepts apply (if the concerning institutional legal facts obtain). These rules are called *consequential rules*, and finally, the law has rules that determine when the phenomena at stake disappear again. These rules are called *terminative rules*; cf. Figure 14.

The figure agrees with our abstract model. Institutional legal facts are then states, the coming into existence and disappearing of which is regulated by causal rules.

†Quotation after Lloyd (1979, p. 625).






FIGURE 14. Institutional legal facts.

Constitutive rules deal with the states of affairs which are constituted by states. As Ross’ discussion shows, claims fit nicely in this picture.

8.2. PROPERTY RIGHTS

The next example is having a property right, such as the ownership of a house. If A owns the house H, it holds that, with the exclusion of everybody else, A is entitled to use, say inhabit, the house. Moreover, A has the power to transfer the ownership. The law may also attach other legal consequences to the ownership of a house. In the Netherlands and in Belgium, owners of houses are, for instance, subject to special taxes. These consequences of ownership are attached by special legal rules to the state of ownership. The rules might have been different, which goes to show that the legal consequences of ownership are not part of the ownership itself, but rather states of affairs which are non-causally connected to ownership.†

The ownership of a house can be acquired in different ways. The most common one is that somebody else was the owner, and transferred his ownership to the new owner. Such a transfer is an event which has the direct effects that the original owner loses his property right, and that the new owner acquires it. The transfer has also indirect effects, because all legal consequences which are attached to ownership disappear for the original owner and come into existence for the new owner.

Another way to acquire the ownership of a house is to build the house on ground which one owns. This event only causes a new ownership to come into existence, not the disappearance of a previous ownership. The passing away of the original owner is a way for an inheritor to acquire ownership. All these different ways of becoming the owner of a house indirectly lead to the legal consequences attached to ownership.

There are also several ways to lose ownership. Transfer is again the most prominent one, but passing away of the owner, devastation of the property, prescription and expropriation are other ways to lose ownership.

As this example about the ownership of a house illustrates, property rights can be treated as ‘empty’ states, the coming into existence, the (legal) consequences and the disappearance of which is governed by rules; cf. Figure 15.

The similarity of Figures 14 and 15 is obvious.

†It may be argued that some consequences of ownership are so essential that if they would not exist, the underlying state would not be ownership anymore, but rather some other state. The discussion of this view falls outside the scope of this paper.




FIGURE 15. Acquisition, consequences and loss of ownership.

### 8.3. HUMAN RIGHTS

Human rights, such as the right of freedom of expression, differ in nature from property rights. Nevertheless, having a human right is also a kind of state, and is in that respect very similar to having a property right. We take a closer look at the freedom of expression.

If *P* has the freedom of expression, this has several consequences. The first and foremost consequence is that *P* is in principle permitted to express his opinion about any issue. (Remember the defeasibility of rule application.) If we follow Dworkin (1978, p. 184f.), having a human right also involves that regulations that infringe these rights are invalid. In other words, for regulations that infringe these rights, the rule that regulations which were validly made contain valid law is not applicable (cf. Hage, 1997, p. 173).

Legal systems usually attribute human rights to all persons on the basis of their being humans. This means that (instances of) human rights come into existence as soon as a human being comes into existence, and end when human beings pass away.

The important thing to note about rights is that, in spite of the different nature of claims, property rights and human rights, the same scheme applies: there are events by which these rights come into existence, and other events by which they disappear again; rules of law determine the legal consequences of the rights. In other words, rights are legal states on which legal consequences supervene (in the sense of the sections 3.1 and 3.2).

## 9. Proof

In Section 3 on states of affairs, we included as examples the states of affairs that, from the point of view of civil law, O.J. Simpson killed his wife, and that, from the point of view of criminal law, O.J. Simpson did not kill his wife. The examples show that the states of affairs in different points of view can be in conflict.

The reason why this seeming inconsistency can obtain is that for many legal purposes it is not the truth that counts, but rather what is proven. The sentences that O.J. killed his wife and that he did not kill his wife cannot both be true, but it can both be true that according to the standards of criminal law, O.J. counts as not having killed his wife (presumption of innocence plus—according to the standards of criminal law—insufficient proof), while according to the standards of criminal law, the counts as having killed his wife (no presumption of innocence plus—according to the standards of criminal law—sufficient proof).

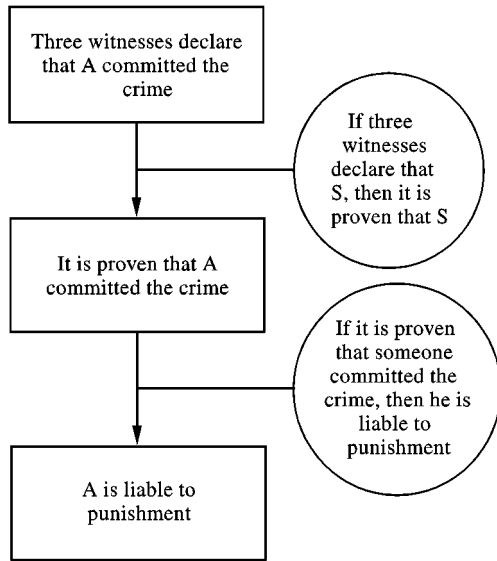


FIGURE 16. Proof.

The state of affairs that something is proven (which is a state of affairs *about* a state of affairs) supervenes on states of affairs that form the proof. The connection between these states of affairs is determined by a rule of proof.† Compare the example in Figure 16.

### 10. Juristic acts

Juristic acts are acts to which the law assigns consequences because of the intention to invoke these consequences by means of the act. For instance, engaging into a contract is a juristic act, to which the law assigns the consequence that a contract exists.

A juristic act supervenes on another act which legally counts as a juristic act. To count as a juristic act, the underlying act must satisfy a number of conditions, such as the condition that the actor is competent to perform the juristic act in question. For instance, to be able to engage into a contract, both parties must have the competence to do so. To make legislation, the actor must have the competence to legislate.

Being competent is a kind of anankastic state of affairs (cf. Section 3.3), which must supervene on another states of affairs. For instance, one must be of age to be competent to engage into a contract.

Figure 17 (from which the rules are left out) depicts a typical juristic act with its preconditions and its consequences. Notice that this figure contains two actions (represented in the dual way of Figure 9), namely signing the sales contract and engaging into

†Rules of proof are constitutive rules for facts of the kind that something counts as proven. They do *not* consume the proven facts. For example, a rule that defines when it is proven that somebody committed a crime constitutes the fact that some crime is *proven*, but it does not constitute the crime itself. Hage (1997, pp. 73, 74) gives more details on the relation between rules of proof (there called “epistemic rules”) and constitutive rules.

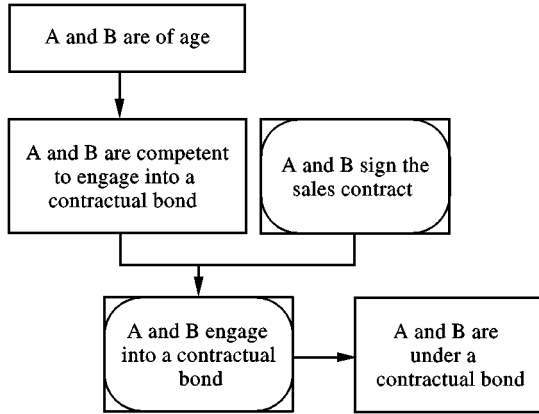


FIGURE 17. A juristic act and its consequences.

a contractual bond. The former counts as a juristic act, because the actor was competent to perform that juristic act. In other words, the juristic act supervenes on its underlying brute action.

Notice moreover that the competence to engage into contracts is itself a state of affairs that supervenes on another state of affairs, namely being of age.

## 11. Validity

In the law, the notion of validity is used for acts, for products and for rules.

If an act satisfies all the conditions which hold for a juristic act, the act is valid as a juristic act. Juristic acts can aim at the creation of a particular product, such as a contract, a license or legislation. If the juristic act is valid, its product is also said to be valid: contracts, licenses and legislation are said to be valid if the acts from which they result are valid as juristic acts.

In the case of legislation, there is still another form of validity. The rules which are created through valid legislation are said to be valid too. This validity is nothing else than the rule's mode of existence (cf. Kelsen, 1979, p. 136). So, in the case of rules based on legislation, we can distinguish three kinds of validity, which supervene upon each other.

- Validity of the legislative act as a juristic act.
- Validity of the legislative product (e.g. the statute).
- Validity of the rules created by means of the legislative product.

Figure 18 gives an example containing the three kinds of validity.

The actions Parliament performs of making a statute lead to the valid making of a statute since Parliament is competent to make statutes. The resulting valid statute leads to the validity of some rule, say about sale contracts. Note that the rule and its validity (i.e. the state of affairs that the rule is valid) are shown in the figure in a dual way similar to the way in which an event and its occurrence are shown (cf. Figure 9). The validity of the rule gives rise to a connection between states of affairs by constitution.

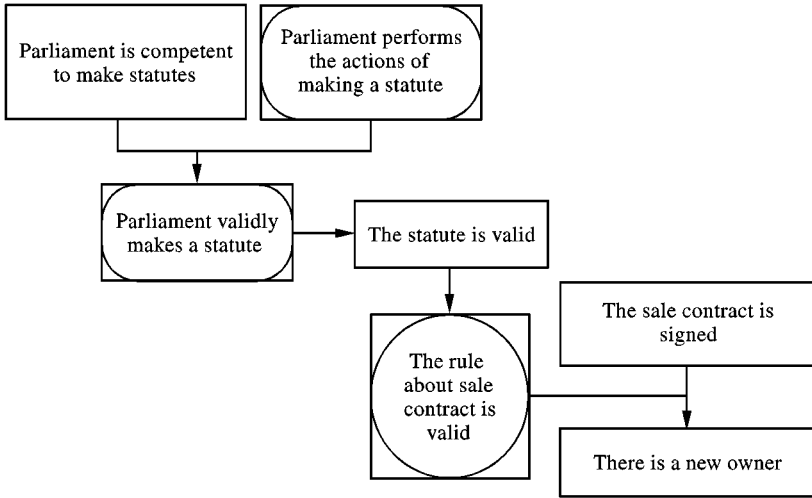


FIGURE 18. The validity of acts, products and rules.

## 12. Juristic facts

Traditionally, continental jurisprudence distinguishes the notions of “juristic fact”, “act”, “bare juristic fact”, ‘juristic act’, and ‘factual act’, which seem to be closely connected to the primitives of our abstract model.

*Juristic facts* are facts to which the law attaches consequences. Examples of juristic facts are sale, theft, death and lapse of time. Possible legal consequences of these examples include the coming about of the vendor’s right to be paid, the liability of the thief to be punished, inheritance and the preclusion of criminal proceedings, respectively. Juristic facts are divided into *acts* (that in the law cannot only be performed by humans, but also, more generally, by juristic persons), such as sale and theft, and *bare juristic facts*, such as death and the passing of time.

Acts are divided in *juristic acts* and *factual acts*. Juristic acts require an intention aimed at legal consequences as manifested by a declaration. Examples of juristic acts are buying a house and recognizing a child. Factual acts are those acts that have legal consequences, but are not meant as such. Examples of factual acts are torts and undue payment.

The traditional categories and their relations are summarized in the scheme in Figure 19.

How do these traditional categories compare to our primitives? The first thing to notice is that the notion of a state of affairs is preferable as a primitive to the notion of a fact. The choice for states of affairs has the advantage that it becomes possible to distinguish between obtaining and non-obtaining states of affairs. This is useful if one wants to deal with connections between hypothetical states of affairs, as in ‘If the state of affairs that John has stolen obtains, then the state of affairs that John is punishable obtains’.

Second, it should be noticed that we distinguish acts (as a kind of events) from the occurrence of acts (as states of affairs). This has the advantages that the changes in the




FIGURE 19. Traditional categories of juristic facts and their relations.

obtaining states of affairs which are caused by acts are appreciated, and that the difference between causation and constitution can be made explicit. In the traditional model sketched above, acts are treated as a sub-category of facts, which seems to be a category mistake. It is therefore better to regard the acts in the traditional model as the corresponding facts that an act occurred.

Just like the traditional view, our abstract model treats juristic acts as a kind of acts.

It is interesting that (intended or unintended) legal consequences of juristic facts are central in the traditional categories. In our abstract model, these correspond to the consequences supervening on a state of affairs by legal rules.

From this brief comparison, it will be clear that our abstract model is richer than the traditional model and is internally more consistent, while remaining on a similar level of abstraction.

### 13. The elements of the abstract model reconsidered

After the discussion of examples in the Sections 6–12, it is time to step back and reconsider the elements of our model. The starting point of the reconsideration is the distinction, familiar from predicate logic, between sentences and terms. This distinction on the linguistic level has an ontological counterpart in the distinction between states of affairs and individuals. This latter distinction (not made explicit until now) plays a central, somewhat complicated role in our abstract model.

The first primitive of our abstract model are states of affairs. By definition, states of affairs play the role of the ontological counterpart of a sentence. As we have seen, states of affairs can be about individuals, as in the state of affairs that John is a thief, and in particular *about states of affairs*, as in the state of affairs that the fact that John is a thief is regrettable. In this sense, states of affair can have a nested structure. As a result, states of affairs are also individuals, and play the role of the ontological counterpart of certain terms, viz. those that denote states of affairs.

In Section 3, we have met other types of internal structure of states of affairs: tense, modality, points of view. Yet other types of internal structure are whether the sentence expressing the state of affairs is a feature-placing sentence (Strawson, 1959, p. 202), has a subject-predicate structure or expresses a relationship. Again another aspect of internal structure is whether the expressing sentence is adverbially qualified. In an action-sentence, the internal structure deals with the identification of the actor and the type of action. All of these additional aspects of internal structure may be relevant for the function of states of affairs in legal arguments.

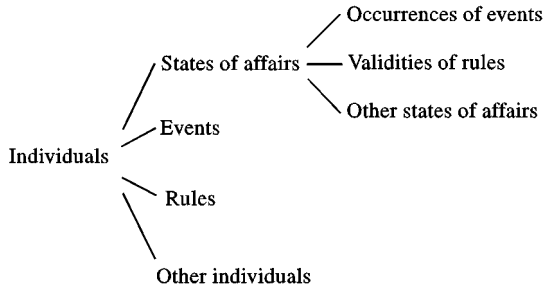


FIGURE 20. Tree of individuals.

The second primitive of our abstract model are events. Events cause changes in the set of obtaining states of affairs. There can be states of affairs about events, one of which plays a special role, namely the state of affairs that an event occurs. As a result, events can bring about states of affair about events, e.g. prohibiting that your opponent check mates you. So also events can have a nested structure.

The third primitive of our abstract model are rules. The conditions and the conclusion of a rule consist of states of affairs. We treat rules as individuals, but these individuals are about states of affairs. Moreover, since rules are treated as individuals, there can be states of affairs about rules. One kind of state of affairs about rules plays a special role in our abstract model, namely the state of affairs that a rule is valid. Since the states of affairs in a rule can be about rules, as in a rule about the validity of rules, also rules can have a nested structure.

The scheme in Figure 20 shows the primitives of our abstract model in a tree of individuals.

#### 14. Heuristic guidelines for legal knowledge representation

As an application of the abstract model of the law, we discuss heuristic guidelines for legal knowledge representation as suggested by the model.

A representation of a legal domain based on our abstract model needs the three primitive elements, that is states of affairs, events and rules. In principle, events and rules can even be represented by corresponding states of affairs of the occurrence of events and the validity of a rule. However, it is wise to distinguish the three primitives, because of the different functions of the three primitives in the abstract model.

The following heuristics for the representation of a legal domain are suggested by our abstract model.

1. *Identify (preliminarily) the types of states of affairs, events and rules occurring in the domain.* These form the skeleton of the representation.
2. *Determine for each state of affairs whether it supervenes on another state of affairs.* Check for every supervening state of affairs whether the rule which connects it with its underlying states of affairs was already identified. Avoid circular connections of states of affairs, where one state of affairs in the end supervenes upon itself. All modal states of affairs and states of affairs which deal with exceptions to rules, validity, or proof, must supervene on other states of affairs.

3. *Identify which states of affairs are states.* Determine for every state which events govern its coming about and disappearing. Check whether these events were already identified as belonging to the domain.
4. *Check for every event whether the rules which govern its effects have already been identified.*
5. *Check for all rules in the broad sense whether they are rules in the strict sense, principles, or goals.*†
6. *Check for all rules whether their application is defeasible or not.* Check for all defeasible rules which states of affairs in the domain may block their application. Check application. Check for every potential exception whether the rule that governs its effects was identified.
7. *Check for every state of affairs whether it must be proven.* For states of affairs that require proof, determine which point of view sets the standards for the proof.

These guidelines for the modeling of legal knowledge domains end the exposition of our abstract model of the law. In the following section we will briefly compare our model with related work.

## 15. Related research

We put our abstract model of the law in perspective by a discussion of related work by McCarty (1989), Valente (1995), and van Kralingen (1995) and Visser (1995).

### 15.1. MCCARTY'S LANGUAGE OF LEGAL DISCOURSE

McCarty's (1989) begins the development of a "deep conceptual model" of the law in his paper on the basic features of a language for legal discourse (LLD). This model takes shape as an abstract description of a knowledge representation language. This description deals first with atomic formulae, including reified relationships, sorts and sub-sorts and the distinction between count terms and mass terms. The section on rules and proofs describes the use of Horn clauses and an extension of it by means of negations and embedded implications, the use of default rules and proofs and reasoning by prototypes and deformations. The section on modalities, finally, deals with time, events and actions and with deontic modalities.

Because McCarty's focus is on a knowledge representation language and what should be incorporated in it, while we focus on the abstract model underlying heuristics for legal knowledge representation, a comparison between McCarty's work and ours is a bit hazardous. Nevertheless, some striking similarities can be found. For instance, in LLD, relations are treated as individuals. McCarty mentions as an example that it is possible to talk about ownership as an individual, e.g. by the use of the word ownership 01. This strongly resembles our treatment of states of affairs as logical individuals. Other similarities are the use of default rules (which was less obvious in 1989 than it is nowadays), and the special attention for events and actions. McCarty also mentions part-whole relations

†Compare the terminology of the end of Section 5.



between elementary and complex events, which are not incorporated in our abstract model.

### 15.2. VALENTE'S FUNCTIONAL ONTOLOGY OF LAW

Valente (1995) has developed a functional ontology of law [also partly described by den Haan (1996)]. This ontology is based on a functional perspective on the legal system, in which it is assumed that the main function of the legal system is to react to social behaviour (Valente, 1995, p. 49).

Valente extends this functional perspective from the legal system as a whole to the elements of the legal system, which he discusses as categories of legal knowledge. He distinguishes six primitive categories of legal knowledge, that is normative knowledge, world knowledge, including classificatory and causal knowledge, responsibility knowledge, reactive knowledge, meta-legal knowledge and creative knowledge.

In our abstract model, Valente's primitive categories of legal knowledge (except creative knowledge) correspond to different kinds of legal rules, where the differences between the kinds of rules is based on different kinds of conclusions of the rules. For instance, normative knowledge would consist of rules with deontic conclusions, while responsibility knowledge would consist of rules in which behaviour is imputed to actors. To the extent that the knowledge categories of Valente correspond to kinds of rules in our terminology, Valente's distinctions can be regarded as a refinement of our abstract model.

Valente's category of creative knowledge cannot be regarded as a kind of rules. The legislator uses, according to Valente, creative knowledge, if he creates some entity that did not exist before in the world. An example would be the creation of a department within the government or a company (Valente, 1995, p. 67). What Valente calls the use of creative knowledge corresponds in our view more or less to the performance of a juristic act.

### 15.3. THE FRAME-BASED CONCEPTUAL MODEL OF VAN KRALINGEN AND VISSER

In two dissertations defended on the same day, van Kralingen (1995) and Visser (1995) have developed a frame-based conceptual model of the law. They distinguish three main types of entities, which can be represented in three corresponding types of frames. The entity types are norms, acts and concepts (van Kralingen, 1995, Chapter 3).

Two types of norms are distinguished, namely norms of conduct and norms of competence. These two types both belong to the category of rules in our terminology. Moreover, they identify eight slots in norm frames, four of which stand for elements of the content of the norm, such as its legal modality and the conditions of application, and four of which stand for other characteristics of the norm.

Acts are discussed primarily from the point of view of legislation which deals with acts. The authors identify six characteristics of acts, but the corresponding frames for acts have 14 slots, three of which deal with auxiliary information about the norm in which the act is mentioned.

Just as acts, concepts are primarily dealt with as elements in rules about the concepts, e.g. in legal definitions. Concept frames have seven slots.

Because of their focus on the elements of norms and rules of meaning, the work of van Kralingen and Visser can be seen as a refinement of and an addition to the minimal theory about the internal structure of states of affairs, events and rules as presented in Section 13 above. (Notice that the structure of a norm can be seen as part of the structure of the states of affairs that a norm is valid.)

#### 15.4. REASON-BASED LOGIC

Although reason-based logic (e.g. Hage & Verheij, 1994; Hage, 1996, 1997; Verheij, 1996) is usually presented as a tool for defeasible reasoning with (legal) rules, its original inspiration was a study in ontology (Hage, 1987). The basic intuition behind reason-based logic is that some facts are reasons for the presence of other facts, and that the former facts derive their status as reasons from rules (in the broad sense). This intuition is elaborated in a formal model of rules and reasons. As a consequence, reason-based logic can be regarded as an ontological theory about the ways in which the (legal) world is structured by means of rules. The work on reason-based logic can be read as a formal elaboration of our abstract model of the law to the extent that it deals with rules.

## 16. Summary and conclusion

We have presented an abstract model of the law. The primitives of the model are states of affairs, event and rules. The model of the law can be summarized as the view of the law as a dynamic system of states of affairs, which are connected by events and rules.

To illustrate the uses of our model, we have given examples of legal topics that can fruitfully be analysed in terms of the model. Moreover, we used the model to suggest heuristic guidelines for legal knowledge representation.

The high level of abstraction of the abstract model allows many additions in which details of the model are refined. In the discussion about related research, we have indicated how the work of McCarty, Valente, Van Kralingen and Visser, and earlier work of the present authors provides such refining additions. Our abstract model can be thought of as a top ontology of the law.

The authors gladly acknowledge the financial support by the Dutch National Programme Information Technology and Law (ITeR) for the research reported in this paper (project number 01437112).

## References

- ALEXY, R. (1985). *Theorie der Grundrechten*. Nomos Verlagsgesellschaft, Baden-Baden.
- BARWISE, J. & PERRY, J. (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.
- BENCH-CAPON, T. J. M. (1989). Deep models, normative reasoning and legal expert systems. *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, pp. 37–45. New York: ACM.
- BENCH-CAPON, T. J. M. & VISSER, P. R. S. (1996). Deep models, ontologies and legal knowledge based systems. In R. W. VAN KRALINGEN, H. J. VAN DEN HERIK, J. E. J. PRINS, M. SERGOT, & J. ZELEZNIKOW, Eds. *Legal Knowledge Based Systems. JURIX '96. Foundations of Legal Knowledge Systems*, pp. 3–13. Tilburg University Press.

- BENCH-CAPON, T. J. M. & VISSER, P. R. S. (1997). Ontologies in legal information systems; the need for explicit specifications of domain conceptualisations. *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pp. 132–141. New York, NY: ACM.
- DWORKIN, R. (1978). *Taking Rights Seriously* (2nd edn.) London: Duckworth.
- GARDNER, A. VON DER LIETH (1987). *An Artificial Intelligence Approach to Legal Reasoning*. Cambridge, (MA): MIT Press.
- GRUBER, T. R. (1995). Toward principles for the design of ontologies used for knowledge sharing. *International Journal of Human-Computer Studies*, **43**, 907–928.
- GUARINO, N. (1995). Formal ontology, conceptual analysis and knowledge representation. *International Journal of Human-Computer Studies*, **43**, pp. 625–640.
- GUARINO, N. (1997). Understanding, building and using ontologies. *International Journal of Human-Computer Studies*, **45**, 293–310.
- GUARINO, N. & POLL, R. (1995). Editorial: the role of formal ontology in the information technology. *International Journal of Human-Computer Studies*, **43**, 623–624.
- HAAN, N. DEN (1996). *Automated legal reasoning*. Dissertation. Amsterdam.
- HAGE, J. C. (1987). *Feiten en betekenis; een verhandeling over ontologie en praktische rede* (Facts and meaning. A treatise on ontology and practical reason). Dissertation, Leiden.
- HAGE, J. C. (1996). A model of legal reasoning and a logic to match. *Artificial Intelligence and Law*, **4**, 199–273. Also in: Prakken, H. & Sartor, G., eds. *Logical Models of Legal Argumentation*. Dordrecht: Kluwer Academic Publishers. pp. 43–117.
- HAGE, J. C. (1997). *Reasoning with Rules*. Dordrecht: Kluwer Academic Publishers.
- HAGE, J. C. & VERHEIJ, B. (1994). Reason-based logic: a logic for reasoning with rules and reasons. *Law, Computers and Artificial Intelligence*, **3**, 171–209.
- VAN HEIJST, G. SCHREIBER, A. TH. & WIELINGA, B. J. (1997a). Using explicit ontologies in KBS development. *International Journal of Human-Computer Studies*, **45**, 183–292.
- VAN HEIJST, G. SCHREIBER, A. TH. & WIELINGA, B. J. (1997b). Rules are not classes: a reply to Nicola Guarino. *International Journal of Human-Computer Studies*, **45**, 311–318.
- HOBBS, J. R. (1995). Sketch of an ontology underlying the way we talk about the world. *International Journal of Human-Computer Studies*, **43**, 819–830.
- JONES, O. R. (1995). Supervenience. Entry in: Honderich, T. ed. (1995). *The Oxford Companion to Philosophy*, p. 860. Oxford: Oxford University Press.
- KELSEN, H. (1979). In K. Ringhofer & R. Walter, Hrsg. *Allgemeine Theorie der Normen*. Wien: Manz-Verlag.
- VAN KRALINGEN, R. W. (1995). *Frame-based conceptual models of statute law*. Dissertation. Leiden.
- LLOYD (1979). *Introduction to Jurisprudence* (4th edn with M. D. A. Freeman as co-editor) London: Stevens and Sons.
- MACCORMICK, D. N. & WEINBERGER, O. (1986). *An Institutional Theory of Law*. Dordrecht: Reidel.
- MCCARTY, L. T. (1989). A language for legal discourse I. Basic features. *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pp. 180–188. New York: ACM.
- POLLOCK, J. L. (1987). Defeasible reasoning. *Cognitive Science*, **11**, 481–518.
- PRAKKEN, H. (1997). *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law*. Dordrecht: Kluwer Academic Publishers.
- RAZ, J. (1975). *Practical Reason and Norms*. London: Hutchinson.
- ROSS, A. (1957). Tū-Tū. *70 Harvard Law Review*, 812.
- SARTOR, G. (1991). The structure of norm conditions and nonmonotonic reasoning in law. *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pp. 155–164. New York, NY: ACM.
- SARTOR, G. (1994). A formal model of legal argumentation. *Ratio Juris*, **7**, 177–211.
- SEARLE, J. R. (1969). *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- SOWA, J. F. (1995). Top-level ontological categories. *International Journal of Human-Computer Studies*, **43**, 669–685.
- STRAWSON, P. F. (1959). *Individuals. An Essay in Descriptive Metaphysics*. London: Methuen.

- VALENTE, A. (1995). *Legal knowledge engineering. A modelling approach*. Dissertation. Amsterdam: IOS Press.
- VERHEIJ, B. (1996). *Rules, reasons, arguments. formal studies of argumentation and defeat*. Dissertation, Maastricht. A summary and table of contents are available on the World-Wide Web at <http://www.metajur.unimaas.nl/~bart/proefschrift/>.
- VERHEIJ, B. & HAGE, J. C. (1997). States of affairs, events, and rules: an abstract model of the law. In A. OSKAMP, R. V. DE MULDER, C. VAN NOORTWIJK, C. A. F. M. GRÜTTERS, K. ASHLEY & T. GORDON, Eds. *Legal Knowledge Based Systems. JURIX: The 10th Conference* pp. 3–20. Nijmegen: The Foundation for Legal Knowledge Systems.
- VERHEIJ, B., HAGE, J. C. & VAN DEN HERIK, H. J. (1998). An integrated view on rules and principles. *Artificial Intelligence and Law*, **6**, 3–26.
- VISSER, P. R. S. (1995). *Knowledge specification for multiple legal tasks*. Dissertation, Leiden.
- VISSER, P. R. S. (1997). Towards distributed legal information systems: a discussion of ontology mismatches. In A. OSKAMP, R. V. DE MULDER, C. VAN NOORTWIJK, C. A. F. M. GRÜTTERS, K. ASHLEY & T. GORDON, Eds. *Legal Knowledge Based Systems. JURIX: The 10th Conference* pp. 43–54. Nijmegen: The Foundation for Legal Knowledge Systems.
- VISSER, P. R. S. & BENCH-CAPON, T. J. M. (1998). A comparison of four ontologies for the design of legal knowledge systems. *Artificial Intelligence and Law*, **6**, 27–57.
- VISSER, P. R. S. BENCH-CAPON, T. J. M. & VAN DEN HERIK, H. J. (1997). A method for conceptualising legal domains. *Artificial Intelligence and Law*, **5**, 207–242.
- VISSER, P. R. S. & WINKELS, R. G. F., Eds. (1997). *Proceedings of the 1st International Workshop on Legal Ontologies. LEGONT '97*, Melbourne.
- VON WRIGHT, G. H. (1963). *Norm and Action. A Logical Enquiry*. London: Routledge & Kegan Paul.

## Appendix A

After a summary of the ontological commitments of the abstract model, we make our abstract model as explicit as possible by providing a formalism for it. We also show how the figures as used in the paper can be translated to the formalism.

### A.1. ONTOLOGICAL COMMITMENTS

The abstract model has three primitives: states of affairs, events and rules.

A state of affairs is a possible part of the world as expressed by a sentence.† Which states of affairs exist depends on a language. Obtaining state of affairs (or facts) are those states of affairs that belong to the actual world.

An event is characterized by a change of the obtaining states of affairs. Each event has a finite number of initial and final states of affairs. The initial states of affairs of an event obtain before it, the final states of affairs after it. (The initial and final states of affairs do not form the whole world, but only the part of the world that is affected by the event.)

A rule is characterized as a directed connection between states of affairs. Each rule has a finite number of states of affairs as conditions, and one state of affairs as conclusion. If the conditions of the rule obtain, its conclusion obtains, or tends to obtain. (In the following, we do not go into defeasibility issues. The reader is referred to Hage (1997) and Verheij (1996).)

†We assume that states of affairs have an implicit time stamp. This has the disadvantage that states of affairs do not exactly correspond to sentences, and the advantage that the world is conceived as a fixed set of states of affairs that stretches over time, and not as a sequence of sets of states of affairs that change over time.

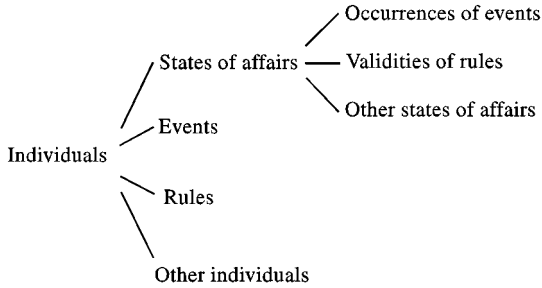


FIGURE A1. Tree of individuals.

Sentences can be about individuals. States of affairs, events and rules are all individuals. As a result, sentences can be about states of affairs, events and rules. There are two types of special states of affairs, one about events and one about rules: the states of affairs that an event occurs and that a rule exists. (Compare the tree of individuals in Figure A1, taken from Section 13.)

A.2. A FORMALISM FOR THE ABSTRACT MODEL

In the following, we provide a formalism for the abstract model. Each formal stipulation is followed by a brief comment.

There are disjoint sets `TERMS` and `LANGUAGE`, the elements of which are *terms* and *sentences*, respectively. The set `LANGUAGE` is the (*representation*) *language*. There is a mapping `TERMS_OF_SENTENCE: LANGUAGE → φ(TERMS)`, where  $\phi(\text{TERMS})$  is the set of all finite subsets of `TERMS`.

Terms denote individuals, sentences express states of affairs. The language can represent a world or a part of the world. The mapping `TERMS_OF_A_SENTENCE` maps each sentence to the set of terms denoting the individuals the sentence is about.

The formalism abstracts from the language as much as possible in order to stress the essentials. A typical example of a language is a language for first-order predicate logic. As an illustration of the formalism, we use a first-order language that has strings beginning with upper-case characters as predicate symbols and strings beginning with lower-case characters as function symbols. In our first-order language, `mary` and `father_of(mary)` are examples of terms, and `It_is_raining` and `Is_father_of(john, mary)` examples of sentences.

In the following, the minimal requirements on the terms and language, as needed for our abstract model, are discussed.

There are disjoint sets `STATES_OF_AFFAIRS ⊆ TERMS`, `EVENTS ⊆ TERMS` and `RULES ⊆ TERMS`.

The elements of `STATES_OF_AFFAIRS`, `EVENTS` and `RULES` are terms denoting states of affairs, events and rules, respectively.

There is a bijective mapping `SENTENCE_TO_TERM: LANGUAGE → STATES_OF_AFFAIRS`.

The mapping `SENTENCE_TO_TERM` maps each sentence to the unique term denoting the states of affairs expressed by the sentence. It should be noted that

SENTENCE\_TO\_TERM (*Sentence*) does not denote the reification of the sentence *Sentence*, but of the state of affairs expressed by *Sentence*. Compare the difference between ‘The sentence “John is a thief” contains 15 characters’ and ‘It is regrettable that John is a thief’.

In our first-order language, a mapping SENTENCE\_TO\_TERM can (for atomic sentences) simply be given by changing the first upper-case character of the sentence to lower case. Then, SENTENCE\_TO\_TERM (*It-is-raining*) is *it-is-raining*, and SENTENCE\_TO\_TERM (*Is-father-of (john, mary)*) is *is-father-of (john, mary)*.

There are mappings INITIAL\_STATES\_OF\_AFFAIRS: EVENTS  $\rightarrow$

$\wp(\text{STATES\_OF\_AFFAIRS})$  and

FINAL\_STATES\_OF\_AFFAIRS: EVENTS  $\rightarrow \wp(\text{STATES\_OF\_AFFAIRS})$ , where

$\wp(\text{STATES\_OF\_AFFAIRS})$  is the set of all finite subsets of STATES\_OF\_AFFAIRS.

INITIAL\_STATES\_OF\_AFFAIRS (*event*) and FINAL\_STATES\_OF\_AFFAIRS (*event*) are the sets of terms denoting the initial and the final states of affairs of the event denoted by the term *event* in *Events*, respectively.

In our first-order language, events can be denoted as structured terms of the form *event(initial\_states\_of\_affairs, final\_states\_of\_affairs)*, where *initial\_states\_of\_affairs* and *final\_states\_of\_affairs* have the form (*State\_of\_affairs*<sub>1</sub>, ..., *state\_of\_affairs*<sub>*n*</sub>) for terms *state\_of\_affairs*<sub>1</sub>, ..., *state\_of\_affairs*<sub>*n*</sub> in STATES\_OF\_AFFAIRS.† The mappings INITIAL\_STATES\_OF\_AFFAIRS and FINAL\_STATES\_OF\_AFFAIRS could then be defined by means of the sets of terms that occur in *initial\_states\_of\_affairs* and *final\_states\_of\_affairs*, respectively. For instance, the set INITIAL\_STATES\_OF\_AFFAIRS (*event({it-is-not-raining}, {it-is-raining})*) then consists of the term *it-is-not-raining*, and FINAL\_STATES\_OF\_AFFAIRS (*event({it-is-not-raining}, {it-is-raining})*) of *it-is-raining*.

There are mappings CONDITIONS: RULES  $\rightarrow \wp(\text{STATES\_OF\_AFFAIRS})$  and

CONCLUSION: RULES  $\rightarrow \text{STATES\_OF\_AFFAIRS}$ .

CONDITIONS (*rule*) is the set of terms denoting the states of affairs that are the conditions of the rule denoted by the term *rule* in RULES. CONCLUSION(*rule*) is the term denoting the state of affairs that is the conclusion of the rule denoted by the term *rule* in RULES.‡

In our first-order language, rules can be denoted as structured terms of the form *rule(conditions, condition)*, where *conditions* has the form {*state\_of\_affairs*<sub>1</sub>, ..., *state\_of\_affairs*<sub>*n*</sub>} for terms *state\_of\_affairs*<sub>1</sub>, ..., *state\_of\_affairs*<sub>*n*</sub> in STATES\_OF\_AFFAIRS, and *condition* is a term in STATES\_OF\_AFFAIRS. The mappings CONDITIONS and CONCLUSION could then be defined by means of *conditions* and *condition*, respectively. For instance, the set CONDITIONS (*rule({john-is-a-thief}, john-is-punishable)*) then consists of the term *john-is-a-thief*, and CONCLUSION (*rule({john-is-a-thief}, john-is-punishable)*) is *john-is-punishable*.

†The first-order language must have special function symbols {., {.,}, {.,...}, ..., etc.

‡For clarity, but unrealistically, the rules in the formalism have states of affairs as conditions and conclusion, and not generic states of affairs. Technically, the use of variables and instantiation can do the trick.

There is a mapping  $\text{OCCURS: EVENTS} \rightarrow \text{LANGUAGE}$ , such that for each element  $event$  of  $\text{EVENTS}$  the set  $\text{TERMS\_OF\_SENTENCE}(\text{OCCURS}(event))$  consists of the term  $event$ .

There is a mapping  $\text{EXISTS: RULES} \rightarrow \text{LANGUAGE}$ , such that for each element  $rule$  of  $\text{RULES}$  the set  $\text{TERMS\_OF\_SENTENCE}(\text{EXISTS}(rule))$  consists of the term  $rule$ .

For terms  $event$  in  $\text{EVENTS}$  and  $rule$  in  $\text{RULES}$ , the sentences  $\text{OCCURS}(event)$  and  $\text{EXISTS}(rule)$  are the special sentences expressing the states of affairs that the event (denoted by the term)  $event$  occurs and that the rule (denoted by the term)  $rule$  exists, respectively.

The notation in the formal stipulation suggests special sentences in our first-order language. Examples are  $\text{OCCURS}(event(\{it\_is\_not\_raining\}, \{it\_is\_raining\}))$  and  $\text{EXISTS}(rule(\{john\_is\_a\_thief\}, john\_is\_punishable))$ .

There is a mapping  $\text{RULES\_OF\_EVENT: EVENTS} \rightarrow \wp(\text{RULES})$ , such that the following hold.

1. If  $rule$  is in  $\text{RULES\_OF\_EVENT}(event)$ , then

$$\begin{aligned} \text{CONDITIONS}(rule) &= \text{INITIAL\_STATES\_OF\_AFFAIRS}(event) \cup \{\text{OCCURS}(event)\} \\ \text{CONCLUSION}(\text{RULES\_OF\_EVENT}(event)) &\in \text{FINAL\_STATES\_OF\_AFFAIRS}(event) \end{aligned}$$

2. For each element  $state\_of\_affairs$  in  $\text{FINAL\_STATES\_OF\_AFFAIRS}(event)$ , there is a (unique) rule in  $\text{RULES\_OF\_EVENT}(event)$ , such that

$$\text{CONCLUSION}(\text{RULES\_OF\_EVENT}(event)) = state\_of\_affairs$$

Here  $\wp(\text{RULES})$  is the set of all finite subsets of  $\text{RULES}$ .

The rules denoted by the terms in  $\text{RULES\_OF\_EVENT}(event)$  are exactly the rules that have the occurrence of the event denoted by  $event$  and its initial states of affairs as conditions and one of the event's final states of affairs as conclusion.

In our first-order language, the set  $\text{RULES\_OF\_EVENT}(event(\{it\_is\_not\_raining\}, \{it\_is\_raining\}))$  consists of the term  $rule(\{it\_is\_not\_raining, occurs(event)\}, it\_is\_raining)$ , where the term  $event$  is equal to  $event(\{it\_is\_not\_raining\}, \{it\_is\_raining\})$ .

We give two definitions of a representation, the first for a language in general, the second for a language with negation. For a language in general, we define:

A *representation* is a subset of Language.

A representation represents the possible world that consists of the states of affairs expressed by the sentences it contains.

For a language with negation (defined below), the definition of a representation is slightly different.

A language Language is a *language with negation* if there is a mapping  $\text{NEGATION: LANGUAGE} \rightarrow \text{LANGUAGE}$ .

For each sentence *Sentence*, the sentence  $\text{NEGATION}(\textit{Sentence})$  is its negation. (For generality, nothing is said about the nature of negation.) For a language with negation, we define:


A subset of Language is *consistent* if there is no sentence in the subset the negation of which is also in the subset. A *representation* is a consistent subset of Language. A representation is *complete* if it is a maximal consistent subset of Language (i.e. any sentence is in the subset, or its negation is).

A representation represents the partial possible world that consists of the states of affairs expressed by the sentences it contains. A complete representation represents a possible world.

A.3. TRANSLATING THE FIGURES TO THE FORMALISM

In the following, we show how the figures, as they are used throughout the paper, can be translated to the formalism. Each figure is interpreted as a partial depiction of some representation  $\text{REPRESENTATION}$  in the formalism.


States of affairs are depicted throughout the paper as rectangular boxes:



In the formalism, the figure corresponds to

$\textit{State\_of\_affairs} \in \text{REPRESENTATION}$

Events are depicted in two ways, as an individual, using a rounded box, and as its corresponding state of affairs:



If an event is shown in a figure, it represents its occurrence. Therefore, both the left and the right-hand sides of the figure correspond to

$\text{OCCURS}(\textit{event}) \in \text{REPRESENTATION}$

Also rules are depicted in two ways, as an individual, using a circle, and as its corresponding state of affairs:




Both the left and the right-hand sides of the figure correspond to

$\text{EXISTS}(\textit{rule}) \in \text{REPRESENTATION}$




The figure below shows an event with one of its initial states of affairs and one of its final states of affairs (cf. Figures 1 and 3).



The figure corresponds to the following in the formalism:

*State\_of\_affairs<sub>1</sub>* ∈ REPRESENTATION  
*State\_of\_affairs<sub>2</sub>* ∈ REPRESENTATION  
 OCCURS ( *event* ) ∈ REPRESENTATION  
*State\_of\_affairs<sub>1</sub>* ∈ INITIAL-STATES-OF-AFFAIRS ( *event* )  
*State\_of\_affairs<sub>2</sub>* ∈ FINAL-STATES-OF-AFFAIRS ( *event* )

Note that not necessarily all initial and final states of affairs of the event are depicted. The next figure depicts the connection between two states of affairs by a rule (cf. Figure 2).




In the formalism, the figure corresponds to the following:

*state\_of\_affairs<sub>1</sub>* ∈ REPRESENTATION  
*state\_of\_affairs<sub>2</sub>* ∈ REPRESENTATION  
 EXISTS ( *rule* ) ∈ REPRESENTATION  
*state\_of\_affairs<sub>1</sub>* ∈ CONDITIONS ( *rule* )  
*state\_of\_affairs<sub>2</sub>* ∈ CONCLUSIONS ( *rule* )

Not necessarily all conditions of the rule are depicted.


There is an alternative way to depict the connection between states of affairs resulting from an event (cf. Figure 5):



In the formalism, the figure corresponds to the following:

*State\_of\_affairs<sub>1</sub>* ∈ REPRESENTATION  
*State\_of\_affairs<sub>2</sub>* ∈ REPRESENTATION  
 OCCURS ( *event* ) ∈ REPRESENTATION  
*State\_of\_affairs<sub>1</sub>* ∈ INITIAL-STATES-OF-AFFAIRS ( *event* )  
*State\_of\_affairs<sub>2</sub>* ∈ FINAL-STATES-OF-AFFAIRS ( *event* )


There is a rule (or there are rules) underlying the connection resulting from an event (cf. Figure 9):



Formally, the figure corresponds to

$State\_of\_affairs_1 \in REPRESENTATION$   
 $State\_of\_affairs_2 \in REPRESENTATION$   
 $OCCURS(event) \in REPRESENTATION$   
 $State\_of\_affairs_1 \in INITIAL\_STATES\_OF\_AFFAIRS(event)$   
 $State\_of\_affairs_2 \in FINAL\_STATES\_OF\_AFFAIRS(event)$   
 $EXISTS(rule) \in REPRESENTATION$   
 $State\_of\_affairs_1 \in CONDITIONS(rule)$   
 $OCCURS(event) \in CONDITIONS(rule)$   
 $State\_of\_affairs_2 \in CONCLUSION(rule)$


The supervenience of events is depicted in two ways (cf. Figure 8):



Formally, both figures correspond to

$OCCURS(event_1) \in REPRESENTATION$   
 $OCCURS(event_2) \in REPRESENTATION$   
 $EXISTS(rule) \in REPRESENTATION$   
 $OCCURS(event_1) \in CONDITIONS(rule)$   
 $OCCURS(event_2) \in CONCLUSIONS(rule)$

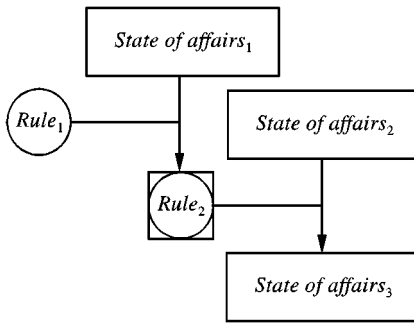
The rule underlying the supervenience of events can have additional conditions (cf. Figure 17):



Formally, the figure corresponds to

$State\_of\_affairs \in REPRESENTATION$   
 $OCCURS (event_1) \in REPRESENTATION$   
 $OCCURS (event_2) \in REPRESENTATION$   
 $EXISTS (rule) \in REPRESENTATION$   
 $State\_of\_affairs_1 \in CONDITIONS (rule)$   
 $OCCURS (event_1) \in CONDITIONS (rule)$   
 $OCCURS (event_2) \in CONCLUSIONS (rule)$

The existence of a rule can be the conclusion of another rule, as in the following figure (cf. Figure 18):



Formally, the figure corresponds to

$State\_of\_affairs_1 \in REPRESENTATION$   
 $State\_of\_affairs_2 \in REPRESENTATION$   
 $State\_of\_affairs_3 \in REPRESENTATION$   
 $EXISTS (rule_1) \in REPRESENTATION$   
 $EXISTS (rule_2) \in REPRESENTATION$   
 $State\_of\_affairs_1 \in CONDITIONS (rule_1)$   
 $EXISTS (rule_2) \in CONCLUSIONS (rule_1)$   
 $State\_of\_affairs_2 \in CONDITIONS (rule_2)$   
 $State\_of\_affairs_3 \in CONCLUSIONS (rule_2)$