A Logic of Classification

with Applications to Linguistic Theory

Vom Fachbereich Informatik
der FernUniversität Hagen
angenommene

Dissertation

zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

von
Dipl.-Math.
Rainer Osswald

Gutachter: Prof. Dr. Hermann Helbig, Prof. Dr. Klaus Weihrauch

Datum der mündlichen Prüfung: 18. September 2002

Hagen, September 2002

Preface

Acknowledgments

In the first place I want to thank Professor Helbig for his continual advice, for stimulating discussions on applications of the framework of this thesis, and for his patience in waiting for its final version. His Chair of Applied Computer Science VII at the University of Hagen has provided me with an exceptionally comfortable working environment. As a member of his research group I have been taking part in the enterprise of developing systems for natural language understanding and knowledge engineering that are based on his MultiNet paradigm. My practical experience in maintaining a lexical knowledge base and implementing a lexicographer's workbench has been both a source of inspiration for my theoretical work and an opportunity to put some of my ideas to practice.

Many thanks to the technical and administrational staff for running everything so smoothly as well as to my colleagues for being always helpful and open to discussion. Especially Sven Hartrumpf has been of invaluable help with every matter ranging from supplying the Scheme compiler with appropriate parameters to discussing the subtleties of default inheritance.

The present form of this thesis has benefited considerably from proof reading and detailed feedback by Sven Hartrumpf, Heike and Tanja Hötte, Ralf Naumann, and Wiebke Petersen, of which of course none is responsible for any remaining errors and fallacies. Special thanks to Mai Gehrke, who called my attention to the work of B. A. Davey and T. P. Speed on the ordered prime spectrum of distributive lattices, and to Eckhard Meyer-Zwiffelhoffer for revealing to me the bibliographic secrets of Abelard's *Editio super Porphyrium*. Finally, I gladly take the opportunity to thank Professor Schnelle and my former colleagues at the Linguistics Department of the University of Bochum for teaching me the basics of general and formal linguistics.

Some Conventions

The following common and less common conventions will be used throughout this thesis: ‘iff’ means ‘if and only if’, ‘$\simeq$’ means ‘is isomorphic to’. A *surjective* function is *onto*, an *injective* function is *into*.

For expository purposes, I occasionally allow *symbol overloading*. The same name is used, for instance, for an algebraic structure and its carrier set, and for an interpretation and its interpretation function. Moreover, I drop indices if possible without causing confusion.

I use *adicity* for predicates and *arity* for relations; so I speak of *dyadic* predicates and *binary* relations. By a *two-place* function or operation, I mean a function that takes two arguments. If R is a binary relation, I say that x *bears* R *to* y or R *is borne by* x *to* y in case $\langle x, y\rangle$ belongs to R. Moreover, I say that a (one-place) function f *takes* x *to* y if $y = f(x)$.

As for the definition of functions as functional relations, I use a convention which is less common nowadays in that I take functional relations as *one-many* relations. So for me, a binary relation is functional if it is borne by no two different things to the same thing. (The reason for this decision is given in (11.4).) As a consequence, the usual definitions of the image and the inverse image of a set by a binary relation have to be interchanged with each other.

Following W. V. Quine’s dictum that “our language is our serious all-purpose language”, I draw no strict line between object language and metalanguage; for example, I use a logical symbol like ‘$\wedge$’ both as a symbol of the language of predicate logic, which I regard as a regimented and formalized piece of natural language, as well as a notational substitute for (logical) ‘and’ in ordinary discourse.

Contents

List of Figures

List of Tables

Introduction

This thesis is about the logic and the formal semantics of classification systems. Though linguistic classification has been its main motivation, the approach presented here is not limited to linguistics at all. Another potential area of application is ontological engineering, as nowadays employed for the Semantic Web. Moreover, the framework is general enough to cover, for instance, the classification of program states in the style of denotational approaches to program semantics.

In what follows I give an overview over the main theses and contributions of this dissertation, indicate its relation to the work of others, list the prerequisites of my ideal reader, and give a brief outline of the chapter contents.

Theses and Contributions

My approach to classification systems is based on the simple observation that classifiers are first and foremost monadic predicates, and that a classification system is essentially a collection of *law-like statements* which have the form of universally quantified conditionals. So the claim is that classification systems are best taken as theories consisting of statements of the schematic form 'everything which is an A is a B', where 'A' and 'B' stand for monadic predicates; in symbols: $\forall x(Ax \rightarrow Bx)$. I will underpin this claim by several examples of classification systems used in linguistics ranging from taxonomic trees to systemic networks.

Furthermore, I take the position that classification consists in making *affirmative assertions*; that is, classifiers, or predicates, can only be affirmed to hold, and not be denied to hold. This means that in order to classify an entity, for instance, as being not animate one has to ascribe the predicate 'not-animate' or, better, 'inanimate' to it. As a consequence, if we think of entities as "bundles" of properties, an absent property does not imply that the entity in question does not satisfy that property.

In contrast to negation, logical conjunction and disjunction of classifiers do not affect the affirmative character of assertions. This is so because an affirmative assertion of a conjunction or disjunction of classifiers is equivalent to a conjunction or disjunction of affirmative assertions, respectively, which is affirmative in turn. The logic of affirmative assertions is sometimes seen as a *logic of observations*, because properties that are observed as being present preserve this status under conjunction and disjunction but not under negation.

For this reason, I assume that the premises and conclusions of the statements defining a classification system are built by finite conjunction and disjunction from primitive monadic predicates but must not include negation. Stressing the observational metaphor, I speak of *observational predicates* and *statements*, respectively. A classification system then is a set of observational statements and will be referred to as an *observational theory*. It is convenient to add two classifiers 'anything' and 'nothing' to the domain specific ones, in order to express, for example, that everything (in the domain of discourse) is an animate being, or that nothing is both animate and inanimate.

When an entity is classified against the background of a classification system, the latter allows to draw inferences about the properties of the entity under consideration. For example, if it is part of the classification system that humans are animate beings and something is classified as human, it can be automatically classified as being animate too. I call a set of properties "closed", if no more inferences can be drawn from these properties by the given classification system.

When viewed from the perspective of a classification system or observational theory, an entity is fully characterized by the set of classifiers (or concepts or properties or predicates) it satisfies. By a *generic entity* I mean any collection of classifiers that is consistent and closed with respect to the given theory. These entities can be seen as sets of properties that can occur in principle, under the assumption that the given theory is true. One can take the set of generic entities as the universe of a model of the theory in the formal sense of first-order predicate logic, where each classifier is satisfied by all generic entities it is a member of; I speak of the *(canonical) generic universe* and the *(canonical) generic model*, respectively. In a sense, the generic universe represents the *ontology* generated by the given theory.

The generic model of a finite observational theory determines the theory up to equivalence. So, if a classification system/observational theory is considered as a knowledge base, the generic model of the theory represents that knowledge base. Besides providing a classification of theories, this representation is of practical use if the canonical universe, whose cardinality can be exponential in the number of primitive predicates, has tractable size. As for classifications in linguistics, I take this to be the case for lexical theories, or at least for considerable large subtheories of those. Suitable candidates are systemic networks, which I will formalize in terms of *choice system theories*, and related formalisms used in feature-based lexical theories. (Syntax is a different thing because here the combinatorial potency of natural language comes into play.)

Due to the lack of negation, there is a natural *specialization* ordering on the generic universe: a generic entity x is specialized by a generic entity y just in case y satisfies every predicate satisfied by x. If specialization is interpreted as an increase of information, the generic universe can be viewed as a *space of information states*, which gives rise to the following *paradigm for information processing*: A cognitive agent processes information about an entity "observed" by applying his classificational knowledge to it, thereby traversing the space of information states defined by the theory in question. Under this perspective, there is a natural distinction between permanent and transient knowledge: The theory provides the background knowledge, which in turn defines the space of possible information states; the actual knowledge about an entity under consideration, in contrast, serves as a "trigger" for traversing this space from less informative to more informative states.

The way generic entities/information states are ordered by specialization strongly affects how new information can be processed. If, for instance, each bounded subset of the generic universe has a supremum then every conjunction of consistent observations has a least satisfier in the generic universe. In this case, there is a unique state which is maximal with respect to the given information. It is thus of interest to know how the ordering structure of the generic universe depends on the type of a classification system/observational theory. A typical result is that the ordered generic universe has suprema for bounded subsets just in case the theory in question is equivalent to a *Horn theory*, that is, to a theory with purely conjunctive statements. I give a systematic analysis of how a theory and its ordered generic universe depend on each other, including a detailed classification of Horn theories and simple inheritance theories.

To summarize the main points stated so far:

- Observational logic, which is the logic of affirmative assertions, provides a unifying approach to formalize and analyze classification systems as theories. This claim is supported by representing several types of classification systems as observational theories.
- The generic model of an observational theory, which determines the theory up to equivalence, can be usefully employed for knowledge representation. For instance, every logical consequence of the theory can be straightforwardly read off from the generic model.
- Observational theories together with their ordered generic universes set up a paradigm for information processing. The theory is regarded as the background knowledge of a cognitive agent, whereas the generic universe is interpreted as a space of information states the agent traverses on the basis of newly acquired information.

Every observational theory/classification system has an associated *Lindenbaum algebra*. An element of this algebra can be thought of as an equivalence class of observational predicates, where equivalence is logical equivalence with respect to the theory in question. The Lindenbaum algebra of an observational theory is a distributive lattice with zero and unit, which I refer to as an *observational algebra*. By definition, an observational theory is determined up to logical equivalence by its Lindenbaum algebra.

The algebraic viewpoint is useful in several respects. For instance, I prove the strong completeness of an inference calculus for observational statements by using a fairly standard algebraic argument which essentially relies on the fact that an observational algebra can be represented as an observational algebra of sets over its *prime spectrum*, i.e. the set of its prime filters. Viewed from another angle, the latter fact says that the generic universe of a theory can be represented by the prime spectrum of the Lindenbaum algebra of that theory. The interrelation between an observational theory, its Lindenbaum algebra, and its generic model is sketched by the diagram of Figure 1.

$$
\begin{tikzcd}
\text{Observational Theory} \arrow[rr, leftrightarrow] \arrow[dr, leftrightarrow] & & \text{Lindenbaum Algebra} \arrow[dl, leftrightarrow] \\
& \text{Generic Model} &
\end{tikzcd}
$$

FIGURE 1 Trinity of theory, model, and algebra

Lindenbaum algebra and prime spectrum representation also prove helpful for studying the relation between observational theories over different sets of primitive predicates. Concretely, I introduce *morphisms* of observational theories, which I define as translations of observational predicates subject to the condition that the target theory entails the translation of the source theory. As indicated by the term ‘morphism’, I make use of the language of *category theory*. In particular, I study the *(contravariant) functor* that assigns to each theory its ordered generic universe and to each morphism of theories an order-preserving function of generic universes that is *continuous* in that it preserves suprema of directed subsets. This functor is shown to take *(quasi-)colimits* of observational theories to *limits* of the respective generic universes. I use this general result to derive an algorithm for constructing the generic universe of the inductive limit of a sequence of theory extensions.

Finally, I apply my approach to *feature-based classification*. Here, a primitive classifier consists of an *attribute* or *feature* term and a *sort* term that determines the value of the respective attribute. For example, the attribute-value classifier COLOR:*red* is true of those entities whose color is red. A detailed logical analysis of this type of classification reveals that one should distinguish between binary features, choice features, and structural features.

I generalize the algebraic approach to feature-based classification systems or *attribute-value theories*. To this end, I introduce the concept of a *feature algebra*, which is an observational algebra together with operations, one for each feature, that preserve zero, meet, and join. I give an algebraic completeness proof for an inference calculus for attribute-value statements. Again, the key is a representation result according to which every feature algebra is representable as an algebra of sets over its prime spectrum, where the operations are realized by so-called *Peirce operations*. In addition, the generic entities of an attribute-value theory are shown to be the feature trees of that theory.

Relation to Other Work

My formalization of classification systems is inspired to a considerable degree by Carpenter and Pollard's (1991) approach to *The Logic of Linguistic Classification* (see also Carpenter 1992, Chap. 2) and by the shortcomings of that proposal. As for systemic networks, though they have been formalized in various ways (Patten and Ritchie 1987, Mellish 1988, Brew 1991, Calder 1999, Penn 2000), apparently no systematic logical and semantic analysis has been given before.

The idea to take an observational theory as a knowledge base and to represent it by its generic universe is also implicit in Oles' (2000) *Application of Lattice Theory to Knowledge Representation*. However, Oles does not seem to recognize the intimate connection between a theory and its generic universe. Typically enough, he sees "no clear connections" (*ibid*, p. 164) between his approach and the role of lattices in the framework of *Formal Concept Analysis* (Ganter and Wille 1999). The connection is this: in the finite case, the lattices of Formal Concepts Analysis are the generic universes of Horn theories without an inconsistency predicate (cf. Section 5.3.4).

The observational metaphor is strongly influenced by Vickers's (1989) *Topology via Logic*, in spite of the fact that I dispense with topological concepts almost completely. My reason for abstaining from topology is that, in contrast to Vickers, I do not take *infinite disjunctions* into account. So, on the algebraic side, I can keep to distributive lattices. (Nor do I use the intuitionistic conditional, which has lead Pollard and Sag 1987 to require complete Heyting algebras.) Another influence of Vicker's book can be found in my treatment of finite specifiability in Chapter 7. His Stone duality theorem for spectral algebraic locales bears much resemblance to my characterization of finitistic theories in terms of their Lindenbaum algebra and their ordered generic universe.

The difference between my approach and logical approaches to domain theory, be it Abramsky's (1991) *Domain Theory in Logical Form*, Zhang's (1991) *Logic of Domains*, or Droste and Göbel's (1990) *Non-Deterministic*

Information Systems and Their Domains, can perhaps best be phrased as follows: Logical approaches to domain theory ask for suitable representations of domains in terms of theories whereas I am interested in domains only insofar as they arise as ordered generic universes of observational theories. So I take a given theory as the primary object of interest; in particular, no deductive closure is presupposed, nor even a normal form.

The category of observational theories I make use of does not seem to correspond to any of those I found in the literature. Often, only primitive morphisms are taken into account (Dimov and Vakarelov 1998, Oles 2000); Coquand and Zhang (2000), in contrast, consider so-called approximable relations, which are too general for my purposes.

My logical analysis of attributive descriptions can be seen as a nominalistic reply to Moshier and Pollard's (1994) *modeling convention* for feature structures. Feature logic (with identity) is treated at length e.g. in Carpenter 1992 and Rounds 1997; the connection to modal logic is emphasized in Blackburn 1993 and Kracht 1995. Although the algebraic approach is well established in modal logic, it has not, to my knowledge, been applied to attribute-value theories before.

Intended Readership and Prerequisites

The intended readership ranges from everybody interested in formal aspects of classification systems to mathematicians and computer scientists interested in applications of domain theory and lattice theory to the logic of classification. As emphasized before, though the main motivation and most of the examples have their roots in linguistics, the presented approach to classification systems is not limited to linguistic applications.

As for mathematical and logical prerequisites, Chapters 1 to 3 and the most part of Chapter 5 presume only acquaintance with basic set theory and first-order logic. The same is true of Chapters 10 and 11. For Chapters 4 and 7, some familiarity with domain theory is helpful though not necessary since all relevant concepts are introduced either in the text or in the Appendix. Chapters 6 to 9 and Chapter 12 make frequent use of distributive lattices, especially their prime spectrum representation and the construction of quotient lattices by congruences. The Appendix covers most of the relevant definitions. In Chapters 8 and 9, the language of category theory is employed to some extent; again, necessary definitions are given in the text or in the Appendix.

Chapter Overview

Chapter 1 presents several examples of classification systems as used in linguistics. In addition, the generic entities of a classification system are informally introduced as sets of classifiers that are closed and consistent with respect to the given classification system.

Part I deals with classification systems without disjunction. In Chapter 2, simple inheritance networks are reanalyzed as classification systems whose statements are either conditionals between primitive concepts or pairwise exclusions of those. The canonical universe of simple inheritance theories is characterized as a subset system over the set of primitive concepts; based on this construction, an inheritance calculus is proved to be complete. In addition, some applications to linguistic classification are given. Chapter 3 deals with classificational theories that allow conjunction but not disjunction – so-called Horn theories. Again their canonical universe is characterized as a subset system, and the completeness of an inference calculus for Horn statements is proved by employing the canonical model. As for the task of formalizing the examples of linguistic classification given in Chapter 1, Horn theories are shown to overcome some of the restrictions of simple inheritance theories but not all. Chapter 4 gives a detailed order-theoretic characterization of the canonical universe of Horn theories and several versions of simple inheritance theories by applying concepts and results from domain theory.

Part II treats observational theories in general. Chapter 5 gives a characterization of their canonical universe in terms of subset systems. The framework is then applied to systemic networks as well as to the problem of inducing theories from instances. Moreover, the status of negation in observational theories is explored to some extent. In Chapter 6, the notion of an observational algebra is introduced, which provides an algebraic approach to observational theories in terms of their Lindenbaum algebras. It is shown that the generic universe of a theory can be represented by the prime spectrum of the Lindenbaum algebra of that theory. In addition, the axiomatic characterization of an observational algebra as a distributive lattice with zero and unit provides a sound and strongly complete inference calculus for observational statements. Chapter 7 addresses the question as to which generic entities are finitely describable, and as to what extent the generic universe is fully determined by entities of this type. An algebraic condition is given in terms of Lindenbaum algebras, and an order-theoretic one in terms of generic universes.

Part III is concerned with translations between theories over different sets of primitives and about ways of constructing new theories from old. Chapter 8 defines a category of theories by introducing an appropriate notion of morphism. The Lindenbaum algebra and the ordered generic universe of an observational theory are then seen as the result of applying appropriate functors. It is shown that a morphism is an equivalence of theories just in case the morphism is conservative and essentially surjective. These two properties and their effect on the induced functions of ordered generic universes are studied in detail. As an application, the translation of systemic networks into Horn theories is considered. Chapter 9 defines (quasi-)coproducts and (quasi-)colimits of observational theories. It is shown that the generic universe functor takes quasi-colimits of theories to limits of their respective generic universes. The special case of a sequence of extensions gives rise to a simple algorithm for constructing the generic universe step-by-step.

Part IV is about feature-based classification, where the primitive classifiers are attributive descriptions. Chapter 10 briefly outlines several examples of feature-based classification in linguistics. In addition, the importance of distinguishing between a classificational theory and its generic universe is exemplified by a case study on the semantic classification of objects by ontological sorts and binary semantic features. Chapter 11 explores more closely the logical structure of attributive descriptions by treating them as regimented natural language expressions that are formalized within first-order predicate logic. This view of attributive descriptions immediately allows to prove logical equivalences between such expressions. Chapter 12 defines feature systems of attribute-value theories as first-order models and introduces algebraic models of attribute-value theories. By using feature algebras in place of observational algebras, an algebraic completeness proof for feature logic is given. Moreover, the generic entities of an attribute-value theory are identified with the feature trees of that theory. Finally, the ordering structure of the generic universe is characterized by reducing feature logic to observational logic.

1

Linguistic Applications

Classification is an elementary method of linguistics and most other sciences. It is prerequisite for revealing regularities in linguistic data. A simple type of regularity has the form of an implication: if certain classifiers or concepts apply to a given linguistic entity then certain others apply as well. By a *classification system* we mean any systematic presentation of a collection of such implicational relations between classifiers or concepts.

In Sections 1.1, 1.2, and 1.3 we introduce three types of classification systems used in linguistics: taxonomic trees, multiple inheritance hierarchies, and systemic networks. In Section 1.4, it is indicated in what sense a given classification system, which is essentially a *theory*, uniquely determines the types of entities it is a theory about. Throughout this chapter the presentation is kept at an informal level.

1.1 Taxonomic Trees

A basic type of relation between concepts is that of a *taxonomic tree*. Early examples can be found in the work of the Phoenician philosopher Porphyry, dating back to the third century A.D., who used tree-based presentations in his introduction to Aristotle's *Categories*.[1] For a truly classical example consider the tree of Figure 2;[2] it expresses the subdivision of bodies into living and nonliving beings, where living beings are further subdivided into animal and vegetal beings, and animals are subclassified as being rational or irrational.

The notion of tree presupposes that there is exactly one *root* concept, the *genus generalissimum*, from which each of the concepts in the tree is reachable by exactly one *path*. (Informally, a path is a sequence of consecutive connections between concepts in the tree.) In particular, for each concept except the root there is exactly one concept immediately superordinate to it. It is custom to draw trees upside down, that is, the root is at the top and the *leaves*, which

[1] See e.g. Eco 1984 for details.
[2] Adapted from Eco 1984, p. 50.

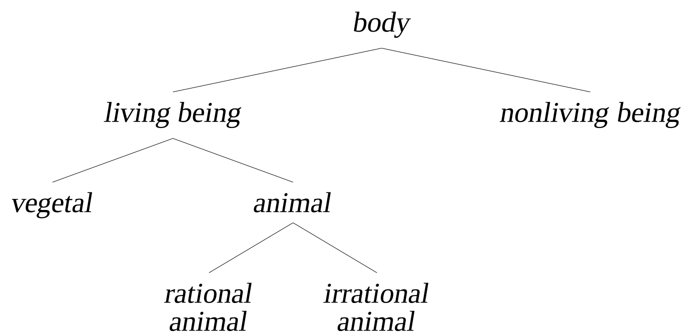

FIGURE 2 Fragment of a Porphyrian tree

are the concepts lacking any subconcept, are at the bottom.

From a logical point of view each concept of a taxonomic tree implies its superordinate concept, its *genus proximum*. For example, according to the taxonomy of Figure 2, *animal* implies *living being*. Furthermore, any two immediate subconcepts of one and the same concept are incompatible – as e.g. *vegetal* and *animal*. Phrased in terms of *structural semantics*,[3] each concept implies its *hyperonym* whereas its *hyponyms* are incompatible, that is, *co-hyponyms* are *heteronymous*.

A taxonomy is called *exhaustive* if each concept implies the disjunction of its immediate subconcepts. In other words, if something belongs to a taxonomic class then it belongs to one of its immediate subclasses. Exhaustiveness is often referred to as *closed world assumption* in the sense that the taxonomy is assumed to exhaustively subclassify everything in the universe of discourse. If the tree of Figure 2 is exhaustive, a body is either a living being or a nonliving being, a living being is either a vegetal or an animal, and an animal is either rational or irrational.

Figure 3 graphically presents a possible taxonomy of German nominal words as it is common in traditional grammar.[4] Strictly speaking, the taxonomic tree of Figure 3 is inconsistent due to the multiple occurrence of *definite* and *indefinite*. For *definite* and *indefinite* would imply both *article* and *pronoun*, which are incompatible by assumption. One can easily resolve this seeming paradox by replacing the two occurrences of *definite* respectively by *definite article* and *definite pronoun*, and correspondingly for *indefinite*. This solution however gives rise to a “combinatorial explosion” of concepts. In addition, one might ask in such cases whether it is not more adequate to reorganize the taxonomic hierarchy to the effect that definite nominals are distinguished from indefinite nominals first and thereafter definite and indefinite

[3] See e.g. Lyons 1977, Chap. 9, 1995, Chap. 4.
[4] Compiled from Eisenberg 1999, Chap. 5.

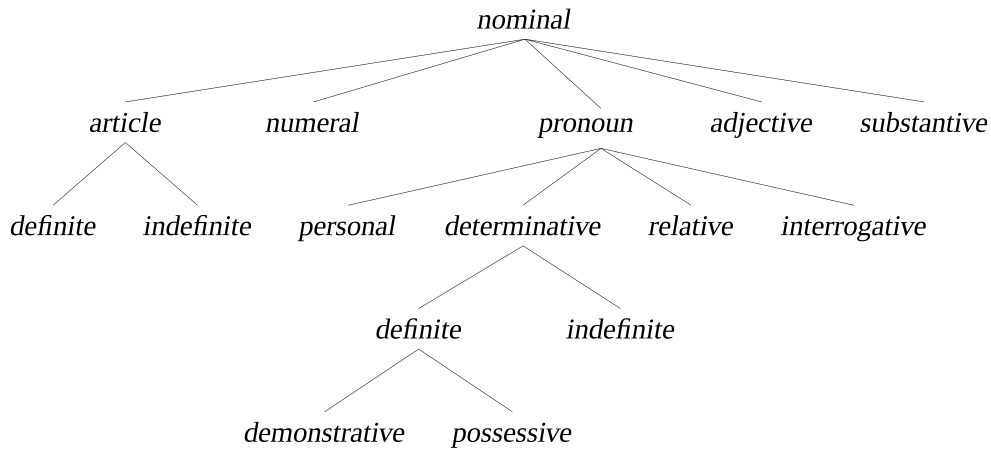

FIGURE 3 Subclassification of German nominals

articles respectively from definite and indefinite pronouns. It is worth mentioning that this type of arbitrariness in building the tree is an objection already put forward by Boethius in the middle of the first millennium A.D. and later by medieval philosophers like Abelard.[5] The more general representational means of Section 1.2 will overcome these problems insofar as they allow to subclassify concepts along several dimensions simultaneously.

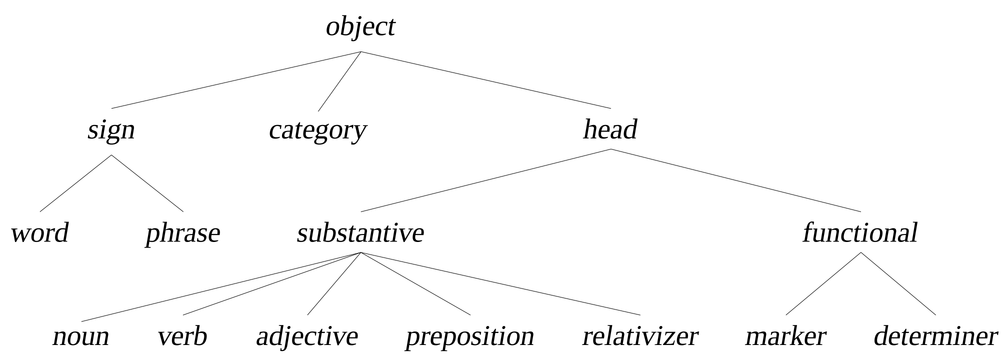

FIGURE 4 Fragment of the HPSG sort hierarchy

Taxonomic trees also play an important role in feature-based grammatical theories like Head-Driven Phrase Structure Grammar (HPSG). A small part of the HPSG *sort hierarchy* is shown in Figure 4.[6] The sort hierarchy is intended to subclassify all linguistic entities the grammar deals with. Its classes obviously are not as self-explaining as those of traditional taxonomies like that of

[5] See Eco, *op cit*, Sect. 2.2.
[6] See Pollard and Sag 1994, pp. 395ff.

Figure 3. For example, one intuitively expects *noun* to be a subclass of *category*. To reveal the intended meaning of the sorts, the so-called *feature declarations* of HPSG have to be taken into account. Slightly simplified, HPSG assumes linguistic entities of sort *sign* to have an attribute CATEGORY, whose value is of sort *category* and has in turn an attribute HEAD with value of sort *head*. The sorts *noun*, *verb*, etc thus classify the values of the attribute HEAD whereas entities of sort *category* have an additional attribute SUBCAT specifying the valence of the sign. See Part IV for a closer examination of this kind of approach.

1.2 Multidimensional Classification

Categories are often subclassified along several *attribute* or *feature dimensions*. German nominal word forms, for example, can be subclassified with respect to case, gender, and number as indicated in Figure 5. In terms of feature-based

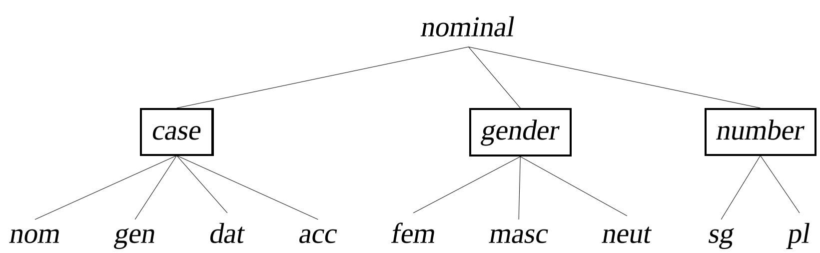

FIGURE 5 Classification of German nominals by case, gender, and number

grammatical theories it is common to say that *case*, *gender*, and *number* are *attributes* of German nominals which can have certain *values*. The attribute *case* for instance can take the values *nominative*, *genitive*, etc. A more natural manner of speaking is that German nominals *are* either nominative, genitive, etc and *are* either singular or plural. The categories *nominative*, *genitive*, etc are thus taken as *subcategories* of *nominal*. *Case* is then a collection of certain subcategories of *nominal*.[7] It is nevertheless possible to come up with a category *case* as part of the classificational hierarchy: just let *case* be satisfied by those expressions which *have* or *carry* a case.

If a category such as *case* is most naturally analyzed as a collection of alternative categories, here $\{nominative, genitive, \ldots\}$, we speak of a *choice category*. The set of alternatives is then referred to as a *choice system*. Choice systems are well known in traditional linguistics as so-called *paradigms*.

Koenig (1999) calls trees like that of Figure 5 AND/OR *trees*, drawing on standard terminology of Artificial Intelligence.[8] Its leaves need not be incom-

[7] See Eisenberg 1999, p. 18 or Schnelle 1973, p. 151 for similar views.

[8] Cf. e.g. Barr and Feigenbaum 1981. Those who are acquainted with AND/OR trees will notice

patible, contrary to the case of taxonomic trees. For example, a nominal expression may simultaneously be nominative, feminine, and singular. The condition on taxonomic trees that every two subconcepts of a given concept are incompatible is thus relaxed for AND/OR trees: choice categories are exempt from it. Furthermore, each concept is required to imply every choice category subordinate to it – e.g., *nominal* implies *case*, *gender*, and *number*, which means that all nominals carry case, gender, and number.

Multidimensional classification is related to *componential analysis*, a well-known method of structural semantics proposed e.g. by Louis Hjelmslev.[9] For example, the concept *stallion* can be said to have the sense components *mature*, *male*, and *horse* – which incidentally resemble the typical dictionary entry for 'stallion'. Choosing the factors *immature*, *female*, and *human* instead gives the componential analysis of *girl*. Such sense factorizations can obviously be seen as members of the Cartesian product of certain choice systems like *gender* and *stage of development*.

The venerable terminology of Aristotle's theory of definition,[10] partly familiar from Section 1.1, is even better suited for the above examples: *horse* is the *genus proximum* of the concept *stallion* whereas *mature* and *male* are its *differentiae specificae*. Figure 6 illustrates how this interplay between taxonomy and decomposition can be represented by combining a taxonomic tree with an AND/OR *net*.[11] AND/OR nets differ from AND/OR trees insofar as they allow concepts to have more than one immediate superordinate concept. Logically, each concept implies all of its superordinate concepts. Needless to say that Figure 6 is only an imperfect sketch – for instance, having a stage of development should not imply to be an animal.

Classification systems as expressed by AND/OR nets are also referred to as *multiple inheritance hierarchies*. The idea is that concepts inherit along several "dimensions" from superordinate concepts in the sense that concepts may have more than one immediate superordinate, each belonging to a different choice system. For a further linguistic application of multiple inheritance consider the AND/OR net of Figure 7, which shows part of the cross-classification of English lexemes with respect to part of speech and argument selection.[12] Notice that the elements at the bottom behave more like *instances* of their superordinates than as subconcepts. This is reflected by the fact that in contrast to Figure 6 it is not

---

the difference in presentation: the standard approach uses arcs to group conjunctive connections; the present notation, in contrast, provides special means for putting together elements into exclusive disjunctions.

[9] See Lyons, *op cit*.

[10] See Eco, *op cit*.

[11] Again Koenig's terminology.

[12] Adapted from Sag and Wasow 1999, p. 361. The terms 'subject raising' and 'subject control' refer to different types of dependencies between the subjects of verbs in infinitival complements and those of the head lexeme; see e.g. *ibid*, Chap. 12 for background.

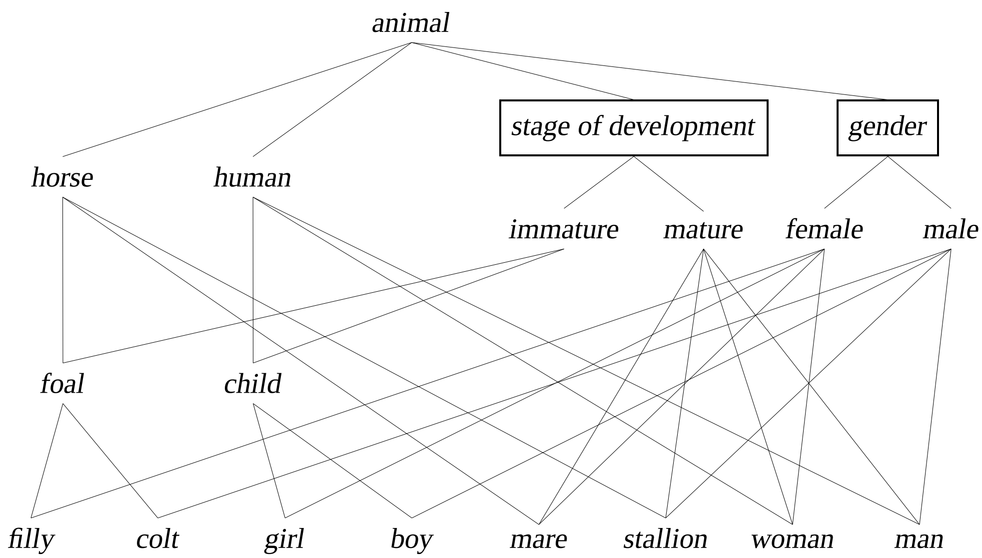

FIGURE 6 Of horses and men

appropriate to assume that the bottom elements are implied by the conjunction of their superordinates because there is usually more than one instance of the conjunctive concept in question.

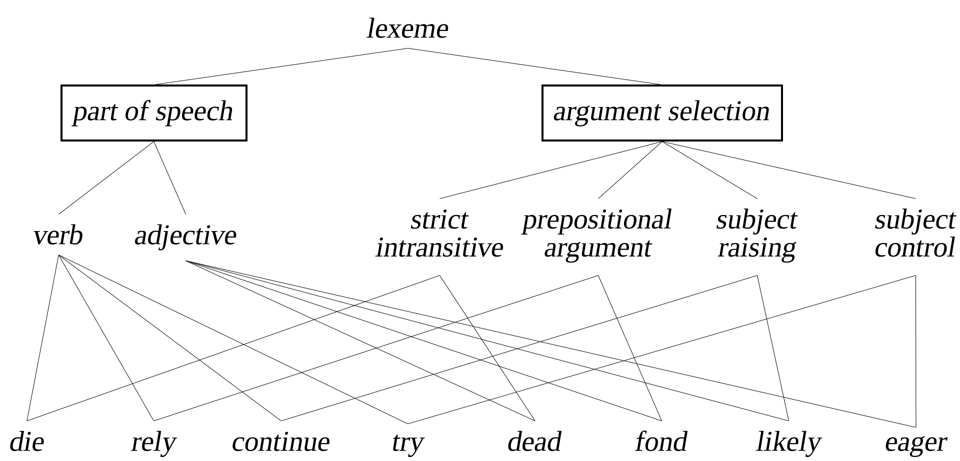

FIGURE 7 Cross-classification by part of speech and argument selection

1.3 Systemic Networks

Systemic networks are a key component of *systemic grammar* (Kress 1976, Halliday 1994). Much like AND/OR nets they express logical dependencies among choice systems. Consider Terry Winograd's well-known example of a

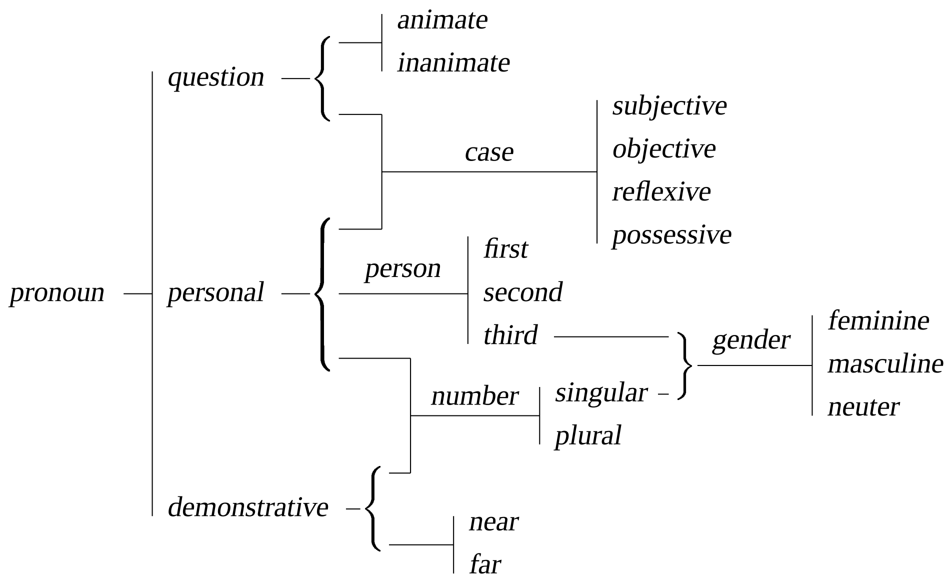

FIGURE 8 Winograd's systemic network for English pronoun classification

systemic network for English pronouns presented in Figure 8.[13] Choice systems correspond to vertical bars with their members aligned in a column to the right of the bar. Each choice system has a so-called *entry condition*, which is built from concepts occurring in the network by conjunction and disjunction. The entry condition of a choice system determines whether or not one of the concepts of the system applies to a given entity. In logical terms, choice disjunction and entry condition imply each other. Graphical correlates of conjunction and disjunction are respectively brace and bar. The network of Figure 9, for instance, determines the entry condition of the choice system P to be $(a \wedge b) \vee c$.

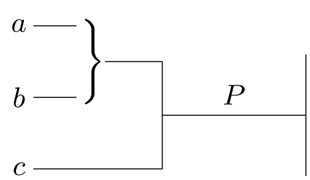

FIGURE 9 Choice system P with entry condition $(a \wedge b) \vee c$

Systemic networks are more expressive than AND/OR nets in that they allow entry conditions of choice systems to consist of arbitrary combinations of conjunction and disjunction. By requiring that entry conditions and choice disjunctions are equivalent, we tacitly assume exhaustiveness. A weaker as-

[13] Winograd 1983, p. 293; unlike Winograd we do not count possessive determiners as pronouns.

sumption is that a choice disjunction implies its entry condition, which means that each concept implies the entry condition of the choice system it is a member of.[14]

1.4 Classification Systems as Constraint Systems

Another perspective on classification systems – of central importance in subsequent chapters and presented here in an informal manner – emphasizes the *constraints* the classification system imposes on the *types of entities* described. According to the constraints given by the classification system of Figure 5, for example, there are nominals which are accusative and neuter but no nominal is both feminine and neuter. Similarly, the AND/OR net of Figure 6 excludes human foals and female colts.

The notion of an entity type, or *generic entity*, can be conceived as follows. Imagine there is no other way to describe an entity than to use the concepts of a given classification system. It is then reasonable to say that an entity is fully characterized by the set of all concepts it satisfies – we can, in a sense, *identify* the entity with the respective concept set. Any set of concepts that can arise that way, given the classification system is true, will now be taken to represent an entity type of the system.

Since, by assumption, a classification system is true of its entity types, it follows that a concept set representing an entity type is "closed" or "saturated" with respect to all classificational constraints of the system. For example, if one of the constraints is that *foal* implies *horse*, then an entity type satisfying *foal* is assumed to satisfy *horse* as well; consequently, every concept set representing an entity type has to satisfy the condition that *horse* belongs to its members whenever *foal* does. Moreover, *inconsistent* concept sets, i.e. sets with incompatible members, do not represent any entity type. For instance, no consistently closed set includes both *female* and *colt* since *colt* implies *male*, which is incompatible with *female*.

When viewed as a system of constraints each classification system uniquely determines a set of (representatives of) entity types. Take (finite) taxonomic trees as a simple example. Disregarding exhaustiveness for the moment, a concept set is closed with respect to the taxonomic constraints if all superordinates of concepts of the set are members of it in turn. Since concepts with common superordinate are incompatible, a closed set is consistent if and only if its members form a linearly ordered subset, i.e. a *chain* of the tree. In particular, every consistently closed concept set either is empty or has a most specific member, of which its other members are direct or indirect superordinates. Because

[14] The non-exhaustive version is used by Carpenter and Pollard 1991 and Carpenter 1992 whereas Patten and Ritchie 1987, Mellish 1988, and Brew 1991 presume exhaustiveness; see also Section 5.3.2.

of this one-to-one correspondence between nonempty consistently closed sets and their most specific members, the entity types determined by a taxonomic tree can equally well be represented by the concepts of the tree itself (plus an additional element for the empty set). Hence, although consistently closed concept sets provide a *canonical* way of representing entity types, it is by far not the only one.

Exhaustiveness makes things slightly more intricate. Recall from Section 1.1 that a taxonomic tree is exhaustive if each concept implies the disjunction of its immediate subconcepts (if there are any). In accordance with the standard meaning of disjunction, an entity type is said to satisfy a disjunction of concepts if it satisfies at least one of them. Consequently, at least one of the concepts occurring in the disjunction is required to be a member of the corresponding closed concept set. (Similarly for conjunction: an entity type satisfies a conjunction of concepts if it satisfies each of them, that is, every concept occurring in the conjunction must be a member of the corresponding closed concept set.) To exemplify the effect of exhaustiveness consider the classification system of Figure 3. One of its constraints is that *article* implies *definite* or *indefinite*. So every closed concept set is subject to the condition that *definite* or *indefinite* belongs to its members whenever *article* does. It is not difficult to see that exhaustiveness enforces the consistently closed sets to coincide with the *maximal* chains (also called *branches*[15]) of the taxonomic tree. Consequently, when using the representation of consistently closed sets by their most specific members, only the *leaves* of the tree (the *infimae species*) remain as representatives of entity types. This gives a first example of the unsurprising fact that more constraints mean less entity types.

For AND/OR nets there is no such direct correspondence between concepts and entity types. Let us confine ourselves to AND/OR *trees*, i.e. to cases like that of Figure 5, or that of Figure 7 except the bottom line. Figure 10 depicts the entity types determined by the latter example when exhaustiveness is not assumed. (The meaning of the edges is explained below. For expository purposes, braces and commas in set names are often omitted.) These entity types are represented by the concept *lexeme*, by the members of the choice systems *part of speech* and *argument selection*, and by the sets having exactly one member of each choice system, i.e. by *strict intransitive verb*, *strict intransitive adjective*, etc. Exhaustiveness preserves only the last mentioned entity types. Notice that by assuming e.g. *strict intransitive verb* to represent an entity type we tacitly introduce a new kind of representative: concept sets which are not necessarily closed but for which there is a *unique closure*, i.e. a unique minimal consistently closed set containing them – as e.g. {*strict intransitive*, *verb*, *lexeme*} is one for {*strict intransitive*, *verb*}. Applied to

[15] See e.g. Landman 1991, Sec. 2.3 for an introduction to this and related notions.

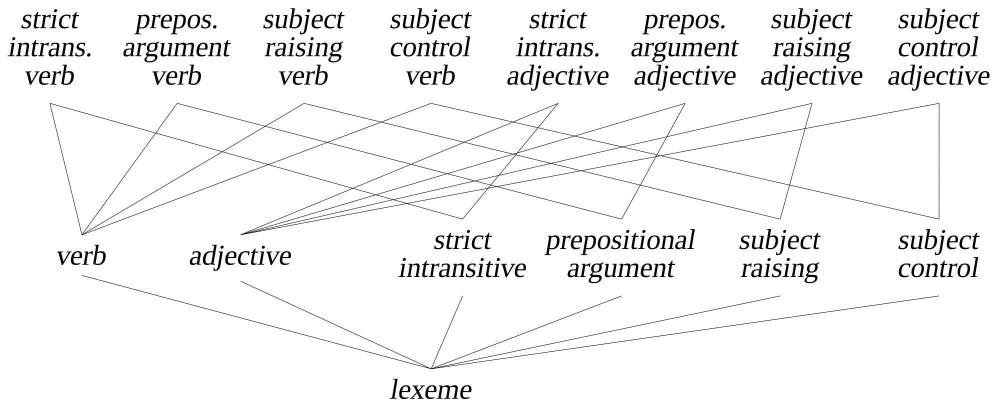

FIGURE 10 Entity types determined by the upper part of Figure 7.

singleton sets, i.e. sets with exactly one member, this resembles the representation of entity types by single concepts.

Just like taxonomic trees and AND/OR nets, a systemic network uniquely determines a set of entity types. Again of course, this set crucially hinges upon the presumed logical dependencies. For example, if exhaustiveness is assumed, Winograd's pronoun network of Figure 8 leads to 44 entity types, witness Figure 11,[16] whereas without exhaustiveness we get 100 entity types (as the dauntless reader may want to convince himself).[17]

The canonical representation of entity types by consistently closed concept sets comes along with set inclusion as a natural (partial) order. The respective order on the set of entity types is commonly called *specialization*. Notice that if x bears the specialization relation to y, we do not say that x specializes y but, on the contrary, we say that x *is specialized by* y; for to specialize x means to satisfy every concept satisfied by x (and possibly some more). Notice further that for other representations of entity types by concept sets, specialization usually is *not* set inclusion. Take taxonomic trees without exhaustiveness: as noted above, entity types can be represented by singleton sets of concepts, of which obviously none is contained in one of the others. Specialization in this case is given by (the converse of) the hierarchical relation between concepts as defined by the taxonomic tree. A concrete example of specialization is provided by Figure 10, where an edge between two entity types indicates that the lower one in the diagram is specialized by the upper one.

We conclude this section by turning once more to structural semantics, for

[16]Partly taken from Brew 1991.

[17]Notice that without any constraints the number of entity types is $2^{20} \approx 10^6$, which is just the cardinality of the power set of the given concept set. If the incompatibility constraints of the choice systems are taken into account, the number of entity types reduces to $2 \cdot 3^3 \cdot 4^3 \cdot 5 = 17{,}280$; cf. (9.2).

question pronoun animate subjective	who
question pronoun animate objective	whom
question pronoun animate reflexive	?
question pronoun animate possessive	whose
question pronoun inanimate subjective	what
question pronoun inanimate objective	what
question pronoun inanimate reflexive	?
question pronoun inanimate possessive	whose
personal pronoun subjective first singular	I
personal pronoun subjective first plural	we
personal pronoun subjective second singular	you
personal pronoun subjective second plural	you
personal pronoun subjective third singular feminine	she
personal pronoun subjective third singular masculine	he
personal pronoun subjective third singular neuter	it
personal pronoun subjective third plural	they
personal pronoun objective first singular	me
personal pronoun objective first plural	us
personal pronoun objective second singular	you
personal pronoun objective second plural	you
personal pronoun objective third singular feminine	her
personal pronoun objective third singular masculine	him
personal pronoun objective third singular neuter	it
personal pronoun objective third plural	them
personal pronoun reflexive first singular	myself
personal pronoun reflexive first plural	ourselves
personal pronoun reflexive second singular	yourself
personal pronoun reflexive second plural	yourselves
personal pronoun reflexive third singular feminine	herself
personal pronoun reflexive third singular masculine	himself
personal pronoun reflexive third singular neuter	itself
personal pronoun reflexive third plural	themselves
personal pronoun possessive first singular	mine
personal pronoun possessive first plural	ours
personal pronoun possessive second singular	yours
personal pronoun possessive second plural	yours
personal pronoun possessive third singular feminine	hers
personal pronoun possessive third singular masculine	his
personal pronoun possessive third singular neuter	its
personal pronoun possessive third plural	theirs
demonstrative pronoun near singular	this
demonstrative pronoun near plural	these
demonstrative pronoun far singular	that
demonstrative pronoun far plural	those

FIGURE 11 English pronouns

which Figure 6 serves as a running example. Without additional assumptions, the concept sets *immature horse* and *foal* represent different entity types. Exhaustiveness, however, is too strong if we want *foal* to represent an entity type. For exhaustiveness has the consequence that *foal* implies the disjunction of its immediate subordinates *filly* and *colt*. So we get two different minimal consistently closed sets with member *foal*, with *filly* belonging to one, and *colt* to the other. What is needed instead is the assumption that a concept is implied by the conjunction of its immediate superordinates. Concretely, the concept *foal* is implied by (and hence equivalent to) the conjunction of *immature* and *horse*. Following Carpenter and Pollard (1991) we call *foal* a *defined* concept.

It is of some interest for structural lexical semantics to distinguish entity types that are representable by single concepts from those which are not. The latter case is often referred to as an "accidental semantic gap" in the lexicon or a *lexical gap*, for short.[18] The entity type *immature horse* for instance is represented by the concept *foal* whereas there is no single concept for *mature horse* or for *immature male animal*.

[18] See Chomsky 1965, p. 170; cf. Lyons 1977, Sect. 9.6 for some critical reflections on this topic.

Part I

Definite Classification

2

Simple Inheritance Theories

The following brief exposition of the theory of simple inheritance networks exemplifies some constructions and arguments that will frequently be used in later sections. In particular, a first formalization of the ideas sketched in Section 1.4 is given.

In Section 2.1, the canonical universe of a simple inheritance network is introduced as the system of consistently closed sets of primitive concepts. We characterize the subset systems that arise that way. In addition, it is shown that under certain conditions the consistently closed sets have a nonredundant representation by maximally specific elements.

In Section 2.2, the links of a simple inheritance network are translated into universally quantified conditionals. This reformulation within (first-order) predicate logic allows us to use the standard notions of interpretation and model. Moreover, we give an inference calculus for inheritance statements that is sound and complete with respect to first-order entailment.

Section 2.3 recapitulates Carpenter and Pollard's (1991) proposal to formalize systemic networks and multiple inheritance hierarchies by translating them into simple inheritance networks. We point out the limitations of this approach and indicate how to overcome them.

2.1 Simple Inheritance Networks

Simple inheritance networks, or IS(NOT)A networks, are usually introduced as collections of links between concept nodes.[1] There are two kinds of links: ISA links express subordination of concepts whereas ISNOTA links indicate incompatibility. More precisely, a (*simple*) *inheritance network* Γ consists of a set Σ_Γ of concepts and two relations ISA_Γ and ISNOTA_Γ on Σ_Γ; ISA_Γ and

[1] Since concepts will be taken to denote sets of entities, we disregard *singular concepts* (or constants or names) for simplicity; otherwise we would need to make sure in one way or another that certain concepts denote singleton sets, which is beyond the scope of the formalism presented in this section.

ISNOTA_Γ links are thus ordered pairs of concepts. (We drop the subscript if reference to Γ is clear from the context.) Notice that we do not require the relations ISA_Γ and ISNOTA_Γ to satisfy any closure properties whatsoever, like reflexivity or transitivity.[2]

For a simple example of an inheritance network see Figure 12, where arrows indicate ISA links and lines terminated by shaded circles indicate ISNOTA links. According to this network, for instance, lemons are yellow and whatever is yellow is not red.

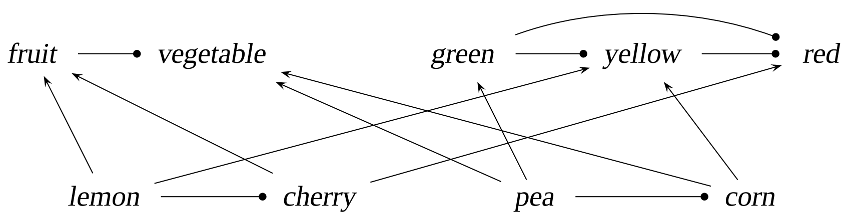

FIGURE 12 Simple inheritance network of some fruits and vegetables

(2.1) Remark *(Nonmonotonic inheritance)* In the following, only *strict* or *monotonic* inheritance is taken into account. Approaches that allow *defeasible* links in inheritance networks, thus giving rise to *nonmonotonic* inheritance, can be found in Horty 1994 and Thomason 1997. See also Osswald 2003a for an outline how to adapt the logic presented in this and later chapters to nonmonotonic reasoning.

2.1.1 *The Canonical Universe*

Every inheritance network Γ determines a set of entity types in the sense of Section 1.4. Again, the basic idea is to represent an entity type by a set of concepts that is consistent and closed with respect to inheritance by Γ. We say that a subset X of Σ is Γ-*closed* if it is the case that

$$\text{if} \quad p \in X \quad \text{and} \quad p \ \text{ISA} \ q \quad \text{then} \quad q \in X.$$

A closed set is called *consistent* if there is no ISNOTA link between any two of its members. Let $\mathrm{C}(\Gamma)$ be the set of consistently Γ-closed subsets of Σ, ordered by set inclusion. We refer to this ordering as *specialization*. The empty set $\varnothing$ is the least element of $\mathrm{C}(\Gamma)$ with respect to specialization; it corresponds to the most general entity type "something". Figure 13 depicts the system of consistently closed sets given by the example network of Figure 12.

[2]The same definition of an inheritance network is used by Carpenter and Pollard 1991, p. 10.

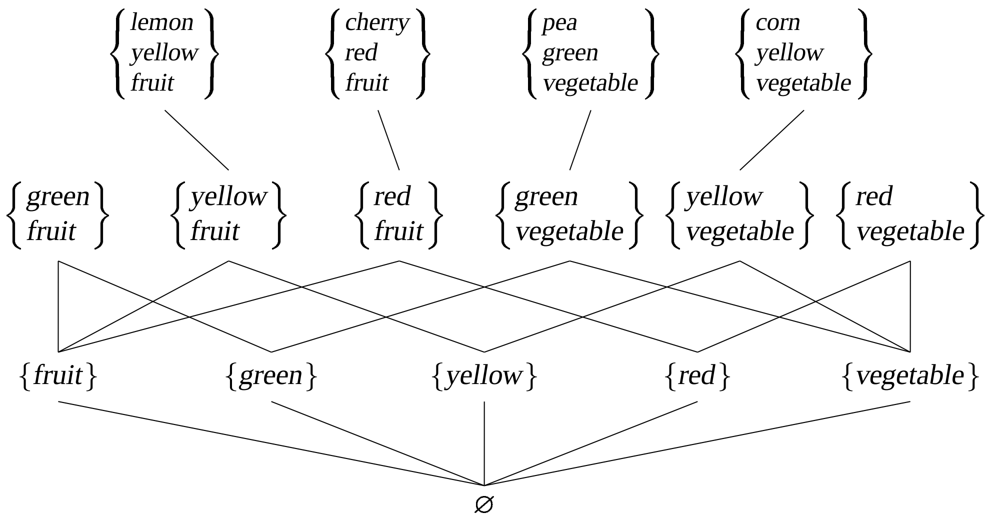

FIGURE 13 Canonical universe of the fruit/vegetable example

Let us call $\mathrm{C}(\Gamma)$ the *canonical* (*generic*) *universe* of Γ. As in Section 1.4, we say that an entity type represented by a member X of $\mathrm{C}(\Gamma)$ *satisfies* a concept p if p belongs to X. Instead of entity types we also speak of *generic entities*. This terminology is motivated by the observation that no two different members of the canonical universe satisfy exactly the same concepts. In the above example, for instance, we have precisely one generic fruit, one generic lemon, etc.

It is not difficult to characterize the structure of $\mathrm{C}(\Gamma)$ as an ordered set, which to a considerable part is revealed by the following simple observation:

(2.2) Lemma The intersection and the union of a system of Γ-closed sets are Γ-closed in turn.[3]

Proof. Suppose $\mathcal{S}$ is a system of Γ-closed sets and p ISA_Γ q. If $p \in \bigcap \mathcal{S}$ then, for every member X of $\mathcal{S}$, $p \in X$ and hence $q \in X$, that is, $q \in \bigcap \mathcal{S}$. If $p \in \bigcup \mathcal{S}$ then there is a member X of $\mathcal{S}$ such that $p \in X$ and thus $q \in X \subseteq \bigcup \mathcal{S}$. $\lrcorner$

Intersection of nonempty systems of Γ-closed sets obviously preserves consistency; union does not: take the network consisting of two concepts a and b, no ISA link, and one ISNOTA link between a and b. Then $\{a\}$ and $\{b\}$ are consistent but $\{a, b\}$ is not. So one has to impose additional constraints on a subset of $\mathrm{C}(\Gamma)$ to ensure that its union belongs to $\mathrm{C}(\Gamma)$ in turn. Call a subset

[3] The intersection of the empty system needs some attention. With $x \in \bigcap \mathcal{S}$ iff $x \in X$ for every member X of $\mathcal{S}$, everything would be a member of $\bigcap \varnothing$. Since we are not interested in everything but only in members of Σ, $\bigcap \varnothing$ is assumed to be Σ in this context.

of $\mathrm{C}(\Gamma)$ *pairwise bounded* if it is nonempty and the union of every two of its members is a subset of some member of $\mathrm{C}(\Gamma)$.

(2.3) Theorem The canonical universe of a simple inheritance network is closed with respect to intersection of nonempty subsets and with respect to union of pairwise bounded subsets.

Proof. By (2.2) and the definition of consistency, $\mathrm{C}(\Gamma)$ is closed with respect to intersection of nonempty subsets. Suppose $\mathcal{S}$ is a pairwise bounded subset of $\mathrm{C}(\Gamma)$. According to (2.2), it remains to check the consistency of $\bigcup\mathcal{S}$. If $p, q \in \bigcup\mathcal{S}$, there are members X and Y of $\mathcal{S}$ such that $p \in X$ and $q \in Y$. Since by assumption $X \cup Y$ is a subset of some member of $\mathcal{S}$, it is not the case that p ISNOTA q. ⌟

The two properties of (2.3), namely to be closed with respect to nonempty intersection and pairwise bounded union, are not only necessary but also sufficient for a subset system $\mathcal{U}$ over Σ to be the canonical universe of some simple inheritance network over Σ. For suppose $\mathcal{U}$ has both properties; consider the network $\Gamma(\mathcal{U})$ over Σ with

(2.4)
$$\begin{aligned} p \ \mathrm{ISA}_{\Gamma(\mathcal{U})}\ q \quad &\text{iff} \quad \forall X \in \mathcal{U}\ (p \in X \rightarrow q \in X), \\ p \ \mathrm{ISNOTA}_{\Gamma(\mathcal{U})}\ q \quad &\text{iff} \quad \nexists X \in \mathcal{U}\ (p \in X \wedge q \in X). \end{aligned}$$

Then the consistently closed sets determined by $\Gamma(\mathcal{U})$ are the members of $\mathcal{U}$:

(2.5)
$$\mathrm{C}(\Gamma(\mathcal{U})) = \mathcal{U}.$$

To see this, observe first that the members of $\mathcal{U}$ are consistently $\Gamma(\mathcal{U})$-closed by definition of $\Gamma(\mathcal{U})$; hence, $\mathcal{U} \subseteq \mathrm{C}(\Gamma(\mathcal{U}))$.

For proving the reverse inclusion, we make use of the least satisfier of a satisfiable concept. We say that a concept p is *satisfiable in* $\mathcal{U}$, if p is satisfied by some member of $\mathcal{U}$, that is, in accordance with the above definition of satisfaction, if p belongs to some member of $\mathcal{U}$. In case p is satisfiable in $\mathcal{U}$, the set $\{Y \in \mathcal{U} \mid p \in Y\}$ is nonempty; its intersection $\mathrm{s}(p)$ therefore belongs to $\mathcal{U}$ since $\mathcal{U}$ is closed with respect to nonempty intersection:

$$\mathrm{s}(p) = \bigcap\{Y \in \mathcal{U} \mid p \in Y\} \in \mathcal{U}.$$

By definition, $\mathrm{s}(p)$ is the least member of $\mathcal{U}$ satisfying p, in short, *the least satisfier of* p *(in* $\mathcal{U}$*)*.

Assume now that $X \in \mathrm{C}(\Gamma(\mathcal{U}))$. We need to show that $X \in \mathcal{U}$. Since X is consistent, every member p of X is satisfiable in $\mathcal{U}$, because otherwise

p ISNOTA$_{\Gamma(\mathcal{U})}$ p, by (2.4). So, every member of X has a least satisfier in $\mathcal{U}$; in addition, by (2.4), p ISA$_{\Gamma(\mathcal{U})}$ q iff $q \in \mathrm{s}(p)$; hence, since X is closed, $\mathrm{s}(p) \subseteq X$ whenever $p \in X$. To sum up,

$$X = \bigcup\{\mathrm{s}(p) \mid p \in X\}.$$

It remains to check that the set $\{\mathrm{s}(p) \mid p \in X\}$ is pairwise bounded to make sure that its union belongs to $\mathcal{U}$. Every two members p and q of X belong to some member Y of $\mathcal{U}$ since X is consistent; hence $\mathrm{s}(p) \cup \mathrm{s}(q) \subseteq Y \in \mathcal{U}$, which concludes the proof of (2.5) and hence of:

(2.6) Theorem Every subset system over a (concept) set Σ that is closed with respect to nonempty intersection and pairwise bounded union is the canonical universe of an inheritance network over Σ.

Notice that, in general, $\Gamma(\mathrm{C}(\Gamma))$ differs from Γ though both networks determine the same set of entity types according to (2.5). In fact, the ISA relation of $\Gamma(\mathrm{C}(\Gamma))$ as defined by (2.4) is reflexive and transitive and includes ISA$_\Gamma$, that is, includes the reflexive and transitive closure of ISA$_\Gamma$.

In view of (2.4) and (2.2) the special case of pure ISA networks, i.e. of inheritance networks lacking ISNOTA links, can be characterized as follows:

(2.7) Corollary The canonical universes of ISA networks over Σ are precisely the *complete subset lattices over* Σ, i.e. those subset systems over Σ that are closed with respect to arbitrary intersection and union.

If, in contrast, there are only ISNOTA links then all concept sets are closed; in particular, every subset of a consistently closed set is closed and consistent in turn. Moreover, (2.4) applied to a subset system with this property obviously does not give rise to any ISA-links besides reflexive ones, i.e. links of the form p ISA p, which can be left out without harm; consequently:

(2.8) Corollary The canonical universes of ISNOTA networks over Σ are precisely the subset systems over Σ that are closed with respect to the formation of subsets and pairwise bounded union.

Subset systems over Σ as described in (2.8) are also known as *coherence spaces over* Σ.[4]

[4]See Girard et al. 1989.

2.1.2 *Representation by Nonredundant Concept Sets*

Let Γ be an inheritance network over Σ. As pointed out in Section 1.4, it is not essential to use consistently closed concept sets for representing the entity types of Γ. The members of any set U do as well, in case there is a one-to-one correspondence between U and $\mathrm{C}(\Gamma)$. In particular, arbitrary subsets of Σ could be used as entity types. More concretely, a member X of $\mathrm{C}(\Gamma)$ can be represented by a subset of X that is uniquely determined by X; the other way around: for each such set Y there must be a unique consistently closed set containing Y. Now observe that if an *arbitrary* subset Y of Σ is included by some member of $\mathrm{C}(\Gamma)$ then, according to (2.3), there is a *least* member of $\mathrm{C}(\Gamma)$ including Y, the ISA-*closure* $\mathrm{cl}(Y)$ *of* Y, which is given by intersection:

$$\mathrm{cl}(Y) \;=\; \bigcap\{X \in \mathrm{C}(\Gamma) \mid Y \subseteq X\}.$$

Thus, what is needed is a function f taking every member X of $\mathrm{C}(\Gamma)$ to a subset of Σ such that $\mathrm{cl}(f(X)) = X$. Then $f(\mathrm{C}(\Gamma))$ is a suitable system of subsets for representing $\mathrm{C}(\Gamma)$.

Suppose we want to choose Y as small as possible with closure X. Then it is reasonable to take the set $\mathrm{m}(X)$ of *maximally specific* members of X, i.e. the set of those members p of X such that there is no member q of X different from p with q ISA p.[5] The maximally specific members are thus the ISA-*initial* ones, that is, the ISA*-*minimal* ones, where ISA* is the reflexive and transitive closure of ISA.

The existence of maximally specific members is not always guaranteed. If e.g. p ISA q and q ISA p then $\{p, q\}$ does not have maximally specific members. Each (nonempty) subset of Σ has maximally specific members just in case the relation ISA minus identity (i.e. $\{\langle p, q\rangle \mid p \text{ ISA } q \wedge p \neq q\}$) is *well-founded*.[6]

(2.9) Lemma If ISA minus identity is well-founded then $\mathrm{cl}(\mathrm{m}(X)) = X$ for every member X of $\mathrm{C}(\Gamma)$.

Proof. Suppose $X \in \mathrm{C}(\Gamma)$. It suffices to show that for every member p of X there is an ISA-initial member q of X such that q ISA* p. But the set $X \cap \{r \mid r \text{ ISA* } p\}$, by assumption, has ISA-initial members, and any of these is such a q. ⌟

[5] This representation is for example taken up in Carpenter 1992, Chap. 2.

[6] A relation R on a set U is said to be *well-founded* if each nonempty subset X of U has a member x to which R is borne by no member y of X. Well-foundedness of R is equivalent to the condition that there is no infinite sequence such that every element of the sequence bears R to its predecessor; see e.g. Wechler 1992, p. 37.

When applied to the fruit/vegetable example of Figure 12, the representation of entity types by nonredundant concept sets is given by Figure 13 with the upmost line replaced by $\{$*lemon*$\}$, $\{$*cherry*$\}$, $\{$*pea*$\}$ and $\{$*corn*$\}$. The set $\{$*lemon*$\}$, for example, is the nonredundant set corresponding to $\{$*lemon*, *yellow*, *fruit*$\}$ because *yellow* and *fruit* belong to every consistently closed set with member *lemon*, as required by the given inheritance network.

To find out what specialization adds up to in the case of nonredundant sets observe that, by definition, X is specialized by Y iff $\mathrm{cl}(X) \subseteq \mathrm{cl}(Y)$, that is, iff $X \subseteq \mathrm{cl}(Y)$, which is the case iff for every member p of X there is a member q of Y such that q ISA* p.

2.2 The Logic of Simple Inheritance

Suppose Σ is a set of concepts. Let us think of a concept as *denoting* a subset of the universe of discourse or as being *satisfied* by certain members of the universe. It does not require much effort to formally subsume the canonical universe $\mathrm{C}(\Gamma)$ of an inheritance network Γ over Σ under this point of view: a subset X of Σ *satisfies* a concept p iff $p \in X$; notation: $X \models p$. By definition, X belongs to $\mathrm{C}(\Gamma)$ iff it holds that $X \models q$ whenever $X \models p$ and p ISA q, and $X \not\models q$ whenever $X \models p$ and p ISNOTA q. Hence $\mathrm{C}(\Gamma)$ could serve as the universe of a model of Γ with satisfaction relation $\models$. What we have done here implicitly is to interpret ISA and ISNOTA in terms of (first-order) predicate logic.[7] In the following we do this explicitly and beforehand.

2.2.1 *Simple Inheritance Theories with Binary Exclusions*

By taking ISA and ISNOTA respectively as marks of subordination and incompatibility, it is straightforward how to rephrase them in predicate logic: typical candidates are regimented phrases like ‘everything which is an A *is a* B’ and ‘everything which is an A *is not a* B’, with the schematic letters ‘A’ and ‘B’ standing for monadic (i.e. one-place) predicates. Formalized in terms of standard predicate logic: $\forall x(Ax \rightarrow Bx)$ and $\forall x(Ax \rightarrow \neg Bx)$, respectively. For convenience, let us abbreviate these statements by ‘A *is-a* B’ and ‘A *is-not-a* B’, thereby introducing two predicate operators ‘*is-a*’ and ‘*is-not-a*’; we refer to such expressions respectively as (*simple*) *inheritance* and (*binary*) *exclusion statements*.

Viewed from this perspective, an inheritance network is a first-order theory Γ consisting of simple inheritance and binary exclusion statements built from members of a certain set Σ of primitive monadic predicates; we speak of a *simple inheritance theory (with binary exclusions)*. One can now *define* relations ISA_Γ and ISNOTA_Γ as follows: p ISA_Γ q iff $\ulcorner p$ *is-a* $q\urcorner \in \Gamma$, and p ISNOTA_Γ q iff

[7] If not otherwise indicated, the terms ‘predicate logic’, ‘quantificational logic’, and ‘first-order logic’ are used interchangeably.

$\ulcorner p$ *is-not-a* $q\urcorner \in \Gamma$, for members p and q of Σ.

(2.10) Remark *(Selective quotation)* A few words about quotation might be in order. The corner quotes indicate *quasi-* or *selective quotation*.[8] If p and q are predicates, i.e. expressions, then $\ulcorner p$ *is-a* $q\urcorner$ is the expression resulting by consecutive concatenation of p, '*is-a*' and q. When selective quotation is employed, it has to be agreed upon by convention that certain letters (or expressions), here 'p' and 'q', *refer to expressions* whereas others, here '*is-a*', *are expressions* to be concatenated. Notice that in contrast to $\ulcorner p$ *is-a* $q\urcorner$ the expression 'p *is-a* q' does *not* consist of the expressions p and q but of the letters 'p' and 'q'. The main purpose of selective quotation is to allow quantification over subexpressions. Keeping the distinction between use and mention in mind we shall omit quotation marks if possible without causing confusion, in order to make the text easier to read. We shall say, for example, that the expression p *is-a* q belongs to a set Γ instead of saying that $\ulcorner p$ *is-a* $q\urcorner$ belongs to Γ. (A more fundamental way to get rid of quotation problems is to introduce appropriate term operations, thereby defining the so-called *term algebra*. We will do this from Chapter 5 onwards.)

2.2.2 *Interpretations and Models*

When viewed as first-order theories, inheritance networks are equipped with the standard set-theoretic semantics: an *interpretation* of Σ consists of a set U, called the *universe*, and an *interpretation* or *denotation function* M taking each member p of Σ to a subset $M(p)$ of U;[9] a *model* of Γ is an interpretation M of Σ such that

$$\begin{aligned} \text{if} \quad \ulcorner p \textit{ is-a } q\urcorner \in \Gamma \quad &\text{then} \quad M(p) \subseteq M(q), \\ \text{if} \quad \ulcorner p \textit{ is-not-a } q\urcorner \in \Gamma \quad &\text{then} \quad M(p) \cap M(q) = \varnothing. \end{aligned}$$

To define such a function M is equivalent to defining a *satisfaction relation* $\models$ that is borne by members of U to those of Σ subject to the following conditions: if $\ulcorner p$ *is-a* $q\urcorner \in \Gamma$ and $x \models p$ then $x \models q$; if $\ulcorner p$ *is-not-a* $q\urcorner \in \Gamma$ and $x \models p$ then $x \not\models q$. In other words, $x \models p$ iff $x \in M(p)$.

The canonical universe $\mathrm{C}(\Gamma)$ of Γ, which was introduced in Section 2.1, can be used as the universe of a model of Γ. The satisfaction relation of this model is defined as indicated at the beginning of Section 2.2: $X \models p$ iff $p \in X$. The corresponding denotation function M_Γ takes each member p of Σ to the set of consistently Γ-closed subsets of Σ with member p:

(2.11)
$$\mathrm{M}_\Gamma(p) \;=\; \{X \in \mathrm{C}(\Gamma) \mid p \in X\}.$$

[8] Cf. e.g. Forbes 1994, Sect. 2.6 or Quine 1951, §§4–6.

[9] We usually refer to an interpretation and its interpretation function by the same symbol.

By definition of $\mathrm{C}(\Gamma)$, this gives a first-order model of Γ, which we call the *canonical model of* Γ.

2.2.3 *A Sound and Complete Inheritance Calculus*

An *inheritance calculus* is a collection of *inference schemes* over the set of inheritance statements. Recall the definition of deductive entailment: a set Γ of statements *deductively entails* a statement α by some inference calculus if there is a *formal proof* of α from Γ by that calculus, i.e. a finite sequence of statements with α as its last element such that each element of the sequence either belongs to Γ or can be inferred from some of its predecessors by one of the inference schemes of the calculus. If an inference scheme has no premises, its conclusion is called an *axiom scheme*.

The following inheritance calculus IC_0 is easily proved to be sound with respect to standard predicate logic.

$$\frac{}{A \textit{ is-a } A}\ \textit{(reflexivity)} \qquad \frac{A \textit{ is-a } B \quad B \textit{ is-a } C}{A \textit{ is-a } C}\ \textit{(transitivity)}$$

$$\frac{A \textit{ is-not-a } B}{B \textit{ is-not-a } A}\ \textit{(symmetry)} \qquad \frac{A \textit{ is-a } B \quad B \textit{ is-not-a } C}{A \textit{ is-not-a } C}\ \textit{(chaining)}$$

That is, if a simple inheritance theory Γ deductively entails an inheritance or exclusion statement α by IC_0 then Γ entails α by any (sound and complete) inference calculus for first-order logic. In particular, for every model M of Γ,

$$\text{if} \quad \Gamma \vdash_{IC_0} p \textit{ is-a } q \quad \text{then} \quad M(p) \subseteq M(q),$$
$$\text{if} \quad \Gamma \vdash_{IC_0} p \textit{ is-not-a } q \quad \text{then} \quad M(p) \cap M(q) = \varnothing,$$

where '$\vdash_{IC_0}$' stands for 'entails by IC_0'.

As usual, the question of completeness is slightly more intriguing. Call an inheritance calculus *(strongly) complete* iff for each simple inheritance theory Γ, any inheritance or exclusion statement first-order provable from Γ is provable from Γ by the calculus in question. The above calculus is easily seen to be *incomplete*. Consider for example the theory over two predicates a and b that consists of the sole exclusion statement a *is-not-a* a. Then both a *is-a* b and a *is-not-a* b are first-order deducible[10] but neither is entailed by Γ. The crucial point here is the presence of a predicate (or concept) which is not satisfiable. This gap can easily be closed by adding the following inference schemes to IC_0 (which are of course instances of *ex falso quodlibet*):

[10] If $\forall(A \to \neg A)$ then $\forall(\neg A \vee \neg A)$ and thus $\forall\neg A$, which implies that $\forall(\neg A \vee B)$ as well as $\forall(\neg A \vee \neg B)$, that is, $\forall(A \to B)$ and $\forall(A \to \neg B)$; see Section 3.1.1 for an explanation of the variable-free notation. In terms of models: if a *is-not-a* $a \in \Gamma$ then $M(a)$ is necessarily the empty set; therefore, $M(a) \subseteq M(b)$ and $M(a) \cap M(b) = \varnothing$.

$$\frac{A \textit{ is-not-a } A}{A \textit{ is-a } B} \qquad \frac{A \textit{ is-not-a } A}{A \textit{ is-not-a } B} \qquad \textit{(inconsistency)}$$

The inheritance calculus *IC* obtained that way turns out to be complete.

One way of proving the completeness of *IC* is to employ the canonical model M_Γ of a simple inheritance theory Γ, which has the following key property:

(2.12)
$$\begin{array}{llll} \text{if} & \mathrm{M}_\Gamma(p) \subseteq \mathrm{M}_\Gamma(q) & \text{then} & \Gamma \vdash_{IC} p \textit{ is-a } q, \\ \text{if} & \mathrm{M}_\Gamma(p) \cap \mathrm{M}_\Gamma(q) = \varnothing & \text{then} & \Gamma \vdash_{IC} p \textit{ is-not-a } q. \end{array}$$

Proof. Let X_p be the set $\{q \mid \Gamma \vdash_{IC} p \textit{ is-a } q\}$. By reflexivity, $p \in X_p$; by transitivity, X_p is Γ-closed. Hence, if X_p is consistent then X_p is the least satisfier $\mathrm{s}(p)$ of p. If X_p is inconsistent, it has members r and s such that $r \textit{ is-not-a } s$ belongs to Γ. Therefore, since Γ entails $p \textit{ is-not-a } r$ and $p \textit{ is-not-a } s$ by definition of X_p, Γ entails $p \textit{ is-not-a } p$ by symmetry and chaining, and thus, by inconsistency, $p \textit{ is-a } q$ and $p \textit{ is-not-a } q$ for every q; in particular, $X_p = \Sigma$. Now suppose $\mathrm{M}_\Gamma(p)$ is empty. Then X_p is inconsistent. So Γ entails $p \textit{ is-a } q$ and $p \textit{ is-not-a } q$ for every q. Correspondingly, if $\mathrm{M}_\Gamma(q)$ is empty then X_q is inconsistent; hence Γ entails $q \textit{ is-not-a } q$ and thus $p \textit{ is-not-a } q$. Finally, if $\mathrm{M}_\Gamma(p)$ is nonempty and $\mathrm{M}_\Gamma(p) \subseteq \mathrm{M}_\Gamma(q)$ then X_p is consistent and $X_p = \mathrm{s}(p) \in \mathrm{M}_\Gamma(q)$, i.e., $q \in X_p$. ⌟

Notice the connection to the canonical inheritance network $\Gamma(\mathrm{C}(\Gamma))$ of $\mathrm{C}(\Gamma)$. By (2.4), $\Gamma(\mathrm{C}(\Gamma))$ has an ISA link from p to q iff $\mathrm{M}_\Gamma(p) \subseteq \mathrm{M}_\Gamma(q)$, that is, by (2.12), iff Γ entails $p \textit{ is-a } q$. Similarly, $\Gamma(\mathrm{C}(\Gamma))$ has an ISNOTA link from p to q iff Γ entails $p \textit{ is-not-a } q$.

(2.13) Theorem Reflexivity, transitivity, symmetry, chaining, and inconsistency constitute a sound and strongly complete inheritance calculus.

Proof. Suppose $p \textit{ is-a } q$ is deducible from Γ by first-order logic. Since M_Γ is a model of Γ, it follows by soundness of first-order logic that $\mathrm{M}_\Gamma(p) \subseteq \mathrm{M}_\Gamma(q)$. Thus, according to (2.12), $\Gamma \vdash_{IC} p \textit{ is-a } q$. Similar for $p \textit{ is-not-a } q$. ⌟

2.2.4 *Discussion*

There has been some debate about the adequacy of translating inheritance networks into first-order theories. Pat Hayes (1979) is a well-known proponent; Thomason, Horty, and Touretzky (1987) on the other hand speak of a "mistaken folk theorem".[11] Their argument is based on a counterexample like ours

[11] See also Carpenter and Thomason 1990 or Horty 1994.

above for showing the incompleteness of the calculus IC_0. It differs insofar as it uses a singular concept, that is, a *name* of single entity, whose classification within the network gives rise to contradictions. According to Thomason, Horty, and Touretzky, the inconsistency rule is inadequate for handling contradictory information.

From our point of view an inheritance network is an assumably true theory about the entities of a certain domain; it can serve as background knowledge for information processing tasks. An inconsistent theory then simply is a false theory, which has to be revised. (Labeling some hypotheses explicitly as defeasible is a different thing.) It is nevertheless possible from this perspective that in the course of an actual application of this knowledge as part of some act of perception, an entity might get classified in an inconsistent manner, due to contradictory evidence (which e.g. could be due to some spoiled reception process). For handling such information, the framework of *relevance logic*,[12] for example, could probably be put to use. It is of course a question of reliability whether to revise the actual perceptual information or the knowledge base.

Another objection might question not the adequacy of the reduction to predicate logic but its necessity. To some it may be more in accordance with standard practice to define a denotational semantics for inheritance statements without explicit translation into an interpreted language: p *is-a* q holds in M iff $M(p) \subseteq M(q)$, and similar for p *is-not-a* q. Soundness and completeness of an inheritance calculus is then defined with respect to this semantics. The "Truth Lemma" (2.12) again does the work: If p *is-a* q holds in every model of Γ then in particular in the canonical model M_Γ, which implies that Γ entails p *is-a* q. Notice though that this definition is not an arbitrary stipulation but does implicitly rest on the intended meaning of *is-a* and *is-not-a*. A reformulation of *is-a* and *is-not-a* within predicate logic simply makes this explicit.

2.3 Inheritance and Linguistic Classification

Let us ask to what degree simple inheritance with binary exclusions is able to cover the examples of linguistic classification presented in Chapter 1. Consider taxonomic trees: simple inheritance corresponds to the hyponym (or daughter) relation whereas binary exclusions allow to express that co-hyponyms (or sisters) are incompatible. Exhaustiveness, on the other hand, escapes a formalization by simple inheritance because disjunction is not available. Notice that (without exhaustiveness) every (nonempty) consistently closed set of a taxonomic tree has exactly one maximally specific member; vice versa, every primitive concept of the tree is the most specific member of a consistently closed set (cf. Section 1.4).

To illustrate these observations by a simple example, consider the subclas-

[12] See e.g. Dunn 1986.

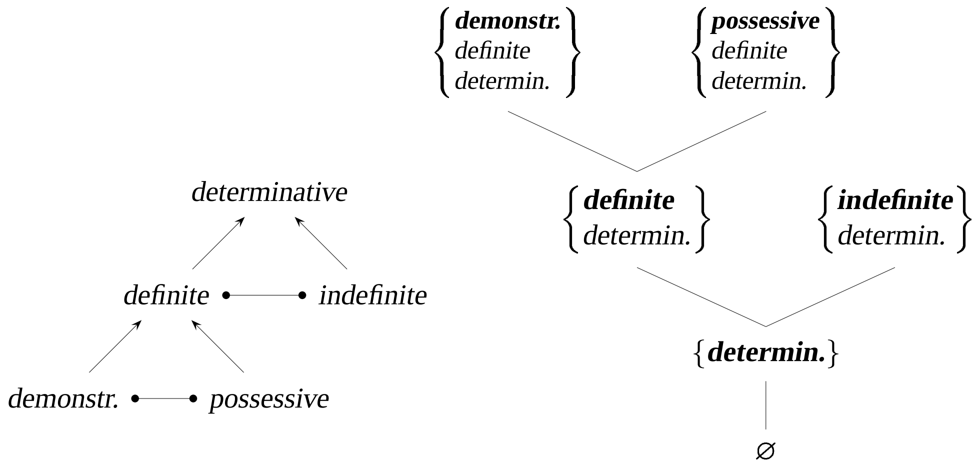

FIGURE 14 Simple inheritance tree and its canonical universe

sification of determinative pronouns in the taxonomic tree of Figure 3. The corresponding inheritance network and its canonical universe are respectively depicted on the left and the right of Figure 14, with maximally specific elements in boldface.

As in the case AND/OR trees and systemic networks, simple inheritance and exclusion statements again allow to express subordination and incompatibility. More precisely, since subordination in systemic networks means that the elements of a choice system P imply the entry condition of P, the latter must not contain disjunctions in order to be expressible by simple inheritance. (Conjunctions, in contrast, pose no problem because constraints with conjunctive conclusion can be split into simple inheritance statements: φ implies $\psi_1 \wedge \psi_2$ just in case φ implies ψ_1 and ψ_2; cf. Section 3.3.1.) Recall from Section 1.3 that an exhaustive classification by a systemic network requires in addition that the disjunction of the elements of P is implied by the entry condition of P. This is clearly beyond the scope of simple inheritance. Due to the lack of conjunction, even the restricted form of exhaustiveness mentioned at the close of Section 1.4 is not available: it is not possible to express by simple inheritance statements that a concept, e.g. *child*, is a defined concept, i.e., is implied by the conjunction of its immediate superordinates, here *immature* and *human* (cf. Figure 6).

In spite of these restrictions, Carpenter and Pollard (1991) consider simple inheritance with binary exclusion as adequate for capturing “the logic of linguistic classification”. In order to cope with defined concepts and closed world reasoning they stipulate additional modifications of the canonical universe of the given simple inheritance theory. We briefly review their approach by means of the following example, which is essentially Carpenter and Pollard’s.

(2.14) Example Suppose Σ is a set of (primitive) concepts a, b, c, d, e, f. Let Γ be the simple inheritance theory over Σ that is graphically depicted on the lower left of Figure 15; the canonical universe $\mathrm{C}(\Gamma)$ of Γ is shown on the right, with maximally specific elements in boldface. In the terminology of Carpenter and Pollard, the elements of $\mathrm{C}(\Gamma)$ are the *conjunctive concepts* of the simple inheritance theory Γ.[13] The framed elements of $\mathrm{C}(\Gamma)$ are the least satisfiers of members of Σ; they can be read off from the specialization diagram of $\mathrm{C}(\Gamma)$ as those elements with exactly one element immediately below them (cf. Section 4.2.1).

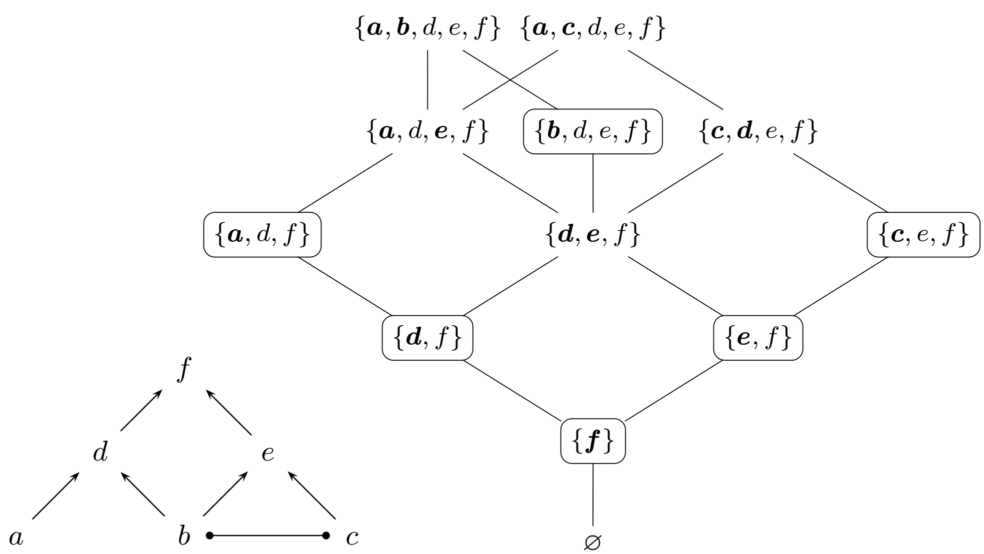

FIGURE 15 Simple inheritance network and its canonical universe

Recall from Section 1.4 that a concept is called *defined* if it is implied by the conjunction of its immediate superordinates. In the above example, the concept b is a potential candidate for a defined concept; to consider b as defined means to assume that b is implied by the conjunction of d and e. (If this is the case then b is *equivalent* to the conjunction of d and e because the implication of d and e by b is already part of the simple inheritance theory.) Consequently, if b is regarded as defined, every entity type satisfying d and e should also satisfy b; hence every consistently closed set containing d and e should also contain b. In other words, taking b as defined disqualifies certain conjunctive concepts as being consistently closed. Concretely, if the conjunction of d and

[13] In Carpenter and Pollard 1991, the specialization relation is borne by more specific elements to less specific ones, which is dual to our definition (and that of Carpenter 1992, Chap. 2).

e implies b, the sets $\{d, e, f\}$, $\{a, d, e, f\}$, $\{c, d, e, f\}$, and $\{a, c, d, e, f\}$ have to be eliminated from the canonical universe displayed in Figure 15.

To take this sort of constraint into account, Carpenter and Pollard introduce an additional condition a conjunctive concept X must satisfy: if p is a defined concept and every q with p ISA q belongs to X then so does p. However, this approach blurs the simple picture that entity types are members of the canonical universe of a certain type of (first-order) theory. We shall see in Chapter 3 how to extend this picture straightforwardly from simple inheritance theories to *Horn theories*, where constraints involving conjunctions are part of the theory. The additional condition given by Carpenter and Pollard will then become a consequence of the definition of the canonical model of a Horn theory.

Let us finally turn to closed world reasoning. Carpenter and Pollard call a concept *closed*, if it implies the disjunction of its immediate subconcepts. In order to handle constraints with disjunctions, they introduce *disjunctive concepts* as sets of conjunctive concepts that are (upwards) closed with respect to specialization. That is, a subset $\mathcal{X}$ of $\mathrm{C}(\Gamma)$ is a disjunctive concept if whenever $X \in \mathcal{X}$ and $X \subseteq Y \in \mathrm{C}(\Gamma)$ then $Y \in \mathcal{X}$. Clearly the disjunctive concepts form a distributive subset lattice over $\mathrm{C}(\Gamma)$. Moreover notice that specialization between disjunctive concepts is *reverse* set inclusion; for disjunction behaves dually to conjunction with respect to specialization.

In the presence of closed concepts, every disjunctive concept $\mathcal{X}$ is required to satisfy the additional condition that, for every $X \in \mathcal{X}$, if $p \in X$ is closed then there is a $q \in X$ such that q ISA p. We illustrate this definition by a simplistic example (which is again Carpenter and Pollard's):

(2.15) Example Consider the minimalistic taxonomy presented on the right of Figure 16 in form of an inheritance network; it says that words and phrases are (linguistic) signs and nothing is both a word and a phrase. Its canonical universe is shown in the middle, with the conjunctive concepts represented by nonredundant concept sets (i.e. with $\{\mathit{word}\}$ and $\{\mathit{phrase}\}$ instead of $\{\mathit{word}, \mathit{sign}\}$ and $\{\mathit{phrase}, \mathit{sign}\}$). The diagram on the right of Figure 16 depicts the corresponding lattice of disjunctive concepts, where the framed elements indicate the embedding of the canonical universe into this lattice. Treating the concept *sign* as closed, in contrast, gives rise to the lattice displayed in the middle of Figure 17. For, according to the above definition, the conjunctive concept $\{\mathit{sign}\}$ is disqualified as a member of a disjunctive concept since $\{\mathit{sign}\}$ contains a closed concept but none of its immediate subconcepts.

Let us notice two things. First, it is obviously not necessary to employ disjunctive concepts in order to specify the effect of disjunctive constraints on conjunctive concepts. Indeed, the condition given by Carpenter and Pollard is a constraint on conjunctive concepts in the first place: $\{\mathit{sign}\}$ does not sat-

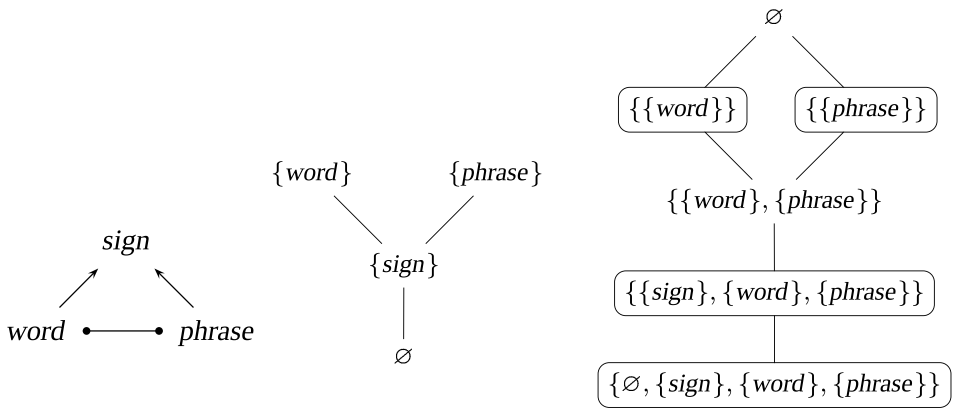

FIGURE 16 Words and phrases

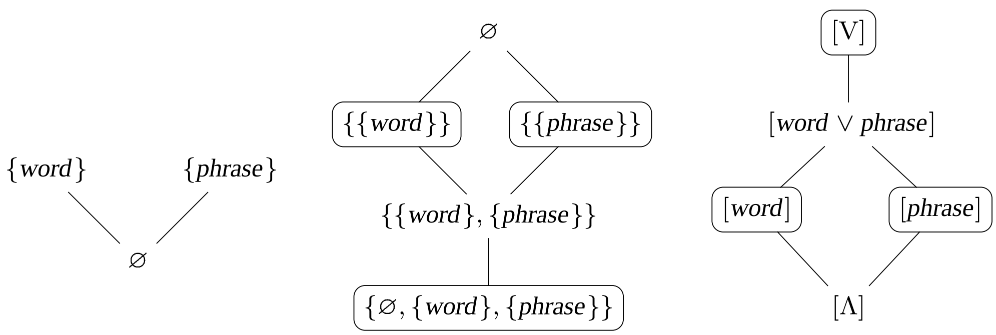

FIGURE 17 More words and phrases

isfy the constraint that *sign* implies the disjunction of *word* and *phrase*. Hence the canonical universe of the simple inheritance theory of Figure 16 extended by this disjunctive constraint is as depicted on the left of Figure 17, whereas the lattice of disjunctive concepts in the middle is just the system of upwards closed sets of this canonical universe. In Chapter 5, we will make precise in what sense a theory with disjunctive statements determines a canonical universe.

The second observation is concerned with the status of disjunctive concepts. As said before, we can do without them if we are interested in the generic entities of a classificational theory, be it a simple inheritance theory or one with conjunctive and disjunctive statements. There are two ways disjunctive concepts can be made use of. One is the modeling of non-determinism: disjunctive concepts represent the sets of entity types that are "informationally" different with respect to the given theory. Since we do not consider non-deterministic processing, we will not pursue this approach further (but see (2.16) below).

The other option is to take disjunctive concepts as *logical* disjunctions of conjunctions of primitive concepts (or predicates). Then $\{\{word\}, \{phrase\}\}$ is for *word* $\vee$ *phrase* (whereas $\{\{word, phrase\}\}$ is for *word* $\wedge$ *phrase*). From this perspective, the distributive lattice of disjunctive concepts can be identified with the order dual of the *Lindenbaum algebra* of the given theory: each disjunctive concept represents a set of predicates which are pairwise equivalent with respect to the theory in question; see the diagram on the right of Figure 17 (where the predicates V and Λ are respectively satisfied by everything and nothing in the universe of discourse). Notice that one can recover the entity types from this lattice by picking out all elements with exactly one element immediately below them. See Chapters 6 and 7 for a detailed account of these matters.

(2.16) Remark *(Upper power domain)* In the previous discussion, the set Σ of primitive concepts was implicitly presumed to be *finite*. Suppose Γ is a set of statements with finite conjunctions and disjunctions (i.e. an observational theory in the sense of Chapter 5). The indicated modeling of non-determinism employs the so-called *upper (or Smyth) power domain* construction. We have seen that for finite Σ, the upper power domain of $\mathrm{C}(\Gamma)$ can be identified with the order dual of the Lindenbaum algebra $\mathrm{L}(\Gamma)$ of Γ, where the latter can be constructed by taking the upwards closed subsets of $\mathrm{C}(\Gamma)$ ordered by set inclusion. For infinite Σ, however, the relation between the power domain and the Lindenbaum algebra is more complicated. If we restrict ourselves to finitistic theories (see Chapter 7), the situation is as follows: The canonical universe $\mathrm{C}(\Gamma)$ is a coherent algebraic domain, the Lindenbaum algebra $\mathrm{L}(\Gamma)$ is given by the upwards closed subsets of $\mathrm{C}(\Gamma)$ that are generated by finite sets of compact elements, and the upper power domain of $\mathrm{C}(\Gamma)$ is the ideal completion of the order dual of $\mathrm{L}(\Gamma)$. (For a general definition of the upper power domain, the reader is referred to Vickers 1989, Gunter and Scott 1990, or Abramsky and Jung 1994. It should be added that in the latter two references, the empty set is omitted from the upper power domain.)

3

Horn Theories

Whereas Carpenter and Pollard (1991) formalize lexical hierarchies and systemic networks by simple inheritance networks, it is our rationale to express the latter as well as the former within the framework of observational logic, which can be seen as a sublanguage of standard predicate logic. Observational logic in general will be treated in Part II; the present chapter is confined to statements without disjunction – so-called Horn statements;

In Section 3.1, we reformulate inheritance statements as Horn statements and introduce a standard model-theoretic semantics for the latter. The canonical model of a Horn theory is defined in Section 3.2. We describe it as a subset system over the set of primitive predicates and characterize the subset systems that arise that way. Section 3.3 presents a sound and complete inference calculus for Horn statements. In Section, 3.4 we apply the framework to defined concepts in simple inheritance theories; in addition, we generalize binary exclusion statements to finitary ones.

3.1 From Inheritance to Horn

Recall from Section 2.2 that simple inheritance theories consist of universally quantified conditionals of the form

$$\forall x(px \rightarrow qx) \qquad \text{and} \qquad \forall x(px \rightarrow \neg qx),$$

where p and q are (primitive) monadic predicates. In Section 2.2, the two predicate operators '*is-a*' and '*is-not-a*' have been introduced for representing both types of statements in a variable free form that resembles closely the original notation of simple inheritance networks.

3.1.1 *Inheritance Statements Reformalized*

In the following we make use of the predicate operators '$\preceq$', '$\neg$', '$\wedge$', '$\rightarrow$', '$\forall$' etc. The operator '$\wedge$', for instance, takes two monadic predicates 'A' and 'B' to the monadic predicate '$\{x \mid Ax \wedge Bx\}$'; correspondingly, '$\neg A$' stands for '$\{x \mid \neg Ax\}$'. The notation '$\{x \mid \ldots x \ldots\}$' indicates *predicate abstraction*; it can be read as 'an x such that $\ldots x \ldots$'.[1] (It is important to keep in mind that predicate abstraction is a purely syntactic issue – though the notation may evoke set-theoretic associations.) The operator '$\forall$' takes a monadic predicate 'A' to the statement '$\forall x(Ax)$', whereas '$A \preceq B$' stands for '$\forall(A \rightarrow B)$', that is, for '$\forall x(Ax \rightarrow Bx)$'.

Using the predicate operators '$\preceq$' and '$\neg$', inheritance statements can be written in the form '$p \preceq q$' and '$p \preceq \neg q$'. Alternatively, one can introduce a "constant" monadic predicate (a "zero-place operator") 'Λ' that denotes nothing in any universe of discourse. Then '$p \wedge q \preceq \Lambda$' can take the place of '$p \preceq \neg q$'. The sole remnant of negation is now incorporated into the predicate 'Λ'. Put differently, 'Λ' is equivalent to '$\neg \mathrm{V}$', where 'V' is a monadic predicate denoting everything of the domain of discourse in question. The attentive reader might object that introducing 'V' and 'Λ' means to lapse back to the attitude criticized in Section 2.2.4 of stipulating uninterpreted expressions. However, one can *define* 'V' as '$\{x \mid x = x\}$', and consequently 'Λ' as '$\{x \mid \neg(x = x)\}$', thereby working within *predicate logic with identity*.[2]

3.1.2 *Horn Statements and Horn Theories*

By taking conjunction as an unrestricted operator, the above reformulation of simple inheritance statements gives rise to statements of the general form $\varphi \preceq \psi$, where φ and ψ are arbitrary *conjunctive* predicates over a given set Σ of *primitive* monadic predicates; that is, φ and ψ are inductively built by finite conjunction from members of $\Sigma \cup \{\mathrm{V}, \Lambda\}$. We speak of such statements as (*universal monadic*) *Horn statements* over Σ.[3] The predicates V and Λ are in the following always assumed not to be members of Σ. A *Horn theory* over Σ is a set of Horn statements over Σ.

For Horn theories we have the standard notions of first-order interpretation and model: a set-theoretic *interpretation* (or interpretation in the category $\mathbf{Set}$ of sets, or $\mathbf{Set}$-valued interpretation) of the vocabulary Σ consists of a set U, the *universe*, and an *interpretation* or *denotation function* M that takes each member of Σ to a subset of U. The function M can be extended to V and Λ by taking them to U and $\varnothing$, respectively, and inductively to conjunctive predicates:

(3.1) $$M(\varphi \wedge \psi) \;=\; M(\varphi) \cap M(\psi).$$

[1] Cf. Quine 1982, Sect. 21.

[2] The abbreviations 'V' and 'Λ' for the predicate abstracts '$\{x \mid x = x\}$' and '$\{x \mid \neg(x = x)\}$' are used e.g. in Quine 1951.

[3] Named after Alfred Horn. It is easily seen that Horn statements are logically equivalent to finite conjunctions of *universal Horn clauses*, i.e. of statements of the form $\forall(\alpha_1 \vee \ldots \vee \alpha_n)$, where each α_i is either of the form p or $\neg p$, with p primitive, and at most one of the α_i's is a negated primitive.

A statement $\varphi \preceq \psi$ *holds* or *is true with respect to* M iff $M(\varphi) \subseteq M(\psi)$. A *model* of a Horn theory is an interpretation of its vocabulary with respect to which each of its statements is true. As mentioned in Section 2.2, an interpretation function can also be defined in terms of *satisfaction*: $x \models \varphi$ iff $x \in M(\varphi)$. In particular, (3.1) becomes:

(3.2) $$x \models \varphi \wedge \psi \quad \text{iff} \quad x \models \varphi \quad \text{and} \quad x \models \psi.$$

Furthermore, every member of the universe satisfies V, and none satisfies Λ.

3.2 The Canonical Universe

3.2.1 *Canonical Model and Universe*

Let Γ be a Horn theory over a set Σ of primitive monadic predicates. Our primary interest is again in the set of entity types determined by Γ. As before, these entities can be thought of as consistently closed sets of primitive predicates. Satisfaction is defined as in Section 2.1: a subset X of Σ satisfies a member p of Σ iff $p \in X$. To repeat the inductive extension of satisfaction to conjunctive predicates: every subset of Σ satisfies V, no subset satisfies Λ, and a subset satisfies $\varphi \wedge \psi$ iff it satisfies φ and ψ. An immediate consequence is that two subsets X and Y of Σ are identical iff they satisfy the same predicates over Σ; in terms of set inclusion:

(3.3) $$X \subseteq Y \qquad \text{iff} \qquad \forall \varphi (X \models \varphi \rightarrow Y \models \varphi).$$

Proof. $X \subseteq Y$ just in case $p \in Y$ whenever $p \in X$, i.e., iff Y satisfies a member p of Σ whenever X satisfies p. The case of V and Λ is trivial since $X \not\models \Lambda$ and $Y \models \mathrm{V}$. We proceed by term induction: Suppose $X \subseteq Y$. If $X \models \varphi \wedge \psi$ then X satisfies φ and ψ, which both, by induction hypothesis, are satisfied by Y; hence $Y \models \varphi \wedge \psi$. ⌟

We call a subset X of Σ *(consistently) closed with respect to* $\varphi \preceq \psi$ if X satisfies ψ whenever X satisfies φ. In particular, X is consistently closed with respect to $\varphi \preceq \Lambda$ iff X does not satisfy φ. Let $\mathrm{C}(\Gamma)$ be the set of those subsets of Σ that are consistently closed with respect to every statement of the theory Γ. It follows from the definition of consistently closed sets that $\mathrm{C}(\Gamma)$ is the universe of a model M_Γ of Γ, whose interpretation function takes each member p of Σ to the set of consistently Γ-closed sets satisfying p, i.e.

$$\mathrm{M}_\Gamma(p) \;=\; \{X \in \mathrm{C}(\Gamma) \mid p \in X\}.$$

The model M_Γ is henceforth called the *canonical model of* Γ, its universe $\mathrm{C}(\Gamma)$ is referred to as the *canonical universe of* Γ. As in Section 2.1.1, we speak of set inclusion on $\mathrm{C}(\Gamma)$ as *specialization*.

(3.4) Remark Satisfaction of V can be handled by membership too. Then, however, only subsets of $\Sigma\cup\{\mathrm{V}\}$ containing V are permitted since every subset is required to satisfy V. Consequently V belongs to every member of $\mathrm{C}(\Gamma)$ in this case.

3.2.2 *A Universal Property*

Suppose Γ is a Horn theory over Σ and M is a model of Γ with universe U. Let ε_M be the function that takes each member x of U to the set of those members of Σ which are satisfied by x:

$$\varepsilon_M(x) \;=\; \{p\in\Sigma \mid x\in M(p)\}.$$

The function ε_M has the following property: for every $x\in U$ and every conjunctive predicate φ over Σ,

(3.5) $$x\in M(\varphi) \qquad \text{iff} \qquad \varepsilon_M(x)\in \mathrm{M}_\Gamma(\varphi).$$

Proof. We first show by term induction that $x\in M(\varphi)$ iff $\varepsilon_M(x)\models\varphi$. In case φ belongs to $\Sigma\cup\{\mathrm{V},\Lambda\}$, there is nothing to show. By (3.1) and induction hypothesis, $x\in M(\varphi\wedge\psi)$ iff $\varepsilon_M(x)$ satisfies φ and ψ and hence $\varphi\wedge\psi$. Now suppose $\varphi\preceq\psi$ belongs to Γ and $\varepsilon_M(x)\models\varphi$. Then $x\in M(\varphi)\subseteq M(\psi)$, and thus $\varepsilon_M(x)\models\psi$. Therefore, $\varepsilon_M(x)$ is consistently Γ-closed. ⌟

In particular, ε_M is a homomorphism of models from M to M_Γ.[4] A further consequence of (3.5) is that the canonical model is *universal* in the sense that if $\mathrm{M}_\Gamma(\varphi)\subseteq\mathrm{M}_\Gamma(\psi)$ then $M(\varphi)\subseteq M(\psi)$ for every model M of Γ; for if $x\in M(\varphi)$ then $\varepsilon_M(x)\in\mathrm{M}_\Gamma(\varphi)\subseteq\mathrm{M}_\Gamma(\psi)$ and hence $x\in M(\psi)$. This universal property is equivalent to the following condition, where $\vdash$ is the entailment relation given by any sound and complete inference calculus for first-order logic with identity:

$$\text{if} \quad \mathrm{M}_\Gamma(\varphi)\subseteq\mathrm{M}_\Gamma(\psi) \quad \text{then} \quad \Gamma\vdash\ \varphi\preceq\psi.$$

Predicates that are coextensive with respect to the canonical model of Γ are thus first-order equivalent modulo Γ; in short, the canonical model satisfies *equivalence of coextensives*.

Let us call a model of Γ *generic* if it is isomorphic to the canonical model of Γ. By (3.3), any two generic entities, i.e. any two members of the generic universe, are identical whenever they satisfy the same set of predicates. So,

[4] A *homomorphism of models* from M to N is a function f from universe to universe such that $f(M(p))\subseteq N(p)$, i.e. if $x\models p$ then $f(x)\models p$, for every member p of Σ.

"the" generic model can be said to satisfy *identity of indiscernibles*. Now observe that, as a consequence of (3.5), a Γ-model M satisfies identity of indiscernibles just in case ε_M is one-to-one, i.e. is an *embedding of models*. In other words, the generic universe of Γ is "as large as possible" subject to the condition that indiscernibles are identical. (Concerning the question as to what degree equivalence of coextensives determines genericity, the reader is referred to Osswald 2003b.)

3.2.3 *Characteristic Functions and Truth Values*

Before we turn to the ordering structure on the canonical universe given by specialization, i.e. by set inclusion, let us notice that there are other canonical ways to define a universe with the desired properties. Because of the one-to-one correspondence between subsets of Σ and their *characteristic functions*, i.e. functions with arguments in Σ and values in a fixed set of two elements, the latter can equally well be employed for representing entity types. In particular, if one takes the two values to be the subsets of a singleton set, that is, to be either the empty set or the singleton set itself, then the characteristic functions are interpretations of Σ with the singleton set as its universe. Members of $\mathrm{C}(\Gamma)$ then correspond to those characteristic functions that are (first-order) models of Γ. A possible view on this representation is that each entity type is the sole inhabitant of a universe of discourse satisfying exactly those predicates that characterize that entity type.

It is probably more in accordance with standard practice to regard the two possible values as *truth values* and the characteristic functions as *truth-valued interpretations* and *models*, respectively. This however means to switch from predicate logic to propositional logic; members of Σ are now treated as statements and not as predicates. Conceptually, the truth-valued models represent different "possible worlds", in which certain members of Σ are true, subject to the statements of Γ. The canonical universe of Γ therefore consists of all possible worlds compatible with Γ. This close connection to propositional models is of course due to the fact that the logic of universal monadic Horn statements coincides with the logic of propositional Horn statements.

3.2.4 *Inductive Intersection Systems and Least Satisfiers*

As with simple inheritance theories, the canonical universe of Horn theories turns out to have the property that every nonempty subset has an infimum, i.e. a greatest lower bound, which is given by intersection. Likewise, suprema do not exist in general because of possible inconsistencies. If existent, however, suprema need not necessarily be given by set union, in contrast to simple inheritance – cf. (2.3). For it can happen that the antecedent of a Horn statement is satisfied by the union $\bigcup \mathcal{S}$ of a set $\mathcal{S}$ of closed sets but not by any member of $\mathcal{S}$. On the other hand, suprema of *upwards directed* sets exist and coincide with union. These facts are consequences of the following observation:

(3.6) Lemma Suppose $\mathcal{S}$ is a subset system over Σ and φ is a conjunctive predicate over Σ.

(i) If $\mathcal{S}$ is nonempty then $\bigcap\mathcal{S} \models \varphi$ iff $X \models \varphi$ for every member X of $\mathcal{S}$.

(ii) If $\mathcal{S}$ is upwards directed then $\bigcup\mathcal{S} \models \varphi$ iff $X \models \varphi$ for some member X of $\mathcal{S}$.

Proof. By term induction. (i) If $\mathcal{S}$ is nonempty then it is not the case that every member of $\mathcal{S}$ satisfies Λ. Thus $\bigcap\mathcal{S} \models \Lambda$ iff every member of $\mathcal{S}$ satisfies Λ. Furthermore, $\bigcap\mathcal{S} \models p \in \Sigma$ iff $p \in \bigcap\mathcal{S}$, that is, iff $p \in X$ for every member X of $\mathcal{S}$. By induction hypothesis, $\bigcap\mathcal{S} \models \varphi \wedge \psi$ iff every member X of $\mathcal{S}$ satisfies φ and ψ and thus $\varphi \wedge \psi$. (ii) $\bigcup\mathcal{S} \models p \in \Sigma$ iff $p \in \bigcup\mathcal{S}$, that is, iff $p \in X$ for some member X of $\mathcal{S}$. By induction hypothesis, $\bigcup\mathcal{S} \models \varphi \wedge \psi$ iff for some members X and Y of $\mathcal{S}$, $X \models \varphi$ and $Y \models \psi$. If $\mathcal{S}$ is upwards directed, there is a member Z of $\mathcal{S}$ including X and Y, which satisfies $\varphi \wedge \psi$ according to (3.3). ⌟

A subset system over Σ that is closed with respect to intersection of nonempty subsets is called an *intersection system over* Σ. A subset system is called *inductive* if it is closed with respect to union of upwards directed subsets.

(3.7) Theorem The canonical universe of a Horn theory over Σ is an inductive intersection system over Σ.

Proof. Let Γ be a Horn theory over Σ and $\mathcal{S}$ a nonempty subset of $\mathrm{C}(\Gamma)$. Suppose $\varphi \preceq \psi$ belongs to Γ. By (3.6)(i), if $\bigcap\mathcal{S} \models \varphi$ then every member X of $\mathcal{S}$ satisfies φ and hence ψ; thus $\bigcap\mathcal{S} \models \psi$. A similar argument applies to directed union. ⌟

Consequently, every nonempty subset of $\mathrm{C}(\Gamma)$ has an infimum in $\mathrm{C}(\Gamma)$. In particular, if $\mathrm{C}(\Gamma)$ is nonempty, it has a least element.[5] Notice that this least element is not necessarily the empty set: in case Γ has statements of the form $\mathrm{V} \preceq \varphi$, every member of $\mathrm{C}(\Gamma)$ is required to satisfy φ. Furthermore, every subset $\mathrm{C}(\Gamma)$ with upper bounds in $\mathrm{C}(\Gamma)$ has a least upper bound in $\mathrm{C}(\Gamma)$, namely the intersection of all upper bounds. It should be stressed once more that in general this least upper bound is not given by set union.

(3.8) Example *(Non-distributivity)* Suppose Σ is a set of three primitive predicates a, b, and c. Let Γ be the theory over Σ consisting of the statements

[5] The canonical universe $\mathrm{C}(\Gamma)$ is empty just in case Γ entails $\mathrm{V} \preceq \Lambda$.

$a \preceq b$ and $b \wedge c \preceq a$ and let Γ' be the theory with statements $a \wedge b \preceq c$, $a \wedge c \preceq b$, and $b \wedge c \preceq a$. The respective canonical universes are depicted in Figure 18. They both are non-distributive subset lattices. Notice that in both cases the least upper bound of $\{\{b\}, \{c\}\}$ is not $\{b, c\}$ but $\{a, b, c\}$.

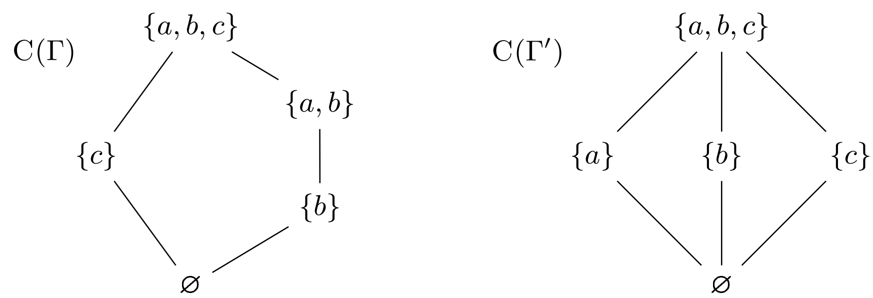

FIGURE 18 Non-distributive canonical universes

We will see below that Horn theories are characterized by (3.7) in much the same way as simple inheritance networks are characterized by (2.3). In order to show this, as well as for later use, we introduce the least satisfier of a satisfiable predicate (cf. Section 2.1). Suppose $\mathcal{U}$ is an inductive intersection system over Σ. If a conjunctive predicate φ over Σ is *satisfiable in* $\mathcal{U}$, i.e., if φ is satisfied by some member of $\mathcal{U}$, the set $\{X \in \mathcal{U} \mid X \models \varphi\}$ is nonempty. Its intersection is therefore a member of $\mathcal{U}$, called the *least satisfier* $\mathrm{s}(\varphi)$ *of* φ *in* $\mathcal{U}$:

(3.9)
$$\mathrm{s}(\varphi) \; = \; \bigcap\{X \in \mathcal{U} \mid X \models \varphi\}.$$

Consequently, for all members X of $\mathcal{U}$,

(3.10)
$$X \models \varphi \qquad \text{iff} \qquad \mathrm{s}(\varphi) \subseteq X.$$

(3.11) Lemma Let $\mathcal{U}$ be an inductive intersection system over Σ. Then, for each subset X of Σ satisfying only predicates that are satisfiable in $\mathcal{U}$, the set $\{\mathrm{s}(\varphi) \mid X \models \varphi\}$ is upwards directed in $\mathcal{U}$ and

$$X \; \subseteq \; \bigcup\{\mathrm{s}(\varphi) \mid X \models \varphi\}.$$

Moreover, X belongs to $\mathcal{U}$ iff $X = \bigcup\{\mathrm{s}(\varphi) \mid X \models \varphi\}$.

Proof. By (3.10) and (3.2), $\mathrm{s}(\varphi \wedge \psi)$ includes both $\mathrm{s}(\varphi)$ and $\mathrm{s}(\psi)$. So the set $\{\mathrm{s}(\varphi) \mid X \models \varphi\}$ is upwards directed. Its union includes X, by (3.3). The last claim of the Lemma follows by (3.10). ⌟

Let us now prove the inverse of (3.7), as promised above. If there is a Horn theory Γ whose canonical universe is identical to a given inductive intersection system $\mathcal{U}$ then, by definition, $\mathrm{M}_\Gamma(\varphi) = \{X \in \mathcal{U} \mid X \vDash \varphi\}$. A Horn statement $\varphi \preceq \psi$ thus holds in M_Γ just in case every member of $\mathcal{U}$ satisfying φ satisfies ψ. In other words, every Horn theory Γ over Σ with canonical universe $\mathcal{U}$ is equivalent to the theory $\Gamma(\mathcal{U})$ over Σ which is defined as follows:

(3.12) $\quad \ulcorner\varphi \preceq \psi\urcorner \in \Gamma(\mathcal{U}) \quad \text{iff} \quad \forall X \in \mathcal{U}\,(X \vDash \varphi \rightarrow X \vDash \psi).$

In case φ is not satisfiable, (3.12) gives rise to statements $\varphi \preceq \psi$ for every conjunctive predicate ψ. In fact, φ is not satisfiable iff $\varphi \preceq \Lambda$ belongs to $\Gamma(\mathcal{U})$. If φ is satisfiable then, by (3.6)(i), we can put (3.12) more succinctly as:

(3.13) $\quad \ulcorner\varphi \preceq \psi\urcorner \in \Gamma(\mathcal{U}) \quad \text{iff} \quad \mathrm{s}(\varphi) \vDash \psi.$

We want to show that every inductive intersection system $\mathcal{U}$ is the canonical universe of the Horn theory $\Gamma(\mathcal{U})$, i.e. $\mathcal{U} = \mathrm{C}(\Gamma(\mathcal{U}))$. Obviously, $\mathcal{U} \subseteq \mathrm{C}(\Gamma(\mathcal{U}))$. Now suppose X is a member of $\mathrm{C}(\Gamma(\mathcal{U}))$. Since X does not satisfy any predicate not satisfiable in $\mathcal{U}$, it follows by (3.11) that $\bigcup\{\mathrm{s}(\varphi) \mid X \vDash \varphi\}$ belongs to $\mathcal{U}$ and includes X. Furthermore, by (3.13) and (3.3), $\mathrm{s}(\varphi) \subseteq X$ whenever $X \vDash \varphi$. Thus, $X = \bigcup\{\mathrm{s}(\varphi) \mid X \vDash \varphi\} \in \mathcal{U}$. All in all, $\mathrm{C}(\Gamma(\mathcal{U})) = \mathcal{U}$. In particular:

(3.14) Theorem Every inductive intersection system over Σ is the canonical universe of a Horn theory over Σ.

3.2.5 *Closure Systems*

An intersection system $\mathcal{U}$ over Σ is called a *closure system* if Σ belongs to $\mathcal{U}$.[6] Let us consider the theory $\Gamma(\mathcal{U})$ determined by an inductive closure system $\mathcal{U}$ as defined by (3.12). Observe that a conjunctive predicate φ over Σ is satisfied by Σ just in case φ is Λ-*free*, i.e. contains no occurrence of Λ. Hence, according to (3.12) and since $\Sigma \in \mathcal{U}$, no statement $\varphi \preceq \psi$ of $\Gamma(\mathcal{U})$ is such that φ but not ψ is Λ-free. Moreover, removing any statement from Γ whose antecedent is not Λ-free does not affect the canonical universe. So every inductive closure system is the canonical universe of a Λ-free Horn theory. On the other hand, Σ belongs to the canonical universe of every Λ-free Horn theory over Σ. Put together:

(3.15) Theorem The canonical universe of a Λ-free Horn theory over Σ is an inductive closure system over Σ, and every such subset system over Σ arises that way.

[6] With $\bigcap \varnothing = \Sigma$, a closure system is a subset system that is closed with respect to arbitrary intersection; cf. footnote 3 on page 25

Suppose $\mathcal{U}$ is a closure system over Σ. Then for each subset Y of Σ there is a least member of $\mathcal{U}$ with subset Y, the *closure* $\mathrm{cl}(Y)$ *of* Y *in* $\mathcal{U}$, which is given by intersection:

$$\text{\textbf{\textit{(3.16)}}} \qquad \mathrm{cl}(Y) \;=\; \bigcap\{X \in \mathcal{U} \mid Y \subseteq X\}.$$

One straightforwardly verifies that cl is a *closure operation* on Σ.[7] It is well known, and easy to show, that the image of each closure operation on Σ is a closure system over Σ, and that this relationship between closure systems and closure operations is one-to-one.[8]

Call a closure operation *inductive* if it corresponds to an inductive closure system. There are several equivalent conditions for a closure operation to be inductive. One is that it preserves unions of directed sets. Another one is that the operation is *finitary* in the sense that the closure of a subset Y of Σ is the union of the closures of all finite subsets of Y. Let us derive this latter condition by using least satisfiers. For every subset Y of Σ, by (3.11), $\bigcup\{\mathrm{s}(\varphi) \mid Y \vDash \varphi\}$ is a member of $\mathcal{U}$ that includes Y and, obviously, it is the least such member; hence:

$$\mathrm{cl}(Y) \;=\; \bigcup\{\mathrm{s}(\varphi) \mid Y \vDash \varphi\}.$$

Now observe that $\mathrm{s}(\varphi)$ is the closure of the finite set $\lfloor\varphi\rfloor$ of all members of Σ occurring in φ.

It is instructive to construct the closure operation cl given by the canonical universe of a Λ-free Horn theory Γ directly in terms of Γ, which can be done as follows. Consider the function d on $\wp(U)$ that takes Y to

$$Y \cup \bigcup\{\lfloor\psi\rfloor \mid \lfloor\varphi\rfloor \subseteq Y \wedge \ulcorner\varphi \preceq \psi\urcorner \in \Gamma\}.$$

Then, by definition, $X \in \mathrm{C}(\Gamma)$ iff $\mathrm{d}(X) = X$. That is, $\mathrm{C}(\Gamma)$ is the set of *fixpoints* of the (extensive and monotone) function d. Hence, since d is finitary, it follows by standard fixpoint arguments that

$$\mathrm{cl}(Y) \;=\; \bigcap\{X \mid Y \subseteq X = \mathrm{d}(X)\} \;=\; \bigcup\{\mathrm{d}^n(Y) \mid n \geqslant 0\}.$$

So, starting with an arbitrary subset of Σ, we inductively arrive at its Γ-closure by closing it successively with respect to all statements of Γ.

[7]A *closure operation* on Σ is a function c on $\wp(\Sigma)$ such that $X \subseteq c(X)$ (*extensivity*), $c(c(X)) \subseteq c(X)$ (*idempotency*), and $c(X) \subseteq c(Y)$ whenever $X \subseteq Y$ (*monotonicity*).

[8]See e.g. Davey and Priestley 1990, Chap. 2 and 3 for a more extensive treatment of these and the following issues.

Finally notice that every inductive intersection system $\mathcal{U}$ over Σ can be made into a closure system by adding $\Sigma \cup \{\Lambda\}$ to $\mathcal{U}$. (Strictly speaking, we switch from a system over Σ to one over $\Sigma \cup \{\Lambda\}$.) From a conceptual point of view, this means to add a single inconsistent element. Satisfaction of Λ is then handled by membership: $\Sigma \cup \{\Lambda\}$ is the only set satisfying Λ. One technical advantage of this move is that $\Sigma \cup \{\Lambda\}$ can serve as the least satisfier of predicates not satisfiable in $\mathcal{U}$. This is compatible with (3.9) if the intersection of the empty set is defined to be $\Sigma \cup \{\Lambda\}$. Consequently (3.13) holds for arbitrary predicates φ.

3.3 Monadic Horn Logic

As in the case of simple inheritance we are interested in an inference calculus for Horn statements that is sound and complete with respect to first-order entailment.

3.3.1 *The Calculus HC*

Consider the set *HC* of inference schemes presented in Figure 19. (R) is *reflexivity*, (T) is *transitivity*, ($\mathrm{I}_\wedge$), ($\mathrm{E}^1_\wedge$), and ($\mathrm{E}^2_\wedge$) are *introduction* and *elimination of conjunction*, (Q) is *de nihilo quodlibet*, and (U) is *universality*. These inference schemes are easily seen to be valid with respect to first-order logic (with identity).

$$\frac{}{A \preceq A}\ (\mathrm{R}) \qquad \frac{}{\Lambda \preceq A}\ (\mathrm{Q}) \qquad \frac{}{A \preceq \mathrm{V}}\ (\mathrm{U})$$

$$\frac{A \preceq B \quad A \preceq C}{A \preceq B \wedge C}\ (\mathrm{I}_\wedge) \qquad \frac{A \preceq B \wedge C}{A \preceq B}\ (\mathrm{E}^1_\wedge) \qquad \frac{A \preceq B \wedge C}{A \preceq C}\ (\mathrm{E}^2_\wedge)$$

$$\frac{A \preceq B \quad B \preceq C}{A \preceq C}\ (\mathrm{T})$$

FIGURE 19 The calculus *HC*

As an illustration, Figure 20 shows a proof of the chaining rule of simple inheritance by *HC*, which says that '$A \preceq B$' and '$B \wedge C \preceq \Lambda$' together deductively entail '$A \wedge C \preceq \Lambda$'. (Notice that the proof goes through for an arbitrary predicate in place of Λ.)

1:	$A \preceq B$	
2:	$B \wedge C \preceq \Lambda$	
3:	$A \wedge C \preceq A \wedge C$	(R)
4:	$A \wedge C \preceq A$	($\mathrm{E}^1_\wedge$) 3
5:	$A \wedge C \preceq C$	($\mathrm{E}^2_\wedge$) 3
6:	$A \wedge C \preceq B$	(T) 4, 1
7:	$A \wedge C \preceq B \wedge C$	($\mathrm{I}_\wedge$) 6, 5
8:	$A \wedge C \preceq \Lambda$	(T) 7, 2

FIGURE 20 Formal proof of ISNOTA chaining by *HC*

3.3.2 *Completeness of HC*

Suppose Γ is a Horn theory over Σ. Since the calculus *HC* is sound with respect to first-order logic, it holds for every model M of Γ that

(3.17) if $\Gamma \vdash_{HC} \varphi \preceq \psi$ then $M(\varphi) \subseteq M(\psi)$.

The key to prove the completeness of *HC* is to show that whenever the least satisfier $\mathrm{s}(\varphi)$ of φ in $\mathrm{C}(\Gamma)$ satisfies some predicate ψ then Γ entails $\varphi \preceq \psi$ by *HC*; that is,

(3.18) if $\mathrm{s}(\varphi) \models \psi$ then $\Gamma \vdash_{HC} \varphi \preceq \psi$.[9]

Proof. Let X_φ be the set of all members p of $\Sigma \cup \{\Lambda\}$ with $\Gamma \vdash_{HC} \varphi \preceq p$. We first show by term induction that $X_\varphi \models \psi$ iff $\Gamma \vdash_{HC} \varphi \preceq \psi$. If ψ belongs to $\Sigma \cup \{\Lambda\}$, there is nothing to show. The case $\psi = \mathrm{V}$ is covered by (U). Now suppose $\psi = \psi_1 \wedge \psi_2$. By induction hypothesis, $X_\varphi \models \psi_1 \wedge \psi_2$ iff Γ entails $\varphi \preceq \psi_1$ and $\varphi \preceq \psi_2$ by *HC*, that is, according to ($\mathrm{I}_\wedge$), ($\mathrm{E}^1_\wedge$), and ($\mathrm{E}^2_\wedge$), iff $\Gamma \vdash_{HC} \varphi \preceq \psi_1 \wedge \psi_2$. Using this result we prove that $X_\varphi = \mathrm{s}(\varphi)$: It follows by (R) that X_φ satisfies φ, and by (T) and (Q) that X_φ either is $\Sigma \cup \{\Lambda\}$ or belongs to $\mathrm{C}(\Gamma)$. Therefore $X_\varphi = \mathrm{s}(\varphi)$. ⌟

By (3.17) and (3.18), and since $\mathrm{s}(\varphi) \models \psi$ iff $\mathrm{M}_\Gamma(\varphi) \subseteq \mathrm{M}_\Gamma(\psi)$,

(3.19) $\mathrm{M}_\Gamma(\varphi) \subseteq \mathrm{M}_\Gamma(\psi)$ iff $\Gamma \vdash_{HC} \varphi \preceq \psi$.

(3.20) Theorem The calculus *HC* is sound and strongly complete with respect to first-order entailment.

[9] Here we take advantage of the modified definition of the least satisfier introduced at the end of Section 3.2.5: if φ is not satisfiable in $\mathrm{C}(\Gamma)$ then $\mathrm{s}(\varphi) = \Sigma \cup \{\Lambda\}$, and vice versa.

Proof. If $\Gamma \vdash \varphi \preceq \psi$ then, by soundness of predicate logic, $\mathrm{M}_\Gamma(\varphi) \subseteq \mathrm{M}_\Gamma(\psi)$. Hence, by (3.19), $\Gamma \vdash_{HC} \varphi \preceq \psi$. ⌟

Since *HC* is complete, one can prove completeness of other sound inference calculi for Horn statements by showing that they are complete with respect to *HC*. In other words, it suffices to show that every inference scheme of *HC* can be replaced by a proof in the calculus in question. For example, consider the following two inference schemes, called *weakening* (W) (or *monotonicity*) and *cut* (C), which are obviously sound with respect to first-order logic.

$$\frac{A \preceq B}{C \wedge A \preceq B}\ \text{(W)} \qquad\qquad \frac{A \preceq B \qquad A \wedge B \preceq C}{A \preceq C}\ \text{(C)}$$

It is immediate that weakening and cut imply transitivity:

$$\cfrac{A \preceq B \qquad \cfrac{B \preceq C}{A \wedge B \preceq C}\,\text{(W)}}{A \preceq C}\,\text{(C)}$$

So, if the inference scheme (T) of *HC* is replaced by (W) and (C), the resulting calculus is again complete.

3.4 Applications

Simple inheritance theories with binary exclusions provide a first example of Horn theories. As shown in Section 3.1.1, every such theory corresponds to a Horn theory over the same set of primitive predicates, with $p \preceq q$ and $p \wedge q \preceq \Lambda$ respectively in place of p *is-a* q and p *is-not-a* q.

3.4.1 *Defined Concepts*

In Sections 1.4 and 2.3, a primitive concept or predicate has been called *defined* if it is implied by the conjunction of its immediate superordinates. Whereas the condition that a concept p is defined cannot be expressed by simple inheritance statements, it clearly poses no problems for Horn statements – at least if the number of immediate superordinates of p is finite: take the statement $\varphi \preceq p$, where φ is the conjunction of all immediate superordinates of p. Obviously, a set X of primitives is closed with respect to $\varphi \preceq p$ just in case X meets the condition of Section 2.3 that Carpenter and Pollard impose on conjunctive concepts with defined concepts.

To give an example, consider the simple inheritance network Γ of (2.14), here repeated on the left of Figure 21. As a Horn theory, Γ consists of the statements

$$a \preceq d, \quad b \preceq d \wedge e, \quad c \preceq e, \quad d \preceq f, \quad e \preceq f, \quad b \wedge c \preceq \Lambda,$$

where we have put the two statements $b \preceq d$ and $b \preceq e$ into one, by introduction of conjunction ($\mathrm{I}_\wedge$). Now suppose the predicates b and f are defined ones. Then the Horn theory Γ has to be extended by the statements $d \wedge e \preceq b$ and $\mathrm{V} \preceq f$. (Since f has no superordinates, to regard f as defined means that f is implied by the conjunction of the empty set, i.e. by V.) The canonical universe of this extended theory is shown on the right of Figure 21.

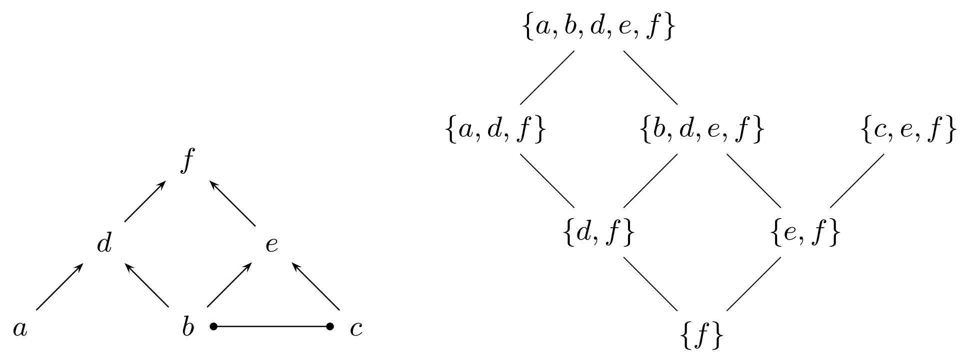

FIGURE 21 Simple inheritance theory with defined concepts

3.4.2 *Finitary Exclusions*

A natural generalization of binary exclusions is to allow finitary ones. The corresponding theory Γ then consists of statements of the form

$$p \preceq q \qquad \text{and} \qquad p_1 \wedge \ldots \wedge p_n \preceq \Lambda.$$

Notice that the only effect of *unary* exclusions, i.e. of statements of the form $p \preceq \Lambda$, with $p \in \Sigma$, is to prevent some of the primitives from appearing in consistently closed sets at all.

(3.21) Proposition The canonical universe of a simple inheritance theory with exclusions is an inductive intersection system which is closed with respect to bounded union.

Proof. Let $\mathcal{S}$ be a bounded subset of $\mathrm{C}(\Gamma)$. We need to show that $\bigcup \mathcal{S} \in \mathrm{C}(\Gamma)$. If $p \preceq q$ belongs to Γ and $p \in \bigcup \mathcal{S}$ then, for some member X of $\mathcal{S}$, $p \in X$ and hence $q \in X \subseteq \bigcup \mathcal{S}$. Suppose $p_1 \wedge \ldots \wedge p_n \preceq \Lambda$ belongs to Γ. Then $\{p_1, \ldots, p_n\} \not\subseteq \bigcup \mathcal{S}$, because otherwise there are members $X_1, \ldots, X_n$ of $\mathcal{S}$, with $p_i \in X_i$, which are by assumption included by some $Y \in \mathrm{C}(\Gamma)$. $\lrcorner$

Since binary exclusion corresponds to subset systems that are closed with respect to pairwise bounded union (see (2.3) and (2.8)), one might expect that the respective property corresponding to finitary exclusion is being closed with respect to union of finitely bounded subsets, where *finitely bounded* means that all nonempty finite subsets are bounded; indeed:

(3.22) Lemma A subset system is closed with respect to finitely bounded union iff it is inductive and closed with respect to bounded union.

Proof. Let $\mathcal{U}$ be a subset system and $\mathcal{S}$ a nonempty subset of $\mathcal{U}$. Observe that $\bigcup\mathcal{S}$ is identical to the union of the directed set $\{\bigcup\mathcal{F} \mid \mathcal{F} \subseteq \mathcal{S} \text{ finite}\}$. Suppose $\mathcal{S}$ is finitely bounded and $\mathcal{U}$ is inductive and closed with respect to bounded union. Then $\bigcup\mathcal{S} \in \mathcal{U}$. Conversely, suppose $\mathcal{U}$ is closed with respect to finitely bounded union. If $\mathcal{S}$ is directed or bounded then all nonempty finite subsets of $\mathcal{S}$ are bounded. Hence $\bigcup\mathcal{S} \in \mathcal{U}$. ⌟

The classification (3.21) again works the other way around as well. Suppose $\mathcal{U}$ is an inductive intersection system over Σ with bounded unions. Consider the theory $\Gamma(\mathcal{U})$ over Σ such that

$$\ulcorner p \preceq q \urcorner \in \Gamma(\mathcal{U}) \quad \text{iff} \quad \forall X \in \mathcal{U}\,(p \in X \rightarrow q \in X),$$
$$\ulcorner p_1 \wedge \ldots \wedge p_n \preceq \Lambda \urcorner \in \Gamma(\mathcal{U}) \quad \text{iff} \quad \forall X \in \mathcal{U}\,(\{p_1, \ldots, p_n\} \not\subseteq X).$$

We prove that $\mathrm{C}(\Gamma(\mathcal{U})) = \mathcal{U}$. Clearly $\mathcal{U} \subseteq \mathrm{C}(\Gamma(\mathcal{U}))$. As for the reverse inclusion, first observe that each member X of $\mathrm{C}(\Gamma(\mathcal{U}))$ is the union of the set $\{\mathrm{s}(p) \mid p \in X\}$ of least satisfiers of members of X; cf. the proof of (2.5). It remains to check that $\{\mathrm{s}(p) \mid p \in X\}$ is finitely bounded. Let $F = \{p_1, \ldots, p_n\}$ be a finite subset of X. Since X is consistent, $p_1 \wedge \ldots \wedge p_n \preceq \Lambda$ does not belong to $\Gamma(\mathcal{U})$. Hence F is included in some member Y of $\mathcal{U}$; thus $\bigcup\{\mathrm{s}(p) \mid p \in F\} \subseteq Y$. So $X = \bigcup\{\mathrm{s}(p) \mid p \in X\}$ belongs to $\mathcal{U}$ because $\mathcal{U}$ is closed with respect to finitely bounded union. To summarize:

(3.23) Theorem The canonical universe of a simple inheritance theory with exclusions is an intersection system that is closed with respect to finitely bounded union, and every such subset system arises that way.

By the same line of reasoning that has led us to (2.8), we can conclude:

(3.24) Corollary The canonical universe of an exclusion theory is closed with respect to the formation of subsets and finitely bounded union, and every such subset system arises that way.

4

Scott Domains

Our characterization of the generic universe determined by an inheritance network or a Horn theory hitherto depends on the representation of entity types by consistently closed sets. It is also possible to give a description in purely order-theoretic terms, that is, a description that does not depend on specific representations but only on properties of the generic universe as a *(partially) ordered set*.

It is well-known that inductive intersection systems are categorically equivalent to Scott domains.[1] We derive this result in Section 4.1 by identifying the least satisfiers of conjunctive predicates as the compact elements of the generic universe. In Section 4.2, several versions of simple inheritance theories are classified with respect to their domains. Restriction to simple inheritance is shown to enforce the corresponding domains to be completely distributive, whereas pure exclusion theories imply atomicity.

4.1 Horn Theories and Scott Domains

4.1.1 *Bounded-Complete Dcpos*

According to (3.7), the generic universe of a Horn theory is a directed-complete (partially) ordered set (*dcpo*), with infima for all nonempty subsets. The latter requirement is equivalent to the condition that every subset with an upper bound has a least upper bound:

(4.1) Lemma An ordered set has infima for all nonempty subsets iff it has suprema for all bounded subsets.

Proof. Suppose P is an ordered set with infima for nonempty subsets. Then the supremum of a subset S of P with upper bounds is the infimum of the set of all upper bounds of S. Vice versa, if P has suprema for bounded subsets then the infimum of a nonempty subset S of P is the supremum of the set

[1] See e.g. Davey and Priestley 1990, Chap. 3.

$\{y \in P \mid \forall x \in S(y \sqsubseteq x)\}$ of all members of P below or equal to every member of S. $\lrcorner$

In particular, every nonempty ordered set with suprema for all bounded subsets has a least element. In the presence of directed completeness, it is enough to require suprema for bounded *finite* subsets to get suprema for all bounded subsets:

(4.2) Lemma If a dcpo has suprema for bounded finite subsets, it has suprema for all bounded subsets.

Proof. Suppose D is a dcpo with suprema for bounded finite subsets. Let S be a bounded subset of D. Then $\bigsqcup F$ exists for every finite subset F of S. Clearly the set $\{\bigsqcup F \mid F \subseteq S \text{ finite}\}$ is directed and its supremum is one of S. $\lrcorner$

We call a dcpo *bounded-complete* if it has suprema for all bounded finite subsets (and thus for all bounded subsets). It is instructive to check that in the finite case, bounded-completeness is all that is needed for an ordered set to be the generic universe of a Horn theory. By (3.14), it suffices to show that such an ordered set is order-isomorphic to an inductive intersection system. Notice that a finite ordered set is necessarily directed-complete since each of its directed subsets has a greatest member.

We can employ the following simple method to represent an arbitrary ordered set P by an intersection system over P. The idea is to represent each member x of P by the subset $\downarrow x = \{y \mid y \sqsubseteq x\}$ of P. Since $\downarrow x \subseteq \downarrow y$ iff $x \sqsubseteq y$, we have that P is order-isomorphic to the subset system $\{\downarrow x \mid x \in P\}$. Moreover, if S is a nonempty subset of P with an infimum in P, then the infimum $\downarrow(\bigsqcap S)$ of $\{\downarrow x \mid x \in S\}$ is given by intersection:

(4.3) $$\{y \mid y \sqsubseteq \bigsqcap S\} \;=\; \{y \mid \forall x \in S(y \sqsubseteq x)\} \;=\; \bigcap\{\downarrow x \mid x \in S\}.$$

So every ordered set P is isomorphic to an intersection system over P.

Suppose P is finite. If S is a directed subset of P then $\bigsqcup S$ belongs to S. It follows that

$$\{y \mid y \sqsubseteq \bigsqcup S\} \;=\; \{y \mid \exists x \in S(y \sqsubseteq x)\} \;=\; \bigcup\{\downarrow x \mid x \in S\}.$$

Hence, for finite P, suprema of directed subsets of $\{\downarrow x \mid x \in P\}$ are given by union. So every finite ordered set with infima for nonempty subsets is isomorphic to an inductive intersection system. Together with (4.1) and (3.14), we have:

(4.4) Proposition Every finite, bounded-complete ordered set P is the generic universe of a Horn theory over P.

For infinite ordered sets, however, this approach fails in two respects: first, there are bounded-complete dcpos that do not represent the generic universe of any Horn theory, and second, the given representation of ordered sets by subset systems may not yield the desired result. The latter type of failure is demonstrated by the following example.

(4.5) Example Suppose Σ is a countable set $\{a_1, a_2, a_3, \ldots\}$ of primitive predicates. Consider the theory Γ over Σ with axioms $a_{n+1} \preceq a_n$ for every n. Let A_n be the set $\{a_m \mid m \leqslant n\}$, with $n \geqslant 0$. Then

$$\mathrm{C}(\Gamma) \;=\; \{A_n \mid n \in \mathbb{N}_0\} \cup \{\Sigma\}.$$

Observe that the subset $\{{\downarrow}A_n \mid n \in \mathbb{N}\}$ of $\{{\downarrow}X \mid X \in \mathrm{C}(\Gamma)\}$ is directed but its union $\{A_n \mid n \in \mathbb{N}_0\}$ differs from its supremum ${\downarrow}\Sigma$. So, suprema of directed subsets of $\{{\downarrow}X \mid X \in \mathrm{C}(\Gamma)\}$ are not necessarily given by union.

The next example indicates why certain bounded-complete dcpos fail to represent the generic universe of a Horn theory.

(4.6) Example Let Σ be $\{a_1, a_2, a_3, \ldots\}$. Consider the subset system over Σ that consists of the sets Σ, $B = \{a_n \mid n > 1\}$, and $A_n = \{a_m \mid m \leqslant n\}$, for all $n \geq 0$; see Figure 22. This bounded-complete ordered set is obviously directed-complete. Now suppose there is an inductive intersection system with an inclusion ordering as depicted in Figure 22 (with Σ, A_n, and B arbitrary otherwise). Then $\bigcup\{A_n \mid n \in \mathbb{N}\} = \Sigma$ and $A_n \cap B = A_0$ for every n, contrary to the assumption that $\Sigma \cap B = B \neq A_0$.

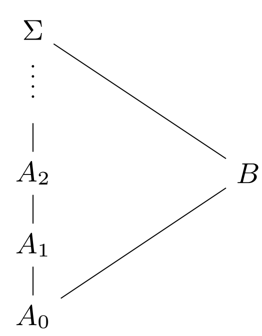

FIGURE 22 A non-algebraic bounded-complete dcpo

4.1.2 *Compact Elements and Algebraicity*

Concerning the representation problem, let us elucidate in general terms what is wrong with Example (4.5). Suppose a dcpo D is an inductive intersection system over Σ. Let K be the set of all elements of D that are least satisfiers of conjunctive predicates over Σ. According to (3.11), every member x of D is the union of the directed set $\{y \in K \mid y \sqsubseteq x\}$ ($= {\downarrow}x \cap K$) of all members of K below or equal to x, that is,

(4.7) $$\downarrow x \cap K \quad \text{is directed} \quad \text{and} \quad x \;=\; \bigsqcup(\downarrow x \cap K).$$

Now, if x is not itself in K, as e.g. Σ in Example (4.5), then x is a member of ${\downarrow}x$ but not of $\bigcup\{{\downarrow}y \mid y \in ({\downarrow}x \cap K)\}$; hence $\downarrow$ does not preserve directed suprema.

The solution is to represent x not by ${\downarrow}x$ but by ${\downarrow}x \cap K$. The resulting subset system $\{{\downarrow}x \cap K \mid x \in D\}$ is order-isomorphic to D because, by (4.7), $x = y$ whenever ${\downarrow}x \cap K = {\downarrow}y \cap K$. To see that this representation of D as a subset system takes suprema of directed subsets to unions, observe first the following consequence of (3.10) and (3.6)(ii): for every member x of K and every directed subset S of D,

(4.8) $$\text{if} \quad x \sqsubseteq \bigsqcup S \quad \text{then} \quad \exists y \in S (x \sqsubseteq y).$$

Consequently,

$$\begin{aligned} \downarrow(\bigsqcup S) \cap K \;&=\; \{x \in K \mid x \sqsubseteq \bigsqcup S\} \\ &=\; \{x \in K \mid \exists y \in S(x \sqsubseteq y)\} \;=\; \bigcup\{\downarrow y \cap K \mid y \in S\}. \end{aligned}$$

Furthermore, since all nonempty subsets of D have infima, it follows in the same way as (4.3) that the infimum of each nonempty subset of $\{{\downarrow}x \cap K \mid x \in D\}$ is given by intersection.

So we succeeded in representing a dcpo D that is known to be isomorphic to an inductive intersection system by an inductive intersection system over the subset K of D. Our aim is to characterize D and K in order-theoretic terms. For that purpose, notice that every member x of D that satisfies (4.8) belongs to K; for by (4.7) and (4.8), there is a member y of ${\downarrow}x \cap K$ such that $x \sqsubseteq y$; hence, $x = y \in K$.

Suppose D is an arbitrary dcpo. A member x of D that satisfies (4.8) for every directed subset S of D is called *compact*. Let $\mathrm{k}(D)$ be the (ordered) set of compact elements of D. The dcpo D is said to be *algebraic* if $\mathrm{k}(D)$ is a *basis* of D in the sense that ${\downarrow}x \cap \mathrm{k}(D)$ is directed with supremum x, for every element x of D.[2] The preceding discussion shows that an algebraic dcpo D is

[2] The general definition of a basis of a dcpo uses the so-called *order of approximation*; cf. Abramsky and Jung 1994, Sect. 2.2. Dcpos with basis are said to be *continuous*.

isomorphic to the subset system $\{{\downarrow}x \cap \mathrm{k}(D) \mid x \in D\}$ over $\mathrm{k}(D)$, where directed suprema are given by union.

Moreover, we have seen, first, that every inductive intersection system is an algebraic dcpo with infima for nonempty subsets, and second, that every such dcpo D is isomorphic to an inductive intersection system over $\mathrm{k}(D)$. Since by (4.1) and (4.2), having infima for nonempty subsets means to be bounded-complete, the ordered sets that can be represented by inductive intersection systems are precisely the bounded-complete algebraic dcpos, which are also known as *Scott domains* (or *Scott-Ershov domains*).[3] To summarize:

(4.9) Theorem

- ***(i)*** Every inductive intersection system is a Scott domain.
- ***(ii)*** If D is a Scott domain then $\{{\downarrow}x \cap \mathrm{k}(D) \mid x \in D\}$ is an inductive intersection system over $\mathrm{k}(D)$ and the function that takes x to ${\downarrow}x \cap \mathrm{k}(D)$ is an order isomorphism.

Reconsider Example (4.6) under this perspective: B is neither compact since $B \sqsubseteq \bigsqcup\{A_n \mid n \in \mathbb{N}\}$ but $B \not\sqsubseteq A_n$ for every n; nor is B the supremum of any directed set of compact elements. Hence Figure 22 depicts a non-algebraic bounded-complete dcpo.

We have thus arrived at a characterization of the generic universe of Horn theories in purely order-theoretic terms. Algebraicity is a necessary and sufficient condition for a bounded-complete dcpo to be the generic universe of a Horn theory:

(4.10) Theorem The generic universe of a Horn theory is a Scott domain and every Scott domain arises that way.

The same line of reasoning can be applied to Λ-free Horn theories and inductive closure systems; cf. Section 3.2.5. The only difference to the previous case is that an inductive closure system has a greatest element. By (4.1) and (4.9), we get:

(4.11) Theorem Every inductive closure system is a complete algebraic lattice. Vice versa, every complete algebraic lattice D is isomorphic to an inductive closure system over $\mathrm{k}(D)$, namely $\{{\downarrow}x \cap \mathrm{k}(D) \mid x \in D\}$.

The generic universes of Λ-free Horn theories are thus precisely the complete algebraic lattices.

---

[3] We allow a Scott domain to be empty in order to cover the case of an empty canonical universe; see footnote 5 on page 44.

4.1.3 *Countability*

An algebraic dcpo with countably many compact elements is said to be *countably based* or ω-*algebraic*.

(4.12) Theorem The generic universe of a Horn theory over a countable set of primitives is a countably based Scott domain and every such domain arises that way.

Proof. The set of conjunctive predicates over a countable set of primitives is countable (because the set of finite sequences over a countable set is countable). Hence the set of least satisfiers of conjunctive predicates is countable in turn. $\lrcorner$

It can be shown that every member of an ω-algebraic domain D is the supremum of an ω-*chain* of compact elements, i.e. the supremum of a countable subset $\{x_n \mid n \in \mathbb{N}\}$ of $\mathrm{k}(D)$ such that $x_n \sqsubseteq x_{n+1}$ for every n.[4] So, by (4.12), every generic entity of a Horn theory with countable vocabulary is the "limit" (the supremum) of a specialization sequence (an ω-chain) consisting of least satisfiers of conjunctive predicates.

4.1.4 *Ideal Completion*

In this Section we briefly review a standard result according to which the representation of a Scott domain D as an inductive intersection system over $\mathrm{k}(D)$ coincides with the ideal completion of $\mathrm{k}(D)$. More generally, every algebraic dcpo D is isomorphic to the ideal completion of $\mathrm{k}(D)$.

An *ideal* of an ordered set P is a nonempty subset of P that is directed and downwards closed. The set $\mathrm{I}(P)$ of ideals of P, ordered by set inclusion, is called the *ideal completion* of P.

(4.13) Proposition If P is an ordered set then $\mathrm{I}(P)$ is an algebraic dcpo. Moreover, $\mathrm{k}(\mathrm{I}(P)) = \{\downarrow x \mid x \in P\} \simeq P$.

Proof. [5] The union of a directed set of ideals is clearly an ideal in turn. So $\mathrm{I}(P)$ is directed complete, and directed suprema are given by set union. Notice that $\downarrow x \in \mathrm{I}(P)$ for all $x \in P$, and that, for every $I \in \mathrm{I}(P)$, the set $\{\downarrow x \mid x \in I\}$ is directed and $I = \bigcup\{\downarrow x \mid x \in I\}$. Consequently, if I is compact in $\mathrm{I}(P)$, there is an $x \in I$ such that $I = \downarrow x$. Moreover, $\downarrow x$ is compact for every $x \in P$, because whenever $\downarrow x \subseteq \bigcup \mathcal{S}$, with $\mathcal{S} \subseteq \mathrm{I}(P)$ directed, then $x \in I$, for some $I \in \mathcal{S}$, and hence $\downarrow x \subseteq I$. So the compact elements of $\mathrm{I}(P)$ are those of the form $\downarrow x$,

[4] See e.g. Vickers 1989, p. 126.
[5] Adapted from Vickers 1989, pp. 116f.

with $x \in P$. In addition, it follows that $\mathrm{I}(P)$ is algebraic. Finally, P is order-isomorphic to $\mathrm{k}(\mathrm{I}(P))$ since the function from P to $\mathrm{I}(P)$ that takes x to $\downarrow x$ is one-to-one. $\lrcorner$

Let D be an algebraic dcpo. Recall from Section 4.1.2 that D can be represented by the subset system $\{\downarrow x \cap \mathrm{k}(D) \mid x \in D\}$ over $\mathrm{k}(D)$. Since the subset $\downarrow x \cap \mathrm{k}(D)$ of $\mathrm{k}(D)$, with $x \in D$, is by definition directed and clearly downwards closed, it is an ideal of $\mathrm{k}(D)$. On the other hand, every ideal I of $\mathrm{k}(D)$ coincides with $\downarrow(\bigsqcup I) \cap \mathrm{k}(D)$. Thus

$$\{\downarrow x \cap \mathrm{k}(D) \mid x \in D\} \;=\; \mathrm{I}(\mathrm{k}(D)).$$

So the representation of D by $\{\downarrow x \cap \mathrm{k}(D) \mid x \in D\}$ is just the ideal completion of $\mathrm{k}(D)$. Consequently:

(4.14) Proposition If D is an algebraic dcpo then $D \simeq \mathrm{I}(\mathrm{k}(D))$.

Since an algebraic dcpo D is determined by the ordered set $\mathrm{k}(D)$ of its compact elements, one can try to characterize properties of the former by properties of the latter. In particular, if D is a Scott domain then the supremum of two compact elements of D, if existent, is compact in turn; that is, $\mathrm{k}(D)$ is a *conditional join semilattice*:

(4.15) Lemma If D is a nonempty Scott domain then $\mathrm{k}(D)$ is a conditional join semilattice with least element.

Proof. One way to prove this simple fact is to employ the order-theoretic definition of compactness. We leave this as an exercise to the reader. Alternatively, one can use the representation of D as an inductive intersection system $\mathcal{U}$ over some set Σ; see Section 4.1.2. Recall that the compact elements of $\mathcal{U}$ are the least satisfiers of conjunctive predicates over Σ. Given two such least satisfiers $\mathrm{s}(\varphi)$ and $\mathrm{s}(\psi)$ which are both included in an element of $\mathcal{U}$, then $\mathrm{s}(\varphi \wedge \psi)$ is the supremum of $\mathrm{s}(\varphi)$ and $\mathrm{s}(\psi)$. $\lrcorner$

In order to show that the ideal completion of a conditional join semilattice with least element is a Scott domain, we need the following fact:

(4.16) Lemma Suppose D is an algebraic dcpo with least element and every bounded two-element subset of $\mathrm{k}(D)$ has a supremum in D. Then D is a Scott domain.

Proof. [6] We need to show that every bounded subset S of D has a supremum in D. Let K be the set $\{y \in \mathrm{k}(D) \mid \exists x \in S(y \sqsubseteq x)\}$. By assumption, every finite subset F of K has a supremum in D. Hence K has a supremum in D, namely the supremum of the directed set $\{\bigsqcup F \mid F \subseteq K \text{ finite}\}$. We show that $\bigsqcup K$ is the supremum of S. If $x \in S$ then $\downarrow x \cap \mathrm{k}(D) \subseteq K$ and consequently $x = \bigsqcup(\downarrow x \cap \mathrm{k}(D)) \sqsubseteq \bigsqcup K$. So $\bigsqcup K$ is an upper bound of S. Since every upper bound of S is also an upper bound of K, it follows that $\bigsqcup K$ is the least upper bound of S. $\lrcorner$

Now suppose P is a conditional join semilattice with least element. We claim that $\mathrm{I}(P)$ is a Scott domain. According to (4.16) and (4.13), it suffices to show that, for all $x, y \in P$, if the set $\{\downarrow x, \downarrow y\}$ is bounded, it has a supremum in $\mathrm{I}(P)$. Suppose $I \in \mathrm{I}(P)$ is an upper bound of $\{\downarrow x, \downarrow y\}$. Then $x \sqcup y \in I$ since I is directed and downwards closed. Hence $\downarrow(x \sqcup y)$ is the supremum of $\{\downarrow x, \downarrow y\}$.

All in all, we have proved the following characterization of Scott domains in terms of their compact elements:

(4.17) Theorem

(i) If D is a nonempty Scott domain then $\mathrm{k}(D)$ is a conditional join semilattice with least element and $D \simeq \mathrm{I}(\mathrm{k}(D))$.

(ii) If P is a conditional join semilattice with least element then $\mathrm{I}(P)$ is a Scott domain and $\mathrm{k}(\mathrm{I}(P)) \simeq P$.

Note that a nonempty downwards closed subset I of a conditional join semilattice P is an ideal of P (as an ordered set) just in case $x \sqcup y \in I$ whenever $\{x, y\} \subseteq I$ is bounded. (In particular, this shows that the ideals of a lattice viewed as an algebra are the ideals of the lattice as an ordered set.)

(4.18) Remark *(Lindenbaum algebra)* Suppose D is the generic universe of a Horn theory Γ over Σ. The compact elements of D correspond to the least satisfiers in D of conjunctive predicates over Σ. Hence each element of $\mathrm{k}(D)$ represents a class of satisfiable predicates that are equivalent with respect to Γ. Put differently, the order dual of $\mathrm{k}(D)$, "lifted" by adjoining a least element for Λ, is a meet semilattice that represents the *Lindenbaum algebra* of Γ. (Beware, the Lindenbaum algebra defined in Chapter 6, which incorporates disjunction too, is the bounded distributive lattice freely generated by this meet semilattice.)

4.2 Domains of Simple Inheritance

We now go on to classify the generic universes of the various versions of simple inheritance theories in order-theoretic terms.

[6] Adapted from Stoltenberg-Hansen et al. 1994, p. 58.

4.2.1 ***Coprimeness and Distributivity***

Let us begin with the general version introduced in Section 3.4.2, which is simple inheritance plus (finitary) exclusion. Recall that the canonical universe $\mathcal{U}$ of such a theory over Σ is an intersection system which is closed with respect to finitely bounded union, that is, by (3.22), an inductive intersection system that is closed with respect to bounded union.

Suppose D is a dcpo isomorphic to $\mathcal{U}$. Let P be the subset of D that corresponds to the set $\{\mathrm{s}(p) \mid p \in \Sigma\}$ of least satisfiers of primitive predicates in $\mathcal{U}$. We know from Section 3.4.2 that every member x of D is the supremum of the set $\{y \in P \mid y \sqsubseteq x\}$; that is:

(4.19) $$x \;=\; \bigsqcup(\downarrow\! x \cap P).$$

Moreover, for every member x of P and every *bounded* subset S of D, we have that

(4.20) $$\text{if} \quad x \sqsubseteq \bigsqcup S \quad \text{then} \quad \exists y \in S\,(x \sqsubseteq y).$$

Proof. Suppose $\mathcal{S}$ is a bounded subset of $\mathcal{U}$ such that $\mathrm{s}(p) \subseteq \bigcup \mathcal{S}$. Then $p \in Y$ for some $Y \in \mathcal{S}$ and hence $\mathrm{s}(p) \subseteq Y$. ⌟

A member x of a dcpo D that satisfies property (4.20) for every bounded subset S of D is called *completely coprime* (or completely $\sqcup$-*prime*). (Notice that the least element $\bot$ of D, if existent, is *not* completely coprime, because $\varnothing$ is bounded and $\bot = \bigsqcup \varnothing$.) Let $\mathrm{p}(D)$ be the set of completely coprime elements of D. A bounded-complete dcpo D is called *coprime-algebraic* if each of its members is the supremum of a subset of $\mathrm{p}(D)$.

We have seen above that the generic universe of a simple inheritance theory with exclusions is a coprime-algebraic Scott domain. The converse also holds because every such domain can be represented by an intersection system that is closed with respect to finitely bounded union:

(4.21) Theorem

- ***(i)*** Every intersection system that is closed with respect to finitely bounded union is a coprime-algebraic Scott domain.
- ***(ii)*** Every coprime-algebraic Scott domain D is isomorphic to the subset system $\{\downarrow\! x \cap \mathrm{p}(D) \mid x \in D\}$ over $\mathrm{p}(D)$, which is closed with respect to nonempty intersection and finitely bounded union.

Proof. We need to verify (ii). Let D be a coprime-algebraic Scott domain. Suppose S is a subset of D such that $\{\downarrow\! x \cap \mathrm{p}(D) \mid x \in S\}$ is finitely bounded. Then S is finitely bounded. Hence $\bigsqcup S \in D$ because D is a Scott domain. By (4.20), it follows that

$$\begin{aligned}\downarrow(\bigsqcup S) \cap \mathrm{p}(D) &= \{x \in \mathrm{p}(D) \mid x \sqsubseteq \bigsqcup S\} \\ &= \{x \in \mathrm{p}(D) \mid \exists y \in S(x \sqsubseteq y)\} \;=\; \bigcup\{\downarrow y \cap \mathrm{p}(D) \mid y \in S\}.\end{aligned}$$

Moreover, since D has infima for all nonempty subsets, $\{\downarrow x \cap \mathrm{p}(D) \mid x \in D\}$ is closed with respect to nonempty intersection. $\lrcorner$

Next we characterize coprime-algebraic Scott domains as ideal completions of the ordered set of their compact elements. An element x of a dcpo D is called *coprime* if it satisfies property (4.20) for all *finite* bounded subsets S of D.

(4.22) Proposition If D is a coprime-algebraic Scott domain then every element of $\mathrm{k}(D)$ is a finite join of coprimes of $\mathrm{k}(D)$.

Proof. According to (4.21)(ii), it suffices to consider the case that D is an inductive intersection system $\mathcal{U}$ over Σ which is closed with respect to bounded union. Recall that every compact element of $\mathcal{U}$ is a least satisfier of a conjunctive predicate $p_1 \wedge \ldots \wedge p_n$ over Σ. But the least satisfier of $p_1 \wedge \ldots \wedge p_n$ in $\mathcal{U}$ is $\mathrm{s}(p_1) \cup \ldots \cup \mathrm{s}(p_n)$. $\lrcorner$

(4.23) Lemma An element of a Scott domain D is completely coprime in D iff it is compact in D and coprime in $\mathrm{k}(D)$.

Proof. Suppose $x \in \mathrm{k}(D)$ is coprime in $\mathrm{k}(D)$, and $x \sqsubseteq \bigsqcup S$ for some bounded $S \subseteq D$. Let K be $\{y \in \mathrm{k}(D) \mid \exists z \in S(y \sqsubseteq z)\}$. Then $\bigsqcup S$ is the supremum of the directed set $\{\bigsqcup F \mid F \subseteq K \text{ finite}\}$; cf. the proof of (4.16). Since x is compact, $x \sqsubseteq \bigsqcup F$ for some finite subset F of K. Hence, since x is coprime in $\mathrm{k}(D)$, there is a $y \in F$ with $x \sqsubseteq y$. So $x \sqsubseteq z$ for some $z \in S$. The rest is obvious. $\lrcorner$

In combination with (4.17)(ii) and the fact that every element of a Scott domain is the supremum of compact ones, we can conclude:

(4.24) Proposition If P is a conditional join semilattice with least element such that every element of P is a finite join of coprimes then $\mathrm{I}(P)$ is a coprime-algebraic Scott domain and $\mathrm{p}(\mathrm{I}(P)) = \{\downarrow x \mid x \in P \text{ coprime}\}$.

(4.25) Example Let Σ be $\{a, b\} \cup \{c_1, c_2, \ldots\}$ and let Γ be the simple inheritance theory over Σ given by the statements

$$a \wedge b \preceq \Lambda, \quad a \preceq c_n, \quad b \preceq c_n, \quad c_{n+1} \preceq c_n \quad (n \geqslant 1).$$

The canonical universe of Γ is depicted by Figure 23. Its only non-compact member is $\{c_1, c_2, \ldots\}$, and every compact element besides $\varnothing$ is completely coprime. This example shows that even in the case of simple inheritance the set of least satisfiers need not be closed with respect to intersection.

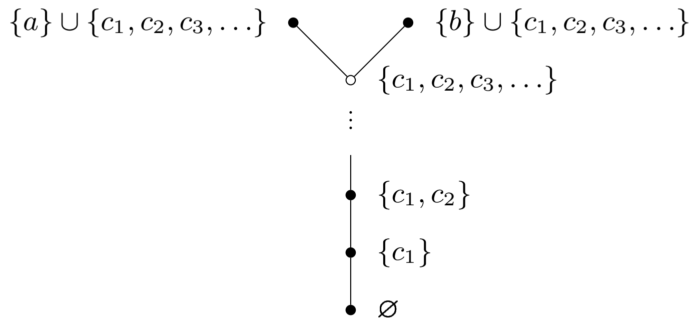

FIGURE 23 Non-compact meet of compact elements

Coprime-algebraic Scott domains correspond to subset systems where infima and suprema, if existent, are given by intersection and union, respectively. Since set union distributes over set intersection and vice versa, it seems reasonable to expect that these domains can be characterized in terms of *distributivity*. We say that a Scott domain D is *completely distributive* if $\downarrow x$ is completely distributive for every $x \in D$. (Notice that $\downarrow x$ is by definition a complete sublattice of D.) For a proof of the following characterization see Winskel (1988).

(4.26) Theorem A Scott domain is coprime-algebraic iff it is completely distributive.

In general, the representation of a coprime-algebraic Scott domain D as a subset system over $\mathrm{p}(D)$ is much more parsimonious than the representation as a subset system over $\mathrm{k}(D)$ via (4.9)(ii). For instance, if $\mathrm{p}(D)$ is a finite antichain then $\mathrm{k}(D)$ ($= D$) is isomorphic to the power set lattice over $\mathrm{p}(D)$ and hence exponential in size.

(4.27) Example Consider the bounded-complete ordered set on the left of Figure 24. It is distributive and hence coprime-algebraic with coprimes x_1, x_2, x_3, and x_5. The diagram on the right shows its representation as a subset system over $\{x_1, x_2, x_3, x_5\}$, which is the canonical universe of the inheritance theory $\{x_3 \preceq x_1,\ x_5 \preceq x_2,\ x_3 \wedge x_5 \preceq \Lambda\}$.

Let us now turn from (finitary) exclusions to simple inheritance with binary exclusions or no exclusions at all. Call a dcpo *pairwise-complete* if it has

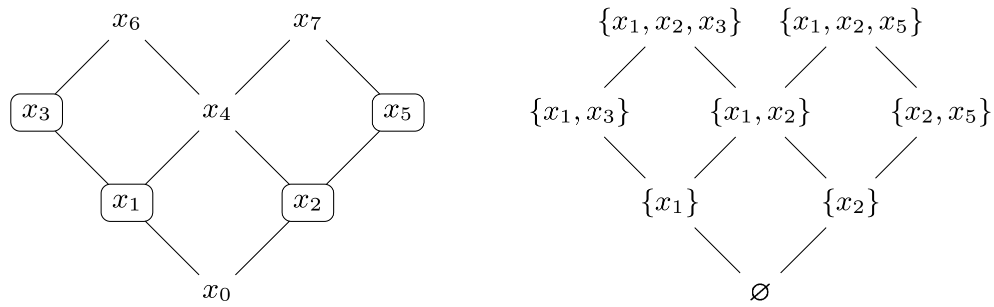

FIGURE 24 Representation of bounded-complete distributive ordered set

suprema for all subsets with pairwise bounded elements.[7] We know from Section 2.1.1 that the generic universe of a simple inheritance theory with binary exclusions is a pairwise-complete coprime-algebraic dcpo. On the other hand, it is easily seen that (4.21)(ii) takes pairwise-complete coprime-algebraic dcpos to subset systems that are closed with respect to nonempty intersection and pairwise bounded union. So the generic universes of simple inheritance theories with binary exclusions are precisely the pairwise-complete completely distributive Scott domains. In the absence of exclusions, finally, we get the complete completely distributive algebraic lattices. The given classification of simple inheritance theories is summarized in rows three to five of Table 1 at the end of this chapter.

(4.28) Remark *(Event structures)* Simple inheritance with (binary) exclusion appears in another guise in Winskel's theory of *event structures*, which are formal models for processes of concurrent computation (see e.g. Winskel 1988, Zhang 1991, Winskel and Nielsen 1995). In this context, members of Σ denote *events*, ISA becomes *causal dependency*, and ISNOTA indicates *conflicts*; elements of the canonical universe are *configurations* and specialization corresponds to a *transition* relation between configurations. In addition, it is assumed that each member of Σ bears the transitive closure of causal dependency only to finitely many other members of Σ. The corresponding domains are the pairwise-complete *dI-domains*, where 'd' means distributivity and 'I' says that there are only finitely many elements below each compact element.

4.2.2 *Atomicity*

Completeness of the generic universe means lack of exclusions. Let us consider the complementary case of exclusions without inheritance. According to (3.24), the canonical universe of an exclusion theory is closed with respect to subsets and finitely bounded union.

[7] Pairwise-completeness is also known as *coherence* (but see footnote 3 on page 128).

Suppose a dcpo D is the generic universe of an exclusion theory. Then an element x of D is completely coprime iff it is an *atom*, that is, iff

$$x \neq \bot \quad \text{and} \quad \text{if} \;\; \bot \neq y \sqsubseteq x \;\; \text{then} \;\; y = x.$$

Proof. Let $\mathcal{U}$ be a subset system over Σ that is closed with respect to subsets and finitely bounded union. An element of $\mathcal{U}$ is completely coprime iff it is the least satisfier $\mathrm{s}(p)$ of some satisfiable member p of Σ. Since $\mathcal{U}$ is closed with respect to subsets, $\mathrm{s}(p) = \{p\}$. But the singleton subsets of $\mathcal{U}$ are precisely the atoms of $\mathcal{U}$. $\lrcorner$

Since the completely coprimes of D are the atoms of D and because of (4.19), every element x of D is the supremum of the set of atoms below or equal to x:

$$x \;=\; \bigsqcup(\downarrow\! x \cap \mathrm{a}(D)),$$

where $\mathrm{a}(D)$ is the set of atoms of D. A dcpo D whose elements have this property is called *atomic*.[8]

(4.29) Theorem

(i) Every subset system that is closed with respect to subsets and finitely bounded union is a bounded-complete atomic dcpo with completely coprime atoms.

(ii) If D is a bounded-complete atomic dcpo with completely coprime atoms, then D is isomorphic to the subset system $\{\downarrow\! x \cap \mathrm{a}(D) \mid x \in D\}$ over $\mathrm{a}(D)$, which is closed with respect to subsets and finitely bounded union.

Proof. It remains to prove (ii). Let D be a bounded-complete atomic dcpo with $\mathrm{a}(D) \subseteq \mathrm{p}(D)$. Then $\mathrm{p}(D) = \mathrm{a}(D)$. For suppose $x \in \mathrm{p}(D)$; since x is the supremum of atoms, there is an atom y such that $x \sqsubseteq y$; hence $x = y$. So D is coprime-algebraic and (4.21)(ii) implies that $\{\downarrow\! x \cap \mathrm{a}(D) \mid x \in D\}$ is closed with respect to finitely bounded union. In order to see that $\{\downarrow\! x \cap \mathrm{a}(D) \mid x \in D\}$ is closed with respect to subsets, it suffices to show that for every two bounded sets X and Y of atoms, whenever $\bigsqcup X = \bigsqcup Y$ then $X = Y$. But if $x \sqsubseteq \bigsqcup Y$ for some $x \in X$ then $x = y$ for some $y \in Y$. $\lrcorner$

As for the compact elements of a bounded-complete atomic dcpo D with completely coprime atoms, we have that every element of $\mathrm{k}(D)$ is a finite join

[8] Some authors use ‘atomistic’ instead of ‘atomic’.

of atoms. Moreover, it follows that the meet of two compact elements is compact in turn and that join and meet in $\mathrm{k}(D)$, if existent, satisfy the distributivity law. So $\mathrm{k}(D)$ is a meet semilattice with finite bounded joins, which is distributive, and every element is a finite join of atoms. Conversely:

(4.30) Proposition Suppose P is a meet semilattice with finite bounded joins, which is distributive, and every element is a finite join of atoms. Then its ideal completion $\mathrm{I}(P)$ is a bounded-complete atomic dcpo with completely coprime atoms.

Proof. By (4.24), it suffices to show that every atom x of P is coprime. Suppose $x \sqsubseteq y \sqcup z$ with $y, z \in P$. Then $x = x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z)$. Hence $x \sqcap y \sqsubseteq x$ and $x \sqcap z \sqsubseteq x$. Since x is an atom, $x \sqcap y = x$ or $x \sqcap z = x$, that is, $x \sqsubseteq y$ or $x \sqsubseteq z$. ⌟

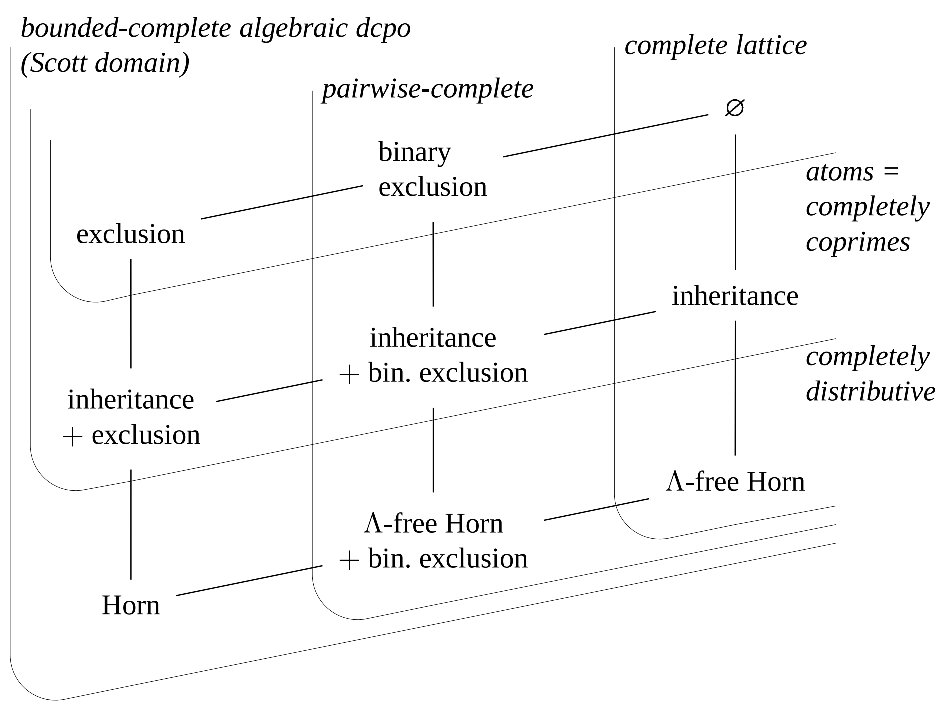

FIGURE 25 Inheritance theories vs. domain properties

When exclusion is restricted to binary exclusion, the corresponding domains are pairwise-complete. In the case of *full binary exclusion*, i.e. if all pairs of primitive predicates are excluded, we get a *flat* domain, that is, every element except the least one is an atom. Table 1 provides an overview of the classification of Horn and inheritance theories given so far. For a condensed version see Figure 25, which illustrates how an increasing restriction of expressivity gives rise to additional properties of the generic universe.

We close this section with another order-theoretic characterization of the generic universe of exclusion theories. A bounded-complete dcpo D is said to satisfy *atomic separation* if for every two of its elements x and y, with $x \not\sqsubseteq y$, there is an atom z of D such that $z \sqsubseteq x$ and $z \not\sqsubseteq y$. The effect of atomic separation becomes more transparent in the following equivalent formulation:

$$\text{if} \quad \forall z \in \mathrm{a}(D)\,(z \sqsubseteq x \rightarrow z \sqsubseteq y) \quad \text{then} \quad x \sqsubseteq y.$$

That is, if ${\downarrow}x \cap \mathrm{a}(D) \subseteq {\downarrow}y \cap \mathrm{a}(D)$ then $x \sqsubseteq y$. In particular, atomic separation implies that $x = \bigsqcup({\downarrow}x \cap \mathrm{a}(D))$, for every $x \in D$, and hence that D is atomic. Conversely, if D is atomic, it clearly satisfies atomic separation. Therefore:

(4.31) Proposition A bounded-complete dcpo is the generic universe of an exclusion theory iff it satisfies atomic separation and has completely coprime atoms.

(4.32) Remark *(Plural lattices)* Atomic separation has been used to axiomatize *plurality*. According to the definition of Link (1998), a *plural lattice* is a bottomless join semilattice which satisfies atomic separation and has completely coprime atoms. That is, a plural lattice is a complete atomic Boolean algebra minus the bottom element. When viewed from the perspective of plurality, the entity types (except the most general one) determined by an exclusion theory are pluralities built of atomic entity types, where the formation of pluralities is subject to the exclusion constraints of the theory. The atomic entity types correspond to the (satisfiable) primitive concepts.

Γ	$\mathrm{C}(\Gamma)$	$D \simeq \mathrm{I}(P) \simeq \mathrm{C}(\Gamma)$	$P \simeq \mathrm{k}(D)$
Horn theory $\varphi \preceq \psi$	nonempty intersection $+$ upwards directed union $=$ inductive intersection system	Scott domain $=$ bounded-complete algebraic dcpo	conditional join semilattice $+$ least element (if nonempty) $=$ ordered set $+$ finite bounded joins
Λ-free Horn theory	intersection $+$ upwards directed union $=$ inductive closure system	complete algebraic lattice	join semilattice $+$ least element $=$ ordered set $+$ finite joins
simple inheritance $+$ exclusions $p \preceq q,\ \varphi \preceq \Lambda$	nonempty intersection $+$ finitely bounded union (bounded $+$ directed union)	completely distributive Scott domain $=$ bounded-complete coprime-algebraic dcpo	ordered set $+$ finite bounded joins $+$ finite coprime decomposition
simple inheritance $+$ binary exclusions $p \preceq q,\ p \wedge q \preceq \Lambda$	nonempty intersection $+$ pairwise bounded union	pairwise-complete completely distributive Scott domain	ordered set $+$ finite pairwise bounded joins $+$ finite coprime decomposition
simple inheritance $p \preceq q$	intersection $+$ union $=$ complete subset lattice	complete coprime-algebraic lattice	ordered set $+$ finite joins $+$ finite coprime decomposition
exclusions $\varphi \preceq \Lambda$	subsets $+$ finitely bounded union	bounded-complete atomic dcpo with completely coprime atoms	meet semilattice $+$ finite bounded joins $+$ finite atom decomposition $+$ distributivity
binary exclusions $p \wedge q \preceq \Lambda$	subsets $+$ pairwise bounded union $=$ coherence space	pairwise-complete atomic dcpo with completely coprime atoms	meet semilattice $+$ finite pairwise bounded joins $+$ finite atom decomposition $+$ distributivity
empty theory	power set lattice	complete atomic Boolean lattice	distributive lattice $+$ finite atom decomposition
full binary exclusion	empty set and singleton sets	flat domain $=$ atoms $+$ bottom	flat ordered set

TABLE 1 Classification of theories by generic universes and vice versa

Part II

Observational Logic

5

Observational Theories

As we have seen at the end of Chapter 3, Horn theories are not expressive enough to fully cover the examples of linguistic classification presented in Chapter 1. We therefore switch to predicates that are built from primitive predicates by finite conjunction *and disjunction*. Following Vickers (1989, 1999) and Smyth (1992) we call such predicates *observational*.[1] The reason for choosing this terminology is essentially the lack of negation: an absent property is not regarded as an observable one, whereas the conjunction and the disjunction of observable properties is observable in turn.

As in Chapters 2 and 3, the main theme of the present chapter is the interplay between a theory and its ordered generic universe. After introducing models of observational theories in Section 5.1, we explore the properties of their canonical universe in Section 5.2. In Section 5.3 the framework is applied to choice system theories and the induction of theories from instances. Finally, in Section 5.4, we reflect on the status of negation and conditionals in observational theories.

5.1 Theories and Models

5.1.1 *Predicates, Statements, and Theories*

Let Σ be a set of primitive predicates with V and Λ not among them. The set $\mathrm{T}[\Sigma]$ of *observational predicates* over Σ is inductively defined to be the least set that consists of all members of Σ, of V and Λ, and of $\varphi \wedge \psi$ and $\varphi \vee \psi$ whenever φ and ψ belong to $\mathrm{T}[\Sigma]$. The predicate operators '$\wedge$' and '$\vee$' have their standard first-order meaning; e.g. '$A \vee B$' stands for the predicate abstract '$\{x \mid Ax \vee Bx\}$'; the predicates 'V' and 'Λ' respectively abbreviate '$\{x \mid x = x\}$' and '$\{x \mid x \neq x\}$', cf. Section 3.1.1.

[1] Alternatively one could call them *(propositional finitary) geometric*, where 'geometric' points to applications in algebraic geometry; see Vickers 1993, 1999, Mac Lane and Moerdijk 1992, Chap. X. It should be added that Vickers and Smyth also count *infinite disjunctions* as observable properties.

In the following we frequently apply the *principle of term induction*, which has already been widely used in Chapter 3. It says that a certain property holds of all observational predicates just in case it holds of all primitives, of V and Λ, and of $\varphi \wedge \psi$ and $\varphi \vee \psi$ whenever it holds of φ and ψ.

Recall from Section 3.1.1 that '$A \preceq B$' stands for '$\forall x(Ax \rightarrow Bx)$'. A (*universal monadic*) *observational statement* over Σ is a statement of the form $\varphi \preceq \psi$ where φ and ψ are observational predicates over Σ. It is sometimes convenient to use (universally quantified) *biconditionals* instead of conditionals. In this case, the predicate operator '$\equiv$' takes the place of '$\preceq$', where '$A \equiv B$' stands for '$\forall x(Ax \leftrightarrow Bx)$'. The logical interdependence between '$\equiv$' and '$\preceq$' is as follows:

(5.1)

$$\begin{gathered} A \equiv B \quad \text{iff} \quad A \preceq B \;\text{ and }\; B \preceq A, \\ A \equiv A \wedge B \quad \text{iff} \quad A \preceq B \quad \text{iff} \quad A \vee B \equiv B. \end{gathered}$$

An *observational theory* over Σ is a set of observational statements over Σ. Given two observational theories Γ and Γ' over Σ, we say that Γ *entails* Γ', in symbols, $\Gamma \vdash \Gamma'$, if $\Gamma \vdash \alpha$ for every $\alpha \in \Gamma'$. The theories Γ and Γ' are said to be *equivalent* iff they entail each other. (Recall that $\vdash$ is entailment by any sound and complete inference calculus for first-order logic with identity.)

5.1.2 *Term Algebra*

Taking up an algebraic perspective one can pair the *operators* $\wedge$, $\vee$, Λ, and V with *operations*, named by the same symbols, on the set $\mathrm{T}[\Sigma]$ of observational predicates over Σ. For example, $\wedge$ is a two-place operation that takes each pair $\langle\varphi,\psi\rangle$ of members of $\mathrm{T}[\Sigma]$ to $\varphi \wedge \psi$. The quadruple $\langle\wedge,\vee,\Lambda,\mathrm{V}\rangle$ of operations thereby provides $\mathrm{T}[\Sigma]$ with the structure of an *algebra of type* $\langle 2,2,0,0\rangle$, the so-called *term algebra over* Σ.

In general, an *algebra* consists of a *carrier set* A and a family $\langle f_i\rangle_{i\in I}$ of *operations* on A. If f_i has arity k_i then $\langle k_i\rangle_{i\in I}$ is called the *(similarity) type* of the algebra. A *homomorphism of algebras* of the same similarity type is a function h of carrier sets, subject to the condition that

$$h(f_i(x_1,\ldots,x_{k_i})) \;=\; f_i(h(x_1),\ldots,h(x_{k_i}))$$

(with corresponding operations of both algebras named by identical symbols).

The term algebra $\mathrm{T}[\Sigma]$ over Σ can be characterized (up to isomorphism) by the property that every homomorphism from $\mathrm{T}[\Sigma]$ to an algebra of the same type is uniquely determined by its effect on Σ. More precisely:

(5.2) Proposition If f is a function from Σ to an algebra A of type $\langle 2,2,0,0\rangle$, then there is a unique algebra homomorphism $\hat{f}$ from $\mathrm{T}[\Sigma]$ to A such that $\hat{f}(p) = f(p)$ for every member p of Σ.

$$\begin{tikzcd} \Sigma \arrow[r, hook] \arrow[dr, "f"'] & \mathrm{T}[\Sigma] \arrow[d, dashed, "\hat{f}"] \\ & A \end{tikzcd}$$

(5.3) Corollary If f is a function from Σ to A and h is a homomorphism of algebras from A to B then $\widehat{h \circ f} = h \circ \hat{f}$.

Given the situation of (5.2), $\hat{f}$ is said to be the *(homomorphic) extension* of f to $\mathrm{T}[\Sigma]$. The property (5.2) is *universal* in the sense that it characterizes the algebra $\mathrm{T}[\Sigma]$ *uniquely up to isomorphism*. It is common to speak of any such algebra as "the" $\langle 2, 2, 0, 0\rangle$-algebra *freely generated* by Σ.

5.1.3 *Normal Forms*

An observational predicate φ over Σ is said to be in *conjunctive normal form* if φ is Λ or V or $\varphi_1 \wedge \ldots \wedge \varphi_m$, where each φ_i is a disjunction of primitive predicates. Similarly, a predicate is in *disjunctive normal form* if it is Λ or V or a disjunction of conjunctions of primitive predicates. The proof of the following fact, which uses essentially term induction and distributivity, is standard and can be adapted from any textbook on classical propositional logic.

(5.4) Proposition Every observational predicate is equivalent to one in conjunctive normal form as well as to one in disjunctive normal form.

In particular, every observational theory Γ is equivalent to one consisting solely of statements $\varphi \preceq \psi$ where φ is in disjunctive normal form and ψ is in conjunctive normal form. Moreover, one can discard statements with $\varphi = \Lambda$ or $\psi = \mathrm{V}$, because such statements are tautologies.

It follows by introduction and elimination of conjunction that if a statement $\varphi \preceq \psi_1 \wedge \psi_2$ of a theory Γ is replaced by the two statements $\varphi \preceq \psi_1$ and $\varphi \preceq \psi_2$, the resulting theory is equivalent to Γ. The same is true, according to introduction and elimination of disjunction, if a statement $\varphi_1 \vee \varphi_2 \preceq \psi$ is replaced by $\varphi_1 \preceq \psi$ and $\varphi_2 \preceq \psi$. Hence, by term induction, every observational theory Γ is equivalent to one that consists solely of statements of the form

$$p_1 \wedge \ldots \wedge p_m \preceq q_1 \vee \ldots \vee q_n,$$

$$\mathrm{V} \preceq q_1 \vee \ldots \vee q_n, \qquad \text{and} \qquad p_1 \wedge \ldots \wedge p_m \preceq \Lambda,$$

with $p_i, q_j \in \Sigma$. In addition, we can assume that each primitive occurs at most once on either side of a statement. An observational theory consisting solely of statements of this type is said to have *normal form*. Moreover, each primitive can be assumed to occur at most once in the whole statement because otherwise the statement follows from reflexivity by weakening. We then speak of a *reduced normal form*.

(5.5) Proposition Every observational theory is equivalent to a theory in (reduced) normal form (over the same set of primitives).

5.1.4 *Interpretations and Models*

Repeating the definitions of Chapter 3, a *(set-valued) interpretation* of the vocabulary Σ consists of a *universe* U and an *interpretation function* M taking members of Σ to subsets of U. Application of (5.2) allows us to state in a precise way what it means to extend M to $\mathrm{T}[\Sigma]$. We simply have to consider the power set $\wp(U)$ of U as an algebra of type $\langle 2, 2, 0, 0\rangle$ with operations $\langle \cap, \cup, \varnothing, U\rangle$. The homomorphic extension of M to $\mathrm{T}[\Sigma]$, here also written as M, satisfies:

$$M(\varphi \wedge \psi) \;=\; M(\varphi) \cap M(\psi),$$
$$M(\varphi \vee \psi) \;=\; M(\varphi) \cup M(\psi),$$
$$M(\Lambda) \;=\; \varnothing \quad \text{and} \quad M(\mathrm{V}) \;=\; U.$$

Following a long tradition in philosophy and logic, we call $M(\varphi)$ the *extension of* φ. Instead of $M(\varphi)$ we also write $[\![\varphi]\!]_M$, or $[\![\varphi]\!]$, for short.

The *satisfaction* relation $\models$ (or $\models_M$) corresponding to M is defined by: $x \models \varphi$ iff $x \in M(\varphi)$. In particular, supplementing (3.2),

(5.6) $$x \models \varphi \vee \psi \quad \text{iff} \quad x \models \varphi \;\; \text{or} \;\; x \models \psi.$$

The definition of models of observational theories can be transferred without changes from that of Horn theories: $\varphi \preceq \psi$ *is true* with respect to an interpretation M iff $M(\varphi) \subseteq M(\psi)$. A *model* of an observational theory Γ over Σ is an interpretation of Σ with respect to which all statements of Γ are true. Note the following immediate consequence of definitions:

(5.7) Proposition Equivalent theories have identical models.

Consider an interpretation of Σ with universe U and satisfaction relation $\models$. Given two members x and y of U we say that x *is specialized by* y (notation: $x \sqsubseteq y$) if y satisfies every member of Σ that is satisfied by x. It follows that

(5.8) $$x \sqsubseteq y \quad \text{iff} \quad \forall \varphi \in \mathrm{T}[\Sigma]\, (x \models \varphi \;\rightarrow\; y \models \varphi).$$

Proof. The straighforward inductive proof runs analogously to that of (3.3). Suppose $x \sqsubseteq y$ and $\varphi \in \mathrm{T}[\Sigma]$. We need to show that if $x \models \varphi$ then $y \models \varphi$. For primitive φ this is just the definition of $\sqsubseteq$. If φ is V or Λ there is nothing to show since $x \not\models \Lambda$ and $y \models \mathrm{V}$. If $x \models \varphi_1 \wedge \varphi_2$ then x satisfies φ_1 and φ_2, both of which are, by induction hypothesis, satisfied by y; hence $y \models \varphi_1 \wedge \varphi_2$. Similarly, if $x \models \varphi_1 \vee \varphi_2$ then x satisfies φ_1 or φ_2; consequently, by induction hypothesis, y satisfies φ_1 or φ_2 and thus $\varphi_1 \vee \varphi_2$. ⌟

Notice that specialization is in general not a partial order but only a preorder because two different members of U can specialize each other without being identical. If this case does not occur, i.e., if $\sqsubseteq$ is antisymmetric and thus a partial ordering, we say that the interpretation satisfies the condition of *identity of indiscernibles*; it then holds that

(5.9) $$x = y \quad \text{iff} \quad \forall \varphi (x \models \varphi \leftrightarrow y \models \varphi).$$

5.2 The Canonical Model

5.2.1 *Canonical Model and Generic Universe*

The following definition of satisfaction of observational predicates over Σ by subsets of Σ straightforwardly extends the definition given in Section 3.2.1 for conjunctive predicates: a subset X of Σ satisfies a member p of Σ iff $p \in X$; now term induction sets in: every subset of Σ satisfies V, no subset satisfies Λ; X satisfies $\varphi \wedge \psi$ iff X satisfies φ and ψ; and finally, in accordance with (5.6), X satsifies $\varphi \vee \psi$ iff X satisfies φ or ψ. It follows by (5.8) that for every two subsets X and Y of Σ,

(5.10) $$X \subseteq Y \qquad \text{iff} \qquad \forall \varphi \in \mathrm{T}[\Sigma]\, (X \models \varphi \rightarrow Y \models \varphi).$$

So two subsets of Σ are identical iff they satisfy the same predicates over Σ.

Suppose Γ is an observational theory over Σ. As in the case of Horn theories we say that a subset X of Σ is *consistently* Γ*-closed* if

$$X \models \varphi \;\rightarrow\; X \models \psi \quad \text{for every } (\varphi \preceq \psi) \in \Gamma.$$

Again, $\mathrm{C}(\Gamma)$ is the system of consistently Γ-closed subsets of Σ. Moreover, the canonical interpretation function M_Γ, which takes p to $\{X \in \mathrm{C}(\Gamma) \mid p \in X\}$, defines a model of Γ with universe $\mathrm{C}(\Gamma)$ – the *canonical model of* Γ. By (5.10):

(5.11) Proposition Specialization on the canonical universe is set inclusion.

Since specialization is thus a partial ordering on $\mathrm{C}(\Gamma)$, the canonical model satisfies *identity of indiscernibles* (5.9).

For each interpretation M of Σ with universe U, let ε_M be the function from U to $\wp(\Sigma)$ such that

(5.12) $$\varepsilon_M(x) \;=\; \{p \in \Sigma \mid x \vDash_M p\}.$$

By definition of specialization, $x \sqsubseteq y$ iff $\varepsilon_M(x) \subseteq \varepsilon_M(y)$. So ε_M is an order embedding of U into $\wp(\Sigma)$ if M satisfies identity of indiscernibles. Moreover, one shows easily by term induction that

(5.13) $$x \vDash_M \varphi \qquad \text{iff} \qquad \varepsilon_M(x) \vDash \varphi.$$

It follows that if M is a model of an observational theory Γ then ε_M is a homomorphism of models from M to M_Γ (cf. Section 3.2.2).

Following the terminology of Section 3.2.2, we call a model of Γ *generic* in case it is isomorphic to M_Γ. We speak of the universe $\mathrm{U}(\Gamma)$ of such a model as “the” *generic universe* of Γ, and of the members of $\mathrm{U}(\Gamma)$ as the *generic entities* determined by Γ. According to (5.13) and (5.11), the generic model of Γ is the “largest” Γ-model satisfying identity of indiscernibles in the sense that every other such model M is embedded in the generic one via ε_M.

Another consequence of (5.13) is that, for every model M of Γ, if $\mathrm{M}_\Gamma(\varphi) \subseteq \mathrm{M}_\Gamma(\psi)$ then $M(\varphi) \subseteq M(\psi)$; hence, since first-order logic is complete,

(5.14) $$\text{if} \quad \mathrm{M}_\Gamma(\varphi) \subseteq \mathrm{M}_\Gamma(\psi) \quad \text{then} \quad \Gamma \vdash \varphi \preceq \psi.$$

The canonical model of an observational theory thus satisfies *equivalence of coextensives*, that is, two observational predicates over Σ are Γ-equivalent if and only if they have identical extension with respect to the generic model of Γ. In Chapter 6 we will prove this fact without making use of completeness of first-order logic.

5.2.2 *The Principle of Duality*

As in classical propositional logic, the *dual* φ^{d} of an observational predicate φ over Σ is defined as the predicate one gets by replacing every occurrence of $\wedge$ and $\vee$ in φ respectively by $\vee$ and $\wedge$. More precisely, duality on $\mathrm{T}[\Sigma]$ is inductively defined as follows: $p^{\mathrm{d}} = p$ if $p \in \Sigma$, $\mathrm{V}^{\mathrm{d}} = \Lambda$, $\Lambda^{\mathrm{d}} = \mathrm{V}$,

$$(\varphi \wedge \psi)^{\mathrm{d}} = \varphi^{\mathrm{d}} \vee \psi^{\mathrm{d}} \qquad \text{and} \qquad (\varphi \vee \psi)^{\mathrm{d}} = \varphi^{\mathrm{d}} \wedge \psi^{\mathrm{d}}.$$

Notice that $(\varphi^{\mathrm{d}})^{\mathrm{d}} = \varphi$, which is easily seen by term induction.

Our goal is to define duality for observational theories in such a way that the canonical universe of the dual theory is “dual” to that of the original theory. To this end, observe that for any subset X of Σ and any predicate φ over Σ,

(5.15) $$\complement X \models \varphi^{\mathrm{d}} \quad \text{iff} \quad X \not\models \varphi,$$

where $\complement X = \Sigma \setminus X$ is the complement of X in Σ.

Proof. By term induction: Suppose $p \in \Sigma$; then $\complement X$ satisfies p^{d} $(=p)$ iff $p \notin X$, that is, iff $X \not\models p$. Moreover, $\complement X \models \mathrm{V}^{\mathrm{d}}$ iff $X \not\models \mathrm{V}$ because both statements are false, and $\complement X \models \Lambda^{\mathrm{d}}$ iff $X \not\models \Lambda$ because both statements are true. In addition, $\complement X \models (\varphi \wedge \psi)^{\mathrm{d}}$ iff $\complement X \models \varphi^{\mathrm{d}}$ or $\complement X \models \psi^{\mathrm{d}}$, which, by induction hypothesis, is the case iff $X \not\models \varphi$ or $X \not\models \psi$, that is, iff $X \not\models \varphi \wedge \psi$. In the same way, the induction hypothesis implies that $\complement X \models (\varphi \vee \psi)^{\mathrm{d}}$ iff $X \not\models \varphi \vee \psi$. ⌟

If we now define $(\varphi \preceq \psi)^{\mathrm{d}}$ as $\psi^{\mathrm{d}} \preceq \varphi^{\mathrm{d}}$ and the *dual theory* Γ^{d} of an observational theory Γ as $\{\alpha^{\mathrm{d}} \mid \alpha \in \Gamma\}$ then, by contraposition and (5.15),

$$\mathrm{C}(\Gamma^{\mathrm{d}}) = \{\complement X \mid X \in \mathrm{C}(\Gamma)\}.$$

Moreover, since $X \subseteq Y$ iff $\complement Y \subseteq \complement X$, duality of theories carries over to duality of ordered generic universes:

(5.16) Proposition The ordered generic universe of Γ is (isomorphic to) the order dual of the generic universe of Γ^{d}.

The generic universe of observational theories, when viewed as a partially ordered set, is thus subject to the *duality principle*: if all of them satisfy a certain order-theoretic property, they satisfy the dual property too.

(5.17) Example Let Γ be the theory over $\{a, b, c\}$ with statements $a \wedge b \preceq \Lambda$ and $b \preceq c$. Its dual theory Γ^{d} consists of $\mathrm{V} \preceq a \vee b$ and $c \preceq b$. Figure 26 depicts the canonical universe of Γ and that of its dual. As shown above, one gets the elements of $\mathrm{C}(\Gamma^{\mathrm{d}})$ by taking the complements of all elements of $\mathrm{C}(\Gamma)$ with respect to $\{a, b, c\}$, which reverses the inclusion order.

5.2.3 *Properties of the Canonical Universe*

It is easy to see that in general neither the intersection nor the union of consistently Γ-closed sets is again consistently Γ-closed. So, infimum (greatest lower bound) and supremum (least upper bound) with respect to specialization on $\mathrm{C}(\Gamma)$, if existent at all, do not necessarily coincide with intersection and union. Moreover, an infimum or supremum of a set of generic entities may not exist even when the set is bounded below or above, respectively. Let us illustrate these facts by two simple examples.

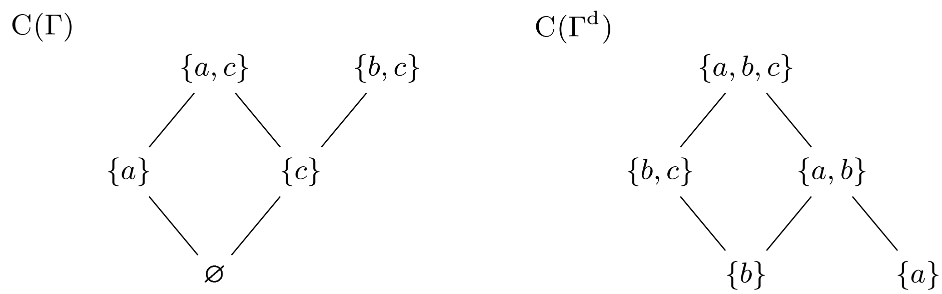

FIGURE 26 Canonical universes of Γ and Γ^{d}

(5.18) Example Take the theory that comedies and thrillers are the only entertaining novels. More explicitly and slightly regimented: something is entertaining and a novel if and only if it is a comedy or a thriller; that is,

$$\textit{entertaining} \wedge \textit{novel} \;\equiv\; \textit{comedy} \vee \textit{thriller}.$$

The theory given by this statement over the set $\{$*entertaining*, *novel*, *comedy*, *thriller*$\}$ of predicates is referred to as '*Fiction*'. The generic entities deter-

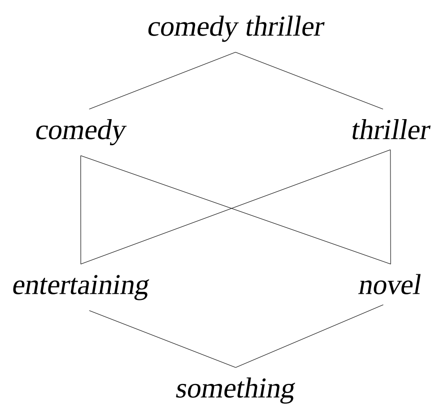

FIGURE 27 Generic entities of *Fiction*

mined by *Fiction* are represented by the *Fiction*-closed subsets of primitive predicates. To list them, there is first the empty set representing "something". Then there are the sets $\{$*entertaining*$\}$ and $\{$*novel*$\}$, that is, the (generic) entertaining thing and the (generic) novel. In addition, we have $\{$*entertaining*, *novel*, *comedy*$\}$ and $\{$*entertaining*, *novel*, *thriller*$\}$ representing respectively the comedy and the thriller, which both are entertaining novels, according to *Fiction*. And last not least, there is the comedy thriller, represented by $\{$*entertaining*, *novel*, *comedy*, *thriller*$\}$. The resulting specialization relation between generic entities is shown in Figure 27. Obviously neither the entertaining thing and the novel have a common least upper bound nor the comedy and the thriller have a common greatest lower bound though in both cases there are upper and lower bounds.

(5.19) Example Consider the taxonomic tree of Figure 28, which classifies lexemes with respect to their inflection type. Assuming subclassifications to be

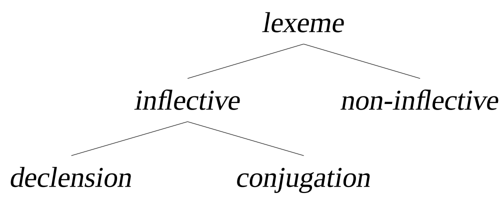

FIGURE 28 Classification of lexemes by inflection type

exhaustive, this taxonomy corresponds to the theory given by the statements

$$\mathit{lexeme} \equiv \mathit{inflective} \vee \mathit{non\text{-}inflective},$$
$$\mathit{inflective} \equiv \mathit{declension} \vee \mathit{conjugation},$$
$$\mathit{inflective} \wedge \mathit{non\text{-}inflective} \equiv \Lambda, \quad \mathit{declension} \wedge \mathit{conjugation} \equiv \Lambda.$$

As indicated in Section 1.4, the canonical universe of such a theory is given by the maximal chains of the tree, i.e. consists of the sets $\{\mathit{non\text{-}infl}, \mathit{lex}\}$, $\{\mathit{decl}, \mathit{infl}, \mathit{lex}\}$, and $\{\mathit{conj}, \mathit{infl}, \mathit{lex}\}$ (with predicates abbreviated appropriately). In addition, the empty set is consistently closed as well since we did not include the statement $\mathit{lexeme} \equiv \mathrm{V}$ into our theory (which would imply that everything in the universe of discourse is a lexeme). So, the canonical universe is as depicted by Figure 29. Now observe that though $\varnothing$ is the infimum of the set of nonempty members of the universe, it is not given by intersection.

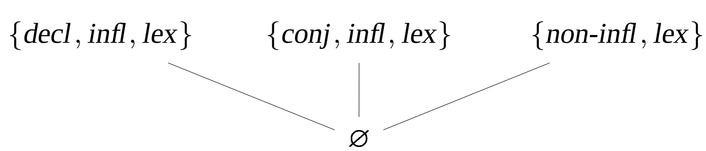

FIGURE 29 Canonical universe of inflection taxonomy

What holds for the canonical universe of observational theories in general is that supremum and infimum exist for upwards and downwards directed subsets, respectively, and coincide with union and intersection, respectively. This is an immediate consequence of the following observation:

(5.20) Lemma Suppose $\mathcal{S}$ is a subset system over Σ and φ is an observational predicate over Σ.

- ***(i)*** If $\mathcal{S}$ is upwards directed then $\bigcup\mathcal{S} \models \varphi$ iff $X \models \varphi$ for some $X \in \mathcal{S}$.
- ***(ii)*** If $\mathcal{S}$ is downwards directed then $\bigcap\mathcal{S} \models \varphi$ iff $X \models \varphi$ for every $X \in \mathcal{S}$.

Proof. We prove (i) by term induction. In light of the proof of (3.6) it remains to verify the induction steps for disjunction. By (5.6) and induction hypothesis, $\bigcup\mathcal{S} \models \varphi \vee \psi$ iff for some members X and Y of $\mathcal{S}$, $X \models \varphi$ or $Y \models \psi$, that is, iff $\varphi \vee \psi$ is satisfied by some member of $\mathcal{S}$ since $\mathcal{S}$ is upwards directed. As for (ii) we apply (5.15) to (i): If $\mathcal{S}$ is downwards directed then $\mathcal{S}' = \{\complement X \mid X \in \mathcal{S}\}$ is upwards directed and $\bigcap\mathcal{S} = \complement(\bigcup\mathcal{S}')$. Hence $\bigcap\mathcal{S} \models \varphi = (\varphi^{\mathrm{d}})^{\mathrm{d}}$ iff $\bigcup\mathcal{S}' \not\models \varphi^{\mathrm{d}}$ iff, for every $X \in \mathcal{S}$, $\complement X \not\models \varphi^{\mathrm{d}}$, i.e. $X \models \varphi$. ┘

(5.21) Proposition The canonical universe of observational theories is closed with respect to the union of upwards directed sets and with respect to the intersection of downwards directed sets.

Notice that in contrast to Horn theories adjoining $\Sigma \cup \{\Lambda\}$ to $\mathrm{C}(\Gamma)$ in general does not turn $\mathrm{C}(\Gamma)$ into a complete lattice since $\mathrm{C}(\Gamma)$ is not necessarily closed with respect to intersection.

(5.22) Corollary If X is a subset of some member of $\mathrm{C}(\Gamma)$ then the set of all members of $\mathrm{C}(\Gamma)$ with subset X has minimal elements.

Proof. The set $\{Y \in \mathrm{C}(\Gamma) \mid X \subseteq Y\}$ is nonempty, by assumption, and downwards directed-complete, by (5.21). Zorn's Lemma therefore guarantees minimal elements. ┘

By the same line of reasoning it follows that every observational predicate φ over Σ satisfiable in the generic universe U of some observational theory over Σ has *minimal satisfiers* in U, namely the minimal members of the nonempty and downwards directed-complete set of satisfiers of φ. Similarly, each subset of U with upper bounds in U has *minimal upper bounds* in U.

Besides being upwards and downwards directed complete, the (ordered) generic universe satisfies the property of *relative atomicity*.[2]

[2] An ordered set P is *relatively atomic* iff, for all $x, y \in P$, if $x < y$ then there are u, v such that $x \leqslant u < v \leqslant y$ and there is no w with $u < w < v$ (i.e. $[u, v] = \{w \mid u \leqslant w \leqslant v\}$ is a *gap* in P).

(5.23) Corollary The generic universe of an observational theory is relatively atomic.

Proof. [3] Suppose $X, Y \in \mathrm{C}(\Gamma)$, $X \subseteq Y$, and $p \in Y \setminus X$. By (5.21) and Zorn's Lemma, there is a maximal element X' of $\{Z \in \mathrm{C}(\Gamma) \mid X \subseteq Z \subseteq Y \setminus \{p\}\}$, and a minimal element Y' of $\{Z \in \mathrm{C}(\Gamma) \mid X' \cup \{p\} \subseteq Y' \subseteq Y\}$. ⊣

In Chapters 3 and 4, considerable effort has been invested to characterize the canonical universe of simple inheritance theories and Horn theories as subset systems and as ordered sets – cumulated in Table 1 at the end of Chapter 4. In the case of observational theories, in contrast, no satisfying order-theoretic characterization is known that picks out all dcpos that arise as generic universes of observational theories.[4] On the positive side there is the result of Speed (1972a) which says that a dcpo D is the generic universe of an observational theory just in case D is the projective limit of a projective system of finite ordered sets; we come back to this characterization in Section 9.3.

Though we are lacking an order-theoretic characterization, it is not difficult to characterize the canonical universes of observational theories as subset systems – which is the topic of the next section.

5.2.4 *Locally Closed Systems*

Let $\mathcal{U}$ be a subset system over Σ. We say that a subset Y of Σ is *locally a member of* $\mathcal{U}$ if for every *finite* subset F of Σ there is a member X of $\mathcal{U}$ such that $Y \cap F = X \cap F$. The system $\mathcal{U}$ is said to be *locally closed* if it contains every subset of Σ which is locally a member of $\mathcal{U}$.[5] Clearly every subset system $\mathcal{U}$ over a *finite* set Σ is locally closed. For if Y is locally a member of $\mathcal{U}$ then, since Σ is finite, there is a member X of $\mathcal{U}$ such that $Y = Y \cap \Sigma = X \cap \Sigma = X$.

Let us first verify that being locally closed is a necessary condition for being the canonical universe of an observational theory.

(5.24) Proposition The canonical universe of an observational theory is locally closed.

Proof. Suppose a subset Y of Σ is locally a member of $\mathrm{C}(\Gamma)$, $\varphi \preceq \psi$ belongs to Γ, and $Y \models \varphi$. We need to show that $Y \models \psi$. Since the set F of all members

[3] Adapted from Droste and Göbel 1990, p. 292.

[4] See Droste and Göbel 1990, p. 307 and the references therein for an example of an upwards and downwards directed-complete ordered set which does not represent the generic universe of an observational theory. Or take Johnstone's (1982, p. 46) example of an upwards (and downwards) directed-complete ordered set for which there is no sober topology inducing the given order. Another example worth mentioning is that of Düntsch (1982), which shows that the ordered set of prime filters of a distributive lattice without unit need not be the generic universe of an observational theory; compare this with (6.12) below.

[5] This definition of locally closed systems is inspired by Davey 1973.

of Σ occurring in the statement $\varphi \preceq \psi$ is finite, there is an $X \in \mathrm{C}(\Gamma)$ such that $X \cap F = Y \cap F$; hence $Y \vDash \psi$. ⌟

Notice that (5.24) allows another simple proof of (5.21). For example, if $\mathcal{S} \subseteq \mathrm{C}(\Gamma)$ is upwards directed then for every finite $F \subseteq \Sigma$ there is an $X \in \mathcal{S}$ such that $\bigcup \mathcal{S} \cap F \subseteq X$; hence $X \cap F = \bigcup \mathcal{S} \cap F$. So, $\bigcup \mathcal{S}$ is locally a member of $\mathrm{C}(\Gamma)$ and thus belongs to $\mathrm{C}(\Gamma)$, by (5.24).

(5.25) Example Suppose Σ is an *infinite* set of primitives. Then there is no observational theory over Σ with canonical universe $\mathcal{U} = \{\{p\} \mid p \in \Sigma\}$. This is so because $\varnothing$ is locally a member $\mathcal{U}$: for every finite subset F of Σ there is a $p \in \Sigma$ such that $p \notin F$; hence $\varnothing \cap F = \{p\} \cap F$.

In order to show that the property of being locally closed characterizes the canonical universe of observational theories, it remains to prove the reverse of (5.24). Let $\mathcal{U}$ be a locally closed system over Σ. We repeat the line of reasoning that precedes (3.12): in case there is an observational theory Γ with canonical universe $\mathcal{U}$, then $\mathrm{M}_\Gamma(\varphi) = \{X \in \mathcal{U} \mid X \vDash \varphi\}$; therefore, an observational statement $\varphi \preceq \psi$ is true in M_Γ iff every member of $\mathcal{U}$ satisfying φ satisfies ψ. Hence Γ, if existent at all, is equivalent to the theory $\Gamma(\mathcal{U})$ over Σ, where

(5.26) $$(\varphi \preceq \psi) \in \Gamma(\mathcal{U}) \quad \text{iff} \quad \forall X \in \mathcal{U}\,(X \vDash \varphi \rightarrow X \vDash \psi).$$

By definition, $\mathcal{U} \subseteq \mathrm{C}(\Gamma(\mathcal{U}))$. As for the reverse inclusion, we first consider the case when Σ is finite.

(5.27) Lemma If Σ is finite then $\mathrm{C}(\Gamma(\mathcal{U})) = \mathcal{U}$.

Proof. Suppose $X \in \mathrm{C}(\Gamma(\mathcal{U}))$. Let φ be the conjunction of all members of X (in any order). Then $\varphi \preceq \Lambda$ does not belong to $\Gamma(\mathcal{U})$ since X satisfies φ. Hence φ is satisfiable in $\mathcal{U}$. Let $Y_1, \ldots, Y_n$ be the minimal satisfiers of φ in $\mathcal{U}$ and let ψ_i be the conjunction of the members of Y_i. Then $\varphi \preceq \psi_1 \vee \ldots \vee \psi_n$ belongs to $\Gamma(\mathcal{U})$. Hence $X \vDash \psi_k$ for some k, that is, $Y_k \subseteq X$. So $X = Y_k \in \mathcal{U}$. ⌟

Given a subset S of Σ, let $\mathcal{U}|_S$ be the system $\{X \cap S \mid X \in \mathcal{U}\}$ of subsets of S – the *restriction of $\mathcal{U}$ to S*. We need the following fact:

(5.28) $$Y \in \mathrm{C}(\Gamma(\mathcal{U})) \;\rightarrow\; Y \cap S \in \mathrm{C}(\Gamma(\mathcal{U}|_S)).$$

Proof. Suppose $Y \in \mathrm{C}(\Gamma(\mathcal{U}))$, $(\varphi \preceq \psi) \in \Gamma(\mathcal{U}|_S)$ and $Y \cap S \vDash \varphi$. We need to show that $Y \cap S \vDash \psi$. Since φ and ψ are predicates over S, they are satisfied by Y iff they are satisfied by $Y \cap S$. Moreover, $(\varphi \preceq \psi) \in \Gamma(\mathcal{U})$. To see this, recall that $\varphi \preceq \psi$ belongs to $\Gamma(\mathcal{U}|_S)$ iff, for every $X \in \mathcal{U}$, $X \cap S \models \varphi \rightarrow X \cap S \models \psi$, that is, according to what has just been said, iff $X \models \varphi \rightarrow X \models \psi$. So $Y \models \psi$, since $Y \cap S \models \varphi$ and $Y \in \mathrm{C}(\Gamma(\mathcal{U}))$; hence $Y \cap S \models \psi$. ⌋

Now suppose Y is a member of $\mathrm{C}(\Gamma(\mathcal{U}))$ and F is a finite subset of Σ. Then, by (5.28), $Y \cap F$ belongs to $\mathrm{C}(\Gamma(\mathcal{U}|_F))$, and hence to $\mathcal{U}|_F$, by (5.27). Since $\mathcal{U}$ is locally closed, it follows that $Y \in \mathcal{U}$. All in all, $\mathrm{C}(\Gamma(\mathcal{U})) = \mathcal{U}$. Consequently:

(5.29) Proposition Every locally closed system over Σ is the canonical universe of an observational theory over Σ.

Suppose $\mathcal{U}$ is a locally closed system over Σ. The theory $\Gamma(\mathcal{U})$ as defined by (5.26) consists of all observational statements over Σ that are true in its canonical model $\mathrm{M}_{\Gamma(\mathcal{U})}$. Since $\Gamma(\mathcal{U})$ is infinite even if Σ is finite, we are interested in more parsimonious theories that are equivalent to $\Gamma(\mathcal{U})$ and thus have the same canonical model. An obvious way to get such a theory is to take the reduced normal form $\Gamma_{\mathrm{rnf}}(\mathcal{U})$ of $\Gamma(\mathcal{U})$, see (5.5). Notice that $\Gamma_{\mathrm{rnf}}(\mathcal{U})$ stays finite if Σ is finite. Moreover, in determining $\Gamma_{\mathrm{rnf}}(\mathcal{U})$ we can do without $\Gamma(\mathcal{U})$: just apply (5.26) only to those statements $\varphi \preceq \psi$, where φ is a conjunction of primitives, ψ is a disjunction of primitives, and no primitive occurs more than once in the statement.

(5.30) Remark *(Nonredundant Basis)* The theory $\Gamma_{\mathrm{rnf}}(\mathcal{U})$ is usually highly redundant. For if $\Gamma_{\mathrm{rnf}}(\mathcal{U})$ contains $\varphi \preceq \psi$ then also $\varphi \wedge p \preceq \psi$ and $\varphi \preceq \psi \vee p$, for every $p \in \Sigma$ not occurring in φ and ψ. The reason is that $\Gamma_{\mathrm{rnf}}(\mathcal{U})$, like $\Gamma(\mathcal{U})$, is deductively closed – this time not with respect to all observational statements but with respect to statements in reduced normal form. For finite Σ, Ganter (1999) shows how to construct a *nonredundant* observational theory over Σ with canonical universe $\mathcal{U}$.[6]

5.2.5 *Some Examples*

The following examples, besides the last one, are primarily counter examples for showing that certain properties, like algebraicity, do not hold for generic universes of observational theories in general.

It is straightforward to see that not every generic universe satisfies the *descending chain condition*; that is, nonempty subsets of a generic universe may lack minimal elements; witness Example (5.31). Consequently, by the principle of duality, also the *ascending chain condition* does not hold in general – which is of course everything but surprising because one of the key aspects of

[6]Ganter's approach generalizes that of Guigues and Duquenne 1986.

Scott domains is that suprema of (ascending) chains of compact elements need not be compact in turn; hence, typically, there are chains of compact elements lacking maximal elements (see Section 4.1).

(5.31) Example *(No descending chain condition)* Suppose Σ is a countable set $\{a_1, a_2, \ldots\}$ of primitives and Γ is the theory $\{a_n \preceq a_{n+1} \mid n \geqslant 1\}$ over Σ. Then $\mathrm{C}(\Gamma) = \{A_1, A_2, \ldots\} \cup \{\varnothing\}$, with $A_n = \{a_m \mid m \geqslant n\}$. Clearly, the descending chain $A_1, A_2, \ldots$ has no minimal elements.

It is not difficult to find an observational theory whose canonical universe is not algebraic:

(5.32) Example *(Non-algebraicity)* Let Γ be the theory over $\Sigma = \{a_1, a_2, \ldots\}$ with statements

$$a_{n+1} \preceq a_n \quad \text{and} \quad a_n \preceq a_1 \vee a_{n+1} \quad (n > 1).$$

Its canonical universe is precisely the *non-algebraic* ordered set of Figure 22, repeated here as Figure 30, with $A_n = \{a_m \mid m \leqslant n\}$ and $B = \{a_n \mid n > 1\}$.

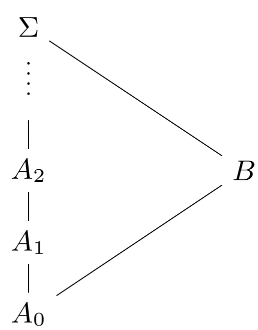

FIGURE 30 Non-algebraic canonical universe

So, we cannot assume that a generic universe has a basis of compact elements. Indeed, there are cases with no compact elements at all, as the following example shows.

(5.33) Example *(No compact elements)* Let Σ be the union of two countable sets $A = \{a_1, a_2, \ldots\}$ and $B = \{b_1, b_2, \ldots\}$, and let Γ be the theory over Σ with statements

$$a_{n+1} \preceq a_n, \quad b_{n+1} \preceq b_n, \quad \mathrm{V} \preceq a_1 \vee b_1, \quad a_n \vee b_n \preceq a_{n+1} \vee b_{n+1},$$

for every $n \geqslant 1$. Then

$$\mathrm{C}(\Gamma) = \{A_n \cup B \mid n \in \mathbb{N}_0\} \cup \{A \cup B_n \mid n \in \mathbb{N}_0\} \cup \{A \cup B\},$$

with $A_n = \{a_m \mid m \leqslant n\}$ and $B_n = \{b_m \mid m \leqslant n\}$; see Figure 31. Now observe that $A_m \cup B \sqsubseteq A \cup B = \bigsqcup\{A \cup B_n \mid n \in \mathbb{N}_0\}$ but $A_m \cup B \not\sqsubseteq A \cup B_n$ for every n. So $A_m \cup B$ is not compact, and neither is $A \cup B_m$, by symmetry.

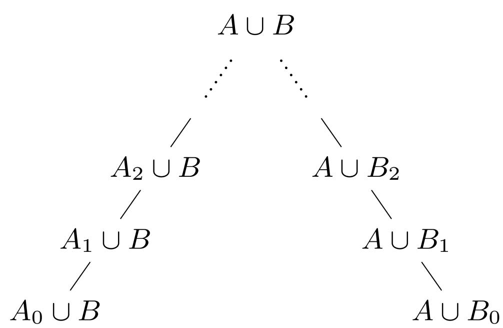

FIGURE 31 Canonical universe without compact elements

The next example shows that even if the generic universe is an algebraic domain D, it need not be the case that every finite subset of $\mathrm{k}(D)$ has a finite set of minimal upper bounds in $\mathrm{k}(D)$. In technical terms, an algebraic generic universe need not be *coherent* algebraic; cf. Section 7.2.

(5.34) Example *(Non-coherence)* Suppose Γ is the theory over $\{a, b\} \cup \{c_0, c_1, \ldots\} \cup \{d_0, d_1, \ldots\}$ that consists of the statements

$$a \wedge b \equiv c_0 \vee d_0, \quad c_n \equiv c_{n+1} \vee d_{n+1}, \quad c_n \wedge d_n \equiv \Lambda \quad (n \geqslant 0).$$

Then $\mathrm{C}(\Gamma)$ consists of $\varnothing$, $A = \{a\}$, $B = \{b\}$, $C = \{a, b\} \cup \{c_0, c_1, c_2, \ldots\}$, and $D_n = \{a, b\} \cup \{c_0, c_1, \ldots, c_{n-1}, d_n\}$, $n \geqslant 0$; see Figure 32. Every member of this ordered set is compact, but the set of minimal upper bounds of $\{A, B\}$ is infinite.

As promised above, our final example is a positive one: it shows that there is an observational theory whose generic universe is an (infinite) binary tree. Of course, we know this already since trees are pairwise-complete, completely distributive Scott domains and therefore can be represented as universes of simple inheritance theories with binary exclusions (Section 4.2.1). Concretely, the primitives of such a theory stand in a one-to-one correspondence to the

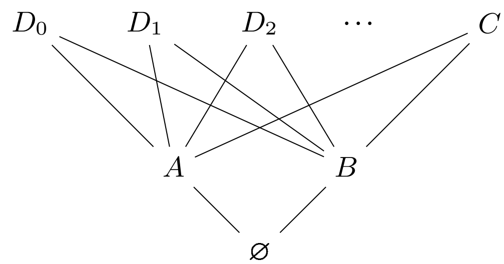

FIGURE 32 Non-coherent algebraic canonical universe

nodes of the tree. The representation used in Example (5.35), in contrast, is based on a much more parsimonious set of primitives, whose cardinality, in the finite case, is logarithmic in the number of nodes.

(5.35) Example *(Binary tree)* Suppose Σ is $\{a_1, a_2, \ldots\} \cup \{b_1, b_2, \ldots\}$ and Γ is given by the statements

$$a_{n+1} \vee b_{n+1} \preceq a_n \vee b_n, \qquad a_n \wedge b_n \preceq \Lambda \qquad (n \geqslant 1).$$

The canonical universe $\mathrm{C}(\Gamma)$ of Γ is an infinite binary tree; witness Figure 33 (where '$a_1 a_2 a_3$' is short for '$\{a_1, a_2, a_3\}$', etc). Notice that $\mathrm{C}(\Gamma)$ is isomorphic to the set of finite strings over $\{a, b\}$ ordered by the prefix relation.

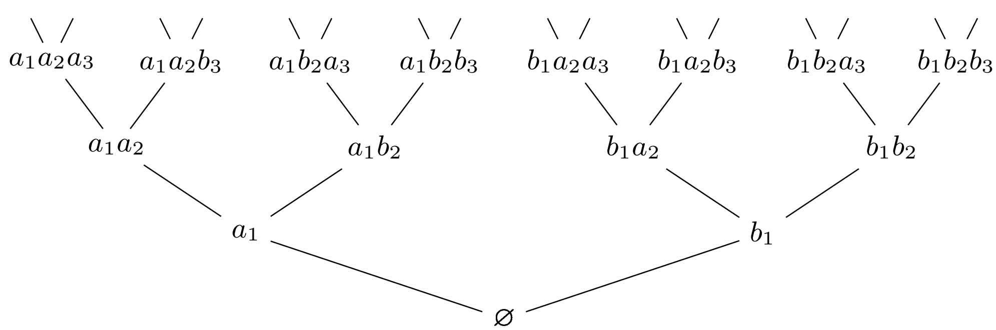

FIGURE 33 An infinite binary tree represented by a canonical universe

5.3 Applications

For a first application of observational theories recall that closed world reasoning as described in Section 2.3 rests on the assumption that a concept implies the disjunction of its immediate subconcepts. Such a condition is obviously expressible by an observational statement. The generic entities (or entity types

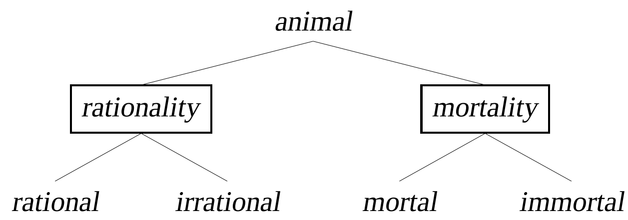

FIGURE 34 Classification of animals by rationality and mortality

or conjunctive concepts) determined by a simple inheritance theory plus closed world assumption are thus given by the canonical universe of an observational theory.

5.3.1 *Trees from Multidimensional Classification*

The taxonomy of German nominals displayed in Figure 3 of Section 1.1 was noticed to involve a certain amount of structural arbitrariness, for the taxonomic tree could be rearranged to the effect that the distinction between definiteness and indefiniteness precedes the subdivision into articles and pronouns.

We also mentioned that this sort of arbitrariness in arranging a taxonomic tree was already observed by Boethius and Abelard in their discussion of Porphyry.[7] So Abelard says in his *Editio super Porphyrium*: "pluraliter ideo dicit genera, quia animal dividitur per rationale animal et irrationale; et rationale per mortale et immortale dividitur; et mortale per rationale et irrationale dividitur."[8] One is thus free to characterize animals first with respect to rationality and thereafter with respect to mortality, or to proceed the other way around. This freedom of choice can be expressed by the AND/OR-tree of Figure 34. With the choice categories *rationality* and *mortality* taken as abbreviations for $\mathit{rational} \vee \mathit{irrational}$ and $\mathit{mortal} \vee \mathit{immortal}$, this AND/OR-tree corresponds to the observational theory Γ consisting of the statements

$$\begin{array}{rclcrcl} \mathit{rational} \wedge \mathit{irrational} & \preceq & \Lambda, & \quad & \mathit{mortal} \wedge \mathit{immortal} & \preceq & \Lambda, \\ \mathit{rational} \vee \mathit{irrational} & \preceq & \mathit{animal}, & \quad & \mathit{mortal} \vee \mathit{immortal} & \preceq & \mathit{animal}, \end{array}$$

which is equivalent to a simple inheritance theory with binary exclusions. Let us add $\mathrm{V} \preceq \mathit{animal}$ to Γ, thereby presuming that everything in the universe of discourse is an animal. The canonical universe $\mathrm{C}(\Gamma)$ of Γ is displayed in Figure 35.

Though $\mathrm{C}(\Gamma)$ is not a tree, it "contains" several trees whose leaves are the maximally specific elements of $\mathrm{C}(\Gamma)$ and whose root is *animal*. One of these

[7] Eco 1984, Sect. 2.2.4.

[8] Cited after Eco, *ibid*, p. 66. (Adelard's *Editio super Porphyrium* is a commentary on Porphyry's *Isagoge*, based on the Latin translation by Boethius. The *Isagoge* in turn is an introduction into Aristotle's *Categoriae*, the first book of his *Organon*.)

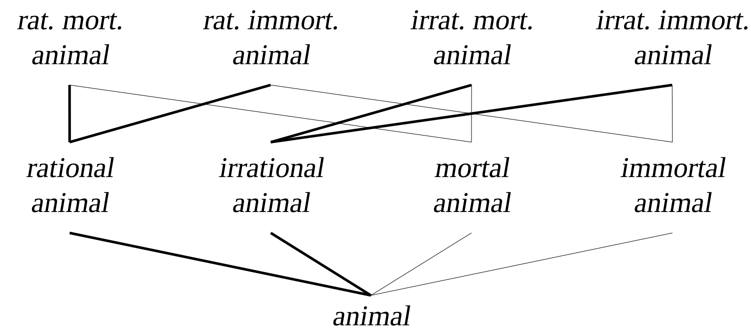

FIGURE 35 Canonical universe given by the AND/OR-tree of Figure 34

trees is indicated in Figure 35 by bold lines. Notice that this tree represents the specialization of animals to rational and irrational ones, of rational animals to mortal and immortal ones, etc, on the level of generic entities. Notice further that the alternative specialization of animals to mortal and immortal ones etc corresponds to the tree given by the non-bold lines of Figure 35.

Since extending the theory Γ by adding statements (not deducible from Γ) has the effect of "deleting" elements of $C(\Gamma)$, we can ask for an extension of Γ whose canonical universe is one of these subtrees. Clearly, adding the statement

$$\mathit{mortal} \vee \mathit{immortal} \preceq \mathit{rational} \vee \mathit{irrational}$$

to Γ yields the bold-line tree (whereas adding the reverse statement gives rise to the complementary tree). The assumption behind this observational statement is that everything which is specified with respect to mortality is also specified with respect to rationality, i.e., there is no generic entity specified with respect to mortality but not with respect to rationality.

It is therefore possible to carve out trees from the generic universe of a multidimensional classification by imposing a linear order on compatible but otherwise independent choice systems. (Notice that (5.35) can be regarded as a general example of this fact.) The precise form of the tree depends on the chosen ordering of choice systems and is thus in a sense arbitrary with respect to the underlying multidimensional classification. All in all, this gives us another illustration of the fact that explicating classification systems in form of observational theories allows us to link properties of the resulting generic universe to specific assumptions inherent in the classification in question.

Let us finally turn to the taxonomy of German nominals mentioned at the beginning of this section. An inheritance hierarchy that is neutral as to whether to first distinguish articles from pronouns or definite from indefinite nominals could be of the form shown in Figure 36. This hierarchy, however, is inadequate because not every nominal word is subject to definiteness, and neither is every pronoun (cf. Figure 3).

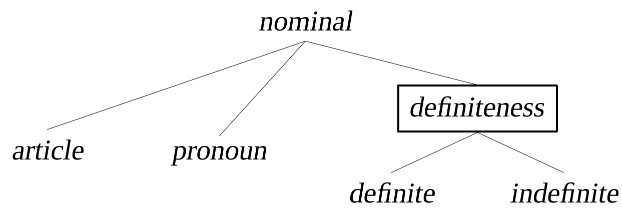

FIGURE 36 Taxonomy of nominal words and definiteness

It is thus reasonable in this case not to regard both choice dimensions as independent but to assume a hierarchy as depicted on the left of Figure 37, where the choice category *definiteness* is meant to imply the *disjunction* of its immediate superordinates (in contrast to the convention of Section 1.2). In terms of observational statements:

$$\mathrm{V} \preceq \mathit{nominal}, \quad \mathit{article} \wedge \mathit{pronoun} \preceq \Lambda, \quad \mathit{definite} \wedge \mathit{indefinite} \preceq \Lambda,$$
$$\mathit{article} \preceq \mathit{nominal}, \quad \mathit{pronoun} \preceq \mathit{nominal}, \quad \mathit{determinative} \preceq \mathit{pronoun},$$
$$\mathit{definite} \vee \mathit{indefinite} \preceq \mathit{article} \vee \mathit{determinative}.$$

The generic universe of this theory is the tree shown on the right of Figure 37. We can conclude that the taxonomic tree of Figure 3 is best seen as the generic universe of an observational theory. Classificational diagrams in linguistics, on the other hand, are primarily used to graphically present such theories (at least implicitly). An example is provided by the diagram on the left of Figure 37; another one by systemic networks, which are the topic of the next section.

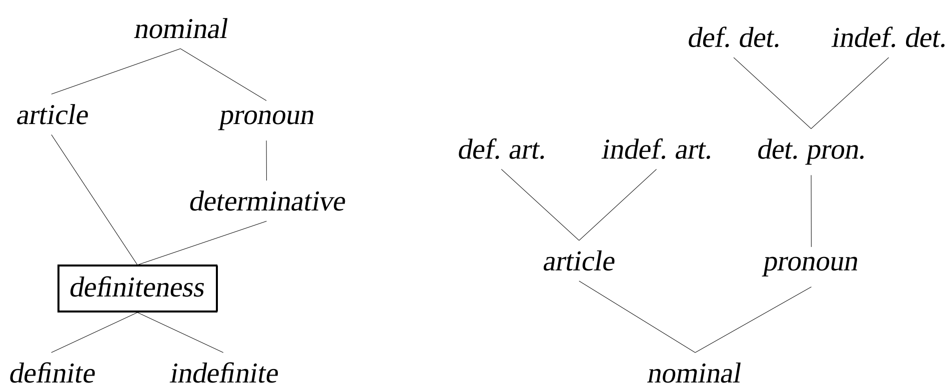

FIGURE 37 Nominal words, definiteness, and canonical universe

5.3.2 *Choice System Theories*

Systemic networks as introduced in Section 1.3 are based on a partition Π of the set Σ of primitive predicates into finite *choice systems*. Each choice system P can have an *entry condition* φ_P, which we assume to be any observational predicate over Σ without occurrences of V and Λ. The intended logical content of such a network is that entry condition and choice system disjunction imply each other and that every two different members of one and the same choice system are incompatible:

$$\varphi_P \;\equiv\; p_1 \vee \ldots \vee p_n \qquad \text{and} \qquad p_i \wedge p_j \;\equiv\; \Lambda \quad (i \neq j),$$

where $P = \{p_1, \ldots, p_n\}$ belongs to Π. Observational theories of this form are henceforth referred to as *choice system theories*. (For a *non-exhaustive* version of choice system theories, in the sense of Chapter 1, one would only require that φ_P is implied by each p_i and thus by $p_1 \vee \ldots \vee p_n$.)

For instance, suppose $P = \{p_1, p_2, p_3\}$ is a choice system of some systemic network and the entry condition of P is graphically presented as follows (cf. Section 1.3):

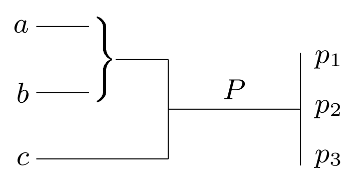

Then the corresponding choice system theory contains the statement

$$(a \wedge b) \vee c \;\equiv\; p_1 \vee p_2 \vee p_3$$

as well as the exclusion statements $p_i \wedge p_j \equiv \Lambda$, for $i \neq j$.

Systemic networks are usually assumed to be *acyclic*. More precisely, what is assumed to be acyclic is the *precedence relation* on Π that is borne by Q to P if some member of Q occurs in the precondition φ_P of P.[9] Figure 38 shows the precedence diagram of Winograd's systemic network of English pronouns as presented in Figure 8 on page 15 (with choice systems named suitably). For possibly infinite Σ, it is convenient to require the precedence relation to be *well-founded*, which implies acyclicity (and is equivalent to acyclicity for finite Σ).

Let us henceforth presume that the precedence relation of a choice system theory is by definition well-founded. Well-foundedness of precedence give us the following identity criterion for generic entities:

[9] Equivalently, one can consider the relation on Σ that is borne by q to p iff q occurs in the precondition of the choice system which p belongs to.

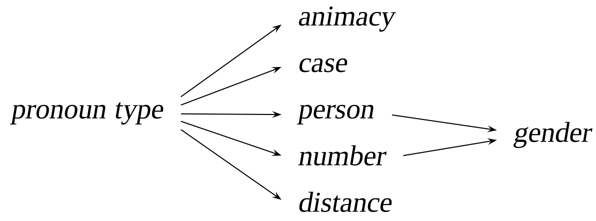

FIGURE 38 Precedence relation of Winograd's pronoun classification

(5.36) Lemma Suppose x and y are generic entities of a choice system theory such that $x \sqsubseteq y$. If, for every *unconditioned* choice system P, x satisfies some member of P whenever y does, then $x = y$.

Proof. Let δ_P be $p_1 \vee \ldots \vee p_n$, for each choice system $P = \{p_1, \ldots, p_n\}$. Suppose $x \sqsubseteq y$, and if P has no entry condition then $y \vDash \delta_P \rightarrow x \vDash \delta_P$. We need to show that $y \sqsubseteq x$. It suffices to show that for every choice system P, if $y \vDash \delta_P$ then $x \vDash \delta_P$; for the members of P are pairwise incompatible and Σ is covered by choice systems. We apply the principle of well-founded (or Noetherian) induction: Assume that $y \vDash \delta_Q \rightarrow x \vDash \delta_Q$ for all Q preceding P. Suppose P has an entry condition and y satisfies δ_P. By induction hypothesis, x satisfies each primitive predicate that occurs in φ_P and is satisfied by y. Hence, by term induction, x satisfies φ_P and thus δ_P. ⌟

The condition described in (5.36) becomes particularly simple if the choice system theory has a *root system*, i.e. a choice system for which no entry condition is defined and which directly or indirectly precedes every other choice system. In other words, the transitive closure of the precedence relation is assumed to have a least element. In this case, every two nontrivial members of the generic universe turn out to be incomparable with respect to specialization:

(5.37) Theorem Rooted choice system theories have a flat generic universe.

Proof. Let Γ be a rooted choice system theory. Then $\varnothing \in \mathrm{C}(\Gamma)$, because V does not occur in any of the preconditions. Suppose $X, Y \in \mathrm{C}(\Gamma)$ with X nonempty and $X \subseteq Y$. Let P_0 be the root system Γ. Since $X \neq \varnothing$, there is a choice system P such that $X \cap P \neq \varnothing$, i.e. $X \vDash \delta_P$ (with δ_P defined as above). If $P \neq P_0$ then $X \vDash \varphi_P$ and hence $X \vDash \delta_Q$ for some Q preceding P. Therefore, by induction, $X \vDash \delta_{P_0}$. Now apply (5.36). ⌟

(5.38) Example *(Exhaustive taxonomic trees)* An *exhaustive taxonomic tree* is a rooted choice system theory where all preconditions are primitive and every choice system except the root is preceded by exactly one choice system;

see (5.19) for a simple example. According to (5.37), exhaustive taxonomic trees induce flat specialization order.

There are variations on this theme: Given a (not necessarily rooted) choice system theory Γ, let us define the *completion* of Γ to be the theory consisting of all statements of Γ plus all statements of the form $\mathrm{V} \equiv p_1 \vee \ldots \vee p_n$, where $\{p_1, \ldots, p_n\}$ is a choice system lacking an entry condition. In other words, each generic entity is enforced to satisfy some member of every unconditioned choice system. Then (5.36) implies:

(5.39) Corollary The specialization order induced by the completion of a choice system theory is discrete, i.e. an antichain.

Suppose Γ is a rooted choice system theory over Σ. Then $\mathrm{C}(\Gamma)$ is flat, by (5.37). Hence Γ is "extensionally equivalent" to a full binary exclusion theory Γ' in the sense that Γ and Γ' have order-isomorphic generic universes. Being not necessarily closed with respect to intersection, $\mathrm{C}(\Gamma)$ is in general not the canonical universe of a Horn theory let alone of an exclusion theory over Σ. So the set Σ' of primitives of Γ' is typically different from Σ. Indeed, the standard method of defining Σ' is to introduce *one primitive predicate for each nonempty member of* $\mathrm{C}(\Gamma)$ whereas Γ' consists of all statements $p \wedge q \equiv \Lambda$ with $p, q \in \Sigma'$ and $p \neq q$. The topic of translating choice system theories into Horn theories will turn up again in Section 8.3.1.

5.3.3 *Closures of Subset Systems*

Though interesting in itself, the main purpose of the following construction is to pave the way for Section 5.3.4. Let $\mathcal{C}$ be a certain class of observational statements over Σ (e.g. the class of Horn statements). We speak of the members of $\mathcal{C}$ as $\mathcal{C}$-*statements* and of the subsets of $\mathcal{C}$ as $\mathcal{C}$-*theories*.

Suppose $\mathcal{U}$ is a subset system over Σ. As in (5.26) and (3.12), we can associate with $\mathcal{U}$ the set $\Gamma_{\mathcal{C}}(\mathcal{U})$ of all $\mathcal{C}$-statements $\varphi \preceq \psi$ such that, for all $X \in \mathcal{U}$, if $X \models \varphi$ then $X \models \psi$. By definition, $\mathcal{U} \subseteq \mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$.

(5.40) Lemma If Γ is a $\mathcal{C}$-theory such that $\mathcal{U} \subseteq \mathrm{C}(\Gamma)$, then $\Gamma \subseteq \Gamma_{\mathcal{C}}(\mathcal{U})$ and $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U})) \subseteq \mathrm{C}(\Gamma)$.

Proof. Suppose $\mathcal{U} \subseteq \mathrm{C}(\Gamma)$ for some $\mathcal{C}$-theory Γ. That is, for all $X \in \mathcal{U}$ and $(\varphi \preceq \psi) \in \Gamma$, if $X \models \varphi$ then $X \models \psi$. In other words, $\Gamma \subseteq \Gamma_{\mathcal{C}}(\mathcal{U})$. Consequently, $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U})) \subseteq \mathrm{C}(\Gamma)$. $\lrcorner$

Let us say that $\mathcal{U}$ is $\mathcal{C}$-*definable* if $\mathcal{U}$ is the canonical universe of a $\mathcal{C}$-theory (which is the case, for instance, if $\mathcal{C}$ is the class of Horn statements and $\mathcal{U}$ is an inductive intersection system). By (5.40):

Statement class $\mathcal{C}$	$\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is closure of $\mathcal{U}$ with respect to
observational	local membership
Horn	nonempty intersection $+$ directed union
Λ-free Horn	intersection $+$ directed union
simple inheritance	intersection $+$ union
exclusion	subsets $+$ finitely bounded union
simple inheritance $+$ exclusion	nonempty intersection $+$ finitely bounded union

TABLE 2 Relationship between $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ and $\mathcal{U}$ depending on $\mathcal{C}$

(5.41) Proposition If $\mathcal{U}$ is $\mathcal{C}$-definable then $\mathcal{U} = \mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$.

Consider the case that $\mathcal{U}$ is not $\mathcal{C}$-definable. The question then is how $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is related to $\mathcal{U}$. By (5.40), it follows that $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is the *least $\mathcal{C}$-definable subset system containing $\mathcal{U}$*. Now suppose we have a characterization of the $\mathcal{C}$-definable subset systems in terms of certain "closure properties", like 'being closed with respect to intersection'. According to what has just been said, $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is the *closure* of $\mathcal{U}$ with respect to these properties. For example, if $\mathcal{C}$ is the class of Horn theories then $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is the closure of $\mathcal{U}$ with respect to nonempty intersection and directed union. For we know from Section 3.2.4 that the Horn-definable subset systems are precisely the inductive intersection systems. Similarly, we can apply this line of reasoning to the various characterizations of Chapter 3 as well as to the characterization given in Section 5.2.4:

(5.42) Theorem Let $\mathcal{U}$ be a subset system. Then the subset system $\mathrm{C}(\Gamma_{\mathcal{C}}(\mathcal{U}))$ is related to $\mathcal{U}$ as indicated in Table 2.

Given an arbitrary observational theory Γ over Σ, we call a $\mathcal{C}$-theory Γ' over Σ a *$\mathcal{C}$-projection* of Γ, if Γ entails Γ' and Γ' entails every other $\mathcal{C}$-theory (over Σ) entailed by Γ. By definition, every two $\mathcal{C}$-projections of Γ are equivalent. Clearly $\Gamma_{\mathcal{C}}(\mathrm{C}(\Gamma))$ is a $\mathcal{C}$-projection of Γ. So, if Γ' is a $\mathcal{C}$-projection of Γ then $\mathrm{C}(\Gamma')$ is the least $\mathcal{C}$-definable subset system containing $\mathrm{C}(\Gamma)$. Moreover, if $\mathcal{C}$ is listed in the first column of Table 2 then $\mathrm{C}(\Gamma')$ is the closure of $\mathrm{C}(\Gamma)$ with respect to the respective properties in the second column. The following simple example illustrates these facts.

(5.43) Example Recall the little inflection taxonomy of Example (5.19). Its canonical universe is depicted in Figure 29. The least Horn-definable subset system (over the same set of primitives) containing this canonical universe is shown in Figure 39; it is the closure of the original system with respect to nonempty intersection.[10] In addition, it is the canonical universe of the Horn-theory

$$\mathit{decl} \preceq \mathit{infl}, \quad \mathit{conj} \preceq \mathit{infl}, \quad \mathit{infl} \preceq \mathit{lex}, \quad \mathit{non\text{-}infl} \preceq \mathit{lex},$$
$$\mathit{infl} \wedge \mathit{non\text{-}infl} \preceq \Lambda, \quad \mathit{decl} \wedge \mathit{conj} \preceq \Lambda,$$

which (consequently) is a Horn-projection of the original theory.

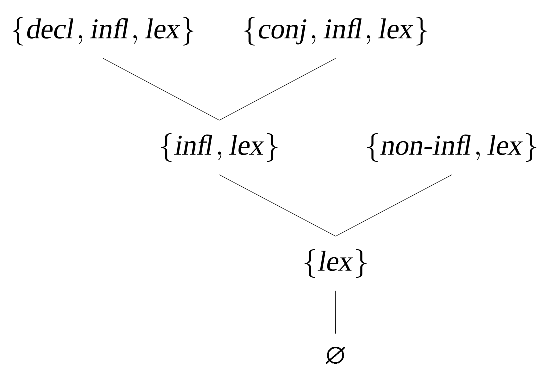

FIGURE 39 Closure of Figure 29 with respect to nonempty intersection

5.3.4 *Induction of Theories by Classifications*

Consider the situation that a certain set U of objects is classified with respect to a set Σ of properties. In other words, we are given a satisfaction relation $\models$ from U to Σ, i.e. an interpretation M of Σ in $\wp(U)$ (see Section 5.1.4). In the terminology of Barwise and Seligman (1997), the triple $\langle U, \Sigma, \models \rangle$ is a *classification*. Given a classification one can ask for a theory that explains the data. To make this precise, we need to fix the type of theory we are interested in. For example, one can ask for a simple inheritance theory with or without exclusions, a Horn theory with or without Λ, or an observational theory in general.

Let $\mathcal{C}$ be a class of observational statements over Σ. We call a $\mathcal{C}$-theory Γ a *complete $\mathcal{C}$-theory of M* if, first, every statement of Γ is true with respect to M, i.e. M is a model of Γ, and, second, Γ entails every $\mathcal{C}$-statement that holds in M, that is, for all $(\varphi \preceq \psi) \in \mathcal{C}$,

[10]Finite subset systems are always closed with respect to directed union.

(5.44) $\qquad$ if $\quad M(\varphi) \subseteq M(\psi) \quad$ then $\quad \Gamma \vdash \varphi \preceq \psi$.

It is an immediate consequence of definitions that a complete $\mathcal{C}$-theory of M is unique up to equivalence. Moreover, there is a trivial way to get a complete theory: take the set $\Gamma_{\mathcal{C},M}$ of all $\mathcal{C}$-statements that are true with respect to M:

$$\Gamma_{\mathcal{C},M} \;=\; \{(\varphi \preceq \psi) \in \mathcal{C} \mid M(\varphi) \subseteq M(\psi)\}.$$

In the remainder of this section we explore more closely the relationship between a given classification and the canonical universe of its complete $\mathcal{C}$-theory. Recall from Section 5.1.4 that a classification, i.e. a satisfaction relation $\models$ from U to Σ, determines a specialization relation $\sqsubseteq$ on U. In addition, there is a (pre)order-preserving function ε_M from U to $\wp(\Sigma)$ that takes $x \in U$ to $\{p \in \Sigma \mid x \models p\}$ (Section 5.2.1). Let $\mathcal{U}_M$ be the image $\{\varepsilon_M(x) \mid x \in U\}$ of U by ε_M. In general ε_M is not one-to-one because there is no guarantee of ***identity of indiscernibles***, i.e. different elements of U may satisfy exactly the same members of Σ. But of course, $\mathcal{U}_M \simeq U/{\sim}$, where $x \sim y$ iff $\varepsilon_M(x) = \varepsilon_M(y)$.

Now notice that the canonical $\mathcal{C}$-theory $\Gamma_{\mathcal{C}}(\mathcal{U}_M)$ associated with $\mathcal{U}_M$ (see Section 5.3.3) coincides with $\Gamma_{\mathcal{C},M}$; for by (5.13), $\varepsilon_M(x) \models \varphi$ iff $x \models \varphi$. So we can apply (5.42) to characterize the canonical universe of a complete $\mathcal{C}$-theory of M. For instance, if Γ is a complete Horn theory of M then $\mathrm{C}(\Gamma)$ is the closure of $\mathcal{U}_M$ with respect to nonempty intersection and directed union; similarly, if Γ is a complete simple inheritance theory of M then $\mathrm{C}(\Gamma)$ is the closure of $\mathcal{U}_M$ with respect to intersection and union.

(5.45) Example Consider again the inflection taxonomy of Example (5.19). Suppose we start off not with a taxonomy but with the classification table of Figure 40. Let M be the associated interpretation function. The corresponding subset system $\mathcal{U}_M$ is depicted on the right of Figure 40. If $\mathcal{C}$ is the class of observational statements, the theory given in (5.19) is a complete $\mathcal{C}$-theory of M with canonical universe $\mathcal{U}_M$. In contrast, the theory stated in (5.43) is a complete Horn-theory of M, whose canonical universe is the closure of $\mathcal{U}_M$ with respect to nonempty intersection, as depicted in Figure 39.

(5.46) Example Let Σ be $\{a, b, c, d, e\}$. Suppose U consists of the seven elements $x_1, x_2, \ldots, x_7$ which are classified by members of Σ according to the classification table of Figure 41. In addition, the figure shows the specialization order on $U/{\sim}$ induced by the given classification (where x_4 and x_6 are indiscernible, i.e. $x_4 \sim x_6$), as well as the corresponding subset system $\mathcal{U}_M$ over Σ. Figure 42 provides an overview of the canonical universes of several complete $\mathcal{C}$-theories of M, with varying $\mathcal{C}$. At the top of the figure, there is the canonical universe of a complete simple inheritance theory of M; it can be determined by

	lexeme	*non-infl*	*infl*	*decl*	*conj*
x_1	X		X		X
x_2					
x_3	X	X			
x_4	X		X	X	

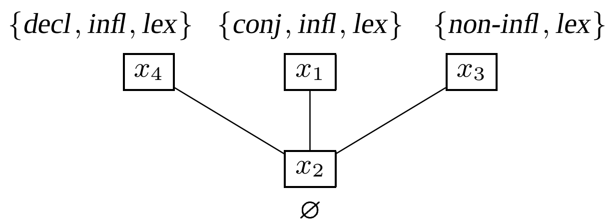

FIGURE 40 Inflection classification and induced specialization order

	a	b	c	d	e
x_1	X	X			X
x_2	X				
x_3		X			
x_4	X	X	X	X	
x_5	X		X		
x_6	X	X	X	X	
x_7	X		X	X	

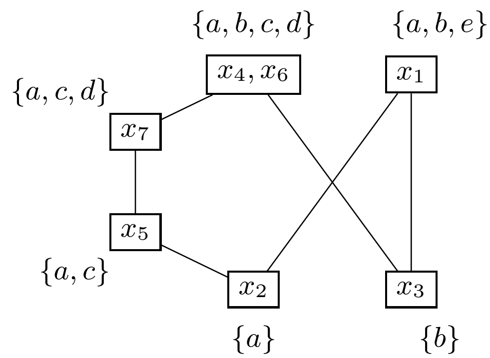

FIGURE 41 Classification table and induced specialization order

closing the subset system $\mathcal{U}_M$ with respect to intersection and union. A (nonredundant) complete simple inheritance theory of M is given by the statements $d \preceq c$, $c \preceq a$, $e \preceq a$, and $e \preceq b$. The diagram below the top on the left depicts the closure of $\mathcal{U}_M$ with respect to intersection of nonempty subsets and union of bounded subsets. It is the canonical universe of the extension of the above simple inheritance theory by the exclusion statement $c \wedge e \preceq \Lambda$. Addition of the Horn statement $b \wedge c \preceq d$ further weakens the closure properties of the associated canonical universe. If the statement $b \wedge c \preceq d$ is added to the simple inheritance theory before the exclusion statement $c \wedge e \preceq \Lambda$, the resulting effect on the respective canonical universes is as depicted by the right branch of Figure 42. Finally, adding the statements $\mathrm{V} \preceq a \vee b$ and $a \wedge b \preceq c \vee e$ leads to a complete observational theory of M, whose canonical universe necessarily is $\mathcal{U}_M$.

Suppose Σ is finite. Consider a Λ-free Horn theory Γ over Σ that is complete with respect to a classification $\langle U, \Sigma, \models\rangle$ (with interpretation function M). Then $\mathrm{C}(\Gamma)$ is a closure system over Σ and hence the image of the closure operation cl on Σ that takes each subset Y of Σ to $\bigcap\{X \in \mathrm{C}(\Gamma) \mid Y \subseteq X\}$; see

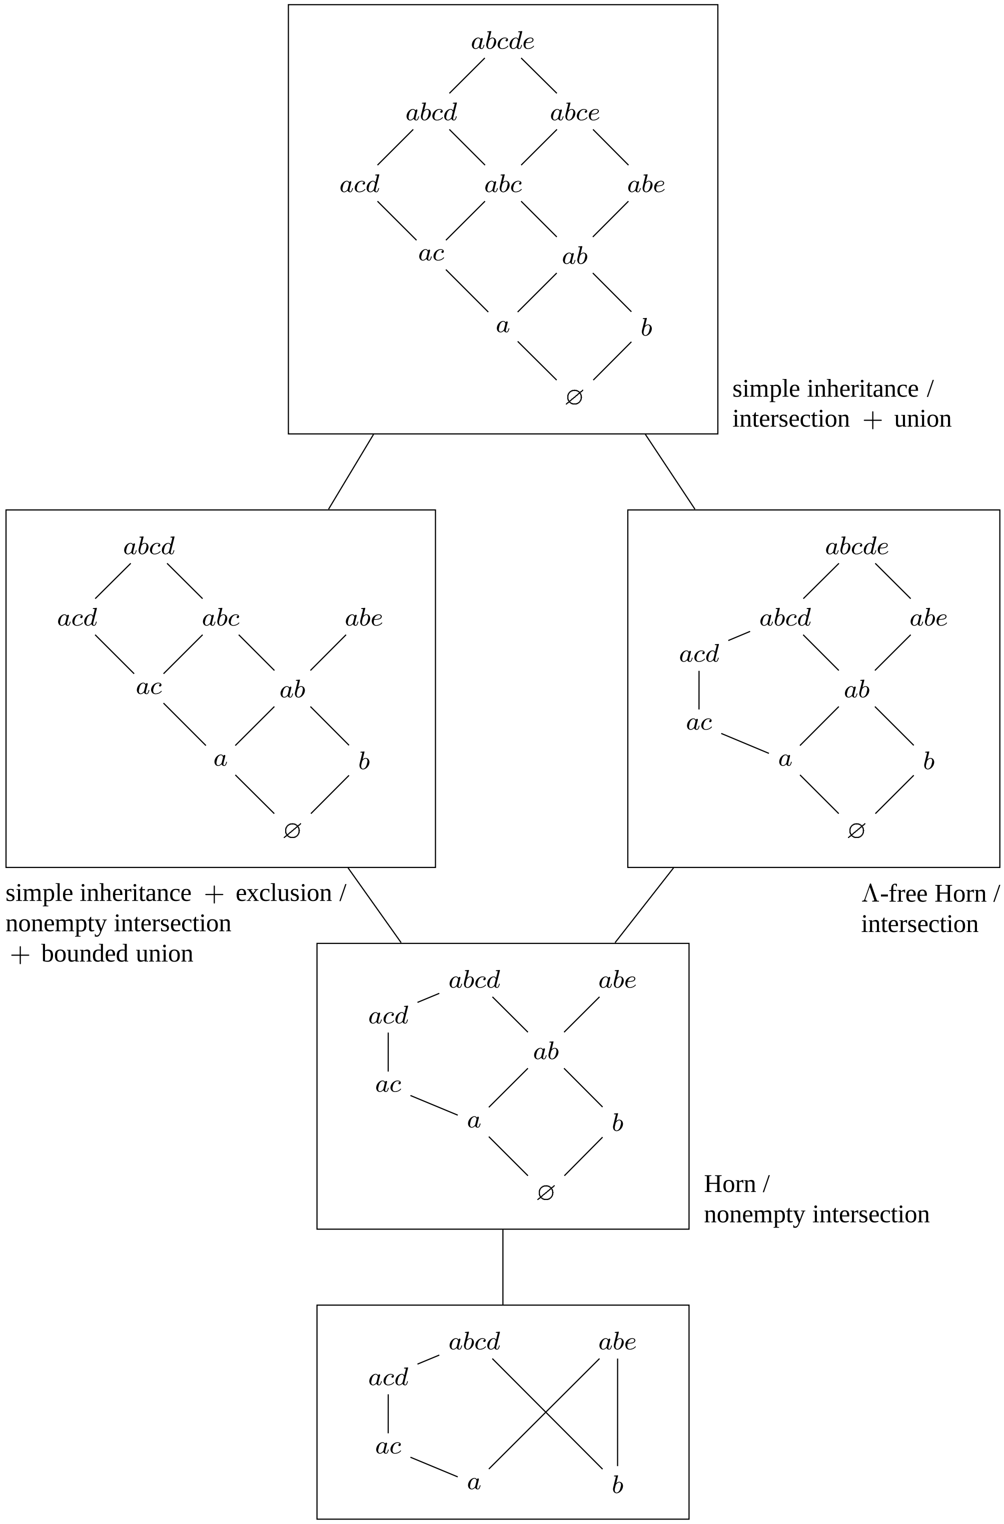

FIGURE 42 Canonical universe of complete $\mathcal{C}$-theory of classification with varying $\mathcal{C}$

(3.16).[11] On the other hand, we know from (5.42) that $\mathrm{C}(\Gamma)$ is the intersection closure of $\mathcal{U}_M$, that is, the members of $\mathrm{C}(\Gamma)$ are the sets of the form $\bigcap\{\varepsilon_M(x) \mid x \in V\}$, with $V \subseteq U$. To combine these two facts we employ the following abbreviations:

$$Y^{\blacktriangleleft} = \{x \in U \mid \forall p \in Y (x \vDash p)\} = \bigcap\{M(p) \mid p \in Y\},$$
$$V^{\blacktriangleright} = \{p \in \Sigma \mid \forall x \in V (x \vDash p)\} = \bigcap\{\varepsilon_M(x) \mid x \in V\},$$

for $Y \subseteq \Sigma$ and $V \subseteq U$. Since $Y^{\blacktriangleleft} = \{x \in U \mid Y \subseteq \varepsilon_M(x)\}$, we have:

(5.47) $$\mathrm{cl}(Y) = (Y^{\blacktriangleleft})^{\blacktriangleright}.$$

So the closure operation associated with the closure system $\mathrm{C}(\Gamma)$ is naturally given by the *Galois connection* associated with $\vDash$.

One can use (5.47) to read off the closure of a subset Y of Σ from a classification table as follows: First take the set $Y^{\blacktriangleleft}$ of all $x \in U$ that satisfy all $p \in Y$, i.e. find all rows in the table that are marked at all columns belonging to Y. Then $\mathrm{cl}(Y)$ $(= (Y^{\blacktriangleleft})^{\blacktriangleright})$ is the set of all $q \in \Sigma$ that are satisfied by every $y \in Y^{\blacktriangleleft}$; that is, take all columns of the table that are marked at every row belonging to $Y^{\blacktriangleleft}$.

(5.48) Remark *(Formal Concept Analysis)* The goal of *Formal Concept Analysis*, introduced by Rudolf Wille in the early 1980s, is to build hierarchical classifications from instances.[12] In the terminology of Formal Concept Analysis, $\langle U, \Sigma, \vDash\rangle$ is a *formal context*, $M(p)$ is the *extent* of $p \in \Sigma$, and $\varepsilon_M(x)$ is the *intent* of $x \in U$. The *concept lattice* of a formal context is essentially the closure system over Σ given by the closure operator (5.47). In the finite case (or, more generally, in the algebraic case), the concept lattice of a formal context is identical to the canonical universe of a complete Λ-free Horn theory of the context in question; see also Osswald and Petersen 2002, Osswald and Petersen 2003.

(5.49) Remark *(Hypothesis space and inductive bias)* It is tempting to apply the terminology of *machine learning* to the problem of inducing theories from classifications.[13] The $\mathcal{C}$-theories constitute the *hypothesis space* $\mathcal{H}$ of the learning problem, whereas the *version space* with respect to $\mathcal{H}$ and M consists of all $\mathcal{C}$-theories with model M. The commitment to statement type $\mathcal{C}$ determines

[11] Since Σ is finite, every closure system over Σ is inductive and every closure operation on Σ is finitary.

[12] See Davey and Priestley 1990, Chap. 11 for a short introduction and Ganter and Wille 1999 for a detailed account. Similar constructions can be found in Barwise and Etchemendy 1990 and Barwise 1992.

[13] Cf. Mitchell 1997.

the *inductive bias*: one can fit the data only as well as $\mathcal{C}$ permits. On the other hand, if $\mathcal{C}$ is too expressive, *overfitting* can occur: the induced theory explains the given data perfectly but does not allow proper generalizations.

5.4 Negation and Conditional

One might argue that the language of observational logic is too restrictive to be useful since it lacks negation. This restriction, however, turns out to be less restrictive than it appears at first glance.

5.4.1 *Spurious Negation (and Conditional)*

We already saw in Section 3.1.1 how to render '*is-not-a*' statements into observational form. Recall that 'A *is-not-a* B' stands for '$\forall x(Ax \to \neg Bx)$', or '$A \preceq \neg B$', for short. In words: 'Everything which is an A is not a B', or, succinctly, 'No As are Bs'. The solution of Section 3.1.1 was to use '$A \wedge B \preceq \Lambda$' instead of '$A \preceq \neg B$', which is justified by the observation that both statements are equivalent with respect to first-order predicate logic with identity. So negation can enter into observational statements via 'Λ', which stands for '$\{x \mid x \neq x\}$' and is hence equivalent to '$\neg \mathrm{V}$'. But also 'V' serves its purpose, for '$\neg A \preceq B$' is equivalent to '$\mathrm{V} \preceq A \vee B$'.

Let us now tackle the problem of eliminating negation in general. Since negation can be eliminated together with conditionals in one sweep, we take the latter into account too. (Recall that '$A \to B$' stands for '$\{x \mid Ax \to Bx\}$', i.e. for 'something which is a B if it is an A'.)

Let Σ be a set of primitive monadic predicates (excluding Λ and V). We refer to monadic predicates that are inductively built from $\Sigma \cup \{\Lambda, \mathrm{V}\}$ by $\wedge$, $\vee$, $\neg$ and $\to$ as *Boolean predicates* over Σ. By a *universal statement* over Σ we mean a statement of the form $\forall\alpha$, where α is a Boolean predicate over Σ. We claim that every universal statement over Σ is logically equivalent to a finite conjunction of observational statements over Σ. Indeed, one just needs to apply the standard transformations of propositional logic to bring every Boolean predicate into conjunctive normal form. Since

$$\begin{aligned}
\neg p_1 \vee \ldots \vee \neg p_m \vee q_1 \vee \ldots \vee q_n &\equiv p_1 \wedge \ldots \wedge p_m \to q_1 \vee \ldots \vee q_n, \\
q_1 \vee \ldots \vee q_n &\equiv \mathrm{V} \to q_1 \vee \ldots \vee q_n, \\
\neg p_1 \vee \ldots \vee \neg p_m &\equiv p_1 \wedge \ldots \wedge p_m \to \Lambda,
\end{aligned}$$

and $\forall(\alpha \wedge \beta) \leftrightarrow \forall\alpha \wedge \forall\beta$, we can conclude:

(5.50) Proposition Every universal statement over a set Σ of primitive monadic predicates is logically equivalent to a finite conjunction of observational statements over Σ.

So it is not a point of expressivity whether we employ negation and conditional as predicate operators or not. From a practical point of view, it can be convenient to use them in the process of formalizing statements given as natural language expressions. Take for example the statement ‘Lions that are not hungry or angry are not dangerous’, which is most naturally formalized by ‘$A \wedge \neg(B \vee C) \preceq \neg D$’, with ‘$A$’ for ‘lion’, ‘$B$’ for ‘hungry’, ‘$C$’ for ‘angry’, and ‘$D$’ for ‘dangerous’. An equivalent observational statement is ‘$A \wedge D \preceq B \vee C$’. The conditional can also be of some use; e.g. ‘Lions are dangerous if they are hungry or angry’ corresponds to ‘$A \preceq (B \vee C \rightarrow D)$’, which is equivalent to ‘$A \wedge (B \vee C) \preceq D$.’

(5.51) Remark *(Traps of logical analysis)* Though convenient in certain cases, the conditional has to be used with care. For example, it is tempting to formalize the statement ‘Animals that are dangerous if they are hungry or angry should be avoided’ by ‘$A \wedge (B \vee C \rightarrow D) \preceq E$’, where ‘$A$’ stand for ‘animal’ and ‘E’ for ‘something that should be avoided’. However, turning this universal statement into observational form gives ‘$(A \preceq B \vee C \vee E) \wedge (A \wedge D \preceq E)$’, which clearly has not the intended meaning of the original statement. A more appropriate logical formalization of the example sentence is: $\forall x(Ax \wedge \forall t(Bxt \vee Cxt \rightarrow Dxt) \rightarrow Ex)$, where ‘$Bxt$’ stands for ‘$x$ is hungry at (time) t’, etc. The given example is therefore beyond the scope of observational logic.

Notice that statements of the form ‘$\neg(A \preceq B)$’, where negation has wide scope over the statement, do not have an observational counterpart; for when negated, universal statements turn into existential ones. From a conceptual point of view, existential assertions as part of a background theory do not fit well into the picture of accumulating information based on observations and the background theory in question. Knowledge about nonexistence, on the other hand, is of course useful because it allows to detect incompatible information; an assertion of nonexistence like ‘$\neg\exists(A \wedge B)$’ is equivalent to ‘$A \wedge B \preceq \Lambda$’, which is an observational statement.

5.4.2 *Predicate Negation and Booleanization*

In the last section we have seen that allowing negation as a predicate operator does not alter the logical expressivity of observational statements. Nevertheless there remains an asymmetry between negated and non-negated predicates if we think of an observational theory as a knowledge base for accumulating *positive* information. This asymmetry towards positive information is reflected in the construction of the canonical universe, whose elements are sets of (positive) predicates and whose ordering signals the increase of positive information.

This asymmetry can be resolved by regarding negated predicates as taking part in *positive* assertions. That is, predications involving negated predicates are not viewed as denials but as assertions. This is the traditional distinction between *predicate denial* and *term negation*, which Horn (1989) traces back to Aristotle; schematically: 'x (is not) A' versus 'x is (not A)' or 'x is not-A'.

To reconcile predicate negation with observational logic one can proceed as follows. Suppose Γ is an observational theory over Σ. Extend Σ by adding primitives $-p$ for all $p \in \Sigma$, and add all statements

(5.52) $$p \wedge -p \preceq \Lambda \qquad \text{and} \qquad \mathrm{V} \preceq p \vee -p,$$

with $p \in \Sigma$. The resulting theory $\overline{\Gamma}$ will be called the *Booleanization* of Γ. Since each member of $\mathrm{C}(\overline{\Gamma})$ contains either p or $-p$, but not both, if follows that no member of $\mathrm{C}(\overline{\Gamma})$ is a proper subset of another member of $\mathrm{C}(\overline{\Gamma})$; therefore:

(5.53) Proposition The ordered generic universe of the Booleanization of an observational theory is an antichain.

(5.54) Example Consider the theory Γ over $\{a, b, c\}$ with statements $a \wedge b \preceq \Lambda$ and $a \vee b \preceq c$. Its canonical universe $\mathrm{C}(\Gamma)$ has the members $\varnothing$, $\{c\}$, $\{a, c\}$, and $\{b, c\}$, whereas the canonical universe $\mathrm{C}(\overline{\Gamma})$ of the Booleanization $\overline{\Gamma}$ of Γ consists of $\{-a, -b, -c\}$, $\{-a, -b, c\}$, $\{a, -b, c\}$, and $\{-a, b, c\}$; see Figure 43.

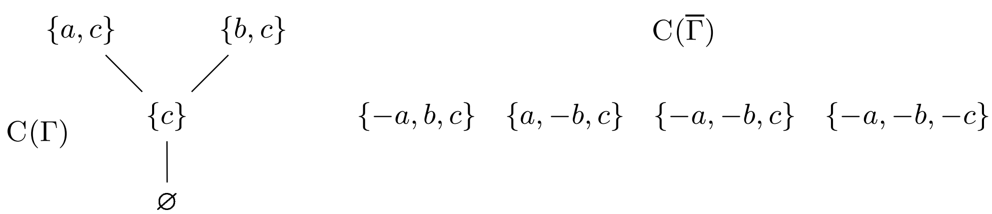

FIGURE 43 Effect of Booleanization on canonical universe

As the previous example indicates, there is a one-to-one correspondence between $\mathrm{C}(\Gamma)$ and $\mathrm{C}(\overline{\Gamma})$ for every observational theory Γ. To see this, consider the function f from $\mathrm{C}(\overline{\Gamma})$ to $\mathrm{C}(\Gamma)$ that takes Y to $Y \cap \Sigma$. Observe that if $X \in \mathrm{C}(\Gamma)$ then the set $\overline{X} = X \cup \{-p \mid p \in \Sigma \setminus X\}$ belongs to $\mathrm{C}(\overline{\Gamma})$ because it is consistently closed with respect to Γ and the statements (5.52). Clearly, $\overline{X} \cap \Sigma = X$ and $\overline{Y \cap \Sigma} = Y$. Hence f is onto and one-to-one:

(5.55) Proposition There is a one-to-one correspondence between the generic universe of an observational theory and that of its Booleanization.

Booleanization means thus loss of specialization while keeping all elements of the generic universe.

Term negation need not be applied to all members of Σ. Given a subset Σ_0 of Σ, let Σ' be $\Sigma \cup \{-p \mid p \in \Sigma_0\}$ and let Γ' be the extension of Γ by all statements of the form (5.52) with $p \in \Sigma_0$. As in the case of Booleanization, there is a one-to-one correspondence between $\mathrm{C}(\Gamma)$ and $\mathrm{C}(\Gamma')$. Specialization on $\mathrm{C}(\Gamma')$, in contrast, is usually nontrivial.

Finally, it should be added that there may be reasons to make use of term negation without admitting the law of excluded middle. That is, $p \wedge -p \preceq \Lambda$ is added to Γ, but not $\mathrm{V} \preceq p \vee -p$. A possible application is the transformation of a taxonomic tree into a binary one.

5.4.3 *Digression: Intuitionistic Negation and Conditional*

Suppose Γ is an observational theory over Σ with ordered generic universe U. By (5.8), it follows that every observational predicate φ over Σ is *persistent* with respect to specialization in the sense that if $x \models \varphi$ and $x \sqsubseteq y$ then $y \models \varphi$, for all $x, y \in U$.

Regarding knowledge as cumulative or persistent is much in accordance with the principles of *intuitionistic logic*. Let us therefore briefly indicate how to incorporate intuitionistic negation and conditional into the present framework. Roughly speaking, they are obtained by making their classical counterparts persistent. That is, when '$\sim$' signals intuitionistic negation, we require that

$$x \models {\sim}\varphi \qquad \text{iff} \qquad \forall y (x \sqsubseteq y \;\rightarrow\; y \not\models \varphi).$$

In words: x satisfies ${\sim}\varphi$ if and only if neither x nor anything more specific than x satisfies φ. Correspondingly for the intuitionistic conditional '$\Rightarrow$'; making the classical conditional persistent gives:

(5.56) $$x \models \varphi \Rightarrow \psi \qquad \text{iff} \qquad \forall y (x \sqsubseteq y \;\rightarrow\; y \models (\varphi \rightarrow \psi)).$$

These interpretations conform with the standard Kripke-style interpretation of intuitionistic logic, where the ordered generic universe serves as a *Kripke-model*.[14] It follows that the laws of intuitionistic logic are extensionally valid; in particular, $[\![\varphi \wedge (\varphi \Rightarrow \psi)]\!] \subseteq [\![\psi]\!]$. Moreover, $[\![{\sim}\varphi]\!] = [\![\varphi \Rightarrow \Lambda]\!]$ and $[\![\varphi \wedge {\sim}\varphi]\!] = [\![\Lambda]\!]$. But it is in general not the case that $[\![\varphi \vee {\sim}\varphi]\!] = [\![\mathrm{V}]\!]$.

(5.57) Example Let Γ be the theory over $\{a, b, c\}$ consisting of the sole statement $c \preceq a$. To bring some life into Γ, say a is 'animal', b is 'black', and c is 'cat'. Consider the extension of the conditional predicate $c \Rightarrow b$ in the

[14] See e.g. van Dalen 1986 or Dunn and Hardegree 2001, Chap. 11.

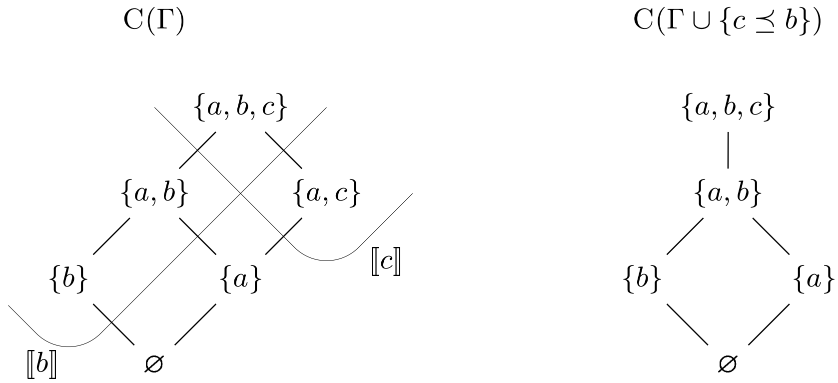

FIGURE 44 Black cats and other animals

canonical universe $\mathrm{C}(\Gamma)$ of Γ. By definition, $c \Rightarrow b$ is satisfied by $X \in \mathrm{C}(\Gamma)$ iff $\uparrow\{X\} \cap \llbracket c \rrbracket \subseteq \llbracket b \rrbracket$. One can easily read off from the diagram of $\mathrm{C}(\Gamma)$ depicted in Figure 44 that $\llbracket c \Rightarrow b \rrbracket = \llbracket b \rrbracket$. (Moreover, $\llbracket {\sim} c \rrbracket = \llbracket {\sim} b \rrbracket = \llbracket {\sim} a \rrbracket = \llbracket \Lambda \rrbracket$.) It is instructive to contrast $c \Rightarrow b$ with the statement $c \preceq b$. First of all, $c \Rightarrow b$ is not a statement but a predicate; $c \Rightarrow b$ is satisfied by those elements of the generic universe of which it is "persistently true" that they are black if they are cats. For example, the generic animal, represented by $\{a\}$, does not satisfy $c \Rightarrow b$ because it can be specialized to a cat that is not necessarily black (represented by $\{a, c\}$). The statement $c \preceq b$, on the other hand, does not exclude non-black non-cats; see Figure 44 for the effect of extending Γ by $c \preceq b$.

Let us ask to what extent $\Rightarrow$ or $\sim$ can be eliminated in the sense of Section 5.4.1. First notice that the interpretation of ${\sim}\varphi$ and $\varphi \Rightarrow \psi$, as introduced above, depends on a given specialization order, which, on the other hand, is induced by a given theory. It is hence far from clear what it means to allow $\Rightarrow$ or $\sim$ as first class predicate operators that can occur in the statements of an observational theory. We therefore restrict ourselves to the question whether the *post hoc* defined predicate $\varphi \Rightarrow \psi$ is extensionally equivalent to an observational predicate over Σ. Notice that (5.56) can be reformulated as follows:

$$\llbracket \varphi \Rightarrow \psi \rrbracket \;=\; \bigcup \{ \uparrow x \mid \uparrow x \cap \llbracket \varphi \rrbracket \subseteq \llbracket \psi \rrbracket \}.$$

Assume that Σ is finite. Then the generic universe U of Γ is finite and every $x \in U$ is the least satisfier of an observational predicate $\varphi \in \mathrm{T}[\Sigma]$, i.e. $\uparrow x = \llbracket \varphi \rrbracket$; see Chapter 7 for details. Consequently, $\llbracket \varphi \Rightarrow \psi \rrbracket$ is the union of the *finite* set of all extensions $\llbracket \chi \rrbracket$ such that $\llbracket \chi \rrbracket \cap \llbracket \varphi \rrbracket \subseteq \llbracket \psi \rrbracket$. It follows that $\varphi \Rightarrow \psi$ is extensionally equivalent to a finite disjunction of observational predicates, which is of course observational in turn. For infinite Σ, however, this need not be the case. Consider for example the full binary exclusion theory $\{p \wedge q \preceq$

$\Lambda \mid p \neq q\}$ over Σ. Then $[\![\sim p]\!] = \Sigma \setminus \{p\}$ is not the extension of an observational predicate over Σ.

(5.58) Remark *(Relative pseudocomplements)* Anticipating some concepts and results of Chapter 6, the problem of eliminating $\Rightarrow$-predicates comes down to the question whether the Lindenbaum algebra $\mathrm{L}(\Gamma)$ of Γ is *relatively pseudocomplemented*, that is, whether for every two elements a and b of $\mathrm{L}(\Gamma)$ there exists an element $a \Rightarrow b$ of $\mathrm{L}(\Gamma)$, the *pseudocomplement of* a *relative to* b, such that

$$c \leqslant a \Rightarrow b \qquad \text{iff} \qquad c \wedge a \leqslant b,$$

for every $c \in \mathrm{L}(\Gamma)$. If $\mathrm{L}(\Gamma)$ is relatively pseudocomplemented, the operation $\Rightarrow$ gives $\mathrm{L}(\Gamma)$ the structure of a *Heyting algebra*. Every *complete* distributive lattice is relatively pseudocomplemented because $a \Rightarrow b$ can be defined by:

$$a \Rightarrow b \;=\; \bigvee \{c \mid c \wedge a \leqslant b\}.$$

In particular, if Σ is finite, $\mathrm{L}(\Gamma)$ is finite too, hence a complete lattice, and thus relatively pseudocomplemented.

6

Logic and Algebra

In Section 6.1, set-theoretic models of observational theories are generalized to algebraic ones, more precisely, to models in distributive lattices with zero and unit, or *observational algebras*, as we will call them. Every observational theory Γ is shown to have a *canonical model in an algebra* $\mathrm{L}(\Gamma)$, which is characterized by a certain universal property. On the other hand, every observational algebra is the canonical algebraic model of an observational theory. This categorical correspondence between theories and algebras will prove useful in Chapters 8 and 9 where translations between theories and constructions of theories are the issue.

In Section 6.2 we show that the generic universe of Γ can be represented by the *prime spectrum* of $\mathrm{L}(\Gamma)$. By a standard application of the Prime Ideal Theorem it then follows that $\mathrm{L}(\Gamma)$ is the *Lindenbaum algebra* of Γ, which, roughly speaking, is obtained by ignoring syntactical differences between observational predicates if they are equivalent with respect to Γ. As a byproduct we get an inference calculus for observational statements that is sound and strongly complete with respect to first-order entailment, which is the topic of Section 6.3.

6.1 Algebraic Models

By an *observational algebra* we mean a distributive lattice with zero and unit. A *homomorphism* of observational algebras is a lattice homomorphism that preserves zero and unit.

6.1.1 *Models in Observational Algebras*

A function m from a set Σ (of primitives) to an observational algebra A will be referred to as an *interpretation of* Σ *with values in* A, or A*-valued interpretation*. The universal property (5.2) of the term algebra $\mathrm{T}[\Sigma]$ over Σ implies that every interpretation m of Σ in A can be uniquely extended to a function $\widehat{m}$ from $\mathrm{T}[\Sigma]$ to A such that $\widehat{m}(\Lambda) = 0$, $\widehat{m}(\mathrm{V}) = 1$,

$$\widehat{m}(\varphi \wedge \psi) = \widehat{m}(\varphi) \wedge \widehat{m}(\psi) \quad \text{and} \quad \widehat{m}(\varphi \vee \psi) = \widehat{m}(\varphi) \vee \widehat{m}(\psi).$$

So every interpretation of Σ in A has a unique extension to an algebra homomorphism (of algebras of type $\langle 2, 2, 0, 0\rangle$) from $\mathrm{T}[\Sigma]$ to A.

Let Γ be an observational theory over Σ. An interpretation m of Σ in A is called an *A-valued model of* Γ if its homomorphic extension $\widehat{m}$ satisfies $\widehat{m}(\varphi) \leqslant \widehat{m}(\psi)$ for all statements $\varphi \preceq \psi$ of Γ. We then also say that $\langle m, A\rangle$ is an *algebraic model* of Γ. This definition can be applied to biconditional statements, i.e. to statements of the form $\varphi \equiv \psi$, by transforming them first into conditional form via (5.1). Alternatively, one could adjust the definition of A-valued models explicitly to biconditional theories: an interpretation m of Σ in A is an A-valued model of Γ if $\widehat{m}(\varphi) = \widehat{m}(\psi)$ for all $(\varphi \equiv \psi) \in \Gamma$. Clearly these two options are equivalent.

(6.1) Example *(Extension algebra)* Every set-valued model M of Γ with universe U is by definition an algebraic Γ-model with values in the power set algebra $\wp(U)$. Moreover, the set of extensions (see Section 5.1.4)

$$\Omega_M \;=\; \{M(\varphi) \mid \varphi \in \mathrm{T}[\Sigma]\}$$

is an *observational algebra of sets over U* since the intersection and the union of every two members of Ω_M belong to Ω_M in turn; in addition, $M(\Lambda)$ $(= \varnothing)$ and $M(\mathrm{V})$ $(= U)$ are respectively zero and unit of Ω_M. We refer to Ω_M as the *algebra of extensions*, or *extension algebra*, of M. It follows that M is an algebraic model of Γ with values in the extension algebra Ω_M.

An algebraic model $\langle m, A\rangle$ of Γ is called *universal* if, for every algebraic model $\langle m', A'\rangle$ of Γ, m' factors uniquely through m by a homomorphism h from A to A' (i.e. $m' = h \circ m$). One easily verifies that a universal model of Γ, if existent, is uniquely determined up to isomorphism by this property; that is, if $\langle m', A'\rangle$ is another universal Γ-model then h is an isomorphism.

As for showing the existence of a universal algebraic model of Γ we can use standard constructions from universal algebra. Let $\cong_\Gamma$ be the least congruence relation on $\mathrm{T}[\Sigma]$ that contains $\langle \varphi, \psi\rangle$ for all $(\varphi \equiv \psi) \in \Gamma$ plus all pairs given by the axioms of a distributive lattice with zero and unit. Define $\mathrm{L}(\Gamma)$ to be the quotient algebra of $\mathrm{T}[\Sigma]$ by $\cong_\Gamma$, that is,

$$\mathrm{L}(\Gamma) \;=\; \mathrm{T}[\Sigma]/{\cong_\Gamma}.$$

The canonical interpretation m_Γ of Σ in $\mathrm{L}(\Gamma)$, which takes p to $[p]_{\cong_\Gamma}$, is by definition a model of Γ in $\mathrm{L}(\Gamma)$. We refer to $\langle \mathrm{m}_\Gamma, \mathrm{L}(\Gamma)\rangle$ as the *canonical (algebraic) model* of Γ.

(6.2) Proposition The canonical algebraic model of an observational theory is universal.

Proof. If $\langle m, A\rangle$ is an algebraic model of Γ, the congruence kernel of $\widehat{m}$ includes $\cong_\Gamma$. Hence there is a unique homomorphism h from $\mathrm{L}(\Gamma) = \mathrm{T}[\Sigma]/{\cong_\Gamma}$ to A such that $\widehat{m}(\varphi) = h([\varphi]_{\cong_\Gamma})$. By (5.2), $\widehat{m}$ is uniquely determined by m. $\lrcorner$

Put differently, (6.2) says that there is a one-to-one correspondence between the set $\mathrm{Mod}(\Gamma, A)$ of models of Γ in A and the set $\mathrm{Hom}(\mathrm{L}(\Gamma), A)$ of homomorphisms from $\mathrm{L}(\Gamma)$ to A, in short,

(6.3) $$\mathrm{Hom}(\mathrm{L}(\Gamma), A) \;\simeq\; \mathrm{Mod}(\Gamma, A),$$

where a homomorphism h from $\mathrm{L}(\Gamma)$ to A is taken to the model $h \circ \mathrm{m}_\Gamma$.

6.1.2 *Two-Valued Models*

It was indicated in Section 3.2.3 that a generic entity of a Horn theory Γ can be viewed as a model of Γ whose universe is a fixed singleton set. We also mentioned that the two possible values of such an interpretation function – the empty set and the singleton set itself – can be regarded as "truth values", thereby leading to truth-valued interpretations and models in the sense of propositional logic. Here we show from an algebraic perspective that the same holds for observational theories in general.

Let $\mathbb{2}$ be an observational algebra consisting of exactly two members, that is, $\mathbb{2} = \{0, 1\}$. There is a one-to-one correspondence between the $\mathbb{2}$-valued interpretations of Σ and the subsets of Σ, where a subset X of Σ is taken to its *characteristic* ($\mathbb{2}$-valued) *function* χ_X, with $\chi_X(p) = 1$ if $p \in X$, and $\chi_X(p) = 0$ otherwise. It follows by term induction that $X \vDash \varphi$ iff $\widehat{\chi}_X(\varphi) = 1$. Hence,

$$X \vDash \varphi \to X \vDash \psi \qquad \text{iff} \qquad \widehat{\chi}_X(\varphi) \leqslant \widehat{\chi}_X(\psi).$$

Consequently, if Γ is an observational theory over Σ then X belongs to $\mathrm{C}(\Gamma)$ iff χ_X is a $\mathbb{2}$-valued model of Γ. Put differently, a $\mathbb{2}$-valued interpretation m of Σ is a Γ-model iff $m^{-1}(1)$ belongs to $\mathrm{C}(\Gamma)$.

Notice that specialization, which is set inclusion on the set of subsets of Σ, is the *pointwise order* on the set of interpretations: $m \sqsubseteq m'$ iff $m(p) \leqslant m'(p)$ for all $p \in \Sigma$. To summarize:

(6.4) Proposition The canonical universe of Γ is order-isomorphic to the set of $\mathbb{2}$-valued models of Γ ordered pointwise; in short, $\mathrm{C}(\Gamma) \simeq \mathrm{Mod}(\Gamma, \mathbb{2})$.

So we can take the set of $\mathbb{2}$-valued Γ-models to represent the generic universe of Γ, where satisfaction is given as follows:

$$\textbf{\textit{(6.5)}} \qquad m \models \varphi \qquad \text{iff} \qquad \widehat{m}(\varphi) = 1,$$

for every $m \in \mathrm{Mod}(\Gamma, \mathbb{2})$ and $\varphi \in \mathrm{T}[\Sigma]$.

6.1.3 *Presentation of Algebras by Theories*

An observational algebra A is said to be *presented* by an observational theory Γ over Σ if A is isomorphic to $\mathrm{L}(\Gamma)$. This is just a reformulation of the standard definition of a *presentation by generators and relations* used in universal algebra, which runs as follows: An algebra A is *generated* by a subset Σ of A if the homomorphic extension of the inclusion of Σ into A is a surjective function from $\mathrm{T}[\Sigma]$ to A. An algebra A is *presented* by a set of *generators* Σ and a relation R on $\mathrm{T}[\Sigma]$ if A is isomorphic to $\mathrm{T}[\Sigma]/{\cong}$, where $\cong$ is the least congruence relation containing R plus all pairs given by the axioms of an observational algebra. If A is presented by an observational theory Γ, the relation $\{\langle\varphi, \psi\rangle \mid (\varphi \equiv \psi) \in \Gamma\}$ takes the role of R.

(6.6) Theorem Every observational algebra is presented by an observational theory.

Proof. Let Σ be a generating subset of an observational algebra A (e.g. $\Sigma = A$) and let ε be the inclusion function from Σ into A. Then, by assumption, the homomorphic extension $\hat{\varepsilon}$ of ε is onto. Therefore A is isomorphic to the quotient $\mathrm{T}[\Sigma]/{\cong}$ of $\mathrm{T}[\Sigma]$ by the congruence kernel $\cong$ of $\hat{\varepsilon}$. Define Γ to be set of all statements $\varphi \equiv \psi$ over Σ such that $\varphi \cong \psi$. Then Γ presents A over Σ. ⌟

Taking (the underlying set of) A itself as the generating subset provides a *canonical* way of presenting A by a theory. Let $\mathrm{Th}(A)$ be the theory over A defined that way. Its statements are given by all identities that hold in A. More precisely, $\varphi \equiv \psi$ is a statement of $\mathrm{Th}(A)$ just in case $\hat{\imath}_A(\varphi) = \hat{\imath}_A(\psi)$, where $\imath_A$ is the identity function on A. The canonical presentation of an observational algebra will turn up again in Section 8.2.1.

Since finite algebras are finitely generated and every observational theory has a normal form, it follows that finite algebras are presentable by finite theories:

(6.7) Corollary Every finite observational algebra is presentable by a finite observational theory over a finite set of primitives.

6.2 Prime Spectrum and Lindenbaum Algebra

6.2.1 *The Prime Spectrum*

The *prime spectrum* $\mathrm{P}(A)$ of an observational algebra A is the set of all prime filters of A ordered by set inclusion.[1] It is an immediate consequence of definitions that $\mathrm{P}(A)$ is closed with respect to upwards directed union and downwards directed intersection. (Incidentally, in the light of (6.12) below, this is just a reformulation of (5.21).)

Let $\mathcal{P}(a)$ be the set $\{F \in \mathrm{P}(A) \mid a \in F\}$ of all prime filters of A with member a. The function $\mathcal{P}$ that takes a to $\mathcal{P}(a)$ is a homomorphism of observational algebras from A to $\wp(\mathrm{P}(A))$. Moreover,

(6.8) $$\mathcal{P}(a) \subseteq \mathcal{P}(b) \qquad \text{iff} \qquad a \leqslant b.$$

Proof. Clearly, if $a \leqslant b$ then $\mathcal{P}(a) \subseteq \mathcal{P}(b)$. Suppose $a \nleqslant b$. Then, by the Prime Ideal Theorem, there is a prime filter F of A such that $\uparrow a \subseteq F$ and $\downarrow b \cap F = \varnothing$. Hence $\mathcal{P}(a) \nsubseteq \mathcal{P}(b)$. ⌟

It follows that $\mathcal{P}$ is an *embedding* of A into $\wp(\mathrm{P}(A))$, that is, A is isomorphic to an observational algebra of sets, namely $\mathcal{P}(A) = \{\mathcal{P}(a) \mid a \in A\}$. We have thus proved the following classical result of Birkhoff and Stone:

(6.9) Theorem Every observational algebra is isomorphic to an observational algebra of sets.

For later use we note the following *compactness* result.

(6.10) Proposition Let a be a member of an observational algebra A. Then every cover of $\mathcal{P}(a)$ by members of $\mathcal{P}(A)$ has a finite subcover.[2]

Proof.[3] Suppose $\mathcal{C} \subseteq \mathcal{P}(A)$ covers $\mathcal{P}(a)$, that is, $\mathcal{P}(a) \subseteq \bigcup \mathcal{C}$. Then $\mathcal{P}(a)$ is covered by the directed set $\mathcal{S} = \{\bigcup \mathcal{F} \mid \mathcal{F} \subseteq \mathcal{C} \text{ finite}\}$. We need to show that $\mathcal{P}(a) \subseteq V$ for some $V \in \mathcal{S}$, i.e. that a belongs to the set

$$I \;=\; \{b \mid \exists V \in \mathcal{S}\,(\mathcal{P}(b) \subseteq V)\},$$

which is clearly an ideal of A. Suppose $a \notin I$. Then, by the Prime Ideal Theorem, there is a prime filter F of A such that $a \in F$ and $I \cap F = \varnothing$. Hence, $F \in \mathcal{P}(a)$, and $F \notin \mathcal{P}(b)$ for all $b \in I$. But this is impossible because $\mathcal{P}(a)$ is covered by $\mathcal{C}$, and $\{b \mid \mathcal{P}(b) \in \mathcal{C}\} \subseteq I$. ⌟

[1] Equivalently one could use the set of prime ideals ordered by reverse inclusion.

[2] If V is a subset of a set U then a system $\mathcal{C}$ of subsets of U *covers* V, or is a *cover* of V, if $V \subseteq \bigcup \mathcal{C}$.

[3] Adapted from Smyth 1992, Sect. 7.2.

Suppose h is a homomorphism of observational algebras from A to $\mathbb{2}$. Then $h^{-1}(0)$ is a prime ideal of A whereas $h^{-1}(1)$ is a prime filter of A. Conversely, given a prime filter F of A (and hence a prime ideal $A \setminus F$), the characteristic function χ_F from A to $\mathbb{2}$ is an algebra homomorphism. Put together, with $\mathrm{Hom}(A, \mathbb{2})$ ordered pointwise:

(6.11) Proposition $\mathrm{P}(A)$ is order-isomorphic to $\mathrm{Hom}(A, \mathbb{2})$.

Assume now that $A = \mathrm{L}(\Gamma)$ for some observational theory Γ over Σ. Combining (6.11) with (6.3) and (6.4) implies that the generic universe of Γ can be represented by the prime spectrum $\mathrm{P}(\mathrm{L}(\Gamma))$ of $\mathrm{L}(\Gamma)$:

(6.12) Proposition The canonical universe of Γ is order-isomorphic to the prime spectrum of $\mathrm{L}(\Gamma)$.

The isomorphism of (6.12) takes each prime filter $F \in \mathrm{P}(\mathrm{L}(\Gamma))$ to the consistently Γ-closed set $\{p \in \Sigma \mid [p] \in F\}$. With prime filters as generic entities, satisfaction therefore becomes: $F \models \varphi$ iff $[\varphi] \in F$. It is instructive to see that this definition of satisfaction, applied to an arbitrary subset F of $\mathrm{L}(\Gamma)$, inevitably leads to prime filters. For we have to require that $[\Lambda] \notin F$, $[\mathrm{V}] \in F$,

$$[\varphi] \wedge [\psi] \in F \quad \text{iff} \quad [\varphi] \in F \text{ and } [\psi] \in F,$$
$$[\varphi] \vee [\psi] \in F \quad \text{iff} \quad [\varphi] \in F \text{ or } [\psi] \in F.$$

It was promised in Section 5.2.1 to prove equivalence of coextensives (5.14) for generic models without making use of completeness of first-order logic. This is the place to do it. With respect to the prime spectrum representation of the generic universe, the extension of $\varphi \in \mathrm{T}[\Sigma]$ is $[\![\varphi]\!] = \mathcal{P}([\varphi])$. By (6.8),

$$[\![\varphi]\!] = [\![\psi]\!] \qquad \text{iff} \qquad [\varphi] = [\psi].$$

Rephrased in terms of the extension algebra – cf. (6.1):

(6.13) Proposition $\mathrm{L}(\Gamma)$ is isomorphic to the extension algebra of the generic model of Γ.

In order to prove equivalence of coextensives it thus remains to show that

$$\text{if} \quad [\varphi] = [\psi] \quad \text{then} \quad \Gamma \vdash \varphi \equiv \psi.$$

By definition of $\mathrm{L}(\Gamma)$, this means to make sure that your favorite calculus of first-order logic entails all those biconditional statements that correspond to the defining equations of a distributive lattice with zero and unit, as e.g. $\varphi \vee \psi \equiv \psi \vee \varphi$ and $\varphi \wedge \mathrm{V} \equiv \varphi$. Certainly, this will cause no problems. The second thing to check is that entailment respects congruence closure, that is, if Γ entails $\varphi \equiv \psi$ and $\varphi' \equiv \psi'$ then also $\varphi \wedge \varphi' \equiv \psi \wedge \psi'$ and $\varphi \vee \varphi' \equiv \psi \vee \psi'$. Again, this is easily seen to be the case. All in all it follows that

$$\text{if} \quad [\![\varphi]\!] = [\![\psi]\!] \quad \text{then} \quad \Gamma \vdash \varphi \equiv \psi.$$

The reverse statement is of course also true since we are working with a first-order model of Γ and first-order logic is sound.

As a further consequence, to be spelled out in more detail in Section 6.3, the foregoing argument shows how to find an inference calculus for observational statements that is sound and (strongly) complete with respect to first-order entailment: choose any calculus that respects congruence closure and entails the aforementioned biconditionals corresponding to the axioms of an observational algebra.

6.2.2 *The Lindenbaum Algebra*

The basic idea behind the Lindenbaum algebra[4] of a theory is to abstract away from syntactical differences between terms whose equivalence is entailed by the theory in question. In the present context, two observational predicates φ and ψ over Σ are said to be *equivalent* with respect to an observational theory Γ just in case they are first-order equivalent with respect to Γ, that is, if Γ entails $\varphi \equiv \psi$ by any sound and complete inference calculus for first-order logic.

The *Lindenbaum* (or *Lindenbaum-Tarski*) *algebra* of Γ is defined to be the quotient algebra of the term algebra $\mathrm{T}[\Sigma]$ modulo (first-order) equivalence. Recall from Section 6.2.1 that

(6.14) $$\Gamma \vdash \varphi \equiv \psi \quad \text{iff} \quad \varphi \cong_\Gamma \psi.$$

Hence the Lindenbaum algebra of Γ is (isomorphic to) $\mathrm{L}(\Gamma)$:

(6.15) Proposition The Lindenbaum algebra of an observational theory is the observational algebra presented by that theory.

The following example illustrates the relationship between canonical universe, extension algebra, Lindenbaum algebra, and prime spectrum. Moreover, it previews part of the content of Chapter 7.

[4] Named after the Polish logician Adolf Lindenbaum, who perished in World War II.

(6.16) Example Let Γ be the theory over $\{a, b, c, d\}$ with statements

$$a \wedge b \preceq c \vee d, \qquad c \wedge d \preceq \Lambda, \qquad c \preceq a, \qquad d \preceq a \wedge b.$$

Its canonical universe $\mathrm{C}(\Gamma)$ and its Lindenbaum algebra $\mathrm{L}(\Gamma)$ are shown in Figure 45. For example, since Γ entails $a \wedge b \preceq (b \wedge c) \vee d$ as well as

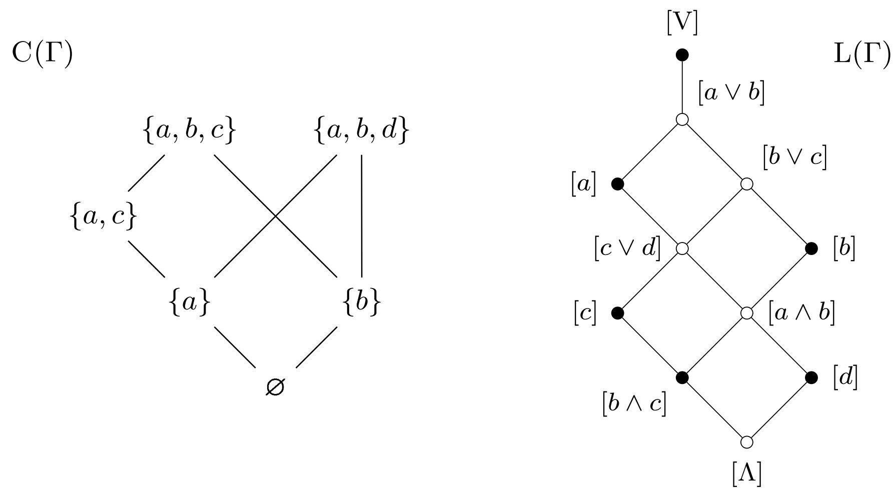

FIGURE 45 Canonical universe and Lindenbaum algebra of Γ

$(b \wedge c) \vee d \preceq a \wedge b$, we have $[(b \wedge c) \vee d] = [a \wedge b]$. The prime spectrum of $\mathrm{L}(\Gamma)$, which according to (6.12) is order-isomorphic to $\mathrm{C}(\Gamma)$, is depicted by Figure 46. Since $\mathrm{L}(\Gamma)$ is finite, all its prime filters are *principal*, i.e. of the form $\uparrow[\varphi] = \{[\psi] \mid [\varphi] \leqslant [\psi]\}$. If $\uparrow[\varphi]$ is a prime filter then $[\varphi]$ is necessarily *join-irreducible*, that is, if $[\varphi] = [\psi \vee \chi]$ then $[\varphi] = [\psi]$ or $[\varphi] = [\chi]$. In the diagram of $\mathrm{L}(\Gamma)$ in Figure 45 the join-irreducible elements are shaded; they can be read off from the diagram as those elements that have precisely one element immediately below them. Figure 47 illustrates the content of (6.13): $\mathrm{L}(\Gamma)$ is isomorphic to the extension algebra of the generic model of Γ. Due to the finiteness of the example, the extension algebra consists of all upwards closed subsets of the ordered generic universe. The doubly framed subsets are those of the form $\uparrow x$; they correspond to the join-irreducible elements of $\mathrm{L}(\Gamma)$.

6.2.3 *Neighborhood Filters*

Suppose M is a model of an observational theory Γ with universe U. Let Ω be the extension algebra of M. In terms of Ω the definition of specialization looks as follows (cf. (5.8)):

$$x \sqsubseteq y \quad \text{iff} \quad \forall V \in \Omega\, (x \in V \to y \in V).$$

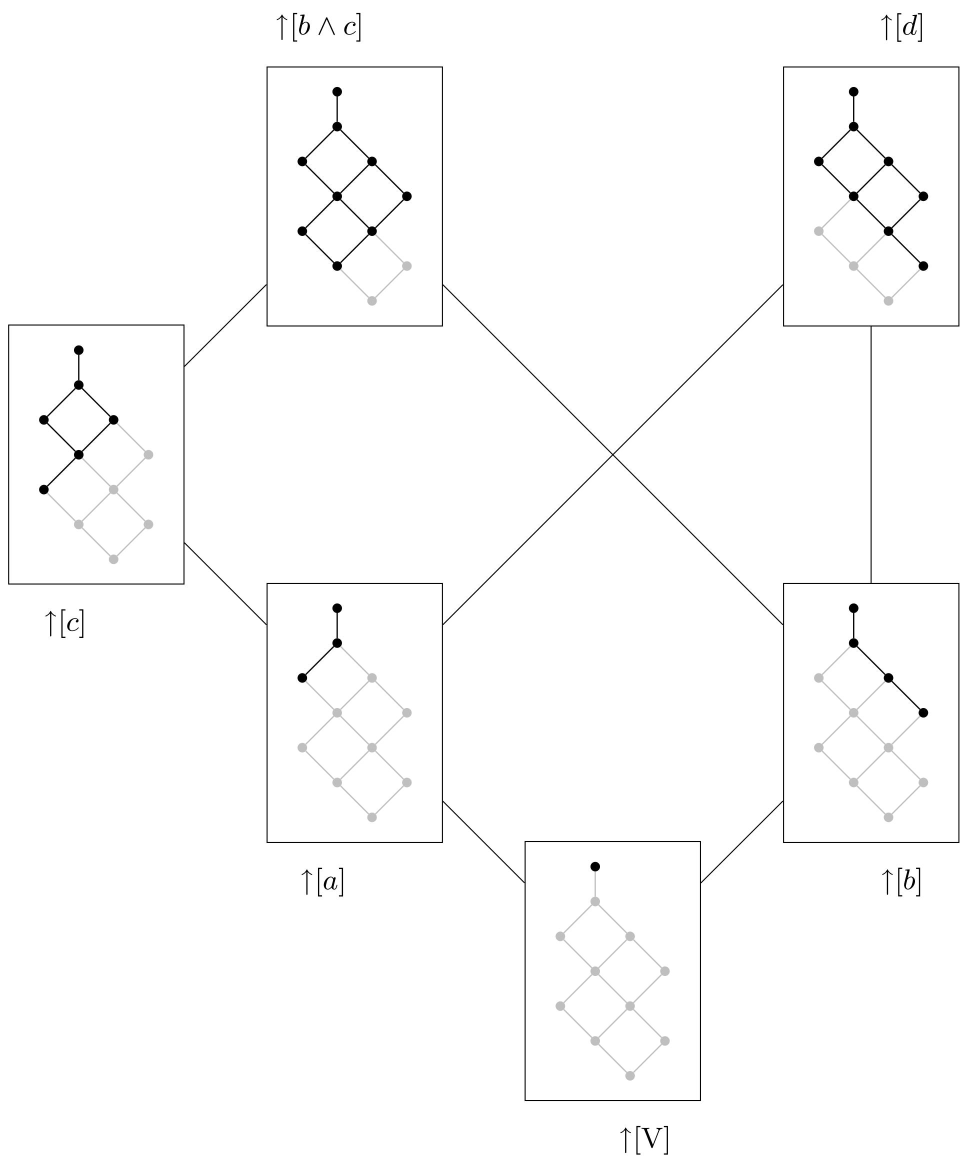

FIGURE 46 The prime spectrum of $\mathrm{L}(\Gamma)$

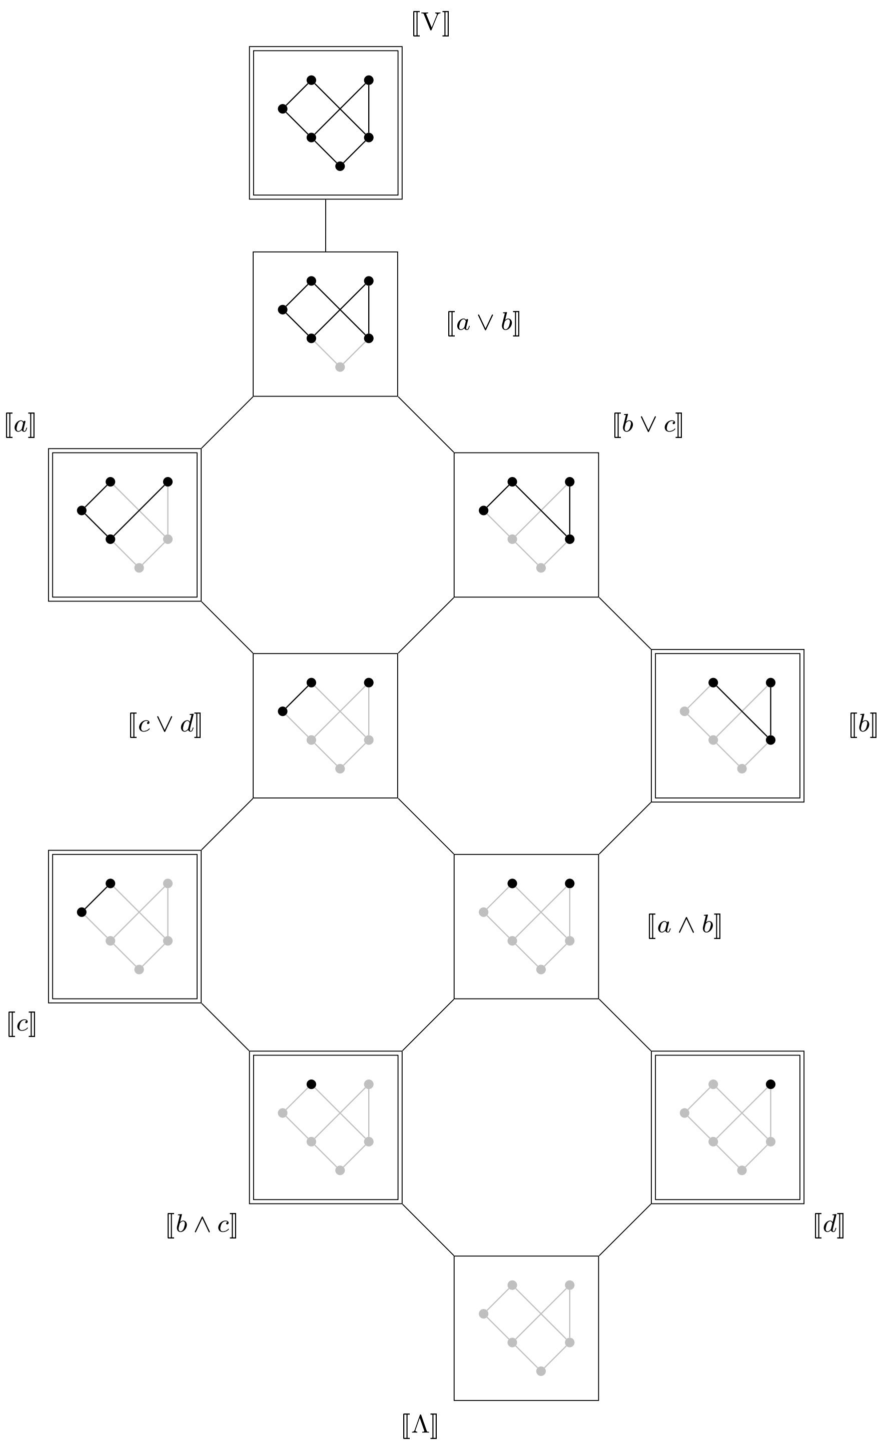

FIGURE 47 Extension algebra of the generic model of Γ

Let $\mathcal{N}(x)$ be the set of extensions of all predicates satisfied by $x \in U$, i.e.

$$\mathcal{N}(x) \;=\; \{[\![\varphi]\!] \mid x \vDash \varphi\} \;=\; \{V \in \Omega \mid x \in V\}.$$

The set $\mathcal{N}(x)$ is upwards closed with respect to inclusion, i.e. if $V \in \mathcal{N}(x)$ and $V \subseteq W$ then $W \in \mathcal{N}(x)$. In addition, $\mathcal{N}(x)$ is closed with respect to finite intersection. Thus $\mathcal{N}(x)$ is a filter of Ω, henceforth called the *neighborhood filter* of x. Moreover, if $V \cup W \in \mathcal{N}(x)$, for any $V, W \in \Omega$, then $V \in \mathcal{N}(x)$ or $W \in \mathcal{N}(x)$. So, since $\varnothing \notin \mathcal{N}(x)$, the neighborhood filter of x is prime.

Assume now that M is a generic model of Γ. Then, by (6.13), Ω is isomorphic to $\mathrm{L}(\Gamma)$. Consequently, by (6.12), U is order-isomorphic to the prime spectrum of Ω. So every element x of U uniquely determines a prime filter of Ω and vice versa. Going through the various isomorphisms shows that this filter is the neighborhood filter of x. Indeed, by (5.12), x is taken to $\{p \mid x \vDash p\}$, which is identical to $m^{-1}(1)$ if m is the $\mathfrak{2}$-valued Γ-model representing x; hence $\{[\varphi] \mid x \vDash \varphi\}$ is the corresponding filter of $\mathrm{L}(\Gamma)$ and $\{[\![\varphi]\!] \mid x \vDash \varphi\}$ is the filter of Ω in question. Vice versa, every prime filter $\mathcal{F}$ of Ω is the neighborhood filter of a unique generic entity x.[5] To sum up:

(6.17) Proposition Suppose Ω is the extension algebra of the generic model of an observational theory. Then the prime filters of Ω are precisely the neighborhood filters of the generic entities.

(6.18) Remark *(Coherent Spaces)* Speaking of neighborhoods suggests that there is a *topological* perspective on extension algebras. Indeed, if we add unions of arbitrary subsets of Ω, we get a *topology* $\tau(\Omega)$ on U.[6] The topological spaces arising that way are precisely the *coherent* (or *spectral*) *spaces*, where coherence means that the compact open sets form a distributive sublattice and a basis of the topology. For assume that U is the spectrum $\mathrm{P}(A)$ of an observational algebra A and Ω ($\simeq A$) consists of the subsets $\{\mathcal{P}(a) \mid a \in A\}$ of $\mathrm{P}(A)$ (cf. Section 6.2.1). By (6.10), it follows that the members of Ω are the *compact open* sets of $\tau(\Omega)$. In addition, it is not difficult to verify that $\tau(\Omega)$ (as a subset lattice) is isomorphic to the ideal completion $\mathrm{I}(A)$ of A. See Johnstone 1982, Vickers 1989, and Smyth 1992 for details.

---

[5] The topological version of this property is known as *sobriety*; cf. e.g. Smyth 1992.

[6] A *topology* on a set U is a subset system τ over U that is closed with respect to finite intersection and arbitrary union. The members of τ are referred to as *open sets*. A *topological space* is a set U together with a topology τ on U. A subset V of U is called *compact* if every cover of V by open sets contains a finite subcover of V.

6.3 Observational Logic

6.3.1 *The Calculus $OC_\equiv$*

At the close of Section 6.2.1 it has been indicated that in order to obtain a sound and complete inference calculus for observational statements we simply have to axiomatize the laws of a distributive lattice with zero and unit as well as the rules for congruence closure. These requirements are obviously met by the inference calculus $OC_\equiv$ whose inference and axiom schemes are stated in Figures 48 and 49.

$$\frac{}{A \equiv A}\ \text{(Refl)} \qquad \frac{A \equiv B}{B \equiv A}\ \text{(Symm)} \qquad \frac{A \equiv B \qquad B \equiv C}{A \equiv C}\ \text{(Trans)}$$

$$\frac{A \equiv B \qquad C \equiv D}{A \wedge C \equiv B \wedge D}\ (\text{Subst}_\wedge) \qquad \frac{A \equiv B \qquad C \equiv D}{A \vee C \equiv B \vee D}\ (\text{Subst}_\vee)$$

FIGURE 48 Inference schemes of $OC_\equiv$

$$\begin{array}{ccr}
(A \wedge B) \wedge C \equiv A \wedge (B \wedge C) & (A \vee B) \vee C \equiv A \vee (B \vee C) & \text{(Assoc)} \\
A \wedge B \equiv B \wedge A & A \vee B \equiv B \vee A & \text{(Comm)} \\
A \wedge A \equiv A & A \vee A \equiv A & \text{(Idemp)} \\
A \wedge (A \vee B) \equiv A & A \vee (A \wedge B) \equiv A & \text{(Absorb)} \\
\multicolumn{2}{c}{A \wedge (B \vee C) \equiv (A \wedge B) \vee (A \wedge C)} & \text{(Distr)} \\
A \wedge \mathrm{V} \equiv A \quad \text{(Unit)} & A \vee \Lambda \equiv A & \text{(Zero)}
\end{array}$$

FIGURE 49 Axiom schemes of $OC_\equiv$

Let Γ be an observational theory over Σ. Recall from Section 6.1.1 that $\cong_\Gamma$ is the least congruence relation on $\mathrm{T}[\Sigma]$ such that $\varphi \cong_\Gamma \psi$ whenever $\varphi \equiv \psi$ belongs to Γ or to the axiom schemes of $OC_\equiv$. By the very definition of $\cong_\Gamma$ and $OC_\equiv$ we have:

$$\varphi \cong_\Gamma \psi \qquad \text{iff} \qquad \Gamma \vdash_{OC_\equiv} \varphi \equiv \psi.$$

Taken together with (6.14) this proves:

(6.19) Theorem The calculus $OC_\equiv$ is sound and strongly complete with respect to first-order entailment.

Figure 50 gives an overview of the various arguments connected to completeness. As said before, (iv) is a mere matter of definitions. (i) is trivial since M_Γ is a model of Γ; (i)$'$ was derived in Section 5.2.1 as an immediate consequence of the definition of M_Γ. The crucial result is (ii), which hinges on the Prime Ideal Theorem, whereas (ii)$'$ is again trivial. (iii) follows from the fact that first-order logic is capable of entailing the rules and axioms of $OC_\equiv$. Finally, (iii)$'$ holds because first-order logic is sound.

$$\begin{tikzcd}
\Gamma \vdash \varphi \equiv \psi \arrow[dr, bend right, "\text{(iii)}'"'] & \varphi \cong_\Gamma \psi \arrow[l, "\text{(iii)}"'] \arrow[r, leftrightarrow, "\text{(iv)}"] \arrow[d, bend left, "\text{(ii)}'"] & \Gamma \vdash_{OC_\equiv} \varphi \equiv \psi \\
& \mathrm{M}_\Gamma(\varphi) = \mathrm{M}_\Gamma(\psi) \arrow[u, bend left, "\text{(ii)}"] \arrow[d, bend left, "\text{(i)}'"] & \\
& \forall M(M(\varphi) = M(\psi)) \arrow[u, bend left, "\text{(i)}"] &
\end{tikzcd}$$

FIGURE 50 Overview of completeness arguments

6.3.2 *The Calculus OC*$_\preceq$

Suppose an observational theory Γ is given in conditional form. If we want to apply the calculus *OC*$_\equiv$ to Γ, the statements of Γ have to be rendered into biconditional form along (5.1). Alternatively, one can ask for an inference calculus that directly applies to conditional statements.

Consider the calculus *OC*$_\preceq$ consisting of the axiom and inference schemes presented in Figure 51. Like the calculus *HC* for Horn statements, see Section 3.3.1, *OC*$_\preceq$ includes *reflexivity* (R), *de nihilo quodlibet* (Q), and *universality* (U), as well as *introduction* ($\mathrm{I}_\wedge$) and *elimination of conjunction* ($\mathrm{E}^1_\wedge$), ($\mathrm{E}^2_\wedge$). In addition, we have *introduction* ($\mathrm{I}_\vee$) and *elimination of disjunction* ($\mathrm{E}^1_\vee$), ($\mathrm{E}^2_\vee$). Finally, there is left and right *weakening* ($\mathrm{W}_\wedge$), ($\mathrm{W}_\vee$), and *cut* (C).

Our goal is to prove the completeness of *OC*$_\preceq$ by showing its equivalence to *OC*$_\equiv$. To this end, every scheme of *OC*$_\equiv$, when put into conditional form via (5.1), must be shown to admit a proof by *OC*$_\preceq$. Take transitivity, for instance, whose conditional form is:

$$\frac{A \preceq B \qquad B \preceq C}{A \preceq C} \qquad \text{(T)}$$

Clearly (T) is provable in *OC*$_\preceq$ by weakening and cut:

$$A \preceq A \quad (\mathrm{R}) \qquad\qquad \Lambda \preceq A \quad (\mathrm{Q}) \qquad\qquad A \preceq \mathrm{V} \quad (\mathrm{U})$$

$$\frac{A \preceq B \qquad A \preceq C}{A \preceq B \wedge C}\;(\mathrm{I}_\wedge) \qquad \frac{A \preceq B \wedge C}{A \preceq B}\;(\mathrm{E}^1_\wedge) \qquad \frac{A \preceq B \wedge C}{A \preceq C}\;(\mathrm{E}^2_\wedge)$$

$$\frac{A \preceq C \qquad B \preceq C}{A \vee B \preceq C}\;(\mathrm{I}_\vee) \qquad \frac{A \vee B \preceq C}{A \preceq C}\;(\mathrm{E}^1_\vee) \qquad \frac{A \vee B \preceq C}{B \preceq C}\;(\mathrm{E}^2_\vee)$$

$$\frac{A \preceq B}{C \wedge A \preceq B}\;(\mathrm{W}_\wedge) \qquad\qquad \frac{A \preceq B}{A \preceq B \vee C}\;(\mathrm{W}_\vee)$$

$$\frac{A \preceq B \vee C \qquad A \wedge B \preceq C}{A \preceq C}\;(\mathrm{C})$$

FIGURE 51 The calculus $OC_{\preceq}$

$$\dfrac{\dfrac{A \preceq B}{A \preceq B \vee C}\,(\mathrm{W}_\vee) \qquad \dfrac{B \preceq C}{A \wedge B \preceq C}\,(\mathrm{W}_\wedge)}{A \preceq C}\,(\mathrm{C})$$

Recall from Section 3.3.2 that weakening and cut, when restricted to conjunctive statements, can be replaced by transitivity. In the presence of disjunction, the situation is more complicated. Weakening still follows from transitivity:

1:	$A \preceq B$		1:	$A \preceq B$	
2:	$C \wedge A \preceq C \wedge A$	(R)	2:	$B \vee C \preceq B \vee C$	(R)
3:	$C \wedge A \preceq A$	$(\mathrm{E}^2_\wedge)$	3:	$B \preceq B \vee C$	$(\mathrm{E}^1_\vee)$
4:	$C \wedge A \preceq B$	(T) 3, 1	4:	$A \preceq B \vee C$	(T) 3, 1

In order to prove cut, however, transitivity is not enough. What is needed in addition is an axiom scheme like (D) that ensures *distributivity* of conjunction and disjunction:

$$A \wedge (B \vee C) \;\preceq\; (A \wedge B) \vee (A \wedge C) \qquad (\mathrm{D})$$

We claim that (D) is provable by $OC_{\preceq}$. To check this, we make use of several intermediate results.[7] For a start, notice that a three times application of elimination of disjunction followed by a two times application of introduction of disjunction gives us the bidirectional inference scheme:[8]

(6.20)
$$\begin{array}{c} (A \vee B) \vee C \;\preceq\; D \\ \hline\hline A \vee (B \vee C) \;\preceq\; D \end{array}$$

Employing elimination and introduction of conjunction instead yields the bidirectional scheme:

(6.21)
$$\begin{array}{c} D \;\preceq\; (A \wedge B) \wedge C \\ \hline\hline D \;\preceq\; A \wedge (B \wedge C) \end{array}$$

Similarly, applying elimination twice and introduction once leads to the following two schemes:

(6.22)
$$\begin{array}{c} A \vee B \;\preceq\; C \\ \hline\hline B \vee A \;\preceq\; C \end{array}$$

(6.23)
$$\begin{array}{c} C \;\preceq\; A \wedge B \\ \hline\hline C \;\preceq\; B \wedge A \end{array}$$

The easy proof of the following four schemes is left as an exercise to the reader.

(6.24)
$$\frac{A \;\preceq\; B}{A \vee C \;\preceq\; B \vee C} \qquad\qquad \frac{A \;\preceq\; B}{C \vee A \;\preceq\; C \vee B}$$

(6.25)
$$\frac{A \;\preceq\; B}{A \wedge C \;\preceq\; B \wedge C} \qquad\qquad \frac{A \;\preceq\; B}{C \wedge A \;\preceq\; C \wedge B}$$

The penultimate step consists in proving the statement scheme

(6.26)
$$A \wedge (B \vee C) \;\preceq\; (A \wedge B) \vee C,$$

which is carried out in Figure 52. Figure 53, finally, presents the desired proof of (D).

Let us now return to the task of proving the "conditionalized" schemes of $OC_{\equiv}$ within $OC_{\preceq}$. Associativity and commutativity follow from the schemes

[7] Basically, we follow Curry 1963, Sect. 4B3.

[8] The double lines indicate inference schemes that are invertible.

1:	$A \wedge (B \vee C) \preceq A \wedge (B \vee C)$	(R)
2:	$A \wedge (B \vee C) \preceq B \vee C$	($\mathrm{E}^2_\wedge$)
3:	$A \wedge (B \vee C) \preceq (B \vee C) \vee (A \wedge B)$	($\mathrm{W}_\vee$)
4:	$(B \vee C) \vee (A \wedge B) \preceq B \vee (C \vee (A \wedge B))$	(R); (6.20)
5:	$A \wedge (B \vee C) \preceq B \vee (C \vee (A \wedge B))$	(T) 3, 4
6:	$C \vee (A \wedge B) \preceq (A \wedge B) \vee C$	(R); (6.22)
7:	$B \vee (C \vee (A \wedge B)) \preceq B \vee ((A \wedge B) \vee C)$	(6.24)
8:	$A \wedge (B \vee C) \preceq B \vee ((A \wedge B) \vee C)$	(T) 5, 7
9:	$A \wedge (B \vee C) \preceq (B \vee C) \wedge A$	(R); (6.23)
10:	$(A \wedge (B \vee C)) \wedge B \preceq ((B \vee C) \wedge A) \wedge B$	(6.25)
11:	$(A \wedge (B \vee C)) \wedge B \preceq (B \vee C) \wedge (A \wedge B)$	(6.21)
12:	$(A \wedge (B \vee C)) \wedge B \preceq A \wedge B$	($\mathrm{E}^2_\vee$)
13:	$(A \wedge (B \vee C)) \wedge B \preceq (A \wedge B) \vee C$	($\mathrm{W}_\vee$)
14:	$A \wedge (B \vee C) \preceq (A \wedge B) \vee C$	(C) 8, 13

FIGURE 52 Proof of $A \wedge (B \vee C) \preceq (A \wedge B) \vee C$

1:	$A \wedge (B \vee C) \preceq A \wedge (B \vee C)$	(R)
2:	$A \wedge (B \vee C) \preceq A$	($\mathrm{E}^1_\wedge$)
3:	$A \wedge (B \vee C) \preceq (A \wedge B) \vee C$	(6.26)
4:	$A \wedge (B \vee C) \preceq A \wedge ((A \wedge B) \vee C)$	($\mathrm{I}_\wedge$) 2, 3
5:	$(A \wedge B) \vee C \preceq C \vee (A \wedge B)$	(R); (6.22)
6:	$A \wedge ((A \wedge B) \vee C) \preceq A \wedge (C \vee (A \wedge B))$	(6.25)
7:	$A \wedge (B \vee C) \preceq A \wedge (C \vee (A \wedge B))$	(T) 4, 6
8:	$A \wedge (C \vee (A \wedge B)) \preceq (A \wedge C) \vee (A \wedge B)$	(6.26)
9:	$A \wedge (B \vee C) \preceq (A \wedge C) \vee (A \wedge B)$	(T) 7, 8
10:	$(A \wedge C) \vee (A \wedge B) \preceq (A \wedge B) \vee (A \wedge C)$	(R); (6.22)
11:	$A \wedge (B \vee C) \preceq (A \wedge B) \vee (A \wedge C)$	(T) 9, 10

FIGURE 53 Proof of (D)

(6.20) to (6.23). Unit and zero are immediate consequences of (U) and (Q). One direction of distributivity is covered by (D). As for the reverse direction, it is not difficult to give proofs of $A \wedge B \preceq A \wedge (B \vee C)$ and $A \wedge C \preceq A \wedge (B \vee C)$. Substitutivity follows from (6.24), (6.25), and (T). This leaves us with absorption and idempotency, each of which is provable in $OC_\preceq$ as the reader will check without problems. Consequently, by (6.19):

(6.27) Theorem The calculus $OC_\preceq$ is sound and strongly complete with respect to first-order entailment.

Notice that if the inference schemes ($\mathrm{W}_\wedge$), ($\mathrm{W}_\vee$), and (C) of $OC_\preceq$ are replaced by (T) and (D), the resulting calculus is again complete.

6.3.3 *Digression: Sequent Structures*

Sequent structures can be seen as a way to represent observational theories in normal form. A *sequent structure* consists of a set Σ of primitives and a binary relation $\preccurlyeq$ on the set of finite subsets of Σ.[9] (We also say that $\preccurlyeq$ is a sequent structure over Σ.) The pairs of finite subsets of Σ belonging to $\preccurlyeq$ are referred to as *sequents*.

Suppose Γ is an observational theory over Σ in normal form. That is, for every statement $\varphi \preceq \psi$ of Γ, φ is either V or a conjunction of primitives and ψ is either Λ or a disjunction of primitives. Let $\preccurlyeq_\Gamma$ be the sequent system over Σ such that

$$\lfloor\varphi\rfloor \preccurlyeq_\Gamma \lfloor\psi\rfloor \quad \text{iff} \quad (\varphi \preceq \psi) \in \Gamma.$$

($\lfloor\varphi\rfloor$ is the set of primitives occurring in an observational predicate φ over Σ.) For example, if Γ contains the statements $a \wedge b \preceq c \vee d$ and $c \wedge d \preceq \Lambda$, then $\{a, b\} \preccurlyeq_\Gamma \{c, d\}$ and $\{c, d\} \preccurlyeq_\Gamma \varnothing$. Clearly, all sequent structures over Σ arise that way from observational theories over Σ.

Since Γ has normal form, a subset X of Σ is consistently Γ-closed iff, for every statement $\varphi \preceq \psi$ of Γ, some element of $\lfloor\psi\rfloor$ belongs to X whenever all elements of $\lfloor\varphi\rfloor$ belong to X. Put differently, X belongs to $\mathrm{C}(\Gamma)$ just in case

$$\text{if} \quad P \subseteq X \quad \text{then} \quad X \cap Q \neq \varnothing$$

for all finite subsets P, Q of Σ with $P \preccurlyeq_\Gamma Q$.

If the calculus $OC_\preceq$ is restricted to statements in normal form, the introduction and elimination rules are obsolete because they either produce or presuppose statements that do not have normal form. The same is true of (Q) and (U), which both are not in normal form. This leaves us with reflexivity, weakening, and cut. In terms of sequents, it is common to state reflexivity, weakening, and cut as follows:

$$\{p\} \preccurlyeq \{p\} \tag{R}$$

$$\text{if} \quad P \preccurlyeq Q \quad \text{then} \quad P \cup R \preccurlyeq Q \cup S \tag{W}$$

$$\text{if} \quad P \preccurlyeq \{p\} \cup Q \quad \text{and} \quad P \cup \{p\} \preccurlyeq Q \quad \text{then} \quad P \preccurlyeq Q \tag{C}$$

for every member p and all finite subsets P, Q, R and S of Σ. Following Barwise (1992), we call a sequent structure *normal* if it is closed with respect

[9] Sequent structures are also called *non-deterministic information system* in Droste and Göbel 1990, *theories* in Barwise and Seligman 1997, and *entailment relations* in Cederquist and Coquand 2000. Our terminology is that of Zhang 1991 and Barwise 1992.

to reflexivity, weakening, and cut. The least normal sequent structure including a given sequent structure $\preccurlyeq$ is referred to as the *normal closure* of $\preccurlyeq$.

Reflexivity, weakening, and cut provide a complete inference calculus for sequents in the following sense:

(6.28) Theorem Suppose $\preccurlyeq_\Gamma$ is the sequent representation of a (normal) observational theory Γ over Σ. Let $\preccurlyeq$ be the normal closure of $\preccurlyeq_\Gamma$. Then

$$P \preccurlyeq Q \quad \text{iff} \quad \forall X \in \mathrm{C}(\Gamma)\,(P \subseteq X \rightarrow X \cap Q \neq \varnothing).$$

A direct proof of this fact can be found for example in Barwise 1992, p. 176. Alternatively, one can employ the completeness of *OC*$_\preceq$. It then remains to show that a proof of α from Γ by *OC*$_\preceq$, with Γ and α in normal form, can be transformed into one that uses only reflexivity, weakening, and cut. We leave this as an exercise to the reader.

(6.29) Remark Cederquist and Coquand (2000) give a construction of the observational algebra presented by an observational theory Γ that directly builds on the sequent structure representation of Γ. They also point out how their approach is related to *cut elimination* – a topic which regrettably had to be left aside in our discussion of the calculus *OC*$_\preceq$.

(6.30) Remark If Γ is a Horn theory in normal form, its associated sequent structure $\preccurlyeq_\Gamma$ consists solely of sequents of the form $P \preccurlyeq_\Gamma Q$, where Q is either empty or a singleton. In this case, the normal closure $\preccurlyeq$ of $\preccurlyeq_\Gamma$ has the property that if $P \preccurlyeq Q$ then $P \preccurlyeq \varnothing$ or $P \preccurlyeq \{q\}$ for some $q \in Q$. Such a sequent structure $\preccurlyeq$ is called *deterministic*. Normal deterministic sequent structures are closely related to so-called *Scott information systems*; see e.g. Zhang 1991, Sect. 3.2.

7

Finite Specifiability

As we saw in Chapter 4, the generic entities that are least satisfiers of (conjunctive) predicates play a key role in establishing the correspondence between Horn theories and Scott domains. Given a Horn theory Γ, the least satisfiers of conjunctive predicates form a basis of the generic universe of Γ, which means that every generic entity is the supremum of a directed set of least satisfiers. Moreover, the least satisfiers are precisely the compact elements of the generic universe.

In Section 7.1 we explore how far the connection between least satisfiers and compact elements carries over to observational theories in general. In Section 7.2 we focus on observational theories whose generic entities are least satisfiers or suprema of directed sets of those. The generic universes of these theories turn out to coincide with the coherent algebraic domains.

7.1 Finite Specifiability and Least Satisfiers

Suppose Γ is an observational theory over Σ. Let us describe what it means that a member x of the generic universe $\mathrm{U}(\Gamma)$ of Γ is the least satisfier of some observational predicate φ over Σ. Suppose x is the least satisfier of φ, that is, by definition,

(7.1) $$\text{if} \quad y \models \varphi \quad \text{then} \quad x \sqsubseteq y, \qquad \text{i.e.} \quad [\![\varphi]\!] = {\uparrow} x.$$

It follows that $y \models \psi$ whenever $y \models \varphi$ and $x \models \psi$; in other words,

(7.2) $$x \models \psi \quad \text{iff} \quad [\![\varphi]\!] \subseteq [\![\psi]\!].$$

Condition (7.2) is not only necessary but also sufficient for x to be the least satisfier of φ. This is so because (7.2) says that if $y \models \varphi$ and $x \models \psi$ then $y \models \psi$, for every ψ; hence, by (5.10), $x \sqsubseteq y$ whenever $y \models \varphi$.[1] Using completeness, we can restate (7.2) in the form:

$$x \models \psi \quad \text{iff} \quad \Gamma \vdash \varphi \preceq \psi.$$

So x is the least satisfier of φ just in case x satisfies φ and whenever x satisfies ψ then Γ entails $\varphi \preceq \psi$.

Since the least satisfier x of φ, if existent, is uniquely determined by φ, we also say that φ *specifies* x, or that x is *finitely specifiable* (because φ is finitely constructed over Σ). A further consequence of (7.2) is that finitely specifiable elements have neighborhood filters which are *principal*, i.e. are generated by singletons; concretely, the neighborhood filter of a finitely specifiable element x is generated by $\uparrow x$. In terms of prime filters of the Lindenbaum algebra $\mathrm{L}(\Gamma)$:

(7.3) Proposition There is a one-to-one correspondence between the finitely specifiable elements of $\mathrm{U}(\Gamma)$ and the principal prime filters of $\mathrm{L}(\Gamma)$.

A principal filter of a distributive lattice is prime iff its generator is join-irreducible. When applied to the extension algebra Ω (of the generic model of Γ), it follows that φ has a least satisfier iff $[\![\varphi]\!]$ is $\cup$-irreducible in Ω; put differently, φ has a least satisfier iff $[\varphi]$ is $\vee$-irreducible in the Lindenbaum algebra, that is, in terms of logic, iff φ is not equivalent to Λ and whenever φ is equivalent to $\psi_1 \vee \psi_2$ then φ is equivalent to ψ_1 or to ψ_2 (where equivalence means equivalence with respect to Γ). Hence:

(7.4) Corollary There is an order-reversing one-to-one correspondence between the finitely specifiable elements of $\mathrm{U}(\Gamma)$ and the join-irreducible elements of $\mathrm{L}(\Gamma)$.

Though the correspondence between finitely specifiable elements and compact elements does not carry over from Horn theories to observational theories in general, it is easy to see that finite specifiability implies compactness. Just notice that, according to (5.20)(i), for every (upwards) directed subset S of $\mathrm{U}(\Gamma)$,

(7.5) $$\bigsqcup S \models \varphi \quad \text{iff} \quad \exists x \in S\, (x \models \varphi).$$

Ergo, if φ specifies x and $x \sqsubseteq \bigsqcup S$ then $\bigsqcup S \models \varphi$ and hence $y \models \varphi$, for some member y of S; hence $x \sqsubseteq y$, by (7.1). Therefore:

[1] Besides unfolding of definitions the argument simply is: $\forall y(y \models \varphi \to \forall \psi(x \models \psi \to y \models \psi))$ iff $\forall y \forall \psi(y \models \varphi \wedge x \models \psi \to y \models \psi)$ iff $\forall \psi(x \models \psi \to \forall y(y \models \varphi \to y \models \psi))$, which is just a way of stating the *Galois connection* between $\wp(\mathrm{U}(\Gamma))$ and $\wp(\mathrm{T}[\Sigma])$ determined by $\models$.

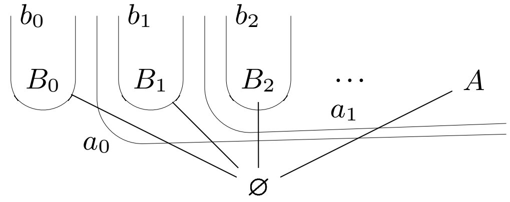

FIGURE 54 Flat canonical universe with non-principal neighborhood filter

(7.6) Proposition Finitely specifiable generic entities are compact with respect to specialization.

The converse of (7.6), however, is false, as the following example shows.

(7.7) Example Let Σ be $\{a_0, a_1, \ldots\} \cup \{b_0, b_1, \ldots\}$, and let Γ be the theory over Σ consisting of the statements

$$a_n \wedge b_n \equiv \Lambda, \qquad a_n \equiv a_{n+1} \vee b_{n+1} \qquad (n \geqslant 0).$$

Then $\mathrm{C}(\Gamma) = \{B_0, B_1, B_2, \ldots\} \cup \{\varnothing, A\}$, with $B_n = \{a_0, a_1, \ldots, a_{n-1}, b_n\}$ and $A = \{a_0, a_1, a_2, \ldots\}$; see Figure 54, which in addition shows the extensions of primitives. Obviously all elements of $\mathrm{C}(\Gamma)$ are compact (with respect to specialization). Now observe that the element A of $\mathrm{C}(\Gamma)$ is not finitely specifiable and hence, by (7.3), corresponds to a non-principal prime filter of $\mathrm{L}(\Gamma)$. For suppose A is the least satisfier of φ. Since $[\varphi]$ is $\vee$-irreducible, we may assume that φ is of the form $p_1 \wedge \ldots \wedge p_m$, with $p_i \in \Sigma$. Then $\{p_1, \ldots, p_n\} \subseteq A$ because A satisfies φ. But φ is also satisfied by infinitely many of the B_n's. Hence A is not the least satisfier of φ and therefore not finitely specifiable. Notice that Γ is a choice system theory in the sense of Section 5.3.2, whose graphical presentation is as follows:

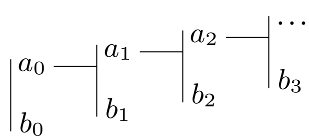

7.2 Finitistic Theories and Coherent Algebraic Domains

7.2.1 *Finitistic Theories*

We call an observational theory Γ *finitistic* if its generic entities are either finitely specifiable or suprema of directed sets of finitely specifiable ones; that is, Γ is finitistic just in case the set of finitely specifiable elements is a basis of

$\mathrm{U}(\Gamma)$. Since each basis of a dcpo includes all compact elements,[2] it follows by (7.6) that Γ is finitistic iff the finitely specifiable generic entities are precisely the compact elements of $\mathrm{U}(\Gamma)$. So, by (7.3), we can note:

(7.8) Proposition An observational theory is finitistic iff its generic universe is an algebraic domain whose compact elements have principal neighborhood filters.

It should be emphasized that an observational theory need not be finitistic to have a generic universe which is an algebraic domain. Non-finitistic theories may even induce flat specialization orders, as Example (7.7) shows. This case of course goes hand in hand with compact elements which are not finitely specifiable. Horn theories, on the other hand, provide positive examples, as we know from Section 4.1.2.

(7.9) Proposition Horn theories are finitistic.

The characterization of finitistic theories given so far refers both to the generic universe and to the Lindenbaum (or extension) algebra. It is also possible to characterize the generic universe of finitistic theories in purely order-theoretic terms, thereby resuming the program of Chapter 4. Likewise, one can try to characterize the Lindenbaum algebra of finitistic theories purely algebraically. Let us pursue the latter task first:

(7.10) Theorem An observational theory is finitistic iff each element of its Lindenbaum algebra is a finite join of join-irreducibles.

Proof. Suppose Γ is a finitistic theory with generic universe U and corresponding extension algebra Ω. By (7.8) and (7.1), the neighborhood filter of every $y \in \mathrm{k}(U)$ is generated by $\uparrow y$, which is $\cup$-irreducible in Ω. Take any $V \in \Omega$. Suppose $x \in V$; since U is algebraic, there is an $S \subseteq \mathrm{k}(U)$ such that $x = \bigsqcup S$; by (7.5), some $y \in S$ belongs to V; hence $x \in \uparrow y \subseteq V$. Therefore, V is the union of the sets $\uparrow y$, with $y \in \mathrm{k}(U) \cap V$. Application of (6.10) now shows that V is a finite union of $\cup$-irreducibles. Conversely, let Γ be an observational theory where every member of Ω is the union of $\cup$-irreducibles. For $x \in U$, let S be the (nonempty) set of those finitely specifiable $y \in U$ such that $y \sqsubseteq x$, i.e. $x \in \uparrow y$. We are finished if we can show that S is directed with supremum x. Suppose $y, y' \in S$. Then $x \in \uparrow y \cap \uparrow y' \in \Omega$. By assumption, x belongs to some $\cup$-irreducible $\uparrow z \subseteq \uparrow y \cap \uparrow y'$; hence $y \sqsubseteq z$ and $y' \sqsubseteq z$. It remains to show that $x = \bigsqcup S$. We have $y \sqsubseteq x$ for each $y \in S$, that is, $\bigsqcup S \sqsubseteq x$. Finally, if $x \in V \in \Omega$

[2] Suppose B is a basis of a dcpo D and $x \in D$; then $x = \bigsqcup S$ for some directed subset S of B. If $x \in \mathrm{k}(D)$ then $x \sqsubseteq y$ for some $y \in S$; hence $x = y \in B$.

then, by assumption, $x \in \uparrow y \subseteq V$ for some finitely specifiable y; thus, for every $V \in \Omega$, if $x \in V$ then $y \in V$, for some $y \in S$, and hence $\bigsqcup S \in V$. Consequently, $x \sqsubseteq \bigsqcup S$. $\lrcorner$

Since the $\cup$-irreducibles of Ω are precisely the sets $\uparrow x$, with x compact, and because $\uparrow S = \bigcup\{\uparrow x \mid x \in S\}$ for every subset S of an ordered set, (7.10) can be reformulated as follows:

(7.11) Corollary An observational theory with generic universe U is finitistic iff the extension algebra consists of the sets $\uparrow F$ with $F \subseteq \mathrm{k}(U)$ finite.

One can therefore recover the Lindenbaum algebra of a finitistic theory from its ordered generic universe.

Since, according to (6.6), every observational algebra A is the Lindenbaum algebra of an observational theory, we can combine our results to derive the following well-known fact about distributive lattices with zero and unit, where $\mathrm{J}(A)$ is the set of join-irreducible elements of A:

(7.12) Theorem Suppose A is an observational algebra whose every element is the join of a finite subset of $\mathrm{J}(A)$. Then $\mathrm{J}(A)^{\mathrm{d}} \simeq \mathrm{k}(\mathrm{P}(A))$. Moreover, A is isomorphic to the algebra $\{\downarrow F \mid F \subseteq \mathrm{J}(A) \text{ finite}\}$ of subsets of $\mathrm{J}(A)$, where $a \in A$ corresponds to $\{x \in \mathrm{J}(A) \mid x \leqslant a\}$, and $\downarrow F$ is taken to $\bigvee F$, for finite $F \subseteq \mathrm{J}(A)$.

7.2.2 *Coherent Algebraic Domains*

We know from (7.8) that the ordered generic universe U of a finitistic theory Γ is an algebraic domain. Let us try to characterize $\mathrm{k}(U)$ with respect to specialization by using (7.10). Recall that the neighborhood filter of each $x \in \mathrm{k}(U)$ is generated by a $\cup$-irreducible extension, namely by $\uparrow x$. Suppose F is a finite subset of $\mathrm{k}(U)$. Then, by (7.10), there is a finite subset M of $\mathrm{k}(U)$ such that

(7.13) $$\bigcap\{\uparrow x \mid x \in F\} \;=\; \bigcup\{\uparrow y \mid y \in M\}.$$

Hence, each $y \in M$ is an upper bound of F in $\mathrm{k}(U)$; moreover, if $z \in \mathrm{k}(U)$ is an upper bound of F then $y \sqsubseteq z$ for some $y \in M$, since $\uparrow z \subseteq \bigcup\{\uparrow y \mid y \in M\}$ implies that $\uparrow z \subseteq \uparrow y$ for some $y \in M$. In addition, we can assume all members of M to be *minimal* upper bounds of F because removing non-minimal upper bounds from M does not affect (7.13). Figure 55 is intended to sketch this situation.

Every finite subset of $\mathrm{k}(U)$ thus has a *finite complete set of minimal upper bounds* in $\mathrm{k}(U)$. This property of $\mathrm{k}(U)$ is sometimes referred to as the *2/3 SFP condition*. An algebraic domain U such that $\mathrm{k}(U)$ satisfies this condition is

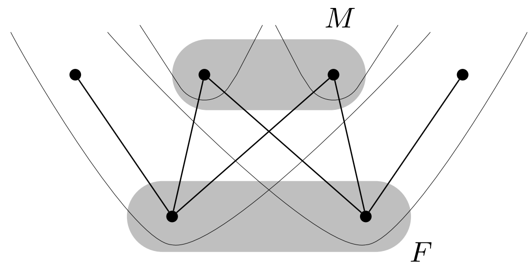

FIGURE 55 Finite complete set M of minimal upper bounds of F

called *coherent*, or a *2/3 bifinite domain*.[3] Notice that the 2/3 SFP condition, when applied to the empty set, implies that $\mathrm{k}(U)$ and thus U can have only finitely many minimal elements.

(7.14) Theorem The generic universe of a finitistic theory is a coherent algebraic domain and every such domain arises that way.

Proof. It remains to show that every coherent algebraic domain U is the generic universe of some finitistic theory. It follows from (7.13) that the system Ω of sets of the form $\uparrow F$, with finite $F \subseteq \mathrm{k}(U)$, is closed with respect to finite intersection and hence, by (7.10), is (isomorphic to) the Lindenbaum algebra of some (finitistic) theory Γ. We still need to show that U is indeed the generic universe of Γ, i.e., that every prime filter of Ω is the neighborhood filter of some member of U. Let $\mathcal{P}$ be a prime filter of Ω. The set $S = \{x \in \mathrm{k}(U) \mid \uparrow x \in \mathcal{P}\}$ is directed; for if F is a finite subset of S and M is the complete set of minimal upper bounds of F then

$$\bigcup\{\uparrow y \mid y \in M\} \;=\; \bigcap\{\uparrow x \mid x \in F\} \;\in\; \mathcal{P};$$

hence $\uparrow y \in \mathcal{P}$ for some $y \in M$, since M is finite, and y is an upper bound for F in S. It follows that $\mathcal{P}$ is the neighborhood filter of $\bigsqcup S \in U$. $\lrcorner$

Since the number of minimal elements of a 2/3 bifinite domain is finite, it is obvious that the property of being finitistic does not dualize in general. For example, the full binary exclusion theory over an infinite set Σ of primitives is finitistic and its generic universe is an infinite flat domain. So the dual theory Γ, which consists of the statements $\mathrm{V} \preceq p \vee q$, with $p \neq q$, has a generic universe with infinitely many minimal elements. Concretely, $\mathrm{C}(\Gamma) = \{\Sigma\} \cup \{\Sigma \setminus \{p\} \mid p \in \Sigma\}$. Clearly $\Sigma \setminus \{p\}$ is not finitely specifiable.

[3] See Abramsky and Jung 1994, Sect. 4.2.3 for a more general definition of coherence, which is independent of algebraicity. (Beware, the use of 'coherent' in domain theory is rather inconsistent. Some authors call a dcpo coherent if it is binary-complete; cf. footnote 7 on page 64.)

A *countable* vocabulary has the same effect on the generic universe of finitistic theories as it has in the case of Horn theories (cf. (4.12)): each generic entity is the limit of a specialization sequence consisting of finitely specifiable generic entities.

(7.15) Theorem The generic universe of a finitistic theory over a countable set of primitives is a countably based coherent algebraic domain, and vice versa.

Proof. If the set of primitives is countable then also the set of observational predicates and hence the set of finitely specifiable generic entities. Now apply (7.14) and the definition of finitistic theories. $\lrcorner$

(7.16) Remark *(Coherent algebraic prelocales)* In order to represent coherent algebraic domains "in logical form", Abramsky (1991) has introduced *coherent algebraic prelocales*, which, roughly speaking, are something halfway between finitistic theories and their Lindenbaum algebras.[4] A coherent algebraic prelocale is a preordered algebra A with zero-place operators 0 and 1, two-place operators $\vee$ and $\wedge$, and a monadic predicate C on A, such that 0 and 1 are respectively least and greatest elements, $\vee$ and $\wedge$ yield suprema and infima, and C is true of those elements that are $\vee$-irreducible. The predicate C is used to axiomatize the requirement that every element of A is a finite join of $\vee$-irreducibles.

In domain theory, one is traditionally interested in categories of domains which are *Cartesian closed*, i.e. closed under function spaces, mainly because this property allows to regard models of the λ-calculus as solutions of a *recursive domain equation*. The category of coherent algebraic (i.e. 2/3 bifinite) domains is known to be *not Cartesian* closed, in contrast to the category of *bifinite domains*. A domain U is called *bifinite* (or *SFP* or *strongly algebraic*) if it is coherent algebraic and every finite subset F of $\mathrm{k}(U)$ has a *finite mub-closure* in $\mathrm{k}(U)$, that is, there is a finite superset M of F in $\mathrm{k}(U)$ such that if $F' \subseteq M$ then M contains a complete set of minimal upper bounds of F' in $\mathrm{k}(U)$. The following standard example of a coherent algebraic domain violating this property has a particularly simple specification in terms of observational statements.

(7.17) Example Let Γ be the theory over $\{a_0, a_1, \ldots\} \cup \{b_0, b_1, \ldots\}$ with statements

$$a_n \wedge b_n \equiv a_{n+1} \vee b_{n+1} \qquad (n \geqslant 0).$$

[4] See also Abramsky and Jung 1994, Sect. 7.3.

The canonical universe $\mathrm{C}(\Gamma)$ of Γ is depicted in Figure 56. Clearly, the set $\{\{a_0\}, \{b_0\}\}$ has an *infinite* mub-closure.

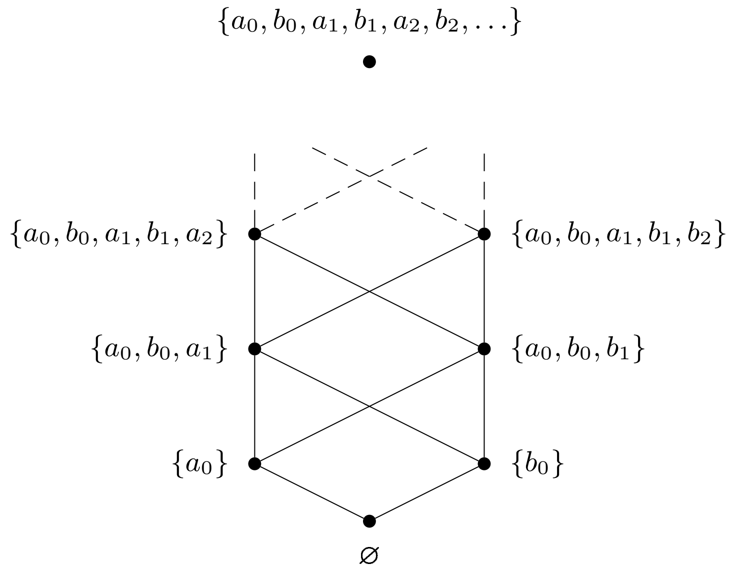

FIGURE 56 A non-bifinite coherent algebraic domain

7.2.3 *Digression: The Descending Chain Condition*

We close the section on finitistic theories by mentioning a classical criterion for an observational algebra to be the Lindenbaum algebra of a finitistic theory. Recall that an ordered set satisfies the *descending chain condition* if every nonempty subset has a minimal element (whereas the *ascending chain condition* requires that every nonempty subset has a maximal element).

Suppose an observational algebra A satisfies the descending chain condition. Then, clearly, every filter of A has a least element and is hence principal. Since, in particular, all elements of $\mathrm{P}(A)$ are principal, A can be presented by a finitistic theory. Consequently, by (7.10):

(7.18) Proposition If an observational algebra satisfies the descending chain condition, then each of its elements is a finite join of join-irreducibles.[5]

So, an observational theory Γ whose Lindenbaum algebra satisfies the descending chain condition is finitistic. Moreover, all members of $\mathrm{U}(\Gamma)$ are finitely specifiable (and hence compact). In addition, it follows by (7.4) that $\mathrm{U}(\Gamma)$ satisfies the ascending chain condition. The converse of the last conclusion is

[5]See e.g. Birkhoff 1967, Chap. VIII, §2 for a direct proof.

of course false: Example (7.7) shows that $\mathrm{U}(\Gamma)$ can satisfy the ascending chain condition without Γ being finitistic. Also notice that $\mathrm{U}(\Gamma)$ may violate the descending chain condition even if $\mathrm{L}(\Gamma)$ satisfies it; witness Example (5.31).

Part III

Translations, Extensions, Constructions

8

Translation of Theories

In this chapter, we study the relationship between observational theories over *different* sets of primitives. A typical problem is to find a translation between two theories given they are "denotationally equivalent" in the sense of having isomorphic generic universes. A related task is to find a translation of a theory into one of a certain type, say a Horn theory, given its generic universe is the generic universe of a theory of the type in question. For instance, every rooted choice system theory is known to have a flat generic universe; see (5.37). Hence such a theory is denotationally equivalent to a Horn theory (and even to a full binary exclusion theory; cf. Table 1).

In Section 8.1, we introduce *morphisms* of observational theories as translations of predicates that respect the theories in question. In the category of theories thus defined, the notion of *equivalence* turns out to be more useful than that of *isomorphism*. Equivalence morphisms between theories are characterized as being *conservative* and *essentially surjective*.

Section 8.2 is concerned with various functors from the category of observational theories to that of observational algebras and of dcpos. The overall goal is to find out in which way morphisms of theories are reflected by (Scott-continuous) functions of their generic universes. Since the "generic universe functor" factors by the "Lindenbaum functor" and the "prime spectrum functor", the latter two are studied in detail. We show, for instance, that the Lindenbaum functor induces an equivalence between the category of observational algebras and the quotient category of observational theories modulo equivalence.

In Section 8.3, this framework is applied to the problem of translating choice system theories into Horn theories. Moreover, we briefly address the problem of translating a (finite) theory into one that uses a minimal number of primitives. Section 8.4 is concerned with *extensions* of theories, which are theory morphisms where the primitives and statements of one theory are included in those of the other. Examples are Booleanization and rule completion.

8.1 Categories of Observational Theories

8.1.1 *Predicate Translations and Theory Morphisms*

Suppose two scientists who have arrived at two different theories about a certain universe of discourse want to decide whether their theories are equivalent. If both agree that they use the same vocabulary in the same way, equivalence means that both theories entail each other, which is the case if every statement of one theory is entailed by the other and vice versa. This is precisely the definition of equivalent theories introduced at the close of Section 5.1.1.

Consider now the case of two theories Γ and Γ' over different sets Σ and Σ' of primitive predicates. Assume that both scientists do not hinge on their chosen predicates (or concepts) but are willing to express them in terms of those of the other. Within the framework of observational logic this means to interpret each primitive predicate of Γ, i.e. each member of Σ, by some member of $\mathrm{T}[\Sigma']$, i.e. by some observational predicate over Σ'.

Let us call a function μ from Σ to $\mathrm{T}[\Sigma']$ a *(predicate) translation*. The homomorphic extension $\hat{\mu}$ of μ is a homomorphism of term algebras from $\mathrm{T}[\Sigma]$ to $\mathrm{T}[\Sigma']$, that is, $\hat{\mu}$ takes predicates over Σ to those over Σ'. Given a predicate φ over Σ we say that $\hat{\mu}(\varphi)$ is the *translation of φ by μ*, or the *μ-translation of* φ; similarly, we speak of μ-translations of statements and theories. If primitives are interpreted by primitives, i.e. if $\mu(\Sigma) \subseteq \Sigma'$, we speak of a *primitive (preserving) translation*. The function μ is then completely determined by a function from Σ to Σ'.

A prerequisite for Γ and Γ' to be equivalent is that Γ' entails the translation of Γ by μ, in short,

(8.1) $$\Gamma' \vdash \hat{\mu}(\Gamma).$$

For the statements of Γ when translated into the language of Γ' must not add information not already deducible from Γ'. If μ satisfies (8.1), we speak of a *morphism of observational theories* from Γ to Γ'.

(8.2) Example *(Extensions)* Suppose Γ and Γ' are observational theories over Σ and Σ', respectively, where $\Sigma \subseteq \Sigma'$ and $\Gamma \subseteq \Gamma'$. Then the inclusion function ε from Σ to Σ' is a (primitive) morphism of theories from Γ to Γ'. We speak in this case of an *extension (morphism)*. Extensions will be investigated in more detail in Section 8.4.

Those readers who ever got in touch with category theory will be familiar with the term “morphism”; those who are not should memorize the slogan that as soon as you consider a certain class of mathematical objects you should be able to say how these objects are related to each other. The concept of *category*

$$\begin{tikzcd}
\Sigma \arrow[d, hook] \arrow[dr, "\mu"] & \Sigma' \arrow[d, hook] \arrow[dr, "\mu'"] & \Sigma'' \arrow[d, hook] \\
\mathrm{T}[\Sigma] & \mathrm{T}[\Sigma'] \arrow[r, dashed, "\widehat{\mu'}"'] & \mathrm{T}[\Sigma'']
\end{tikzcd}$$

$$\begin{tikzcd}
\Gamma \arrow[dr, mapsto] & \Gamma' \arrow[d, phantom, "\top"] \arrow[dr, mapsto] & \Gamma'' \arrow[d, phantom, "\top"] \\
& \hat{\mu}(\Gamma) \arrow[dr, mapsto] & \widehat{\mu'}(\Gamma') \arrow[d, phantom, "\top"] \\
& & \widehat{\mu'}(\hat{\mu}(\Gamma))
\end{tikzcd}$$

FIGURE 57 Composition of theory morphisms

incorporates this idea in terms of morphisms between objects:[1] a *category* consists of a class of *objects* and a class of *morphisms*, where each morphism has a unique "source" and "target" object. In addition, there is an identity morphism $\imath_A$ from each object A to itself. Moreover, for every two morphisms f from A to B and g from B to C there is a morphism $g \circ f$ from A to C, the *composite* of g and f, subject to the condition that $\circ$ is associative and $f \circ \imath_A = f = \imath_B \circ f$.

Our goal is to define a category $\mathbf{Th}$ of observational theories and theory morphisms. More precisely, the objects of $\mathbf{Th}$ are the pairs $\langle \Sigma, \Gamma \rangle$ where Γ is an observational theory over Σ. Keeping this in mind, we often call the pair $\langle \Sigma, \Gamma \rangle$ an observational theory and write 'Γ' instead of '$\langle \Sigma, \Gamma \rangle$' in case Σ is clear from the context or irrelevant. To make $\mathbf{Th}$ into a category, it remains to define appropriate composites of morphisms as well as identity morphisms. The *identity morphism* $\imath_\Gamma$ of Γ is the canonical inclusion function of Σ into $\mathrm{T}[\Sigma]$. (Notice that identity morphisms of different theories over Σ are identical as functions but differ as theory morphisms with respect to source and target.) The composite of two theory morphisms μ from Γ to Γ' and μ' from Γ' to Γ'' is the composite function $\widehat{\mu'} \circ \mu$ from Σ to $\mathrm{T}[\Sigma'']$; see the upper diagram of Figure 57. In order to show that this definition does indeed give rise to a morphism of theories we make use of the following observation:

(8.3) Lemma Suppose μ is a function from Σ to $\mathrm{T}[\Sigma']$, Γ is a theory over Σ, α is a statement over Σ, and $\Gamma \vdash \alpha$. Then $\hat{\mu}(\Gamma) \vdash \hat{\mu}(\alpha)$.

Proof. Since $\hat{\mu}$ is a homomorphism of term algebras, it preserves all inference schemes of the (complete) calculus $\mathit{OC}_{\equiv}$. $\lrcorner$

[1] Standard references on category theory are Mac Lane 1971 and Schubert 1972. Texts with applications to computer science are Poigné 1992 and Barr and Wells 1995. For a modest introduction see Lawvere and Schanuel 1997 or Magnan and Reyes 1994.

(8.4) Proposition The composite of two theory morphisms is a theory morphism in turn. Moreover, composition is associative with identity morphisms as units.

Proof. Suppose μ and μ' are theory morphisms as above. We need to show that $\Gamma'' \vdash \widehat{\widehat{\mu'} \circ \mu}(\Gamma)$. Since μ is a morphism, $\Gamma' \vdash \hat{\mu}(\Gamma)$. By (8.3), it follows that $\widehat{\mu'}(\Gamma') \vdash \widehat{\mu'}(\hat{\mu}(\Gamma))$, as indicated by lower diagram of Figure 57. Hence, because μ' is a morphism and entailment is transitive, $\Gamma'' \vdash \widehat{\mu'}(\hat{\mu}(\Gamma))$. Now apply (5.3). Associativity of composition is also a consequence (5.3), since $(\mu'' \circ \mu') \circ \mu$ is $\widehat{\widehat{\mu''} \circ \mu'} \circ \mu$ and $\mu'' \circ (\mu' \circ \mu)$ is $\widehat{\mu''} \circ \widehat{\mu'} \circ \mu$. Finally, identity morphisms are obviously units with respect to composition. $\lrcorner$

The categorical viewpoint gives us the notion of isomorphic theories for free: a morphism is an *isomorphism* iff it has a two-side inverse with respect to composition. Thus, two observational theories Γ and Γ' are isomorphic iff there are two translations μ from Σ to $\mathrm{T}[\Sigma']$ and μ' from Σ' to $\mathrm{T}[\Sigma]$ such that $\widehat{\mu'} \circ \mu = \imath_\Gamma$, $\hat{\mu} \circ \mu' = \imath_{\Gamma'}$, $\Gamma' \vdash \hat{\mu}(\Gamma)$, and $\Gamma \vdash \widehat{\mu'}(\Gamma')$. Unfortunately, this notion of isomorphism is too restrictive for our purposes. For the empty theory over $\{a\}$ should surely count as equivalent to the theory $\{b \equiv c\}$ over $\{b, c\}$. But if $\hat{\mu}(\mu'(b)) = b$ then necessarily $\mu'(b) = a$ (because $\hat{\mu}$ takes non-primitive terms to non-primitive terms). Similarly $\mu'(c) = a$, and therefore $\hat{\mu}(\mu'(c)) = b$. So, there is no isomorphism between the two theories.

8.1.2 *Equivalent Theories, Equivalent Morphisms*

Before we turn more closely to the problem of characterizing the equivalence of theories, we point out that there is a useful notion of equivalence for theory morphisms too. Recall that a theory morphism μ from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$ gives rise to a translation of predicates over Σ into predicates over Σ' which takes Γ-equivalent predicates to Γ'-equivalent ones. Suppose ν is a second morphism from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$. Resuming our little scientific discourse from the last section, we may ask in which cases the holders of the theories Γ and Γ' are willing to admit that choosing μ or ν is a mere matter of taste. The decision obviously depends upon whether the holder of Γ' regards the μ-translation of each predicate φ over Σ as Γ'-equivalent to the ν-translation of φ. It is thus reasonable to say that μ and ν are *equivalent morphisms* (notation: $\mu \sim \nu$) iff

(8.5)
$$\forall p \in \Sigma \; (\Gamma' \vdash \mu(p) \equiv \nu(p)),$$

which implies that Γ' entails $\hat{\mu}(\varphi) \equiv \hat{\nu}(\varphi)$ for all $\varphi \in \mathrm{T}[\Sigma]$.

Let us now return to the question as to when a morphism μ establishes an equivalence of theories. As mentioned at the close of Section 8.1.1, it is too restrictive to require that μ has an inverse in the category $\mathbf{Th}$. The preceding discussion suggests that it might suffice if there is a morphism ν from Γ' to Γ such that

$$\nu \circ \mu \sim \imath_\Gamma \qquad \text{and} \qquad \mu \circ \nu \sim \imath_{\Gamma'}.$$

We then say that ν is "*quasi-inverse*" to μ. To make sure that the existence of a quasi-inverse is indeed an appropriate criterion for equivalence, it is helpful to consider the problem from another angle. Roughly speaking, we require for equivalence, first, that everything expressible in one theory is expressible in the other, and, second, that every regularity captured by one theory is also captured by the other. If μ is a morphism from Γ to Γ' then, by (8.1) and (8.3), $\Gamma' \vdash \hat{\mu}(\alpha)$ whenever $\Gamma \vdash \alpha$. The second requirement can therefore be expressed as follows:

(8.6) $$\Gamma \vdash \alpha \quad \text{iff} \quad \Gamma' \vdash \hat{\mu}(\alpha).$$

As for the first condition, we take 'identical expressivity' to mean that for every $\varphi' \in \mathrm{T}[\Sigma']$ there is a $\varphi \in \mathrm{T}[\Sigma]$ such that $\hat{\mu}(\varphi)$ is equivalent to φ' with respect to Γ'; it suffices to require this condition for primitives:

(8.7) $$\forall p \in \Sigma' \; \exists \varphi \in \mathrm{T}[\Sigma] \; (\Gamma' \vdash p \equiv \hat{\mu}(\varphi)).$$

The reader might want to convince herself that (8.6) and (8.7) have the desired effect. Every predicate φ over Σ has a translation over Σ' by μ, namely $\hat{\mu}(\varphi)$; this is what the morphism μ provides anyway. In addition, by (8.7), every predicate over Σ' is Γ'-equivalent to the μ-translation of a predicate over Σ. Similar for statements: every statement α' over Σ' that is entailed by Γ' is Γ'-equivalent to the μ-translation of a statement α over Σ, which, by (8.6), is entailed by Γ.

In case a morphism μ satisfies (8.6) we say that μ is *conservative*; if μ satisfies (8.7), it will be called *essentially surjective*.

(8.8) Theorem A theory morphism μ has a quasi-inverse if and only if μ is conservative and essentially surjective.

Proof. Suppose ν is quasi-inverse to μ. Then $\Gamma' \vdash \varphi' \equiv \hat{\mu}(\hat{\nu}(\varphi'))$, for every $\varphi' \in \mathrm{T}[\Sigma']$, which shows that μ is essentially surjective. If $\Gamma' \vdash \hat{\mu}(\varphi) \equiv \hat{\mu}(\psi)$ then $\Gamma \vdash \hat{\nu}(\hat{\mu}(\varphi)) \equiv \hat{\nu}(\hat{\mu}(\psi))$, because ν is a morphism. Moreover, Γ entails $\varphi \equiv \hat{\nu}(\hat{\mu}(\varphi))$ and $\psi \equiv \hat{\nu}(\hat{\mu}(\psi))$. Hence $\Gamma \vdash \varphi \equiv \psi$; so, μ is conservative. Conversely, assume μ is conservative and essentially surjective. Then, by (8.7) (and the axiom of choice), one can choose $\nu(p') \in \mathrm{T}[\Sigma]$, for every $p' \in \Sigma'$, such that $\Gamma' \vdash p' \equiv \hat{\mu}(\nu(p'))$. Thus $\mu \circ \nu \sim \imath_{\Gamma'}$. In particular, for every $p \in \Sigma$, it holds that $\Gamma' \vdash \mu(p) \equiv \hat{\mu}(\hat{\nu}(\mu(p)))$. Hence $\Gamma \vdash p \equiv \hat{\nu}(\mu(p))$, by (8.6); that is, $\nu \circ \mu \sim \imath_{\Gamma}$. ⌟

In the light of this theorem, we call a morphism μ from Γ to Γ' an *equivalence (morphism)* if μ satisfies the conditions of (8.8); the theories Γ and Γ' are then said to be *equivalent*; in symbols: $\Gamma \sim \Gamma'$

(8.9) Example Let Σ be $\{a_0, a_1, \dots a_k\} \cup \{b_0, b_1, \dots b_k\}$, where k is finite, and let Γ be the theory over Σ with statements

$$a_n \equiv a_{n+1} \vee b_{n+1} \qquad \text{and} \qquad a_n \wedge b_n \equiv \Lambda \qquad (0 \leqslant n < k).$$

Notice that Γ is a rooted choice system theory in the sense of Section 5.3.2 and hence, according to (5.37), has a flat generic universe. Concretely, $\mathrm{C}(\Gamma)$ consists of the sets $\varnothing$, $\{a_0, a_1, \dots, a_k\}$, and $\{a_0, a_1, \dots, a_{n-1}, b_n\}$ for all $n \leqslant k$. The generic universe of Γ is therefore isomorphic to the generic universe of the full binary exclusion theory $\Gamma' = \{c_m \wedge c_n \equiv \Lambda \mid m \neq n\}$ over $\Sigma' = \{c_0, c_1, \dots, c_{k+1}\}$. We shall see below in Section 8.2.4 that finite theories are equivalent whenever they have isomorphic generic universes. For the moment, however, we are content with finding an equivalence morphism from Γ to Γ'. Let μ be the function from Σ to $\mathrm{T}[\Sigma']$ with

$$\mu(a_n) = c_{n+1} \vee \dots \vee c_{k+1} \qquad \text{and} \qquad \mu(b_n) = c_n \qquad (0 \leqslant n \leqslant k).$$

Clearly Γ' entails $\hat{\mu}(\Gamma)$ (since $(c_{n+1} \vee \dots \vee c_{k+1}) \wedge c_n \equiv (c_{n+1} \wedge c_n) \vee \dots \vee (c_{k+1} \wedge c_n) \equiv \Lambda$ for all $n \leqslant k$). So μ is a theory morphism from Γ to Γ'. Now consider the function ν from Σ' to $\mathrm{T}[\Sigma]$ such that $\nu(c_{k+1}) = a_k$ and $\nu(c_n) = b_n$ for every $n \leqslant k$. We want to show that ν is a morphism from Γ' to Γ. To this end, observe that Γ (put into conditional form via (5.1)) entails $b_m \preceq a_n$ and hence $b_m \wedge b_n \preceq \Lambda$ for all $m > n$. In addition, Γ entails $a_m \preceq a_n$ if $m \geqslant n$, from which it follows that $a_m \wedge b_n \equiv a_m \wedge a_n \wedge b_n \equiv \Lambda$ whenever $m \geqslant n$. All in all, this proves that Γ entails $\hat{\nu}(\Gamma')$. It remains to check that ν is quasi-inverse to μ. The only nontrivial claim is that Γ entails $\nu(\mu(a_n)) \equiv a_n$, i.e. that $\Gamma \vdash a_n \equiv b_{n+1} \vee \dots \vee b_k \vee a_k$, which follows easily by induction. Figure 58 illustrates the situation for $k = 1$ in terms of Lindenbaum algebras and canonical universes (decorated with extensions of primitives). The shaded circles in the diagrams of $\mathrm{L}(\Gamma)$ and $\mathrm{L}(\Gamma')$ correspond to join-irreducible elements, which, by (7.4), stand in an (order-reversing) one-to-one correspondence to the elements of $\mathrm{C}(\Gamma)$ and $\mathrm{C}(\Gamma')$. Clearly, the equivalence morphism μ from Γ to Γ' induces an isomorphism from $\mathrm{L}(\Gamma)$ to $\mathrm{L}(\Gamma')$. We shall see in Section 8.2.1 that this holds in general.

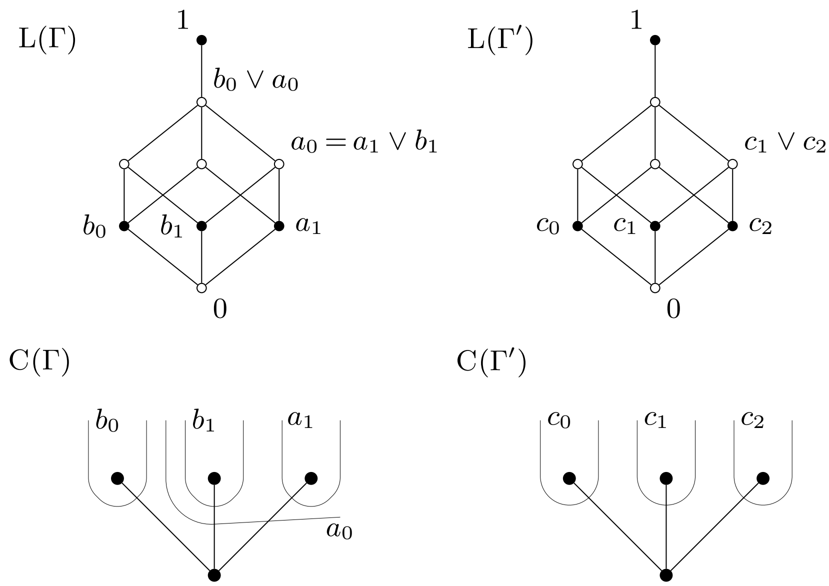

FIGURE 58 The case $k = 1$

(8.10) Remark *(Theories as categories)* Readers with background in category theory will recognize (8.8) as an instance of the fact that a functor F from a category $\mathbf{C}$ to a category $\mathbf{D}$ is an equivalence of categories iff F is full and faithful and every object of $\mathbf{D}$ is isomorphic to $F(c)$ for some $c \in \mathbf{C}$.[2] An observational theory Γ over Σ determines a category as follows: the objects are the members of $\mathrm{T}[\Sigma]$ and the morphisms are the pairs $\langle \varphi, \psi \rangle$ with $\Gamma \vdash \varphi \preceq \psi$. Since there is at most one morphism between any two objects, the category thus defined is a *preorder*. One easily checks that a morphism of theories is a functor between the respective categories. Such a functor is trivially faithful since the categories in question are preorders. Moreover, the functor is full iff it is conservative as a theory morphism. Finally, a theory morphism from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$ is essentially surjective just in case every object $\varphi' \in \mathrm{T}[\Sigma']$ is isomorphic, i.e. Γ'-equivalent, to $\hat{\mu}(\varphi)$ for some $\varphi \in \mathrm{T}[\Sigma]$.

8.2 A Bunch of Functors

The functorial perspective provides the appropriate formal machinery to study how morphisms between observational theories are reflected by functions of their generic universes. Again, the algebraic representation of theories turns out to be a valuable tool.

[2]See e.g. Mac Lane 1971, Sect. IV.4.

8.2.1 *The Functors* $\mathbf{L}$ *and* $\mathbf{Th}$

Our first goal is to show that every morphism of theories gives rise to a homomorphism of the respective Lindenbaum algebras and that this assignment is functorial. As an intermediate step, let us see how theory morphisms act on models:

(8.11) Lemma Let μ be a morphism of observational theories from Γ to Γ'. If m is a model of Γ' in an observational algebra A then $\widehat{m} \circ \mu$ is a model of Γ in A.

Proof. Suppose $(\varphi \equiv \psi) \in \Gamma$. Then $\Gamma' \vdash \hat{\mu}(\varphi) \equiv \hat{\mu}(\psi)$, by assumption. Hence $\widehat{m}(\hat{\mu}(\varphi)) = \widehat{m}(\hat{\mu}(\psi))$, for every model m of Γ'. ⌟

Recall from Section 6.1.1 that Γ' has a universal algebraic model $\mathrm{m}_{\Gamma'}$ in $\mathrm{L}(\Gamma')$. Applying (8.11) to this model shows that every morphism μ of observational theories from Γ to Γ' determines a model $\widehat{\mathrm{m}}_{\Gamma'} \circ \mu$ of Γ in $\mathrm{L}(\Gamma')$. By (6.2), on the other hand, every model m of Γ in $\mathrm{L}(\Gamma')$ uniquely determines a homomorphism h from $\mathrm{L}(\Gamma)$ to $\mathrm{L}(\Gamma')$ such that $m = h \circ \mathrm{m}_\Gamma$:

(8.12)
$$\mathrm{Mod}(\Gamma, \mathrm{L}(\Gamma')) \;\simeq\; \mathrm{Hom}(\mathrm{L}(\Gamma), \mathrm{L}(\Gamma')).$$

Taken together, every morphism μ from Γ to Γ' gives rise to a homomorphism $\mathrm{L}(\mu)$ from $\mathrm{L}(\Gamma)$ to $\mathrm{L}(\Gamma')$ such that $\widehat{\mathrm{m}}_{\Gamma'} \circ \mu = \mathrm{L}(\mu) \circ \mathrm{m}_\Gamma$.

$$\begin{tikzcd}
\Sigma \arrow[r, "\mu"] \arrow[d, "\mathrm{m}_\Gamma"'] & \mathrm{T}[\Sigma'] \arrow[d, "\widehat{\mathrm{m}}_{\Gamma'}"] \\
\mathrm{L}(\Gamma) \arrow[r, "\mathrm{L}(\mu)"'] & \mathrm{L}(\Gamma')
\end{tikzcd}$$

Let $\mathbf{Alg}$ be the category of observational algebras and homomorphisms.

(8.13) Proposition L is a functor from $\mathbf{Th}$ to $\mathbf{Alg}$.

Proof. Obviously L takes identity morphisms to identity homomorphisms. If μ and ν are morphisms from Γ to Γ' and from Γ' to Γ'', respectively, then

$$\mathrm{L}(\nu \circ \mu) \circ \mathrm{m}_\Gamma \;=\; \widehat{\mathrm{m}}_{\Gamma''} \circ \hat{\nu} \circ \mu \;=\; \mathrm{L}(\nu) \circ \mathrm{L}(\mu) \circ \mathrm{m}_\Gamma.$$

Hence $\mathrm{L}(\nu \circ \mu) = \mathrm{L}(\nu) \circ \mathrm{L}(\mu)$, by (8.12). ⌟

Since the Lindenbaum construction, roughly speaking, turns equivalence of predicates into identity, one might expect that this holds for equivalence of theories and morphisms too; indeed:

(8.14) Proposition Suppose μ and ν are theory morphisms from Γ to Γ'. Then $\mathrm{L}(\mu) = \mathrm{L}(\nu)$ iff $\mu \sim \nu$. In particular, $\mathrm{L}(\Gamma) \simeq \mathrm{L}(\Gamma')$ iff $\Gamma \sim \Gamma'$.

Proof. By (8.5) and (6.14), $\mu \sim \nu$ iff, for every $\varphi \in \mathrm{T}[\Sigma]$, $\hat{\mu}(\varphi) \cong_{\Gamma'} \hat{\nu}(\varphi)$, that is, iff $\mathrm{L}(\mu) \circ \mathrm{m}_\Gamma = \mathrm{L}(\nu) \circ \mathrm{m}_\Gamma$. By (8.12), the latter condition implies that $\mathrm{L}(\mu) = \mathrm{L}(\nu)$. ⌟

By (8.14) and (8.13), it follows that $\mathrm{L}(\mu)$ is an isomorphism iff μ has a quasi-inverse, that is, by (8.8), iff μ is conservative and essentially surjective. This fact splits into two parts:

(8.15) Proposition Let μ be a morphism of theories from Γ to Γ'.

(i) $\mathrm{L}(\mu)$ is one-to-one iff μ is conservative.

(ii) $\mathrm{L}(\mu)$ is onto iff μ is essentially surjective.

Proof. (i) $\mathrm{L}(\mu)$ is one-to-one iff $\cong_\Gamma$ is the congruence kernel of $\widehat{\mathrm{m}}_{\Gamma'} \circ \hat{\mu}$, that is, iff, for all $\varphi, \psi \in \mathrm{T}[\Sigma]$, $\varphi \cong_\Gamma \psi$ just in case $\hat{\mu}(\varphi) \cong_{\Gamma'} \hat{\mu}(\psi)$. According to (6.14) and (8.6), this is precisely what conservativity of μ means. (ii) $\mathrm{L}(\mu)$ is onto iff $\widehat{\mathrm{m}}_{\Gamma'} \circ \hat{\mu}$ is onto, which is the case iff for every $\varphi' \in \mathrm{T}[\Sigma']$ there is a $\varphi \in \mathrm{T}[\Sigma]$ such that $\varphi' \cong_{\Gamma'} \hat{\mu}(\varphi)$, that is, by (6.14) and (8.7), iff μ is essentially surjective. ⌟

(8.16) Example Suppose $\Sigma = \{a\}$, $\Sigma' = \{a, b\}$, Γ is empty, and Γ' consists of $a \wedge b \equiv \Lambda$ and $a \vee b \equiv \mathrm{V}$. Consider the (primitive) morphism μ of theories from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$ that takes a to a. Then $\mathrm{L}(\Gamma)$ and $\mathrm{L}(\Gamma')$ are as depicted in the upper part of Figure 59. Since $\mathrm{L}(\mu)$ is an embedding of $\mathrm{L}(\Gamma)$ into $\mathrm{L}(\Gamma')$, μ is conservative. Moreover, μ is not essentially surjective since $\mathrm{L}(\mu)$ is not onto; in terms of condition (8.7), there is no $\varphi \in \mathrm{T}[\Sigma]$ such that $\Gamma' \vdash b \equiv \hat{\mu}(\varphi)$. (In Figure 59, the framed elements are the join-irreducibles of $\mathrm{L}(\Gamma)$ and $\mathrm{L}(\Gamma')$. The lower part of the figure shows the function of canonical universes induced by μ; see also (8.32) below.)

(8.17) Remark A word of caution on the distinction between epimorphisms and surjective functions should be given. In category theory, a morphism f from A to B is called an *epimorphism* if, for every two morphisms g and h from B to C, $g = h$ whenever $g \circ f = h \circ f$. Notice that in the category $\mathbf{Set}$ of sets a morphism, i.e. a function, is an epimorphism iff it is onto. In $\mathbf{Alg}$, however, this in not the case: the homomorphism $\mathrm{L}(\mu)$ of (8.16) is an *epimorphism of observational algebras that is not onto*. For suppose g and h are homomorphisms from $\mathrm{L}(\Gamma')$ to an observational algebra A such that $g \circ \mathrm{L}(\mu) = h \circ \mathrm{L}(\mu)$. Then $g(a) = h(a)$ and hence $g(b) = h(b)$ because in a distributive lattice an element can have at most one complement.

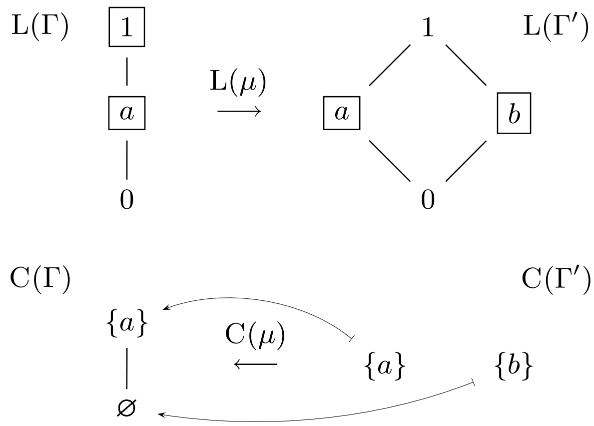

FIGURE 59 Epimorphism $\mathrm{L}(\mu)$ of algebras which is not onto

(8.18) Proposition The functor L from $\mathbf{Th}$ to $\mathbf{Alg}$ is full and every object A of $\mathbf{Alg}$ is of the form $\mathrm{L}(\Gamma)$ for some object Γ of $\mathbf{Th}$.

Proof. Since the second part of the statement is just a reformulation of (6.6), it is enough to verify that L is full. Let h be an algebra homomorphism from $\mathrm{L}(\Gamma)$ to $\mathrm{L}(\Gamma')$. We need to find a morphism μ from Γ to Γ' with $h = \Gamma(\mu)$. For every $p \in \Sigma$ choose $\mu(p) \in \mathrm{T}[\Sigma']$ such that $[\mu(p)]_{\cong_{\Gamma'}} = h([p]_{\cong_\Gamma})$. All we need to show is that μ is a morphism, i.e. that $\Gamma' \vdash \hat{\mu}(\Gamma)$. Suppose $(\varphi \equiv \psi) \in \Gamma$. Then $[\varphi]_{\cong_\Gamma} = [\psi]_{\cong_\Gamma}$ and thus $[\hat{\mu}(\varphi)]_{\cong_{\Gamma'}} = h([\varphi]_{\cong_\Gamma}) = h([\psi]_{\cong_\Gamma}) = [\hat{\mu}(\varphi)]_{\cong_{\Gamma'}}$. Hence $\Gamma' \vdash \hat{\mu}(\varphi) \equiv \hat{\mu}(\psi)$. ┘

So every homomorphism h of observational algebras is *presented* by some morphism μ of theories in the sense that $h = \mathrm{L}(\mu)$.

According to (8.13) and (8.14), L is a functor from $\mathbf{Th}$ to $\mathbf{Alg}$ such that $\mathrm{L}(\mu) = \mathrm{L}(\nu)$ iff $\mu \sim \nu$. It follows that $\sim$ is a *congruence relation* with respect to composition (that is, if $\mu \sim \nu$ and $\mu' \sim \nu'$ then $\mu \circ \mu' \sim \nu \circ \nu'$). The so-called *quotient category* $\mathbf{Th}/{\sim}$ of $\mathbf{Th}$ by $\sim$ has the same objects as $\mathbf{Th}$ whereas its morphisms are the equivalence classes of morphisms of $\mathbf{Th}$ modulo $\sim$.[3] Let $Q_\sim$ be the quotient functor from $\mathbf{Th}$ to $\mathbf{Th}/{\sim}$, that is, $Q_\sim(\Gamma) = \Gamma$ and $Q_\sim(\mu) = [\mu]_\sim$. It is straightforward to show that L factors uniquely by $Q_\sim$ and a faithful functor from $\mathbf{Th}/{\sim}$ to $\mathbf{Alg}$. In combination with (8.18), we get:

(8.19) Theorem The categories $\mathbf{Th}/{\sim}$ and $\mathbf{Alg}$ are equivalent.

(8.20) Remark Another option for defining a category $\mathbf{Th}'$ which is equivalent to $\mathbf{Alg}$ and whose objects are those of $\mathbf{Th}$ is to take the models of Γ in

[3] For more on quotient categories see e.g. Barr and Wells 1995, Sect. 3.5, Mac Lane 1971, Sect. II.8, or Tokizawa and Kanki 1985.

$\mathrm{L}(\Gamma')$ as morphisms from Γ to Γ'. The equivalence of $\mathbf{Th}'$ and $\mathbf{Alg}$ follows from (8.12).

Recall from Section 6.1.3 that every algebra A is presented by the theory $\mathrm{Th}(A)$ over A, that is, $\mathrm{L}(\mathrm{Th}(A)) \simeq A$. In particular, $\mathrm{L}(\mathrm{Th}(\mathrm{L}(\Gamma))) \simeq \mathrm{L}(\Gamma)$ for every theory Γ, and hence $\mathrm{Th}(\mathrm{L}(\Gamma)) \sim \Gamma$, by (8.14). Moreover, an algebra homomorphism from A to B, when viewed as a function from A to B, trivially determines a primitive theory morphism $\mathrm{Th}(h)$ from $\mathrm{Th}(A)$ to $\mathrm{Th}(B)$. Clearly Th is a functor from $\mathbf{Alg}$ to $\mathbf{Th}$.

Let $\mathbf{Th_p}$ be the subcategory of $\mathbf{Th}$, whose objects are those of $\mathbf{Th}$ and whose morphisms are the primitive theory morphisms. $\mathbf{Th_p}$ corresponds to the standard category of presentations by generators and relations as used e.g. by Oles (2000). Since Th by definiton takes algebra homomorphisms to primitive morphisms, it can be regarded as a functor from $\mathbf{Alg}$ to $\mathbf{Th_p}$.

(8.21) Theorem The functor L from $\mathbf{Th_p}$ to $\mathbf{Alg}$ is left adjoint to Th.

Proof. Let Γ be an observational theory over Σ and A an observational algebra. By definition of $\mathrm{Th}(A)$, a primitive morphism μ from Γ to $\mathrm{Th}(A)$ is a function from Σ to A such that $\hat{\mu}(\Gamma) \subseteq \mathrm{Th}(A)$. The latter condition, on the other hand, says precisely that μ is model of Γ in A. So there is a bijection between primitive morphisms from Γ to $\mathrm{Th}(A)$ and models of Γ in A. Together with (6.3) it follows that there is a bijection between primitive morphisms from Γ to $\mathrm{Th}(A)$ and homomorphisms from $\mathrm{L}(\Gamma)$ to A. $\lrcorner$

A consequence of (8.21) is that L as a functor from $\mathbf{Th_p}$ to $\mathbf{Alg}$ preserves colimits, whereas its right adjoint Th preserves limits (cf. Section 9.1.2).

(8.22) Remark *(Subobjects)* A *subobject* of an object A of an arbitrary category $\mathbf{C}$ is usually defined to be an equivalence class of monomorphisms with target A, where two monomorphisms f from B to A and f' from B' to A are said to be equivalent if there is an isomorphism g from B to B' such that $f = f' \circ g$. Since monomorphisms in $\mathbf{Alg}$ are one-to-one, and vice versa, (8.15)(i) says that *subobjects* (or, better, *quasi-subobjects*) in $\mathbf{Th}$ are given by conservative morphisms, with ‘isomorphism’ relaxed to ‘equivalence morphism’.

8.2.2 *The Functors* P *and* $\mathbf{Hom_2}$

Recall that the prime spectrum $\mathrm{P}(A)$ of an observational algebra A is the set of all prime filters of A ordered by set inclusion. It is easy to show that inverse images of prime filters under homomorphisms of observational algebras are prime filters in turn:

(8.23) Lemma If h is an algebra homomorphism from A to B and $F \in \mathrm{P}(B)$ then $h^{-1}(F) \in \mathrm{P}(A)$.

Proof. Suppose $P \in \mathrm{P}(A)$. Then $1 \in h^{-1}(P)$ because $h(1) = 1 \in P$. In addition $0 \notin h^{-1}(P)$ since $h(0) = 0 \notin P$. Moreover, $a \in h^{-1}(P)$ and $b \in h^{-1}(P)$ iff $h(a) \in P$ and $h(b) \in P$ iff $h(a \wedge b) \in P$ iff $a \wedge b \in h^{-1}(P)$. Similarly, $a \vee b \in h^{-1}(P)$ iff $a \in h^{-1}(P)$ or $b \in h^{-1}(P)$. $\lrcorner$

So, an algebra homomorphism h from A to B gives rise to a function $\mathrm{P}(h)$ from $\mathrm{P}(B)$ to $\mathrm{P}(A)$ such that $\mathrm{P}(h)(F) = h^{-1}(F)$. Clearly $\mathrm{P}(h)$ is order-preserving. Moreover, $\mathrm{P}(h)$ can be shown to preserve suprema of directed sets. An order-preserving function of dcpos which preserves suprema of directed sets is said to be *Scott-continuous*.

(8.24) Proposition If h is an algebra homomorphism from A to B then $\mathrm{P}(h)$ is a Scott-continuous function of dcpos from $\mathrm{P}(B)$ to $\mathrm{P}(A)$.

Proof. If $\mathcal{S}$ is an upwards directed subset of $\mathrm{P}(B)$ then $\bigcup \mathcal{S} \in \mathrm{P}(B)$. By definition, $a \in \mathrm{P}(h)(\bigcup \mathcal{S})$ iff $h(a) \in \bigcup \mathcal{S}$, i.e. iff $h(a) \in F$ and hence $a \in \mathrm{P}(h)(F)$ for some $F \in \mathcal{S}$. Therefore, $\mathrm{P}(h)(\bigcup \mathcal{S}) = \bigcup\{\mathrm{P}(h)(F) \mid F \in \mathcal{S}\}$. $\lrcorner$

Let $\mathbf{Dcpo}$ be the category of dcpos and Scott-continuous functions. Since $\mathrm{P}(g \circ h) = \mathrm{P}(h) \circ \mathrm{P}(g)$ and because P takes identity homomorphisms to identity functions, we have:

(8.25) Corollary P is a contravariant functor from $\mathbf{Alg}$ to $\mathbf{Dcpo}$.

Our primary reason for studying the functors L and P is that their composite $\mathrm{P} \circ \mathrm{L}$ is a (contravariant) functor from $\mathbf{Th}$ to $\mathbf{Dcpo}$ which takes each observational theory to its generic universe, see (6.12), and each morphism of theories to a Scott-continuous function.

Suppose, for example, one aims at characterizing those functions of generic universes that are induced by conservative morphisms and by essentially surjective ones. Then, according to (8.15), it suffices to characterize $\mathrm{P}(h)$ if h is one-to-one or onto. In order to do so, we need the following description of the prime filters of a subalgebra.

(8.26) Proposition For every subalgebra B of an observational algebra A,

$$\mathrm{P}(B) \;=\; \{F \cap B \mid F \in \mathrm{P}(A)\}.$$

Proof. [4] Application of (8.23) to the inclusion of B into A shows that $F \cap B$ belongs to $\mathrm{P}(B)$ for every $F \in \mathrm{P}(A)$. Suppose now that $F' \in \mathrm{P}(B)$. We need to show that there is an $F \in \mathrm{P}(A)$ such that $F \cap B = F'$. Observe that $\uparrow F'$ and $\downarrow(B \setminus F')$ are respectively filter and ideal in A with empty intersection. Therefore, by the Prime Ideal Theorem, there is a prime filter F of A such that $\uparrow F' \subseteq F$ and $F \cap \downarrow(B \setminus F') = \varnothing$. Consequently, $F' \subseteq F \cap B$ and $F \cap (B \setminus F') = \varnothing$; hence $F' = F \cap B$. $\lrcorner$

(8.27) Proposition Let h be a homomorphism of observational algebras.

(i) h is one-to-one iff $\mathrm{P}(h)$ is onto.

(ii) h is onto iff $\mathrm{P}(h)$ is an order embedding.

Proof. [5] Suppose h is an algebra homomorphism from A to B. (i) The forward direction is a consequence of (8.26). As for the reverse direction suppose h is not one-to-one, that is, there are $a, b \in A$ such that $a \neq b$ and $h(a) = h(b)$. We may assume that $a \nleq b$. According to the Prime Ideal Theorem, there is a prime filter F of A such that $a \in F$ and $b \notin F$. Since F is not the inverse image by h of any subset of B, $\mathrm{P}(h)$ is not onto. (ii) If h is onto then $h(h^{-1}(X)) = X$ for every $X \subseteq B$; it follows that $\mathrm{P}(h)$ is an order embedding. Conversely, suppose h is not onto. Then there is an $a \in B \setminus B'$ with $B' = h(A)$. Applying the Prime Ideal Theorem to the ideal $\downarrow a$ and the filter $\uparrow(\uparrow a \cap B')$ gives us a prime filter F of B such that $a \notin F$ and $\uparrow a \cap B' \subseteq F$. A second application of the theorem to $\downarrow(B' \setminus F)$ and $\uparrow a$ yields a prime filter F' of B such that $a \in F'$ and $F' \cap (B' \setminus F) = \varnothing$; see Figure 60. Since $F' \cap B' \subseteq F \cap B'$ but $F' \nsubseteq F$, it follows that $\mathrm{P}(h)$ is not an order embedding. $\lrcorner$

According to (6.11), there is an order isomorphism τ_A from the pointwise ordered set $\mathrm{Hom}(A, \mathbb{2})$ to $\mathrm{P}(A)$ such that $\tau_A(g) = g^{-1}(1)$. Now consider the contravariant Hom-functor $\mathrm{Hom}_{\mathbb{2}}$ from $\mathbf{Alg}$ to $\mathbf{Set}$, which takes every algebra A to $\mathrm{Hom}(A, \mathbb{2})$ and every homomorphism h from A to B to the function that takes $g \in \mathrm{Hom}(B, \mathbb{2})$ to $g \circ h \in \mathrm{Hom}(A, \mathbb{2})$.[6] Straightforward unraveling of definitions shows that

$$\tau_A \circ \mathrm{Hom}_{\mathbb{2}}(h) \;=\; \mathrm{P}(h) \circ \tau_B.$$

[4] Adapted from Balbes and Dwinger 1974, p. 74, Theorem 5.

[5] Partly based on Davey and Priestley 1990, p. 217, Exercise 10.6(iii).

[6] More generally, if C is an observational algebra then the *contravariant Hom-functor* Hom_C from $\mathbf{Alg}$ to $\mathbf{Set}$ takes every algebra A to the set $\mathrm{Hom}(A, C)$ and every homomorphism h from A to B to the function which takes $g \in \mathrm{Hom}(B, C)$ to $g \circ h \in \mathrm{Hom}(A, C)$. It is essentially a restatement of definitions that Hom_C takes epimorphisms to one-to-one functions and colimits in $\mathbf{Alg}$ to limits in $\mathbf{Set}$ (cf. Section 9.2.2).

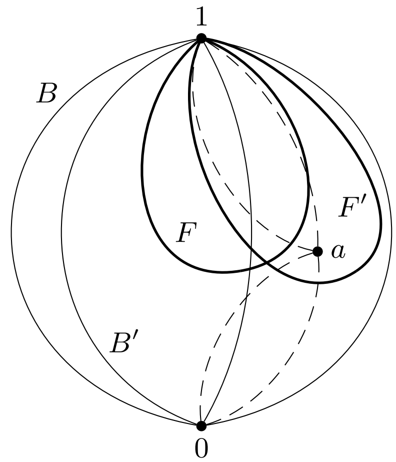

FIGURE 60 The Prime Ideal Theorem at work

Hence τ is a natural isomorphism of functors from $\mathrm{Hom}_{\mathbb{2}}$ to P, with P taken as a functor from $\mathbf{Alg}$ to $\mathbf{Set}$. Moreover, since τ_A is an order isomorphism, it follows by (8.24) that $\mathrm{Hom}_{\mathbb{2}}(h)$ is Scott-continuous. Hence $\mathrm{Hom}_{\mathbb{2}}$ is a functor from $\mathbf{Alg}$ to $\mathbf{Dcpo}$; moreover:

(8.28) Proposition The functors P and $\mathrm{Hom}_{\mathbb{2}}$ from $\mathbf{Alg}$ to $\mathbf{Dcpo}$ are naturally isomorphic.

We are now prepared to define a functor that takes theories to their ordered generic universe and theory morphisms to Scott-continuous functions.

8.2.3 *The Functors* $\mathbf{Mod_{2}}$, C, ... *and* U

Recall from (6.4) that the generic universe of a theory Γ can be represented by the set $\mathrm{Mod}_{\mathbb{2}}(\Gamma)$ of $\mathbb{2}$-valued models of Γ with pointwise order. In order to extend $\mathrm{Mod}_{\mathbb{2}}$ to a functor we make use of the general fact that (8.11) gives rise to a contravariant functor Mod_A from $\mathbf{Th}$ to $\mathbf{Set}$, for every observational algebra A, where $\mathrm{Mod}_A(\Gamma)$ is the set $\mathrm{Mod}(\Gamma, A)$ of A-models of Γ and $\mathrm{Mod}_A(\mu)$, with μ a theory morphism from Γ to Γ', is the function from $\mathrm{Mod}_A(\Gamma')$ to $\mathrm{Mod}_A(\Gamma)$ that takes m to $\widehat{m} \circ \mu$.

(8.29) Proposition The functors Mod_A and $\mathrm{Hom}_A \circ \mathrm{L}$ from $\mathbf{Th}$ to $\mathbf{Set}$ are naturally isomorphic.

Proof. By (6.3), $\mathrm{Hom}(\mathrm{L}(\Gamma), A) \simeq \mathrm{Mod}(\Gamma, A)$, for every theory Γ; concretely, a homomorphism h from $\mathrm{L}(\Gamma)$ to A corresponds to the A-valued Γ-model $h \circ \mathrm{m}_\Gamma$. One easily checks that these data define a natural isomorphism from $\mathrm{Hom}_A \circ \mathrm{L}$ to Mod_A. For let μ be a morphism of theories from Γ to Γ'. Then

$(h \circ \mathrm{L}(\mu)) \circ \mathrm{m}_\Gamma = (h \circ \widehat{\mathrm{m}}_{\Gamma'}) \circ \mu$, for every $h \in \mathrm{Hom}(\mathrm{L}(\Gamma'), A)$; cf. the diagram below (8.12). ⌟

So $\mathrm{Mod}_{\mathbb{2}}$ is naturally isomorphic to $\mathrm{Hom}_{\mathbb{2}} \circ \mathrm{L}$. Moreover, since the isomorphism between $\mathrm{Hom}(\mathrm{L}(\Gamma), \mathbb{2})$ and $\mathrm{Mod}(\Gamma, \mathbb{2})$ is an order isomorphism (with respect to pointwise order), we can conclude:

(8.30) Corollary The functors $\mathrm{Mod}_{\mathbb{2}}$ and $\mathrm{Hom}_{\mathbb{2}} \circ \mathrm{L}$ from $\mathbf{Th}$ to $\mathbf{Dcpo}$ are naturally isomorphic.

Next we extend the function that takes Γ to $\mathrm{C}(\Gamma)$ to a functor C from $\mathbf{Th}$ to $\mathbf{Dcpo}$. According to (6.4), $\mathrm{C}(\Gamma)$ is order-isomorphic to $\mathrm{Mod}_{\mathbb{2}}(\Gamma)$, where each member X of $\mathrm{C}(\Gamma)$ corresponds to its characteristic function χ_X. Employing this correspondence twice yields a Scott-continuous function $\mathrm{C}(\mu)$ from $\mathrm{C}(\Gamma')$ to $\mathrm{C}(\Gamma)$, for every morphism μ from Γ to Γ'. Concretely, if Y belongs to $\mathrm{C}(\Gamma')$ then

$$\mathrm{C}(\mu)(Y) \;=\; \{p \in \Sigma \mid \widehat{\chi}_Y(\mu(p)) = 1\}.$$

In case μ is primitive, i.e. if $\mu(\Sigma) \subseteq \Sigma'$, the effect of $\mathrm{C}(\mu)$ on $Y \in \mathrm{C}(\Gamma')$ simplifies to:

(8.31)
$$\mathrm{C}(\mu)(Y) \;=\; \{p \in \Sigma \mid \mu(p) \in Y\} \;=\; \mu^{-1}(Y).$$

For later use, the special case of extensions is stated separately:

(8.32) Proposition If ε is an extension of theories from $\langle \Sigma, \Gamma \rangle$ to $\langle \Sigma', \Gamma' \rangle$ then $\mathrm{C}(\varepsilon)(Y) = Y \cap \Sigma$. In particular, $\mathrm{C}(\varepsilon)$ is an order embedding if $\Sigma = \Sigma'$.

We have thus defined a functor C from $\mathbf{Th}$ to $\mathbf{Dcpo}$ which, by definition, is naturally isomorphic to $\mathrm{Mod}_{\mathbb{2}}$. Figure 61 gives an overview of the relevant functors introduced so far.

$$\begin{tikzcd}
\mathbf{Th} \arrow[rr, shift left, "\mathrm{L}"] \arrow[dr, "\mathrm{C} \simeq \mathrm{Mod}_{\mathbb{2}}"'] & & \mathbf{Alg} \arrow[ll, shift left, "\mathrm{Th}"] \arrow[dl, "\mathrm{P} \simeq \mathrm{Hom}_{\mathbb{2}}"] \\
& \mathbf{Dcpo} &
\end{tikzcd}$$

FIGURE 61 Overview of relevant functors

Starting with Section 1.4 it has been repeatedly emphasized that the concrete representation of the generic universe of a theory does not matter as long as only satisfaction and specialization are of interest. To take this into account, we have introduced the generic name '$\mathrm{U}(\Gamma)$' in Section 5.2.1. Meanwhile we can do better: Let U be any functor from $\mathbf{Th}$ to $\mathbf{Dcpo}$ that is naturally isomorphic to C and hence to Mod_2, $\mathrm{Hom}_2 \circ \mathrm{L}$, and $\mathrm{P} \circ \mathrm{L}$. So U is a contravariant functor from $\mathbf{Th}$ to the category $\mathbf{Dcpo}$ of dcpos and Scott-continuous functions; it takes objects of $\mathbf{Th}$ to objects of $\mathbf{Dcpo}$ and morphisms of $\mathbf{Th}$ from Γ to Γ' to Scott-continuous functions from $\mathrm{U}(\Gamma')$ to $\mathrm{U}(\Gamma)$.

To see how U interacts with satisfaction, consider its representation by Mod_2. Suppose μ is a morphism of theories from Γ to Γ'. By definition of $\mathrm{Mod}_2(\mu)$, we have $(\mathrm{Mod}_2(\mu)(m))(\varphi) = \widehat{m}(\hat{\mu}(\varphi))$ for all predicates φ over Σ. Hence, by (6.5),

(8.33) $$x \models \hat{\mu}(\varphi) \qquad \text{iff} \qquad \mathrm{U}(\mu)(x) \models \varphi,$$

for every member x of $\mathrm{U}(\Gamma')$ and every predicate φ over Σ.

Since $\mathrm{U} \simeq \mathrm{P} \circ \mathrm{L}$, it follows by (8.14) that equivalent theories have isomorphic generic universes. But beware, theories may have order-isomorphic generic universes without being equivalent:

(8.34) Example Consider the theory $\langle \Sigma, \Gamma \rangle$ of Example (7.7), that is, Σ is $\{a_0, a_1, \ldots\} \cup \{b_0, b_1, \ldots\}$ and Γ is $\{a_n \wedge b_n \equiv \Lambda,\ a_n \equiv a_{n+1} \vee b_{n+1} \,|\, n \geqslant 0\}$. Its generic universe $\mathrm{U}(\Gamma)$ is repeated on the left of Figure 62, with extensions of primitives added. Since $\mathrm{U}(\Gamma)$ is flat, it is (isomorphic to) the generic universe of the full exclusion theory $\langle \Sigma', \Gamma' \rangle$, with $\Sigma' = \{c_0, c_1, \ldots\} \cup \{c_\omega\}$ and $\Gamma' = \{c_m \wedge c_n \preceq \Lambda \mid m \neq n\}$; see again Figure 62. It was shown in (7.7) that $\mathrm{U}(\Gamma)$ contains a compact element that is not finitely specifiable and that $\mathrm{L}(\Gamma)$ consequently has a non-principal prime filter. On the other hand, we know by (7.3) and (7.9) that all prime filters of $\mathrm{L}(\Gamma')$ are principal. Hence $\mathrm{L}(\Gamma)$ is not isomorphic to $\mathrm{L}(\Gamma')$ and thus Γ is not equivalent to Γ', by (8.14). If we compare this result with Example (8.9), we see that the equivalence of theories defined there does not carry over to the infinite case although the generic universes remain isomorphic.

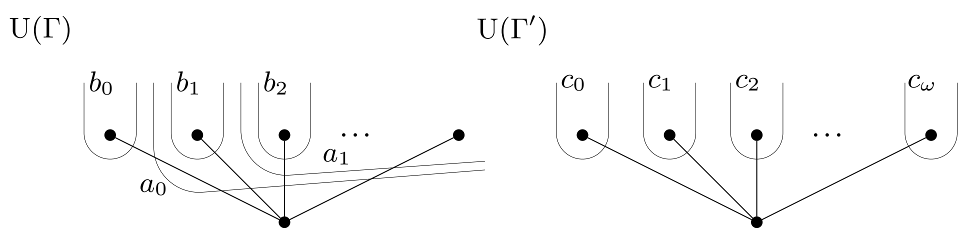

FIGURE 62 Non-equivalent theories with isomorphic generic universes

Finally, we can combine (8.27) and (8.15) to characterize the behavior of conservative and essentially surjective morphisms under U:

(8.35) Proposition Let μ be a morphism of observational theories.

(i) μ is conservative iff $\mathrm{U}(\mu)$ is onto.

(ii) μ is essentially surjective iff $\mathrm{U}(\mu)$ is an order embedding.

For example, if μ is a primitive morphism then, according to (8.31), $\mathrm{C}(\mu)$ takes $Y \in \mathrm{C}(\Gamma')$ to $\mu^{-1}(Y)$. If μ is conservative, it follows by (8.35)(i) that every $X \in \mathrm{C}(\Gamma)$ is of the form $\mu^{-1}(Y)$, with $Y \in \mathrm{C}(\Gamma')$.

8.2.4 *Finitistic Theories and Coherent Algebraic Domains*

As we saw earlier, the connection between theories (and algebras) on the one hand and dcpos on the other is not as tight as it is desirable. Neither is the functor U full nor is every dcpo (isomorphic to) the generic universe of an observational theory. In particular, two observational theories with isomorphic generic universes need not be equivalent (in which case they have non-isomorphic Lindenbaum algebras).

Finitistic theories behave much better in this respect since they are uniquely determined up to equivalence by their generic universe. Recall from Section 7.2 that the generic universes of finitistic theories are precisely the coherent algebraic domains. Moreover, the extension algebra of the generic model of a finitistic theory Γ consists of all sets of the form $\uparrow F$, with $F \subseteq \mathrm{k}(\mathrm{U}(\Gamma))$ finite.

Even when restricted to finitistic theories and coherent algebraic domains, the functor U is not full, that is, not every Scott-continuous function from $\mathrm{U}(\Gamma')$ to $\mathrm{U}(\Gamma)$, with Γ and Γ' finitistic, is of the form $\mathrm{U}(\mu)$, for some morphism μ from Γ to Γ'. For, by (8.33), we have that

(8.36)
$$\mathrm{U}(\mu)^{-1}(\llbracket\varphi\rrbracket) \;=\; \llbracket\hat{\mu}(\varphi)\rrbracket.$$

So, for every $x \in \mathrm{k}(\mathrm{U}(\Gamma))$, it follows that $\mathrm{U}(\mu)^{-1}(\uparrow x)$ is of the form $\uparrow F$, with $F \subseteq \mathrm{k}(\mathrm{U}(\Gamma'))$ finite. Let us call a Scott-continuous function f of algebraic domains from D to D' *coherent* if, for every $x \in \mathrm{k}(D')$, $f^{-1}(\uparrow x)$ is of the form $\uparrow F$, with $F \subseteq \mathrm{k}(D)$.[7] As indicated in Figure 63, there are non-coherent Scott-continuous functions of coherent algebraic domains. (Intuitively, μ must take a primitive of Γ to an *infinite* disjunction in order to give rise to the function sketched in the figure.)

Suppose Γ and Γ' are finitistic and f is a coherent Scott-continuous function from $\mathrm{U}(\Gamma')$ to $\mathrm{U}(\Gamma)$. Let Ω and Ω' be the corresponding extension algebras over $\mathrm{U}(\Gamma)$ and $\mathrm{U}(\Gamma')$. Since f is coherent, $f^{-1}(V)$ belongs to Ω' for every

[7] Topologically speaking, f is required to preserve compactness of Scott-open sets under inverse image.

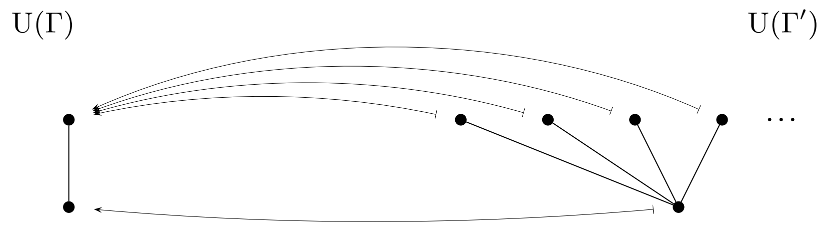

FIGURE 63 Non-coherent Scott-continuous function of coherent algebraic domains

member V of Ω. Hence the function f^* that takes V to $f^{-1}(V)$ is an algebra homomorphism from Ω to Ω'. Recall that $\mathrm{U}(\Gamma) \simeq \mathrm{P}(\Omega) = \{\mathcal{N}(y) \,|\, y \in \mathrm{U}(\Gamma)\}$; cf. (6.17). We claim that $\mathrm{P}(f^*)(\mathcal{N}(x)) = \mathcal{N}(f(x))$, for every $x \in \mathrm{U}(\Gamma')$. But this is just as consequence of definitions: $V \in \mathrm{P}(f^*)(\mathcal{N}(x)) = (f^*)^{-1}(\mathcal{N}(x))$ iff $f^*(V) = f^{-1}(V) \in \mathcal{N}(x)$ iff $x \in f^{-1}(V)$ iff $f(x) \in V$ iff $V \in \mathcal{N}(f(x))$. It follows that f is of the form $\mathrm{U}(\mu)$ for some morphism μ from Γ to Γ'.

In other words, the restriction of U to a functor from the full subcategory of finitistic theories to that of coherent algebraic domains and coherent Scott-continuous functions is full. In addition, this functor is faithful up to equivalence since $\mu \sim \nu$ whenever $\mathrm{U}(\mu) = \mathrm{U}(\nu)$. Hence:

(8.37) Theorem The functor U induces a dual equivalence between the quotient category of finitistic theories modulo equivalence and the category of coherent algebraic domains and coherent Scott-continuous functions.

(8.38) Corollary The functor U induces a dual equivalence between the quotient category of finite theories modulo equivalence and the category of finite ordered sets (and order-preserving functions).

Recall that an observational theory is finitistic if and only if every element of its Lindenbaum algebra is a finite join of join-irreducibles. Consequently, the functor P restricts to a dual equivalence between the full subcategory of those observational algebras whose elements are finite joins of join-irreducibles and the category of coherent algebraic domains and coherent Scott-continuous functions.

(8.39) Remark In domain theory one is primarily interested in *full* subcategories of $\mathbf{Dcpo}$. For instance, one asks for characterizations of the category of coherent algebraic domains and *all* Scott-continuous functions. As seen before, the functor P from $\mathbf{Alg}$ to this category is not full and hence no dual equivalence. While we have restricted the class of Scott-continuous functions to coherent ones in order to get a dual equivalence, one can also extend the class of morphisms of observational algebras from homomorphisms to so-called *join-approximable relations*; see e.g. Abramsky and Jung 1994, Sect. 7.2.

8.3 Applications

An important application of the categorical viewpoint is the construction of (co)products and (co)limits, which is the topic of Chapter 9.

8.3.1 *Horn Translation of Choice System Theories*

According to (5.37), every rooted choice system theory has a flat generic universe. In addition, we know that every flat ordered set is the generic universe of a Horn theory (and even of a full binary exclusion theory). However, not every rooted choice system theory is equivalent to a Horn theory; witness Example (8.34). In the case of finitistic theories, on the other hand, this cannot happen, as shown in Section 8.2.4. In particular:

(8.40) Proposition Every finite rooted choice system theory is equivalent to a Horn theory.

The practical use of this general result is limited since a Horn theory with a flat generic universe of cardinality $n+1$ has at least n primitives. To see this, recall that the canonical universe $\mathrm{C}(\Gamma)$ of a Horn theory Γ over Σ is closed with respect to intersection. Suppose $\mathrm{C}(\Gamma)$ is flat and $|\mathrm{C}(\Gamma)| = n+1$. Then $\mathrm{C}(\Gamma)$ consists of n subsets $X_1, \ldots, X_n$ of Σ that are pairwise incomparable with respect to inclusion, and a subset X of Σ such that $X = X_i \cap X_j$ for all $i \neq j$. Let Y_i be $X_i \setminus X$. Since the sets $Y_1, \ldots, Y_n$ are pairwise disjoint, Σ has at least n elements, for otherwise, by the pigeonhole principle, two of the Y_i's would have an element in common.

So the most parsimonious Horn theory with flat generic universe of cardinality $n+1$ is indeed the full binary exclusion theory $\{p \wedge q \preceq \Lambda \mid p \neq q\}$ over a set of n primitives. Consequently, given a rooted choice system theory Γ over Σ, the "best" we can do to get a Horn theory Γ' equivalent to Γ is to use "brute force": determine $\mathrm{C}(\Gamma)$ and let Γ' be the full exclusion theory over a set Σ' of cardinality $|\mathrm{C}(\Gamma)| - 1$. Equivalence morphisms can then be defined as follows: Suppose $X_1, \ldots, X_n$ are the nonempty members of $\mathrm{C}(\Gamma)$ and Σ' is $\{q_1, \ldots, q_n\}$. Clearly the function μ from Σ' to $\mathrm{T}[\Sigma]$ that takes q_i to the conjunction of all members of X_i is a morphism from Γ' to Γ. Moreover, the function ν that takes each $p \in \Sigma$ to the disjunction of all q_i's with $p \in X_i$ is a morphism from Γ to Γ' that is quasi-inverse to μ.

Exponential growth is a serious obstacle for brute force. Consider for example the rooted choice system theory Γ over $\Sigma = \{a_1, b_1, \ldots a_n, b_n, c\}$ with choice systems $\{a_i, b_i\}$ and $\{c\}$ and statements $c \equiv a_i \vee b_i$ and $a_i \wedge b_i \equiv \Lambda$. ($\{c\}$ is the root system.) The canonical universe of Γ consists of the empty set and all subsets of Σ that contain c and either a_i or b_i, for all i. Hence $|\mathrm{C}(\Gamma)| = 2^n + 1$.

(8.41) Remark Calder (1999) proposes a translation method for systemic networks that works "locally" on the precondition structure of the network and keeps close to the original set of primitives. In certain, well-defined cases, his method yields a Horn translation. In addition, Calder presents a technique for deciding whether a given network allows a Horn translation by his approach. He suggests that the positive cases are essentially the ones with tractable satisfiability problem.

8.3.2 *Minimal Representations*

Let P be a finite ordered set. Recall that P can be represented by a subset system over P, where $x \in P$ corresponds to $\downarrow x$. Of course, P may have a representation by a subset system over a set Σ with lower cardinality than P. For instance, we saw in Section 4.2.1 that if P is bounded-complete and distributive, the set of coprime elements of P can be chosen for Σ (cf. Example (4.27)). A different type of example is given by (5.35), which shows that a full binary tree P of height $k+1$, which has $2^{k+1}-1$ nodes and 2^k leaves, can be embedded in $\wp(\Sigma)$ with $|\Sigma| = 2k$. (Notice that all elements of a tree besides the root are coprime.)

The *representation problem* arising here has the following general form: Given a finite ordered set P, find the least number n such that P can be represented as a subset system over a set of cardinality n. The solution of this problem is of practical importance because representations of ordered sets as subset systems can be easily implemented on a computer via *bit-vector* encoding. However, it can be shown:

(8.42) Theorem The representation problem is NP-complete.

For a proof, see e.g. Caseau et al. 1999, where also an algorithm using graph coloring is given. In terms of theories and translations the representation problem can be rephrased as follows: Given an observational theory Γ over a finite set Σ of primitives, find a set Σ' of primitives with minimal cardinality such that Γ is equivalent to a theory Γ' over Σ'.

We close this section with an application of an elementary result from combinatorics, which is known as Sperner's Lemma (see e.g. Trotter 1995):

(8.43) Proposition If $|\Sigma| = n$ then the width of $\wp(\Sigma)$ is $\binom{n}{\lfloor \frac{n}{2} \rfloor}$.

(The *height* and the *width* of an ordered set P is respectively the cardinality of a maximal chain and a maximal antichain of P.)

If the generic universe of an observational theory Γ is flat or unordered, Sperner's Lemma immediately gives us the minimal number of primitives needed for a theory equivalent to Γ. For example, take Winograd's systemic network for English pronouns presented in Section 1.3. Its (flat) canonical universe has 44 elements plus the bottom element. Since $\binom{7}{3} = 35$ and $\binom{8}{4} = 70$, any observational theory Γ equivalent to this choice system theory has at least 8 primitive predicates, whereas Winograds network uses 20 primitives. However, such a theory Γ will presumably be rather unsatisfying from a linguistic point of view, because its primitives are complex combinations of the original linguistic attributes.

8.4 Theory Extensions

Extension morphisms have been introduced in (8.2) as theory morphisms from $\langle\Sigma, \Gamma\rangle$ to $\langle\Sigma', \Gamma'\rangle$ where the underlying function from Σ to Σ' is an inclusion, and $\Gamma \subseteq \Gamma'$. We then say that $\langle\Sigma', \Gamma'\rangle$ is an *extension* of $\langle\Sigma, \Gamma\rangle$. Notice that the Scott-continuous function of ordered generic universes induced by an extension is in general neither one-to-one nor onto, as the following simple example shows; so extensions need be neither essentially surjective nor conservative.

(8.44) Example Let Γ be the empty theory over $\{a, b\}$ and let ε be its extension to Γ' by a single primitive c and statements $a \preceq c$ and $c \preceq b$. The induced function $\mathrm{C}(\varepsilon)$ from $\mathrm{C}(\Gamma')$ to $\mathrm{C}(\Gamma)$ is depicted by Figure 64; it is neither one-to-one nor onto.

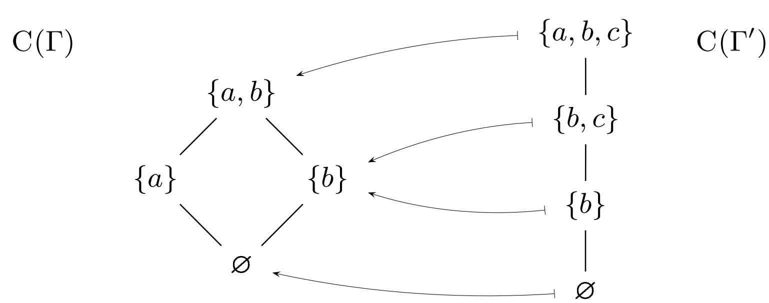

FIGURE 64 Extension ε where $\mathrm{C}(\varepsilon)$ is neither one-to-one nor onto

8.4.1 *Conservative Extensions and Rule Extensions*

Suppose Γ and Γ' are observational theories over Σ such that $\Gamma \subseteq \Gamma'$. Then Γ' is called a *rule extension* of Γ. The corresponding extension morphism from Γ to Γ' is the identity function on Σ. Notice that a rule extension is trivially essentially surjective and recall from (8.32) that the induced function of canonical universes is an inclusion; in particular, $\mathrm{C}(\Gamma') \subseteq \mathrm{C}(\Gamma)$. The definition of the canonical universe of a theory $\langle\Sigma, \Gamma\rangle$ given in Section 5.2 can be seen as an example of this fact: $\mathrm{C}(\Gamma)$ is included in the powerset $\wp(\Sigma)$ of Σ, which is the canonical universe of the empty theory over Σ.

An extension $\langle\Sigma',\Gamma'\rangle$ of $\langle\Sigma,\Gamma\rangle$ is called *conservative* if the corresponding extension morphism is conservative in the sense of (8.6), that is, if $\Gamma' \vdash \alpha$ just in case $\Gamma \vdash \alpha$, for every predicate α over Σ. In other words, conservative extensions do not entail additional statements over the base vocabulary. It follows by (8.32) and (8.35)(i) that $\mathrm{C}(\Gamma) = \{Y \cap \Sigma \mid Y \in \mathrm{C}(\Gamma')\}$.

(8.45) Proposition Let h be a homomorphism of observational algebras.

- ***(i)*** h is one-to-one iff h is presentable by a conservative extension.
- ***(ii)*** h is onto iff h is presentable by a rule extension.

Proof. (i) It suffices to consider the case where h is the inclusion of a subalgebra A into B. Choose a subset $\Sigma \subseteq A$ that generates A and a subset $\Sigma' \subseteq B$ that generates B such that $\Sigma \subseteq \Sigma'$ (e.g. $\Sigma = A$ and $\Sigma' = B$). Then there are congruence relations $\cong$ and $\cong'$ on $\mathrm{T}[\Sigma]$ and $\mathrm{T}[\Sigma']$, respectively, such that $A \simeq \mathrm{T}[\Sigma]/{\cong}$ and $B \simeq \mathrm{T}[\Sigma']/{\cong'}$. Clearly $\cong$ is the congruence kernel of the composite homomorphism $\mathrm{T}[\Sigma] \hookrightarrow \mathrm{T}[\Sigma'] \to \mathrm{T}[\Sigma']/{\cong_{\Gamma'}}$. Moreover, $\cong$ is the restriction of $\cong'$ to $\mathrm{T}[\Sigma]$. Let Γ consist of all statements $\varphi \equiv \psi$ over Σ such that $\varphi \cong \psi$. Define Γ' analogously. Then $\langle\Sigma,\Gamma\rangle$ presents A, $\langle\Sigma',\Gamma'\rangle$ presents B, and the inclusion morphism from $\langle\Sigma,\Gamma\rangle$ to $\langle\Sigma',\Gamma'\rangle$ is a conservative extension that presents h. The reverse direction follows from (8.15)(i). (ii) Suppose A is presented by $\langle\Sigma,\Gamma\rangle$ and there is a homomorphism from $A \simeq \mathrm{T}[\Sigma]/{\cong_\Gamma}$ onto B. By the homomorphism theorems of universal algebra, there is a congruence relation $\cong$ on $\mathrm{T}[\Sigma]$ such that ${\cong_\Gamma} \subseteq {\cong}$ and $B \simeq \mathrm{T}[\Sigma]/{\cong}$. Let Γ' consist of all statements $\varphi \equiv \psi$ such that $\varphi \cong \psi$. Then $\langle\Sigma,\Gamma'\rangle$ presents B. The reverse direction is a special case of (8.15)(ii). ⌟

Suppose h is a homomorphism from $\mathrm{L}(\Gamma)$ to $\mathrm{L}(\Gamma')$. Looking through the proofs of (8.45)(i) and (ii) tells us that one can choose in both cases an extension of Γ to present h. Together with (8.15) we can therefore conclude:

(8.46) Theorem Let μ be a theory morphism from Γ to Γ'.

- ***(i)*** If μ is conservative, Γ' is equivalent to a conservative extension of Γ.
- ***(ii)*** If μ is essentially surjective, Γ' is equivalent to a rule extension of Γ.

8.4.2 *Examples: Booleanization and Rule Completion*

Let Γ be an observational theory over Σ. Recall from Section 5.4.2 that the Booleanization $\overline{\Gamma}$ of Γ is defined as follows: Σ is extended by a disjoint copy $\{-p \mid p \in \Sigma\}$ of Σ, and Γ is extended by all statements

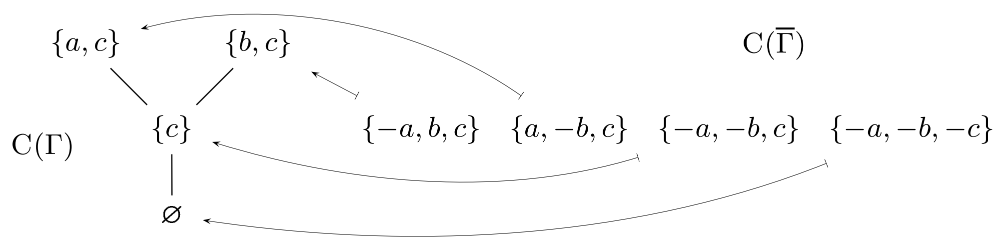

FIGURE 65 Function of canonical universes induced by Booleanization

(8.47) $\qquad p \wedge -p \preceq \Lambda \qquad$ and $\qquad \mathrm{V} \preceq p \vee -p,$

with $p \in \Sigma$. Let ε be the extension morphism from Γ to $\overline{\Gamma}$. By (8.32), $\mathrm{C}(\varepsilon)(Y)$ takes $Y \in \mathrm{C}(\overline{\Gamma})$ to $Y \cap \Sigma$. We have seen in Section 5.4.2 that $\mathrm{C}(\varepsilon)$ is onto and one-to-one; in particular, by (8.35), Booleanization is conservative. To give an illustration, consider the theory Γ over $\{a, b, c\}$ with statements $a \wedge b \preceq \Lambda$ and $a \vee b \preceq c$; cf. Example (5.54). Figure 65 depicts the function $\mathrm{C}(\varepsilon)$ from $\mathrm{C}(\overline{\Gamma})$ to $\mathrm{C}(\Gamma)$.

Since $\overline{\Gamma}$ contains the statements (8.47), it follows that the Lindenbaum algebra of $\overline{\Gamma}$ is *complemented*, i.e. $\mathrm{L}(\overline{\Gamma})$ is a *Boolean lattice*. Since $\mathrm{C}(\varepsilon)$ is onto, $\mathrm{L}(\varepsilon)$ is an embedding of $\mathrm{L}(\Gamma)$ into $\mathrm{L}(\overline{\Gamma})$, by (8.27)(i). It is not difficult to prove the following universal property of $\mathrm{L}(\varepsilon)$ and $\mathrm{L}(\Gamma)$: for every homomorphism h from $\mathrm{L}(\Gamma)$ to a complemented observational algebra A, there exists a unique homomorphism h' from $\mathrm{L}(\overline{\Gamma})$ to A such that $h = h' \circ \mathrm{L}(\varepsilon)$. (This universal characterization of $\mathrm{L}(\overline{\Gamma})$ can be seen as an additional justification of the term 'Booleanization'.)

Besides Booleanization there is a second route to naturally extend an observational theory $\langle \Sigma, \Gamma \rangle$ to a theory whose generic universe is an antichain. The idea is to find an extension of Γ whose generic universe is the set of maximal elements of $\mathrm{U}(\Gamma)$. (Since $\mathrm{U}(\Gamma)$ is directed-complete, it has maximal elements by Zorn's Lemma.) Now observe that if such an extension of Γ exists at all, it can be realized by a rule extension; see (8.35) and (8.46). We then speak of a *rule completion* of Γ. In the case of the above Example, a possible rule completion of Γ is $\Gamma' = \Gamma \cup \{\mathrm{V} \preceq a \vee b\}$. Then $\mathrm{C}(\Gamma')$ consists of $\{a, c\}$ and $\{b, c\}$. Another example of a rule completion is the completion of choice system theories introduced in Section 5.3.2.

Rule completion, however, is not always possible. A simple counterexample is provided by the full binary exclusion theory $\Gamma = \{p \wedge q \preceq \Lambda \mid p \neq q\}$ over an *infinite* set Σ. Recall that $\mathrm{C}(\Gamma)$ is $\{\varnothing\} \cup \{\{p\} \mid p \in \Sigma\}$. According to (5.25), there is no observational theory Γ' over Σ such that $\mathrm{C}(\Gamma') = \{\{p\} \mid p \in \Sigma\}$. Hence there is no rule completion of Γ. (Intuitively, what is needed is the statement $\mathrm{V} \preceq \bigvee \Sigma$, which employs an *infinite disjunction*.) Notice that the notorious example (7.7), which has the same generic universe as Γ, behaves much better in this respect: addition of $\mathrm{V} \preceq a_0 \vee b_0$ leads to rule completion. (Incidentally, since (7.7) is a choice system theory and $\{a_0, b_0\}$ is the only choice system without entry condition, this is just a completion in the sense of Section 5.3.2.)

8.4.3 *Digression: Ultraconservative Morphisms*

In domain theory, *embedding-projection pairs* play an important role in connection with subdomains and so-called bilimits;[8] cf. (9.18) below. Suppose D and E are dcpos, e is a Scott-continuous function from D to E and p is a Scott-continuous function from E to D. Then $\langle e, p\rangle$ is called an *embedding-projection pair* from D to E if $p \circ e = \imath_D$ and $e \circ p \sqsubseteq \imath_E$.

Let us briefly characterize the theory morphisms that give rise to embedding-projection pairs of generic universes. For that purpose it is convenient to define a preorder $\lesssim$ on the morphism sets as follows: if μ and ν are theory morphisms from $\langle\Sigma, \Gamma\rangle$ to $\langle\Sigma', \Gamma'\rangle$ then

$$\mu \lesssim \nu \qquad \text{iff} \qquad \forall p \in \Sigma\ (\Gamma' \vdash \mu(p) \preceq \nu(p)).$$

By definition, $\mu \sim \nu$ iff $\mu \lesssim \nu$ and $\nu \lesssim \mu$. It is not difficult to see that if $\mu \lesssim \nu$ then $\mathrm{L}(\mu) \leqslant \mathrm{L}(\nu)$ and $\mathrm{U}(\mu) \sqsubseteq \mathrm{U}(\nu)$, where $\leqslant$ and $\sqsubseteq$ are pointwise orderings. (For instance, if $\mathrm{L}(\mu) \leqslant \mathrm{L}(\nu)$ then $h \circ \mathrm{L}(\mu) \leqslant h \circ \mathrm{L}(\nu)$ for all $h \in \mathrm{Hom}(\mathrm{L}(\Gamma'), \mathbb{2})$; hence $\mathrm{Hom}_{\mathbb{2}}(\mathrm{L}(\mu)) \leqslant \mathrm{Hom}_{\mathbb{2}}(\mathrm{L}(\nu))$.)

Let us call a theory morphism μ from $\langle\Sigma, \Gamma\rangle$ to $\langle\Sigma', \Gamma'\rangle$ *ultraconservative* if there is a morphism ν from $\langle\Sigma', \Gamma'\rangle$ to $\langle\Sigma, \Gamma\rangle$ such that

$$\nu \circ \mu \sim \imath_\Gamma \qquad \text{and} \qquad \mu \circ \nu \lesssim \imath_{\Gamma'}.$$

Then $\langle \mathrm{U}(\nu), \mathrm{U}(\mu)\rangle$ is an embedding-projection pair from $\mathrm{U}(\Gamma')$ to $\mathrm{U}(\Gamma)$. By (8.35), it follows that μ is conservative and ν is essentially surjective. (So ultraconservatives are conservative, as expected.)

(8.48) Example Let ε be the theory extension from $\langle\Sigma, \Gamma\rangle$ to $\langle\Sigma', \Gamma'\rangle$, where Γ is the empty theory over $\Sigma = \{a, b\}$, Σ' is $\Sigma \cup \{c, d\}$, and Γ' consists of the single statement $a \wedge b \equiv c \vee d$. We claim that ε is conservative but *not* ultraconservative. This is easily seen by considering the induced function $\mathrm{C}(\varepsilon)$ of canonical universes, which is depicted in Figure 66. Obviously there is no function f from $\mathrm{C}(\Gamma)$ to $\mathrm{C}(\Gamma')$ such that $\mathrm{C}(\varepsilon)(f(X)) = X$ and $f(\mathrm{C}(\varepsilon)(Y)) \subseteq Y$ for all $X \in \mathrm{C}(\Gamma)$ and $Y \in \mathrm{C}(\Gamma')$.

[8] See e.g. Abramsky and Jung 1994.

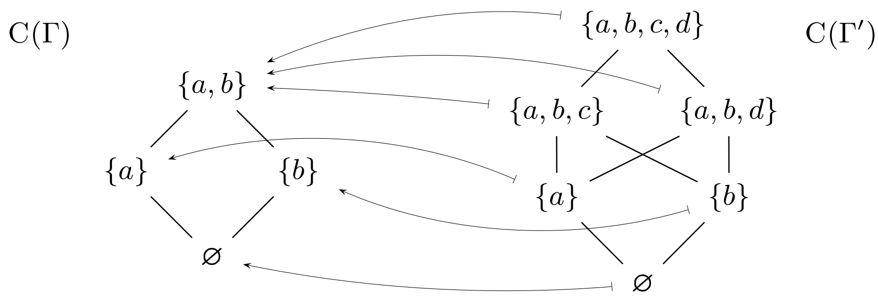

FIGURE 66 A conservative extension that is not ultraconservative

Given an ultraconservative morphism μ from Γ to Γ', it follows by (8.46)(ii) that Γ is equivalent to a rule extension Γ'' of Γ'. Hence there is an ultraconservative morphism μ' from Γ'' to Γ' such that $\mathrm{C}(\Gamma'') \subseteq \mathrm{C}(\Gamma')$. The following example illustrates this fact.

(8.49) Example Let Γ be the theory over $\{a, b\}$ consisting of the sole statement $a \wedge b \preceq \Lambda$. The extension ε of Γ to the theory $\Gamma' = \Gamma \cup \{a \preceq c\}$ over $\{a, b, c\}$ is ultraconservative. For let ν be the morphism from Γ' to Γ with $\nu(a) = \nu(c) = a$ and $\nu(b) = b$. Then $\nu \circ \varepsilon \sim \imath_\Gamma$ and $\varepsilon \circ \nu \lesssim \imath_{\Gamma'}$. The effect of $\mathrm{C}(\nu)$ on elements of $\mathrm{C}(\Gamma)$ can be read off from Figure 67 by reversing all solid arrows. Moreover, one can easily read off a rule extension Γ'' of Γ' that is equivalent to Γ. Since $\mathrm{C}(\Gamma'')$ must consist of $\varnothing$, $\{a, c\}$, and $\{b\}$, one can choose e.g. $\Gamma'' = \Gamma' \cup \{c \preceq a\}$.

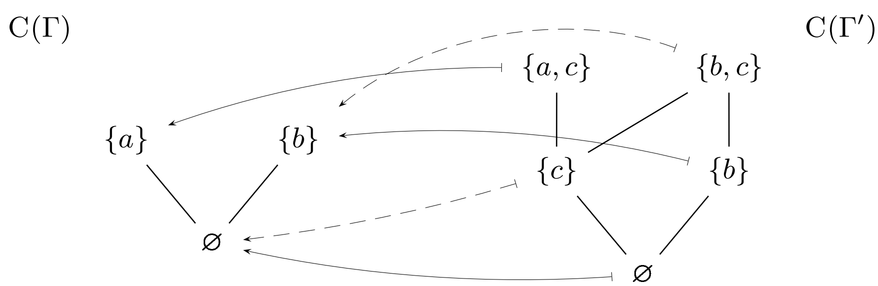

FIGURE 67 An ultraconservative extension

9

Constructions

In order to determine the generic universe of an observational theory it is important – both for theoretical and practical purposes – to understand how the generic universe is affected by an extension or combination of theories. Since we have established observational theories as objects of a category, the natural way for combining theories is to make use of standard categorical constructions like (co)products and (co)limits. The task then is to study how these constructions behave under the functor U from $\mathbf{Th}$ to $\mathbf{Dcpo}$.

In Section 9.1 we first consider the disjoint union of theories, which, categorically speaking, is a coproduct in $\mathbf{Th}$. Its generic universe is the product of the respective generic universes in $\mathbf{Dcpo}$. We have seen in Chapter 8 that it is more useful to characterize theories up to equivalence than up to isomorphism because equivalent theories have isomorphic generic universes. To take this into account, we generalize the standard notion of a colimit in $\mathbf{Th}$ to that of a *quasi-colimit*. We give an explicit construction of quasi-colimits of arbitrary diagrams in $\mathbf{Th}$.

The key result of Section 9.2 is that U takes quasi-colimits in $\mathbf{Th}$ to limits in $\mathbf{Dcpo}$. So the generic universe of a quasi-colimit of a diagram in $\mathbf{Th}$ can be determined (up to isomorphism) by taking the limit of the corresponding diagram of generic universes. In Section 9.3 we focus on *inductive limits* of theories. As an application we show that the generic universes of observational theories are precisely the *profinite* ordered sets, i.e. the projective limits of finite ordered sets. In Section 9.4, we present a simple algorithm for constructing the canonical universe of an observational theory step-by-step.

9.1 Coproducts and Colimits

9.1.1 *Disjoint Union of Theories*

Consider the task of combining two observational theories $\langle \Sigma, \Gamma \rangle$ and $\langle \Sigma', \Gamma' \rangle$. Let us assume that Σ and Σ' and hence Γ and Γ' are disjoint. The most obvious way to combine the two theories into one is to take their *disjoint union*

$$\langle\Sigma,\Gamma\rangle \uplus \langle\Sigma',\Gamma'\rangle \;=\; \langle\Sigma\cup\Sigma',\Gamma\cup\Gamma'\rangle,$$

which is also known as their *direct sum*. Correspondingly, one can define the direct sum of an arbitrary family $\langle\langle\Sigma_i,\Gamma_i\rangle\rangle_{i\in I}$ of theories, given that the Σ_i's are pairwise disjoint:

$$\biguplus_{i\in I}\langle\Sigma_i,\Gamma_i\rangle \;=\; \langle\bigcup_{i\in I}\Sigma_i,\bigcup_{i\in I}\Gamma_i\rangle.$$

It is not difficult to describe the canonical universe of $\biguplus_i \Gamma_i$ in terms of those of the Γ_i's. Let ε_i be the extension morphism from $\langle\Sigma_i,\Gamma_i\rangle$ to the direct sum $\langle\bigcup_i\Sigma_i,\bigcup_i\Gamma_i\rangle$. According to (8.32), the (Scott-continuous) function $\mathrm{C}(\varepsilon_i)$ from $\mathrm{C}(\biguplus_i\Gamma_i)$ to $\mathrm{C}(\Gamma_i)$ takes X to $X\cap\Sigma_i$. Moreover, if $X_i\in\mathrm{C}(\Gamma_i)$ for every i, then $\bigcup_i X_i\in\mathrm{C}(\biguplus_i\Gamma_i)$; hence

$$\mathrm{C}(\biguplus_i\Gamma_i)=\{\bigcup_i X_i \mid X_i\in\mathrm{C}(\Gamma_i)\}.$$

It follows that the canonical universe of $\biguplus_i\Gamma_i$ is order-isomorphic to the Cartesian product $\prod_i\mathrm{C}(\Gamma_i)$ ordered by *coordinatewise inclusion*:

(9.1) $$\mathrm{C}(\biguplus_i\Gamma_i) \;\simeq\; \prod_i\mathrm{C}(\Gamma_i).$$

In other words, $\mathrm{C}(\biguplus_i\Gamma_i)$ together with the family $\langle\mathrm{C}(\varepsilon_i)\rangle_i$ of *projections* is a *(direct) product* of the family $\langle\mathrm{C}(\Gamma_i)\rangle_i$ in the category $\mathbf{Dcpo}$. (One easily checks that the Cartesian product is a product in the category $\mathbf{Dcpo}$; see Section 9.2.1 below.)

(9.2) Example *(Direct sum of choice systems)* Suppose $\langle\Sigma,\Gamma\rangle$ is the direct sum of choice systems, i.e. of full binary exclusion theories. That is, $\langle\Sigma,\Gamma\rangle=\biguplus_i\langle\Sigma_i,\Gamma_i\rangle$, where $\Gamma_i=\{p\wedge q\preceq\Lambda \mid p,q\in\Sigma_i, p\neq q\}$. Then $\mathrm{C}(\Gamma_i)$ is flat and $|\mathrm{C}(\Gamma_i)|=|\Sigma_i|+1$. Hence $|\mathrm{C}(\Gamma)|=\prod_i(|\Sigma_i|+1)$.

(9.3) Example *(Extension by primitives)* Given an observational theory Γ over Σ and a set Σ' disjoint to Σ, we call $\langle\Sigma\cup\Sigma',\Gamma\rangle$ the *extension of* $\langle\Sigma,\Gamma\rangle$ *by primitives* Σ'. Since $\langle\Sigma\cup\Sigma',\Gamma\rangle$ is identical to $\langle\Sigma,\Gamma\rangle\uplus\langle\Sigma',\varnothing\rangle$, it follows by (9.1) that $\mathrm{C}(\langle\Sigma\cup\Sigma',\Gamma\rangle)\simeq\mathrm{C}(\langle\Sigma,\Gamma\rangle)\times\wp(\Sigma')$. In particular, $|\mathrm{C}(\langle\Sigma\cup\Sigma',\Gamma\rangle)|=|\mathrm{C}(\langle\Sigma,\Gamma\rangle)|\cdot 2^{|\Sigma'|}$.

(9.4) Example *(Free primitives)* Let $\langle\Sigma,\Gamma\rangle$ be an observational theory. Call a primitive $p\in\Sigma$ *free with respect to* Γ if p does not occur in any statement of Γ. Let $\sigma(\Gamma)$ be the set of primitives occurring in at least one of the statements of Γ, that is, $\Sigma'=\Sigma\setminus\sigma(\Gamma)$ is the set of free primitives. Then $\langle\Sigma,\Gamma\rangle$ is the direct sum of $\langle\sigma(\Gamma),\Gamma\rangle$ and $\langle\Sigma',\varnothing\rangle$ (and $\langle\Sigma,\Gamma\rangle$ is the extension of $\langle\sigma(\Gamma),\Gamma\rangle$ by primitives Σ'). Hence $\mathrm{C}(\langle\Sigma,\Gamma\rangle)\simeq\mathrm{C}(\langle\sigma(\Gamma),\Gamma\rangle)\times\wp(\Sigma')$.

It is not difficult to see that the disjoint union of a family of theories is a coproduct of that family in the category $\mathbf{Th}$. By definition, a *coproduct* of a family $\langle \Gamma_i \rangle_{i\,\in\, I}$ of theories is a pair $\langle \Gamma, \langle \mu_i \rangle_{i\,\in\, I} \rangle$ consisting of a theory Γ and a family of morphisms μ_i from Γ_i to Γ, also called *injections*, such that for every such pair $\langle \Gamma', \langle \mu'_i \rangle_{i\,\in\, I} \rangle$ there is a unique morphism μ from Γ to Γ' with $\mu'_i = \mu \circ \mu_i$ for every $i \in I$.

Since the Σ_i's need not be disjoint in general, we have to be a bit more careful than before in defining the disjoint union of a family $\langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in I}$ of theories. Let ι_i be the canonical injection of Σ_i into the disjoint union $\biguplus_i \Sigma_i$ of sets. Then ι_i is a morphism from $\langle \Sigma_i, \Gamma_i \rangle$ to the theory $\biguplus \Gamma = \langle \biguplus_i \Sigma_i, \bigcup_i \hat{\iota}_i(\Gamma_i) \rangle$.

(9.5) Proposition If $\langle \Gamma_i \rangle_i$ is a family of theories then $\langle \biguplus_i \Gamma_i, \langle \iota_i \rangle_i \rangle$ is a coproduct of that family in $\mathbf{Th}$.

Proof. Let $\langle \Sigma, \Gamma \rangle$ be a theory and suppose μ_i is a morphism from $\langle \Sigma_i, \Gamma_i \rangle$ to $\langle \Sigma, \Gamma \rangle$, for every i. We need to show that there is a unique morphism μ from $\biguplus_i \Gamma_i$ to Γ such that $\mu_i = \mu \circ \iota_i$. Since every element of $\biguplus_i \Sigma_i$ is of the form $\iota_i(p)$, where i and $p \in \Sigma_i$ are uniquely determined, the only function μ from $\biguplus_i \Sigma_i$ to $\mathrm{T}[\Sigma]$ satisfying the desired condition takes $\iota_i(p)$ to $\mu_i(p)$. It remains to check that μ is a morphism, i.e. that $\Gamma \vdash \hat{\mu}(\hat{\iota}_i(\alpha))$ for all $\alpha \in \Gamma_i$. But $\Gamma \vdash \hat{\mu}_i(\alpha)$, by assumption, and $\hat{\mu} \circ \hat{\iota}_i = \hat{\mu}_i$, by definition of μ. ⌟

The coproduct of theories given by disjoint union will be henceforth referred to as the *canonical coproduct*. As the reader will verify without problems, $\langle \mathrm{C}(\biguplus_i \Gamma_i), \langle \mathrm{C}(\iota_i) \rangle_i \rangle$ is a (direct) product of $\langle \mathrm{C}(\Gamma_i) \rangle_i$ in $\mathbf{Dcpo}$. Indeed, this is a special case of the general result (9.13) below.

9.1.2 *Colimits and Quasi-Colimits*

Colimits are to arbitrary directed graphs what coproducts are to index sets. Let G be a directed graph with vertex set I and edge set E, together with two functions s and t from E to I, where $s(e)$ and $t(e)$ are respectively source and target of the edge e. A *diagram* $\mathcal{D}$ of shape G in the category $\mathbf{Th}$ is a pair consisting of a family $\langle \Gamma_i \rangle_{i \in I}$ of theories and a family $\langle \mu_e \rangle_{e \in E}$ of morphisms such that μ_e is a morphism from $\Gamma_{s(e)}$ to $\Gamma_{t(e)}$. A *cocone* of $\mathcal{D}$ consists of a theory Γ and a family $\langle \mu_i \rangle_{i \in I}$, where μ_i is a morphism from Γ_i to Γ and $\mu_{s(e)} = \mu_{t(e)} \circ \mu_e$ for all $e \in E$. A cocone is said to be a *colimit* if for every cocone $\langle \Gamma', \langle \mu'_i \rangle_i \rangle$ of $\mathcal{D}$ there is exactly one morphism μ from Γ to Γ' such that $\mu'_i = \mu \circ \mu_i$; see Figure 68. In particular, coproducts are colimits of diagrams over graphs without edges. *Cone* and *limit* of a diagram are defined dually by reversing the direction of all arrows of Figure 68 that point downwards.

Since we are more interested in characterizing theories *up to equivalence* than up to isomorphism, it will prove useful to consider "quasi-cocones" and

$$
\begin{tikzcd}
\Gamma_{s(e)} \arrow[rr, "\mu_e"] \arrow[dr, "\mu_{s(e)}"'] \arrow[ddr, bend right, "\mu'_{s(e)}"'] & & \Gamma_{t(e)} \arrow[dl, "\mu_{t(e)}"] \arrow[ddl, bend left, "\mu'_{t(e)}"] \\
& \Gamma \arrow[d, dashed, "\mu"] & \\
& \Gamma' &
\end{tikzcd}
$$

FIGURE 68 Cocone and colimit

"quasi-colimits" of diagrams in $\mathbf{Th}$ instead of cocones and colimits. By a *quasi-cocone* $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ of a diagram $\langle \langle \Gamma_i \rangle_{i \in I}, \langle \mu_e \rangle_{e \in E} \rangle$ of theories we mean a theory Γ and a family of morphisms μ_i from Γ_i to Γ such that $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$ (instead of $\mu_{s(e)} = \mu_{t(e)} \circ \mu_e$, as in the definition of cocones). A quasi-cocone $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ will be called a *quasi-colimit* if for every quasi-cocone $\langle \Gamma', \langle \mu'_i \rangle_i \rangle$ of the diagram there is a morphism μ from Γ to Γ', which is unique up to equivalence, such that $\mu'_i \sim \mu \circ \mu_i$.

In spite of this somewhat involved definition, quasi-colimits of arbitrary diagrams in $\mathbf{Th}$ can be straightforwardly constructed by disjoint union and rule extension:[1]

(9.6) Theorem Let $\langle \langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in I}, \langle \mu_e \rangle_{e \in E} \rangle$ be a diagram of theories. Then the theory

$$\bigcup_i \hat{\iota}_i(\Gamma_i) \;\cup\; \{\iota_{s(e)}(p) \equiv \hat{\iota}_{t(e)}(\mu_e(p)) \mid e \in E \,\wedge\, p \in \Sigma_{s(e)}\}$$

over $\biguplus_i \Sigma_i$, together with the family $\langle \iota_i \rangle_i$ of injections, is a quasi-colimit of the diagram.

Proof. Let $\langle \Sigma, \Gamma \rangle$ be the theory defined in the theorem. Clearly $\langle \Gamma, \langle \iota_i \rangle_i \rangle$ is a quasi-cocone since, by definition, $\iota_{s(e)} \sim \iota_{t(e)} \circ \mu_e$; see (8.5). Suppose $\langle \langle \Sigma', \Gamma' \rangle, \langle \mu_i \rangle_i \rangle$ is a quasi-cocone of the given diagram, i.e. $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$ for all $e \in E$. We claim that there is a morphism μ from Γ to Γ' with $\mu_i = \mu \circ \iota_i$. As in the proof of (9.5), we define μ to take $\iota_i(p)$ to $\mu_i(p)$, where $p \in \Sigma_i$. Now check that μ is a morphism: Γ' entails $\hat{\mu}_i(\alpha)$ and thus $\hat{\mu}(\hat{\iota}_i(\alpha))$, for all $\alpha \in \Gamma_i$; see again the proof of (9.5). This leaves us to check that Γ' entails $\hat{\mu}(\iota_{s(e)}(p)) \equiv \hat{\mu}(\hat{\iota}_{t(e)}(\mu_e(p)))$, i.e. $\mu_{s(e)}(p) \equiv \hat{\mu}_{t(e)}(\mu_e(p))$, for every $p \in \Sigma_{s(e)}$. But this is just the assumption that $\mu_{s(e)} \sim \mu_{t(e)} \circ \mu_e$; so μ is a

[1] Readers with background in category theory will recognize the similarity to the canonical construction of colimits by coproducts and coequalizers.

morphism from Γ to Γ'. It remains to show that if ν is another morphism from Γ to Γ', with $\mu_i \sim \nu \circ \iota_i$, then $\mu \sim \nu$. Since $\mu \circ \iota_i = \mu_i \sim \nu \circ \iota_i$, it follows that $\Gamma' \vdash \hat{\mu}(\iota_i(p)) \equiv \hat{\nu}(\iota_i(p))$ for all $p \in \Sigma_i$. Hence $\mu \sim \nu$, by definition. $\lrcorner$

It is a direct consequence of definitions that quasi-colimits in $\mathbf{Th}$ correspond to colimits in the quotient category $\mathbf{Th}/{\sim}$. More precisely, the quotient functor $Q_\sim$ from $\mathbf{Th}$ to $\mathbf{Th}/{\sim}$ takes quasi-colimits to colimits. Since, according to (8.19), the functor L from $\mathbf{Th}$ to $\mathbf{Alg}$ factors by $Q_\sim$ and an equivalence of categories, we have:

(9.7) Theorem The functor L takes quasi-colimits in $\mathbf{Th}$ to colimits in $\mathbf{Alg}$.

Recall from (8.21) that the restriction of L to a functor from $\mathbf{Th_p}$ to $\mathbf{Alg}$ has a right adjoint and hence preserves colimits.[2] Colimits need not exist for all diagrams in $\mathbf{Th_p}$ (because $\mathbf{Th_p}$ has not enough morphisms). If, however, a diagram $\mathcal{D}$ in $\mathbf{Th_p}$ has a colimit $\langle \Gamma, \langle \mu_i \rangle_i \rangle$ in $\mathbf{Th_p}$ then this cocone is also a quasi-colimit of $\mathcal{D}$ in $\mathbf{Th}$. For $\langle \mathrm{L}(\Gamma), \langle \mathrm{L}(\mu_i) \rangle_i \rangle$ is a colimit of the diagram $\mathrm{L}(\mathcal{D})$ in $\mathbf{Alg}$. In addition, by (9.6), $\mathcal{D}$ has a quasi-colimit $\langle \Gamma', \langle \mu'_i \rangle_i \rangle$ in $\mathbf{Th}$, which, by (9.7), is taken to another colimit of $\mathrm{L}(\mathcal{D})$. Hence $\mathrm{L}(\Gamma) \simeq \mathrm{L}(\Gamma')$ and thus $\Gamma \sim \Gamma'$, by (8.14). We can conclude:

(9.8) Theorem Colimits in $\mathbf{Th_p}$ are quasi-colimits in $\mathbf{Th}$.

The canonical coproduct of a family of theories is a quasi-coproduct of that family. So L preserves coproducts and therefore takes coproducts of theories to coproducts of algebras (also known as *free distributive products*).

(9.9) Remark It follows by (9.6) that the category $\mathbf{Th}/{\sim}$ is *cocomplete*, i.e. has colimits of arbitrary diagrams. By (8.19), the same is true of $\mathbf{Alg}$. Of course, this is well-known since $\mathbf{Alg}$ is an equationally presentable algebraic category and thus complete and cocomplete.

9.2 The Universe in the Limit

The goal of this section is to show that the functor U ($\simeq \mathrm{C} \simeq \mathrm{Mod}_{\mathbb{2}}$) takes quasi-colimits in $\mathbf{Th}$ to limits in $\mathbf{Dcpo}$. In other words, the ordered generic universe of the quasi-colimit of a diagram of theories is the limit of the corresponding diagram of generic universes. Because of (8.29) and (9.7), it suffices to show that $\mathrm{Hom}_{\mathbb{2}}$ ($\simeq \mathrm{P}$) takes colimits in $\mathbf{Alg}$ to limits in $\mathbf{Dcpo}$.

[2] See e.g. Mac Lane 1971, Sect. V.5.

9.2.1 *Limits in* $\mathbf{Dcpo}$

Limits in $\mathbf{Dcpo}$ can be constructed as canonical limits in $\mathbf{Set}$, i.e. as subsets of Cartesian products, with elements ordered coordinatewise. Concretely, if $\langle\langle D_i\rangle_{i\in I}, \langle f_e\rangle_{e\in E}\rangle$ is a diagram in $\mathbf{Dcpo}$ over a graph with vertices I and edges E then the coordinatewise ordered set

$$\{\langle x_i\rangle_i \in \textstyle\prod_{i\in I} D_i \mid \forall e\in E\,(x_{t(e)} = f_e(x_{s(e)}))\}$$

together with the canonical projections is a limit cone of the diagram. The straightforward proof is left to the reader.[3]

We now give a criterion in which cases a limit of a diagram of dcpos (under the forgetful functor) in $\mathbf{Set}$ is actually a limit in $\mathbf{Dcpo}$. To this end, we make use of the fact that the forgetful functor from $\mathbf{Dcpo}$ to $\mathbf{Set}$ takes canonical limits to (canonical) limits.

(9.10) Lemma Suppose $\langle D, \langle f_i\rangle_{i\in I}\rangle$ is a cone of a diagram in $\mathbf{Dcpo}$ which is taken to a limit cone of the diagram in $\mathbf{Set}$ by the forgetful functor. Then the cone is a limit of the diagram in $\mathbf{Dcpo}$ iff, for every two $x, y\in D$, it holds that $x \leqslant y$ whenever $f_i(x) \leqslant f_i(y)$ for all $i\in I$.

Proof. Let $\langle D', \langle p_i\rangle_i\rangle$ be the canonical limit cone in $\mathbf{Dcpo}$ of the diagram in question. Then there is a unique Scott-continuous function f from D to D' such that $f_i = p_i \circ f$. By assumption, f is one-to-one because $\langle D', \langle p_i\rangle_i\rangle$, under the forgetful functor, is a limit cone of the diagram in $\mathbf{Set}$. Since f is Scott-continuous and one-to-one, it is enough to check that f is an order embedding to make sure that f is an isomorphism in $\mathbf{Dcpo}$. Suppose $f(x) \leqslant f(y)$, for $x, y\in D$. Then $f_i(x) = p_i(f(x)) \leqslant p_i(f(y)) = f_i(y)$ for every i, since f and p_i are order preserving. Hence $x \leqslant y$, by assumption. ⌟

9.2.2 *Colimits under* $\mathrm{Hom}_{\mathbf{2}}$

Suppose $\langle A, \langle h_i\rangle_i\rangle$ is a colimit of a diagram $\langle\langle A_i\rangle_i, \langle h_e\rangle_e\rangle$ in $\mathbf{Alg}$ over some graph $\langle I, E\rangle$. This means, by definition, that for every observational algebra B, the set $\mathrm{Hom}(A, B)$ is isomorphic to

(9.11) $$\{\langle g_i\rangle_i \in \textstyle\prod_i \mathrm{Hom}(A_i, B) \mid \forall e\in E\,(g_{s(e)} = g_{t(e)} \circ h_e)\},$$

where a homomorphism h from A to B is taken to the tuple $\langle h \circ h_i\rangle_i$. Observe that (9.11) is the canonical limit of the diagram $\langle\langle \mathrm{Hom}_B(A_i)\rangle_i, \langle \mathrm{Hom}_B(h_e)\rangle_e\rangle$ in $\mathbf{Set}$; therefore $\langle \mathrm{Hom}_B(A), \langle \mathrm{Hom}_B(h_i)\rangle_i\rangle$ is another limit of this diagram

[3] See e.g. Abramsky and Jung 1994, p. 45, Theorem 3.3.1.

in $\mathbf{Set}$. We have thus proved the general fact that contravariant Hom-functors transform colimits into limits in $\mathbf{Set}$.

So $\mathrm{Hom}_{\mathbb{2}}$ takes colimits in $\mathbf{Alg}$ to limits in $\mathbf{Set}$. In order to show that these limits are actually limits in $\mathbf{Dcpo}$, we employ Lemma (9.10).

(9.12) Theorem The functor $\mathrm{Hom}_{\mathbb{2}}$ takes colimits in $\mathbf{Alg}$ to limits in $\mathbf{Dcpo}$.

Proof. Let $\langle A, \langle h_i \rangle_i \rangle$ be a colimit of a diagram in $\mathbf{Alg}$. We know that $\mathrm{Hom}_{\mathbb{2}}$ takes colimits in $\mathbf{Alg}$ to limits in $\mathbf{Set}$. Consider two members v and w of $\mathrm{Hom}_{\mathbb{2}}(A)$ with $\mathrm{Hom}_{\mathbb{2}}(h_i)(v) \leqslant \mathrm{Hom}_{\mathbb{2}}(h_i)(w)$ for every i. Then $v(h_i(a)) \leqslant w(h_i(a))$ for all i and $a \in A_i$. Since A is generated by $\bigcup\{h_i(A_i) \mid i \in I\}$, it follows inductively that $v(a) \leqslant w(a)$ for every $a \in A$; so $v \leqslant w$. Now apply (9.10). ⌟

Taken together with (9.7), we have proved the promised result:

(9.13) Theorem The functor U takes quasi-colimits in $\mathbf{Th}$ to limits in $\mathbf{Dcpo}$.

9.3 Inductive and Projective Limits

In this section we consider diagrams of theories where the underlying graph can be identified with a directed ordering on the vertex set. Let $\langle I, \leqslant \rangle$ be an ordered set which is directed, that is, for every two $i, j \in I$ there is a $k \in I$ such that $i \leqslant k$ and $j \leqslant k$. Suppose for each $i \in I$ there is an observational theory Γ_i and for all $i, j \in I$ with $i \leqslant j$ there is a morphism μ_{ij} from Γ_i to Γ_j such that $\mu_{ii} \sim \imath_{\Gamma_i}$ and $\mu_{jk} \circ \mu_{ij} \sim \mu_{ik}$ for all $i \leqslant j \leqslant k$. If this is the case, we speak of an *inductive* (or *directed*) *quasi-system* of observational theories over I. A quasi-colimit of such a diagram is also called an *inductive* (or *direct*) *quasi-limit*.

9.3.1 *Inductive Limits of Extensions*

We know by (9.6) that every inductive quasi-system in $\mathbf{Th}$ has an inductive quasi-limit. Here we focus on inductive systems of *extensions*, whose inductive quasi-limits turn out to allow a simple construction via set union. Since extensions are primitive morphisms, we can construct inductive limits in the category $\mathbf{Th_p}$ and then apply (9.8).

Let I be a directed ordered set and suppose $\langle\langle \Sigma_i, \Gamma_i \rangle\rangle_{i \in I}$ is an inductive system of theory extensions over I with extension morphisms ε_{ij}, for $i \leqslant j$. Then the theory $\bigcup_i \Gamma_i$ over $\bigcup_i \Sigma_i$ together with the extensions ε_i from Γ_i to $\bigcup_i \Gamma_i$ forms a quasi-cocone of the given inductive system.

(9.14) Proposition If $\langle \Gamma_i \rangle_i$ is an inductive system of theory extensions, then $\langle \bigcup_i \Gamma_i, \langle \varepsilon_i \rangle_i \rangle$ is an inductive quasi-limit of that system.

Proof. According to (9.8), it suffices to show that $\langle \bigcup_i \Gamma_i, \langle \varepsilon_i \rangle_i \rangle$ is an inductive limit of $\langle \Gamma_i \rangle_i$ in $\mathbf{Th_p}$. Suppose there are primitive morphisms μ_i from $\langle \Sigma_i, \Gamma_i \rangle$ to a theory $\langle \Sigma, \Gamma \rangle$ such that $\mu_i = \mu_j \circ \varepsilon_{ij}$ for $i \leqslant j$. Let μ be the function from $\bigcup_i \Sigma_i$ to Σ that takes $p \in \Sigma_i$ to $\mu_i(p)$; μ is well defined because if $p \in \Sigma_i \cap \Sigma_j$, there is a k with $i, j \leqslant k$ and thus $p \in \Sigma_k$; so $\mu_i(p) = \mu_k(p) = \mu_j(p)$. Notice that μ is the only function satisfying $\mu_i = \mu \circ \varepsilon_i$ for all i. It remains to check that μ is a morphism, i.e. that $\Gamma \vdash \hat{\mu}(\alpha)$ for every $\alpha \in \bigcup_i \Gamma_i$. But if $\alpha \in \Gamma_i$ then $\hat{\mu}(\alpha) = \hat{\mu}_i(\alpha)$, and μ_i is a morphism from Γ_i to Γ. ⌟

This result has the following straightforward but useful application. Suppose $\langle \Sigma, \Gamma \rangle$ is an observational theory. For each subset S of Σ, let $\Gamma|_S$ be the set of all statements of Γ whose primitives belong to S. Let $\mathcal{F}$ be the directed set of finite subsets of Σ, ordered by set inclusion. Then the family $\langle \langle F, \Gamma|_F \rangle \rangle_{F \in \mathcal{F}}$ together with the extensions from $\Gamma|_F$ to $\Gamma|_{F'}$, whenever $F \subseteq F'$, is an inductive system of theories, whose inductive limit is $\langle \Sigma, \Gamma \rangle$, by (9.14). Consequently, since every theory over a finite set of primitives is equivalent to a finite theory:

(9.15) Proposition Every observational theory is an inductive quasi-limit of finite observational theories.

It is of course not necessary to take all finite subsets of Σ. Clearly any directed system $\mathcal{F}$ such that $\bigcup \mathcal{F} = \Sigma$ will do. For example, suppose Σ is countable, i.e. $\Sigma = \{p_0, p_1, p_2, \ldots\}$. Let Σ_i be $\{p_0, p_1, \ldots, p_i\}$ and Γ_i be $\Gamma|_{\Sigma_i}$. Then $\langle \Sigma, \Gamma \rangle$ is the inductive limit of the inductive system $\langle \langle \Sigma_i, \Gamma_i \rangle \rangle_{i \in \omega}$, with extensions from Γ_i to Γ_j for all $i \leqslant j$. Hence:

(9.16) Corollary Every observational theory over a countable set of primitives is an inductive quasi-limit of a sequence of finite observational theories.

9.3.2 *Profiniteness*

An ordered set is said to be *profinite* if it is the projective limit of a projective system of *finite* ordered sets. Since, by (9.13), U takes inductive quasi-limits in $\mathbf{Th}$ to projective limits in $\mathbf{Dcpo}$, it follows by (9.15) that the generic universe of an observational theory is profinite. Moreover, being profinite characterizes the generic universes of observational theories:

(9.17) Theorem The ordered generic universe of an observational theory is profinite, and every profinite ordered set arises that way.

Proof. It remains to check the second part. Let $\langle \langle P_i \rangle_i, \langle f_{ij} \rangle_{i,j} \rangle$ be a projective system of finite ordered sets with projective limit P. Because of (8.38), there is an inductive quasi-system $\langle\langle\Gamma_i\rangle_i, \langle\mu_{ij}\rangle_{i,j}\rangle$ of finite observational theories such that $P_i = \mathrm{U}(\Gamma_i)$ and $f_{ij} = \mathrm{U}(\mu_{ij})$. Now apply (9.13). $\lrcorner$

In combination with (6.12), we get the well-known result of Speed (1972a) that an ordered set P is isomorphic to the prime spectrum of a distributive lattice with zero and unit if and only if P is profinite.[4] Speed's proof uses the fact that every observational algebra is an inductive limit of the inductive system of its finite subalgebras. Consequently, by (8.15), every observational theory Γ is equivalent to an inductive limit of an inductive system consisting of *conservative* extensions of finite theories.

Recall from Example (8.34) that non-equivalent theories may have order-isomorphic generic universes. Therefore, though each projective system of finite ordered sets determines an inductive system of theories uniquely up to equivalence, it can happen that two projective systems of finite ordered sets with isomorphic projective limits correspond to two inductive quasi-systems of theories whose inductive quasi-limits are *not* equivalent.

(9.18) Remark *(Bilimits)* If an observational theory Γ is an inductive quasi-limit of a sequence $\langle\langle\Gamma_i\rangle_i, \langle\varepsilon_{ij}\rangle_{i\leqslant j}\rangle$ of *ultraconservative* extensions of finite theories, then $\mathrm{U}(\Gamma)$ is a countably based bifinite domain (and every such domain arises that way). Moreover, Γ is a projective quasi-limit of a sequence $\langle\langle\Gamma_i\rangle_i, \langle\nu_{ij}\rangle_{i\leqslant j}\rangle$, where the ν_{ij}'s are essentially surjective; cf. Section 8.4.3. In particular, Γ is a projective quasi-limit of a sequence of rule extensions of finite theories.

9.4 Building the Universe Step-by-Step

The results of Section 9.3 allow a straightforward algorithmic formulation which gives us a simple method for constructing the generic universe of a finite (or countably infinite) observational theory step-by-step by successive extensions. Extensions, on the other hand, can be broken down as follows.

9.4.1 *Extensions Decomposed*

An extension of theories from $\langle\Sigma,\Gamma\rangle$ to $\langle\Sigma',\Gamma'\rangle$ can be decomposed into an extension of primitives from $\langle\Sigma,\Gamma\rangle$ to $\langle\Sigma',\Gamma\rangle$ followed by a rule extension from $\langle\Sigma',\Gamma\rangle$ to $\langle\Sigma',\Gamma'\rangle$. Recall from (9.3) that

$$\begin{aligned} \mathrm{C}(\langle\Sigma',\Gamma\rangle) \;&=\; \{X \cup Y \mid X \in \mathrm{C}(\langle\Sigma,\Gamma\rangle) \,\wedge\, Y \subseteq \Sigma \setminus \Sigma'\} \\ &\simeq\; \mathrm{C}(\langle\Sigma,\Gamma\rangle) \times \wp(\Sigma \setminus \Sigma'). \end{aligned}$$

[4] See also Speed 1972b and Johnstone 1982, Chap. VI, Sect. 3.

Moreover, $\mathrm{C}(\langle\Sigma', \Gamma'\rangle)$ consists of all members of $\mathrm{C}(\langle\Sigma', \Gamma\rangle)$ that are consistenly closed with respect to $\Gamma' \setminus \Gamma$. So we can construct $\mathrm{C}(\langle\Sigma', \Gamma'\rangle)$ from $\mathrm{C}(\langle\Sigma, \Gamma\rangle)$ by taking first the product of $\mathrm{C}(\langle\Sigma, \Gamma\rangle)$ and $\wp(\Sigma \setminus \Sigma')$ and then deleting those elements that are not consistently closed with respect to $\Gamma' \setminus \Gamma$.

(9.19) Example Let Γ' be the theory $\{a \wedge b \preceq \Lambda,\ a \preceq c\}$ over $\Sigma' = \{a, b, c\}$. Viewed as an extension of the theory $\{a \wedge b \preceq \Lambda\}$ over $\{a, b\}$, the construction of Γ' by product and deletion is as depicted by the upper row of Figure 69. In case Γ' is viewed as an extension of the theory $\{a \preceq c\}$ over $\{a, c\}$, the construction runs as shown in the lower row of the figure. (The shaded elements are subject to deletion because they are not consistently closed with respect to $\Gamma' \setminus \Gamma$.)

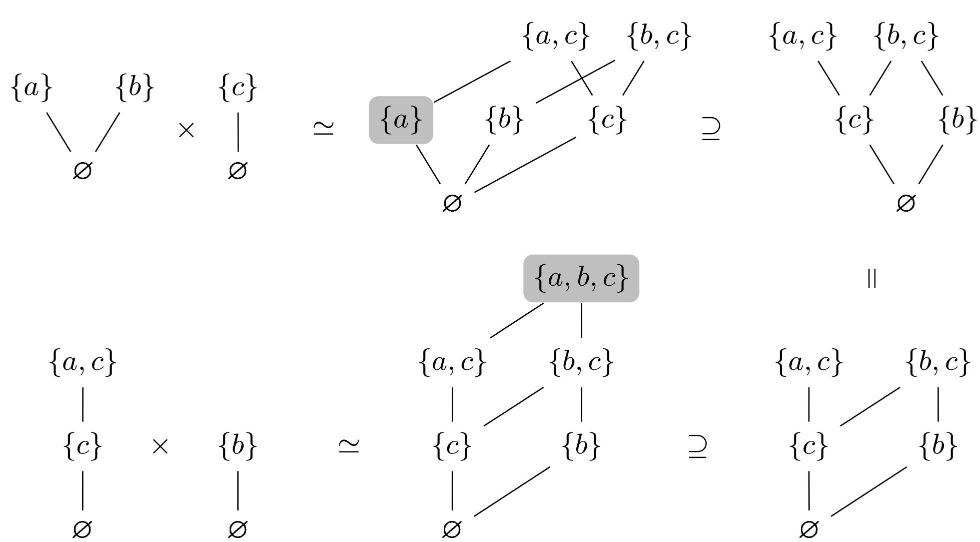

FIGURE 69 Universe of extension as sub-universe of product

9.4.2 *A Simple Construction Algorithm*

Suppose Γ is an observational theory over a finite or countably infinite set Σ of primitives. (Without loss of generality, we can assume that Γ is countable and that $\Gamma|_F$ is finite for every finite subset F of Σ.) In the finite case, there is a trivial way to construct $\mathrm{C}(\Gamma)$ along its very definition: just take all consistently Γ-closed subsets of Σ. This method is of course highly impracticable since it means to check each member of the power set of Σ against every statement of Γ; in addition, it is of no use for approximating the infinite case.

We can do better by using the fact that $\langle\Sigma, \Gamma\rangle$ is the inductive (quasi-)limit of a sequence of (finite) theories and extensions. Suppose Σ is finite.

Let $\Sigma_0, \Sigma_1, \ldots, \Sigma_n$ be a strictly increasing sequence of sets, with $\Sigma_0 = \varnothing$ and $\Sigma_n = \Sigma$, and let Γ_i be $\Gamma|_{\Sigma_i}$. Then $\mathrm{C}(\Gamma)$ can be constructed from $\mathrm{C}(\Gamma_0)$ by applying the two-step method of Section 9.4.1 iteratively to the extension from Γ_{i-1} to Γ_i, for every i. Notice that the canonical universe of $\Gamma_0 = \Gamma|_\varnothing$ is either $\varnothing$ or $\{\varnothing\}$ – concretely, $\mathrm{C}(\Gamma_0)$ is empty iff $\Gamma|_\varnothing$ entails $\mathrm{V} \preceq \Lambda$, in which case $\mathrm{C}(\Gamma)$ is empty too. As we saw in Section 9.4.1, $\mathrm{C}(\Gamma_i)$ can be constructed from $\mathrm{C}(\Gamma_{i-1})$ by taking all sets of the form $X \cup Y$, with $X \in \mathrm{C}(\Gamma_{i-1})$ and $Y \subseteq \Sigma_i \setminus \Sigma_{i-1}$, such that $X \cup Y$ is consistently closed with respect to $\Gamma_i \setminus \Gamma_{i-1}$.

Figure 70 presents an algorithmic formulation of the iteration scheme developed so far. (The variables F and Σ' play the role of $\Sigma_i \setminus \Sigma_{i-1}$ and Σ_i, respectively.) Notice that with $F = \Sigma$ we get the single step construction considered in the first place. So in order to circumvent the combinatorial problems mentioned above, it is reasonable to prefer singletons for F (which leads to up to $|\Sigma|$ iterations of the **while**-loop). However, one cannot be sure that an iterative construction is superior to the single step approach. For if every member of Σ occurs in all statements of Γ, the given algorithm first determines the powerset of Σ, irrespective of the choices of F, before it selects all subsets of Σ belonging to $\mathrm{C}(\Gamma)$.

Calculating the canonical universe of the i-th extension requires to check $|\mathrm{C}(\Gamma_{i-1})| \cdot 2^{k_i}$ sets against $|\Gamma_i \setminus \Gamma_{i-1}|$ statements, with $k_i = |\Sigma_i \setminus \Sigma_{i-1}|$. So in order for the algorithm to be useful in practice, $\mathrm{C}(\Gamma_i)$ should be of considerably low cardinality (i.e. more in the order of $|\Sigma_i|$ than of $2^{|\Sigma_i|}$) and k_i should be near to one. This means not only that the final result $\mathrm{C}(\Gamma)$ is of tractable size but also that the partition of $\Gamma \setminus \Gamma_0$ into the sets $\Gamma_i \setminus \Gamma_{i-1}$ ($1 \leqslant i \leqslant n$) has the effect of keeping $|\mathrm{C}(\Gamma_i)|$ small during the construction process.

Notice that a nonredundant theory is not necessarily the best choice. For though less rules reduce the number of tests a single candidate set of primitives has to undergo (in the subroutine $cc?$ of Figure 70), the total number of sets to be tested during an extension step can increase because additional (redundant) statements would have pruned $\mathrm{C}(\Gamma_i)$ at an earlier stage of the construction.

Suppose Γ has reduced normal form. Then the statements of Γ can be represented by pairs of finite subsets of Σ, i.e. by *sequents* in the sense of Section 6.3.3. Recall that a statement $\varphi \preceq \psi$ of Γ is represented by a sequent $\langle P, Q\rangle$ if φ is the conjunction of the members of P and ψ is the disjunction of the members of Q (with V and Λ arising respectively by conjunction and disjunction of nothing). Moreover $P \cap Q = \varnothing$ since the normal form of Γ is reduced. In terms of the sequent representation of Γ, the **if**-statement of the subroutine $cc?$ becomes an elementary condition on sets:

if $P \subseteq X$ **and** $X \cap Q = \varnothing$ **then** **return** (false);

We close this section with some remarks on implementational issues. *Bit-*

```
function C (Σ: set; Γ: theory): system of sets;
begin
    if not cc? (∅, Γ|∅) then
        C := ∅
    else begin
        C := {∅};
        Σ′ := ∅;
        while Σ ≠ ∅ and C ≠ ∅ do begin
            F := any nonempty subset of Σ;
            Σ′ := Σ′ ∪ F;
            Γ′ := Γ|Σ′;
            C′ := ∅;
            foreach X ∈ C, Y ⊆ F do begin
                X′ := X ∪ Y;
                if cc? (X′, Γ′) then C′ := C′ ∪ {X′}
            end;
            Σ := Σ \ F;
            Γ := Γ \ Γ′;
            C := C′
        end
    end
end;

function cc? (X: set; Γ: theory): boolean;
{ true if X is consistently Γ-closed, false otherwise }
begin
    foreach (φ ⪯ ψ) ∈ Γ do
        if X ⊨ φ and X ⊭ ψ then return (false);
    return (true)
end;
```

FIGURE 70 A simple algorithm for constructing the canonical universe *vectors* of length $|\Sigma|$ provide a simple and effective way to encode subsets of Σ. (Notice that bit-vectors of length $|\Sigma|$ can be regarded as direct instantiations of $\mathbb{2}$-valued interpretations of Σ.) In addition, the algorithm of Figure 70 calls for a method to represent subset systems over Σ. For that purpose, ordered lists of bit-vectors are appropriate.

The specialization ordering on $\mathrm{C}(\Gamma)$ corresponds to the bitwise ordering on the set of bit-vector encodings (which are essentially the $\mathbb{2}$-valued models of Γ). In the following, we refer to bit-vector encodings of members of $\mathrm{C}(\Gamma)$ briefly as *codes*. Consider the task of finding the set M of minimal upper bounds of a set S of codes, which includes the special case of finding the supremum of S, if existent. If Γ is a simple inheritance theory with exclusions, the supremum of each bounded set S exists and can be determined by applying the bitwise logical or-operation to S; for $\mathrm{C}(\Gamma)$ is closed with respect to bounded union (see Chapter 2). In case Γ is a Horn theory, the supremum of a bounded set S of codes still exists, but $\mathsf{or}(S)$ may not be a code in turn; the supremum of S is then the least code above $\mathsf{or}(S)$. For observational theories in general, where suprema need not exist, one has to determine the set of all minimal codes above $\mathsf{or}(S)$ instead.

The determination of minimal upper bounds can be sped up by an explicit representation of the specialization order, e.g. in form of a directed acyclic graph.[5] Such a representation of the ordering structure can be maintained already during the process of constructing the canonical universe. It is not difficult to modify the above algorithm accordingly. Assume that $|\Sigma_i \setminus \Sigma_{i-1}|$ is a singleton, say $\{p_i\}$, for all i. Then $\mathrm{C}(\Gamma_i)$ together with its ordering structure can be constructed from that of $\mathrm{C}(\Gamma_{i-1})$ by taking the direct product $\mathrm{C}(\Gamma_{i-1}) \times \{\varnothing, \{p_{i-1}\}\}$ of ordered sets and deleting all elements that are not consistently closed with respect to $\Gamma_i \setminus \Gamma_{i-1}$; see also Section 9.4.1. (The two operations on directed acyclic graphs involved here are taking products of graphs and deleting nodes.)

9.4.3 *Applications*

In case the canonical universe $\mathrm{C}(\Gamma)$ of a theory Γ is of tractable size, it can be usefully employed for various inference tasks. This includes, for example, the task of finding the minimal (consistent) Γ-closures of a given set of primitives, and hence the task of finding the minimal satisfiers of a conjunctive predicate. Another type of problem easily solvable with the help of $\mathrm{C}(\Gamma)$ is to determine whether a given statement $\varphi \preceq \psi$ is entailed by Γ or not; for Γ entails $\varphi \preceq \psi$ just in case no member X of $\mathrm{C}(\Gamma)$ satisfies φ but not ψ.

The construction algorithm of Figure 70 can be easily adapted to allow an on-line introduction of new primitives and statements. A similar approach is

[5] More on encoding techniques for partial orders can be found e.g. in Aït-Kaci et al. 1989, Habib and Nourine 1994, Fall 1996, Caseau et al. 1999.

that of Oles (2000), who derives the representation of an observational theory (which he calls a knowledge base) by its generic universe (which he does not seem to recognize as such) via a somewhat tedious detour through Lindenbaum algebras. For expository purposes, Oles uses a “simple knowledge representation language” consisting of the three constructs

$$\texttt{emptyKB}, \quad \texttt{newconcept}\, I, \quad \text{and} \quad \texttt{subconcept}\, C_1 C_2,$$

where I is an identifier “given denotations by the knowledge base” and the C_i’s are inductively built from identifiers plus $\top$ and $\perp$ by finite conjunction and disjunction. Obviously, the identifiers are our primitive predicates whereas ‘`subconcept` C_1C_2’ means ‘$C_1 \preceq C_2$’. The `newconcept` construct is for explicitly introducing new primitives, and `emptyKB` is for initialization.

Part IV

Application to Attribute-Value Theory

10

Linguistic Applications

So far, classification was based on primitive monadic predicates. In modern linguistics, classification is often feature-based. Predicates then typically have inner structure in that they consist of a feature, say, GENDER, and a value of that feature, say, *feminine*. Ascribing GENDER:*feminine* to a linguistic entity means to state that the gender of that entity is feminine. In this chapter, we study several paradigmatic examples of feature-based classification in a semi-formal manner; a thorough explication of the logical form of feature-based descriptions will follow in Chapter 11.

Section 10.1 starts with a brief review of binary features in phonology and semantics. Another class of features is shown to be closely related to choice systems. Moreover, we introduce a standard method to impose constraints on feature-value pairs by so-called feature declarations. Finally, it is emphasized that the genuine contribution of feature-based descriptions lies in their capacity to express structural information.

In Section 10.2, we discuss a concrete application of feature-based descriptions to the semantic classification of objects. The example is taken from the current implementation of the MultiNet language understanding system, but the discussion should be general enough to apply to other cases as well. It is pointed out that a hand-crafted set of admissible feature-value combinations usually has defects, and that these defects can be brought to light by checking the feature-value statements that are valid with respect to the predefined set of combinations.

10.1 Classification by Features and Sorts

10.1.1 *Feature Bundles and Binarism*

The basic idea of Chapter 1 was to classify linguistic entities by ascribing certain properties to them. As explained at length in earlier chapters, it is reasonable from the viewpoint of a given classificational theory to *identify* the entities that can be characterized by the theory with the sets of properties that are consistent and closed with respect to the theory. A *generic entity* of the theory, as it was called, "is" thus the set of all its properties. By speaking of properties as *attributes* or *features*, which is surely not uncommon in everyday usage, we can therefore regard generic entities as "*feature bundles*".

In modern linguistics, the characterization of linguistic entities as bundles of *distinctive features* has been employed most prominently in phonology.[1] The phoneme /p/, for instance, can be seen as the result of combining the phonetic features *voiceless*, *plosive*, *bilabial*, etc. The features are distinctive in that changing one of them gives rise to a different phoneme; e.g. replacing *voiceless* by *voiced* leads to /b/. Phonetic features are usually assumed to occur in *opposition pairs*, like *voiced* vs. *voiceless*, *plosive* vs. *non-plosive*, and so on. To take account of this systematic dichotomy, it is common to introduce *binary features* VOICE, PLOSIVE, etc, with possible values $+$ and $-$. Then *voiced* and *voiceless* become $\text{VOICE}{:}+$ and $\text{VOICE}{:}-$, respectively. Notice the twist in terminology: now VOICE is said to be a feature, whereas $\text{VOICE}{:}-$ is a *feature-value pair* (see also Section 11.2.1 below). The phoneme /p/ is then given by the set $\{\text{VOICE}{:}-, \text{PLOSIVE}{:}+, \ldots\}$ of feature-value pairs, for which the following *matrix notation* is in use:

$$\begin{bmatrix} \text{VOICE} & - \\ \text{PLOSIVE} & + \\ \vdots & \end{bmatrix}$$

Feature-based classification, which is closely related to componential analysis (cf. Section 1.2), has been widely employed in semantics as well.[2] Typical semantic opposition pairs are *animate* vs. *inanimate* and *male* vs. *female*. As Lyons (1977, Sect. 9.9) points out, the use of binary features confronts us with the problem of marking one of the equipollent opposites *male* and *female* as negative by choosing either MALE or FEMALE as a binary feature.[3]

Another question raised by Lyons (*ibid*) is concerned with the distinction between concepts like *horse*, for which the feature MALE (or FEMALE) is appropriate but unspecified, and concepts like *house*, to which the feature does not apply at all. The obvious solution is to restrict the feature MALE to animate objects and to allow binary features to have a third value, say *boolean*, that generalizes $+$ and $-$. Notice that if there is no need to distinguish unspecified from inappropriate features, the use of binary features can be straightforwardly simulated by term negation.[4]

---

[1] Cf. Jakobson and Halle 1956, Chomsky and Halle 1968.

[2] E.g. Katz and Fodor 1963.

[3] For a more fundamental critique of linguistic classification based on binary features see e.g. Taylor 1989.

[4] See also Lyons 1995, Sect. 4.2.

10.1.2 *Features versus Choice Systems*

One way to circumvent the problem of opposite but equipollent features like *male* and *female* is to treat them not as binary features but as dichotomous *values* of a single feature SEX.[5] Then SEX:*female* and SEX:*male* can take the place of MALE:$-$ and MALE:$+$ (or FEMALE:$+$ and FEMALE:$-$). In order to express that the feature SEX is appropriate but unspecified, we can introduce a feature value *sex* which subsumes *male* and *female*. This approach is not restricted to opposition pairs but can be applied to arbitrary property paradigms as well. For instance, the properties *red*, *blue*, *green*, etc can be taken as the values of the feature COLOR.

Ascribing COLOR:*red* to an object basically says that the object is red and that red is a color (see Section 11.2.2 below). It is thus tempting to treat *red*, *blue*, etc as members of a choice system and not as values of a feature. Let us speak of a *choice feature* if the values of a feature constitute a choice system, i.e. a system of mutually incompatible properties (or predicates). The choice feature COLOR, for example, corresponds to the choice system $\{red, blue, green, \ldots\}$.

It is of course also possible to turn choice systems into choice features. Consider the AND/OR-tree presented in Figure 5 of Section 1.2, that shows the subclassification of German nominal word forms with respect to case, gender, and number. With choice features CASE, GENDER, and NUMBER, a nominal word form is then said to satisfy, for instance, CASE:*nominative*, GENDER:*feminine*, and NUMBER:*singular* instead of just satisfying *nominative*, *feminine*, and *singular*.

10.1.3 *Feature Declarations and Other Feature-Value Constraints*

More often than not, a feature (or attribute) is not appropriate to all entities of a given domain of discourse but only to a certain subclass of them. The feature CASE, for instance, is appropriate to nominal expressions but not to adverbs. In order to express such appropriateness conditions by observational statements we need to be a bit more explicit on the logical form of feature-value pairs.

Without going into details – see Chapter 11 for a full account – let us adopt the position that features or, better, feature symbols, are functional dyadic predicates and that their "values" are monadic predicates, henceforth referred to as *sorts* or, better, sort symbols (or predicates). If F is a feature symbol and s is a sort symbol then the expression F:s is taken as a monadic predicate that is satisfied by those entities "whose F is of sort s". The above appropriateness constraint for CASE, which says that a (linguistic) entity is a nominal expression whenever it admits a case value, can then be expressed by the observational statement CASE:$\mathrm{V} \preceq$ *nominal*.

[5] Cf. Lyons 1977, p. 325.

A related type of constraint is that all entities of a certain sort have certain features whose values are restricted in a certain way. For example, each linguistic entity of sort *nominal* bears the feature CASE, whose value is restricted to a sort *case* with subsorts *nominative*, *accusative*, etc. The systematic listing of all features appropriate to a sort, together with their respective value restrictions, is also called the *feature declaration* of that sort; see Pollard and Sag 1994 or Sag and Wasow 1999 for extensive linguistic examples. A feature declaration of the form

$$\langle s, \begin{bmatrix} \mathrm{F}_1 & t_1 \\ & \vdots \\ \mathrm{F}_n & t_n \end{bmatrix} \rangle$$

signifies that for $1 \leqslant i \leqslant n$, the feature F_i is appropriate for entities of sort s, and the value such an entity has with respect to F_i is of sort t_i.[6] Notice that this usage of 'value' differs from the foregoing one, where the sort (symbol) occurring in a feature-value pair was taken as the value of the respective feature.

Suppose the sorts which are subject to feature declarations are arranged in a taxonomic hierarchy. It is then presumed that a sort *inherits* all features of its supersorts, that is, a feature that is defined for a sort is also defined for all subsorts of that sort. Suppose a feature-value pair $\mathrm{F}{:}t$ occurs in the feature declaration of a sort s. If F is inherited from a supersort of s, we can assume that t is more specific than the value restriction of F at that supersort, because otherwise $\mathrm{F}{:}t$ would be redundant in the declaration of s. If F is not inherited, it is said to be *introduced* at s. In this case, it is implicitly assumed that bearing the feature F implies to be of sort s.

We can express the meaning of feature declarations in terms of observational statements as follows. If F is introduced at s, then $\mathrm{F}{:}\mathrm{V} \preceq s$. The condition that every object has all features declared for its sort corresponds to the statements $s \preceq \mathrm{F}{:}t$, for all feature-value pairs $\mathrm{F}{:}t$ in the feature declaration of s. In particular, $s \equiv \mathrm{F}{:}t$ in case F is introduced at s.[7] The assumption that sorts inherit the feature-value restrictions of their supersorts is a logical consequence: if s' is a subsort of s, i.e. $s' \preceq s$, then $s' \preceq \mathrm{F}{:}t$ whenever $s \preceq \mathrm{F}{:}t$.[8]

Feature declarations alone are of course insufficient to specify a full grammatical theory, or even a lexical theory. This calls for more general observational statements over feature-value pairs, like e.g. the *feature cooccurrence restrictions* used in Generalized Phrase Structure Grammer.[9]

[6] See also Pollard and Sag 1994, pp. 395ff.

[7] Since $t \preceq \mathrm{V}$ implies $\mathrm{F}{:}t \preceq \mathrm{F}{:}\mathrm{V}$; see Section 11.3

[8] For more on appropriateness and feature introduction see Carpenter 1992, Chap. 6 and Penn 2000.

[9] Gazdar et al. 1985, Sells 1985, Chap. 3.

10.1.4 *From Features to Structures*

As mentioned at the close of Section 10.1.1, the effect of binary features can essentially be simulated by term negation. Choice features like GENDER can be discarded too, for *feminine* does the job of GENDER:*feminine*. And there is no danger of confusing the values of two different choice features as long as the underlying choice systems are pairwise disjoint. So the feature-based descriptions given so far can dispense with features altogether (see also Section 11.2 below).

The situation is different when it comes to structural information. Consider for example the description of the valency frame of a verb in the lexicon. Since the subject of the verb and its objects are specified with respect to the same type of syntactic and semantic information, like part of speech and semantic sort, it is inevitable to explicitly mark the role of the complement to be described. A simple solution is to introduce features SUBJ, OBJ1, OBJ2, etc.

Alternatively, one can characterize the list of complements explicitly as a *list structure*. The standard way to do so is to define a sort *list* with subsorts *nonempty list* (*ne-list*) and *empty list* (*e-list*), where the sort *ne-list* introduces features FIRST and REST with values restricted to V and *list*, respectively.[10] A list of three elements of sort a, b, and c can then be characterized by the following feature-value pairs:

$$
\begin{array}{l}
\textit{ne-list},\\
\text{FIRST}:a,\\
\text{REST}:\textit{ne-list},\\
\text{REST}:(\text{FIRST}:b),\\
\text{REST}:(\text{REST}:\textit{ne-list}),\\
\text{REST}:(\text{REST}:(\text{FIRST}:c)),\\
\text{REST}:(\text{REST}:(\text{REST}:\textit{e-list})).
\end{array}
$$

Clearly features are indispensable in this case. The novel aspect here is that values of features can be feature-value pairs in turn. A convenient notational device for sets or conjunctions of such complex feature-value descriptions is provided by nested attribute-value matrices. The matrix shown in Figure 71, for example, represents the preceding set of feature-value pairs.

Readers interested in feature-based descriptions of syntactic phenomena are referred to Pollard and Sag 1994, Borsley 1996, and Sag and Wasow 1999, all of which employ variants of Head Driven Phrase Structure Grammar. It has to be added that these approaches make use of an additional type of attribute-value predicate, which allows to state that feature values are identical. We briefly touch this topic in Section 12.4.1 below.

[10] Cf. e.g. Carpenter 1992.

$$\begin{bmatrix} \textit{ne-list} & \\ \text{FIRST} & a \\ \text{REST} & \begin{bmatrix} \textit{ne-list} & \\ \text{FIRST} & b \\ \text{REST} & \begin{bmatrix} \textit{ne-list} & \\ \text{FIRST} & c \\ \text{REST} & \textit{e-list} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

FIGURE 71 Attribute-value matrix representing a list of three elements

10.2 Case Study: Sorts and Features in MultiNet

10.2.1 *The MultiNet Paradigm*

The *MultiNet* (short for *Multilayered Extended Semantic Network*) paradigm is an elaborate framework for knowledge representation, which allows to represent the semantics of natural language expressions in a fine grained and cognitively adequate way. To this end MultiNet provides representational means to distinguish, for example, intensional from extensional aspects of meaning, actual from hypothetical entities, and immanent from situative knowledge. There is also a variety of representational means for quantification and definiteness phenomena. For a detailed description of the MultiNet paradigm the reader is invited to consult Helbig 2001.[11]

The core component of MultiNet is a semantic network formalism. A semantic network is basically a labeled directed graph whose nodes represent concepts and whose edges represent semantic relations. If an edge is labeled by a member of a predefined set of relation symbols, the pair of nodes connected by the edge satisfies the *semantic relation* named by the label. Nodes are labeled by symbols that indicate the *ontological sort* of the represented concepts; furthermore, they are specified with respect to several so-called *layer features* that encode e.g. genericity and definiteness information.[12]

Besides providing a *formalism* for semantic representation, the MultiNet paradigm is an integral part of a natural language processing *system*, henceforth referred to as the *MultiNet system*. The MultiNet system includes a syntactico-semantic analyzer that automatically transforms natural language expressions into semantic representations.[13] The transformation of an expression into its semantic representation rests to a large extent on lexical information about the words occurring in that expression. In addition to morphological and syntactic information, the lexicon of the MultiNet system contains detailed semantic information.[14]

[11] See also Helbig and Schulz 1997, Helbig and Gnörlich 2002.
[12] See Helbig 2001, Chap. 10, Hartrumpf and Helbig 2002.
[13] See Helbig and Hartrumpf 1997, Hartrumpf 2002, Chap. 3.
[14] See Schulz 1999, Helbig 2001, Chap. 12.

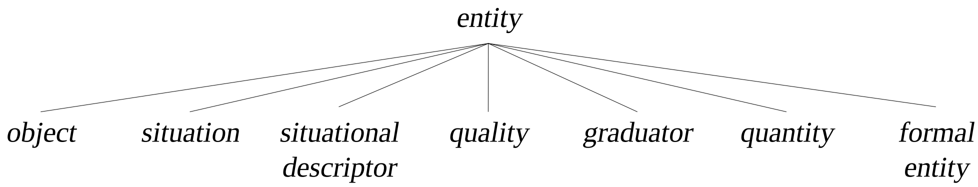

FIGURE 72 Top level of the MultiNet sort hierarchy

In particular, a lexical entry specifies the semantic relations the denoted entity bears to its participants (if there are any) as well as the ontological sort of these entities. For example, the verb 'kill' denotes an entity of sort *action* that bears the relation AGT to its agent and the relation AFF to the entity directly affected by the action. In addition, the entity denoted by the lexeme, as well as its participants, are specified with respect to a predefined set of binary *semantic features*. The entity affected by a killing action, for instance, is assumed to satisfy ANIMATE:$+$. The primary purpose of this classification by semantic features is to efficiently provide the syntactico-semantic analyzer with semantic selectional restrictions; for representing these restrictions within the semantic network formalism proper would call for significantly more costly inference processes. In what follows we take a closer look at the interdependence between ontological sorts and semantic features.

10.2.2 *Ontological Sorts and Semantic Features*

The MultiNet framework presumes a tree-shaped hierarchy of 45 ontological sorts. Figure 72 shows all immediate subsorts of the most general sort *entity*; for a full account of the MultiNet sort hierarchy see Helbig 2001, Sect. 17.1. In addition, there are 16 binary features that allow a more fine grained semantic differentiation (see *op cit*, Sect. 12.2), of which 13 are listed in Table 3, with positive and negative examples.

In the following, we focus on the three ontological sorts *concrete object* (*co*), *discrete object* (*d*), and *substance* (*s*) (see Figure 73) and their semantic differentiation by the features listed in Table 3. If we assume the value of each of these features as being either *boolean* or one of its subsorts $+$ and $-$, we get $3 \cdot 3^{13} = 4{,}782{,}969$ possible entity types. Among these are $2 \cdot 2^{13} = 16{,}384$ maximally specific entity types, that is, types whose ontological sort is either *d* or *s* and whose semantic features have value $+$ or $-$.

The number of possible entity types reduces by several orders of magnitude if the semantic interdependence between feature-value pairs is taken into account. For example, humans and animals are animate beings and potential agents; rendered into logical form:

$$\text{HUMAN:}+ \vee \text{ANIMAL:}+ \;\preceq\; \text{ANIMATE:}+ \wedge \text{POTAG:}+,$$

Feature		+	−
ANIMAL		*bird*	*human*
ANIMATE		*tree*, *animal*	*stone*
ARTIF	(artifact)	*car*, *house*	*animal*
AXIAL		*bottle*, *tree*	*ball*, *snowflake*
GEOGR	(geographical object)	*street*, *mountain*	*animal*
HUMAN		*student*	*animal*, *stone*
INFO	(information)	*newspaper*	*animal*, *stone*
INSTIT	(institution)	*company*	*student*
INSTRU	(instrument)	*knife*, *hammer*	*mountain*
LEGPER	(legal person)	*student*, *company*	*animal*, *stone*
MOVABLE		*table*, *apple*	*mountain*
POTAG	(potential agent)	*horse*, *motor*	*stone*
SPATIAL		*building*	*idea*, *number*

TABLE 3 Subset of semantic features used in MultiNet

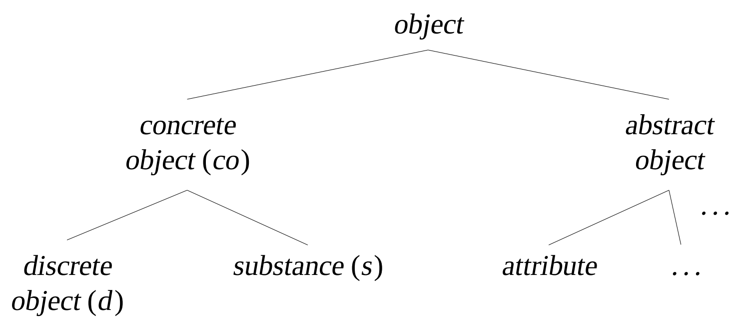

FIGURE 73 Ontological subsorts of *object*

which is logically equivalent to the conjunction of four simple inheritance statements, namely HUMAN:$+ \preceq$ ANIMATE:$+$, HUMAN:$+ \preceq$ POTAG:$+$, etc.

10.2.3 *Encoding of Feature-Value Combinations by Semantic Sorts*

Within the MultiNet system, the differentiation of ontological sorts by semantic features is realized by introducing sort symbols for all possible feature-value combinations.[15] We refer to these additional sorts as *semantic sorts*. The semantic sorts are technically kept apart from the ontological sorts by defining a feature SORT for the former, whose value belongs to the latter. There is, for instance, a semantic sort *human-object* with SORT:d, HUMAN:$+$ (unsurprisingly), ANIMAL:$-$, ANIMATE:$+$, ARTIF:$-$, etc; see the attribute-value matrix on the right of Figure 74.

Some of the semantic sorts bear semantic features that are unspecified in the sense of having the value *boolean*. The semantic sort *animate-object*,

[15]Cf. Schulz 1999, Helbig 2001, Hartrumpf 2002, Hartrumpf et al. 2003.

	object	*con-object*	*con-potag*	*animate-object*	*human-object*
SORT	*o*	*co*	*d*	*d*	*d*
ANIMAL	*b*	*b*	*b*	*b*	−
ANIMATE	*b*	*b*	*b*	+	+
ARTIF	*b*	*b*	*b*	−	−
AXIAL	*b*	*b*	+	+	+
GEOGR	*b*	*b*	−	−	−
HUMAN	*b*	*b*	*b*	*b*	+
INFO	*b*	*b*	−	−	−
INSTIT	*b*	−	−	−	−
INSTRU	*b*	*b*	*b*	−	−
LEGPER	*b*	*b*	*b*	*b*	+
MOVABLE	*b*	*b*	*b*	*b*	+
POTAG	*b*	*b*	+	+	+
SPATIAL	*b*	+	+	+	+

FIGURE 74 Stepwise restriction of feature values

for example, is unspecified with respect to the features ANIMAL, HUMAN, LEGPER, and MOVABLE. The purpose of such non-maximal sorts is to capture "natural" generalizations. (A possible criterion for introducing a non-maximal sort could be that there are lexical concepts of this sort.) Figure 74 shows a specialization chain of semantic sorts, starting with the sort *object*, all of whose semantic features have value *boolean* (abbreviated by *b*); Figure 75 depicts the hierarchy of all subsorts of *con-object* (short for *concrete object*) as used in the MultiNet system (see Helbig 2001, p. 386) – with features only displayed if their value is more restricted than that of the immediate supersort.

The current lexical component of the MultiNet system explicitly encodes the semantic sort hierarchy, accompanied by feature declarations as indicated in Figure 75. Since all semantic features listed in Table 3 are defined for the semantic sort *object*, the feature declarations of *con-object* and its subsorts do not introduce any features but solely consist of value restrictions. Feature declarations alone, however, do not suffice to single out all and only the feature-value combinations that correspond to the semantic sorts. This is so because the logic behind feature declarations does not prevent an entity of a given sort to bear features whose values are more specific than stated in the declaration of that sort. For example, an entity of sort *animate-object* can satisfy MOVABLE:+ without being of sort *human-object* or *animal-object*. Moreover, entities satisfying *con-object*, ANIMATE:−, and HUMAN:+ are excluded neither, contrary to what one might expect. Notice that if the goal is to single out only the maximally specific semantic sorts, it is enough to assume exhaustiveness of the sort hierarchy, which means that each sort implies the disjunction of its immediate subsorts.

According to Schulz (1999, Sect. 3.2.2), the hierarchy of semantic sorts used in the MultiNet system has been determined "by hand" on the basis of lexicographic practice, with an eye on semantic dependencies between feature

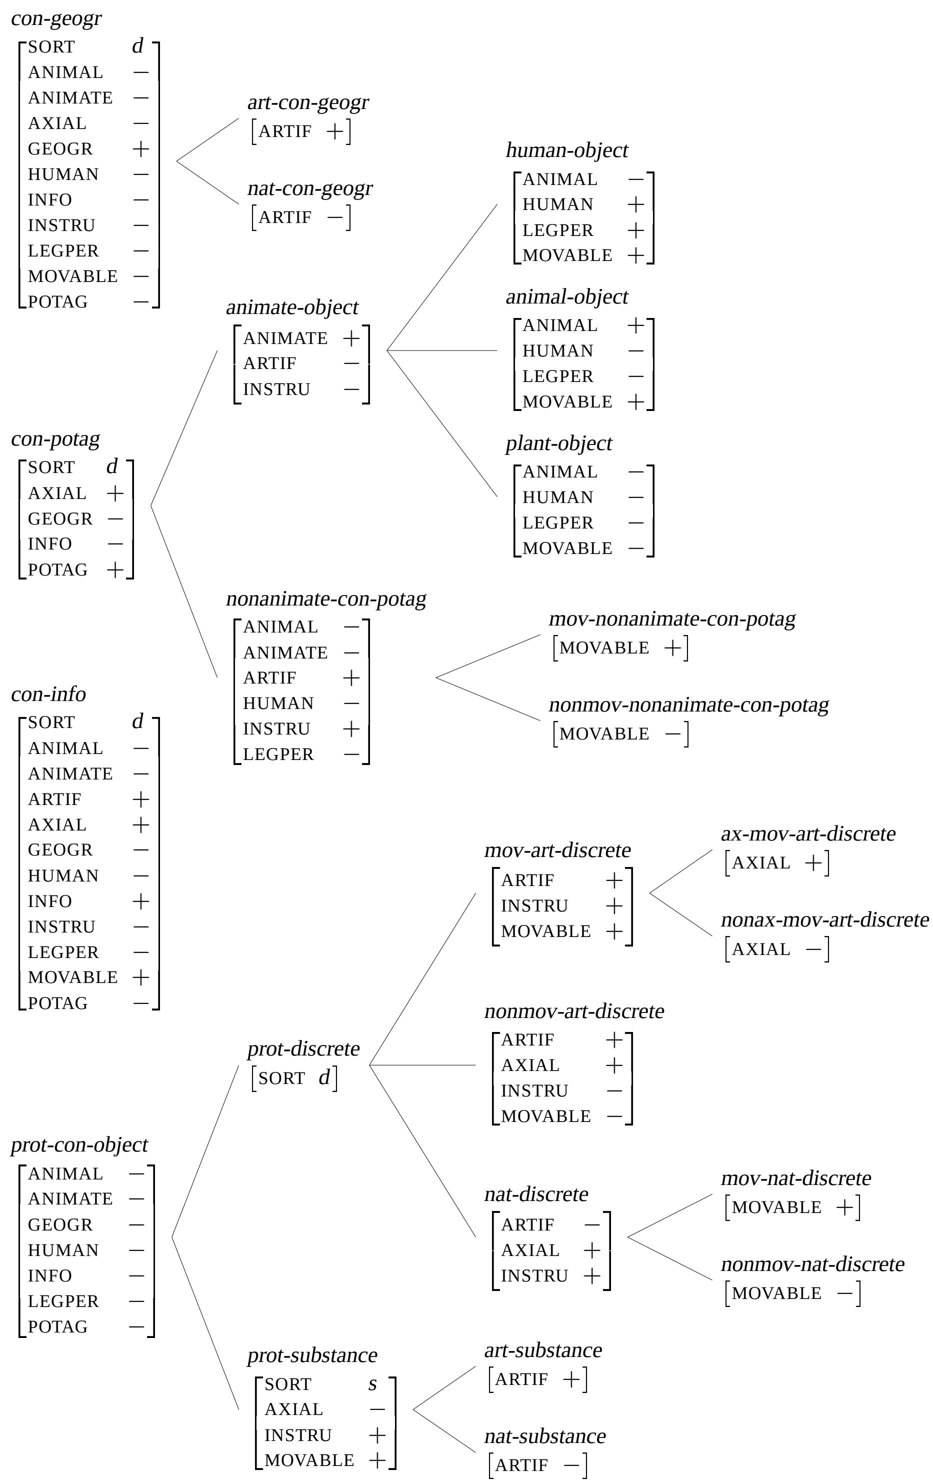

FIGURE 75 Semantic subsorts of *con-object* used in the MultiNet system

values. So, neither has a set of feature-value statements been employed systematically to constrain the possible feature-value combinations nor has any attempt been made to reveal the feature-value statements valid for the chosen system of semantic sorts. In other words, the interdependence between sets of feature-value statements (theories) and systems of sets of feature-value pairs (canonical universes) has been neglected for the most part.

It is hence not very surprising that the provisional set of feature-value statements given in Helbig 2001, Sect. 12.2 is everything but complete with respect to the system of semantic sorts. There is even no guarantee of validity, witness the feature-value statements $\text{INSTR}{:}+ \preceq \text{MOVABLE}{:}+$ and $\text{INFO}{:}+ \preceq \text{SPATIAL}{:}-$ (*ibid*, p. 276). The first statement is incompatible with the feature-value matrix of the sort *nonmov-nonanimate-con-potag*, the second one with that of *con-info*, both of which occur in Figure 75 and are thus semantic sorts of the MultiNet system (*op cit*, p. 386).[16]

Moreover, several of the statements valid with respect to the system of Figure 75 turn out to be rather questionable. For example, the statement

$$\text{ARTIF}{:}- \wedge \text{ANIMATE}{:}- \wedge \text{GEOGR}{:}- \wedge \text{SORT}{:}d \;\preceq\; \text{AXIAL}{:}+ \wedge \text{INSTRU}{:}+$$

holds for each of the semantic sorts. Assuming the given semantic subclassification of concrete objects as complete thus implies that all natural, non-animate, non-geographic discrete objects are typically used as instruments and have an axis. Obviously this is wrong: take snowflakes, for instance.

A manually crafted hierarchy of semantic sorts therefore may implicitly contain assumptions about feature-value dependencies that are not intended by the designer of the hierarchy. This is highly problematic because, once adopted, the hierarchy gains normative impact when applied in classification tasks, say, of assigning semantic sorts to lexemes, where the user is forced to choose from the given set of sorts. Mnemonic designators for semantic sorts, like *nat-discrete*, make it even more difficult to detect defects in the hierarchy, since they hide the actual distribution of feature values.

10.2.4 *Feature-Value Combinations via Feature-Value Statements*

The discussion of the previous section has shown that a definition of a system of semantic sorts, i.e. of feature-value combinations, should be systematically related to a set of feature-value statements. This is not to say that lexicographic examples are of no use in determining the semantic sorts. On the contrary, a reasonably representative set of example lexemes, classified with respect to all semantic features, can serve as a first approximation of possible feature-value

[16] We tacitly assume each statement to hold without exception. The situation is of course different if defaults are allowed.

combinations. (Such a set of classified examples is basically a classification in the sense of Section 5.3.4.)

Since the example set cannot be expected to be complete *a priori*, it is vital to check whether the feature-value statements that hold for the example set are semantically valid. Such a statement is acknowledged as being invalid by giving a *counter example*, which is then added to the example set. The effect is, first, that there is a new semantic sort and, second, that the invalid statement cannot be inferred anymore.[17] For instance, we have observed above that 'snowflake' is a counter example to the statement $\textsc{artif}{:}- \wedge \textsc{animate}{:}- \wedge \textsc{geogr}{:}- \wedge \textsc{sort}{:}d \preceq \textsc{axial}{:}+ \wedge \textsc{instru}{:}+$, which holds for all semantic sorts of Figure 75. Consequently, one has to add a subsort of *prot-discrete* with $\textsc{artif}{:}-$, $\textsc{axial}{:}-$, $\textsc{instru}{:}-$, and $\textsc{movable}{:}+$.

Notice that in order for this approach to be feasible in practice, we need a strategy to present the statements to be checked in a nonredundant way. This task is closely related to the problem of finding a nonredundant basis of observational statements for a given canonical universe; cf. (5.30). The interested reader is referred to Ganter 1999.

Leaving aside for the moment the problem of constructing a hierarchy of semantic sorts by introducing non-maximal sorts, let us assume that all example lexemes give rise to maximally specific feature-value matrices, i.e. each semantic feature has value $+$ or $-$ and the feature $\textsc{sort}$ has value d or s. In other words, we require the statements

$$\textsc{sort}{:}co \;\preceq\; \textsc{sort}{:}d \vee \textsc{sort}{:}s \qquad \text{and} \qquad F{:}b \;\preceq\; F{:}+ \vee F{:}-$$

to be valid, for all semantic features F. Since we only consider concrete objects, for which all semantic features of Table 3 are defined, the statements $\mathrm{V} \preceq \textsc{sort}{:}co$ and $\textsc{sort}{:}co \preceq F{:}b$ can be taken as valid too.

In addition to these exhaustiveness assumptions, the hierarchy of the ontological sorts (see Figure 73) and the fact that $+$ and $-$ are incompatible subsorts of *boolean* give rise to the statements

$$\textsc{sort}{:}d \vee \textsc{sort}{:}s \preceq \textsc{sort}{:}co, \qquad \textsc{sort}{:}d \wedge \textsc{sort}{:}s \preceq \Lambda,$$
$$F{:}+ \vee F{:}- \preceq F{:}b, \qquad \text{and} \qquad F{:}+ \wedge F{:}- \preceq \Lambda,$$

for all semantic features F. (As for the logical deduction of these statements from $d \vee s \preceq co$, etc see Section 11.3.)

The canonical universe of the feature-value theory defined so far is the set of $2 \cdot 2^{13} = 16{,}384$ maximally specific semantic sorts that are possible on

[17] This iterative procedure is simliar to the method of *attribute exploration* used in Formal Concept Analysis; see Ganter 1999, Ganter and Wille 1999.

(1)	ANIMATE:$+$ $\preceq$ ARTIF:$-$ $\wedge$ INSTRU:$-$ $\wedge$ POTAG:$+$
(2)	ANIMAL:$+$ $\preceq$ ANIMATE:$+$ $\wedge$ MOVABLE:$+$
(3)	HUMAN:$+$ $\preceq$ ANIMAL:$-$ $\wedge$ ANIMATE:$+$ $\wedge$ MOVABLE:$+$ $\wedge$ LEGPER:$+$
(4)	GEOGR:$+$ $\preceq$ AXIAL:$-$ $\wedge$ INSTRU:$-$
(5)	INFO:$+$ $\preceq$ INSTRU:$-$ $\wedge$ POTAG:$-$
(6)	SORT:*co* $\preceq$ SPATIAL:$+$ $\wedge$ INSTIT:$-$
(7)	SORT:*s* $\preceq$ MOVABLE:$+$ $\wedge$ AXIAL:$-$
(8)	SORT:*co* $\wedge$ INFO:$+$ $\wedge$ AXIAL:$-$ $\preceq$ Λ
(9)	SORT:*co* $\wedge$ INFO:$+$ $\wedge$ MOVABLE:$-$ $\preceq$ Λ
(10)	SORT:*co* $\wedge$ POTAG:$+$ $\wedge$ AXIAL:$-$ $\preceq$ Λ
(11)	SORT:*co* $\wedge$ HUMAN:$-$ $\wedge$ LEGPER:$+$ $\preceq$ Λ
(12)	SORT:*co* $\wedge$ ANIMATE:$-$ $\wedge$ INSTRU:$-$ $\wedge$ POTAG:$+$ $\preceq$ Λ
(13)	ANIMATE:$-$ $\wedge$ ARTIF:$-$ $\wedge$ POTAG:$+$ $\preceq$ Λ
(14)	ANIMAL:$-$ $\wedge$ ANIMATE:$+$ $\wedge$ HUMAN:$-$ $\wedge$ MOVABLE:$+$ $\preceq$ Λ
(15)	AXIAL:$-$ $\wedge$ GEOGR:$-$ $\wedge$ MOVABLE:$-$ $\preceq$ Λ
(16)	ARTIF:$+$ $\wedge$ INSTRU:$+$ $\wedge$ MOVABLE:$-$ $\wedge$ POTAG:$-$ $\preceq$ Λ
(17)	ANIMATE:$-$ $\wedge$ INFO:$-$ $\wedge$ INSTRU:$-$ $\wedge$ MOVABLE:$+$ $\preceq$ Λ
(18)	SORT:*d* $\wedge$ ARTIF:$-$ $\wedge$ ANIMATE:$-$ $\wedge$ AXIAL:$-$ $\wedge$ GEOGR:$-$ $\preceq$ Λ
(19)	SORT:*d* $\wedge$ ARTIF:$-$ $\wedge$ ANIMATE:$-$ $\wedge$ GEOGR:$-$ $\wedge$ INSTRU:$-$ $\preceq$ Λ

TABLE 4 Valid statements for the semantic sorts of Figure 75

combinatorial grounds. Our aim is to reduce this set to the 15 maximal sorts of Figure 73 by adding appropriate feature-value statements. Table 4 shows a nonredundant set of statements with the desired property.[18] The first block of statements is equivalent to a set of simple inheritance statements consisting of ANIMATE:$+$ $\preceq$ ARTIF:$-$ etc. The second block consists of exclusion statements. Notice that the above exhaustiveness assumptions allow us to state exclusion statements in several equivalent ways. For example, statement (13) is equivalent to ANIMATE:$-$ $\wedge$ POTAG:$+$ $\preceq$ ARTIF:$+$ as well as to POTAG:$+$ $\preceq$ ARTIF:$+$ $\vee$ ANIMATE:$+$.

Some of these statements are rather questionable. As already noted, snowflakes provide a counter example for (18) and (19); moreover they invalidate (17) too. Another problematic example is (16): in contrast to what is presumed there, it seems reasonable to regard a dam or dike as a non-movable, artificial instrument that is not a potential agent – at least in view of the fact that the semantic sort *nonmov-nat-discrete* legitimates non-movable, natural instruments. We can conclude that the systematic relation of feature-value statements to admissible feature-value combinations helps to reveal the shortcomings of a hand-crafted system of semantic sorts and at the same time provides a method

[18]This statement set has been determined semi-automatically by checking that it generates the desired canonical universe, which has been calculated by the algorithm presented in Section 9.4.2.

to improve that system.

Let us turn to the problem of constructing a *hierarchy* of admissible feature-value combinations. To this end, we need to drop the exhaustiveness assumptions $\textsc{sort}{:}co \preceq \textsc{sort}{:}d \vee \textsc{sort}{:}s$ and $F{:}b \preceq F{:}+ \vee F{:}-$. Since the remaining set of statements is equivalent to a simple inheritance theory with exclusions, the resulting ordered set of entity types is a (finite) distributive Scott domain (cf. Chapter 4). Now notice, first, that the maximally specific elements of this hierarchy remain the same as before, and, second, that the hierarchy consists of far more feature-value combinations than just the subsorts of *con-object* defined in Section 10.2.3.

In view of Section 5.3.1, it is nevertheless possible to single out precisely the tree-shaped hierarchy of Figure 75 by imposing an appropriate specification ordering on the semantic features. A typical statement to be added is $\textsc{movable}{:}+ \vee \textsc{movable}{:}- \preceq \textsc{animate}{:}+ \vee \textsc{animate}{:}-$, which ensures that everything specified with respect to MOVABLE is specified with respect to ANIMATE. The question arises on what grounds the tree structure of Figure 75 is to be preferred over other ones. As mentioned before, a possible criterion for stipulating a non-maximal semantic sort is the existence of lexical concepts of that sort. It is however more than questionable whether the system of Figure 75 consists of all and only the feature-value combinations meeting this condition. Another option is to treat the tree of semantic sorts as a *decision tree*. But then the criteria for building decision trees should be taken into account – as, for instance, maximal information gain per attribute specification.[19]

If we dispense with trees altogether, a natural way to construct a hierarchy of feature-value combinations is to take the canonical universe of a complete Horn theory determined by the given set of maximally specific semantic sorts, cf. Section 5.3.4. (Alternatively, one can take the canonical universe of a complete simple inheritance theory with exclusions; but notice that the theory listed in Table 4 is not complete because without exhaustiveness it fails to entail, for instance, the valid statement $\textsc{animal}{:}+ \preceq \textsc{human}{:}-$.) Since the canonical universe of a Horn theory is bounded-complete, every consistent feature-value combination has a least satisfier in the hierarchy. This situation can be argued to be optimal from the viewpoint of information processing because new information can be deterministically processed at once.

[19] See e.g. Mitchell 1997, Chap. 3.

11

Attributive Descriptions

In the foregoing chapter, feature-value pairs have been treated as monadic predicates in a somewhat *ad hoc* fashion. Here, the logical form of feature-based descriptions will be explicated systematically.

Section 11.1 can be seen as an exercise in logical analysis. The logical structure of attributive descriptions, which we regard first and foremost as natural language expressions of a certain type, is revealed by rendering these expressions into regimented and logical form. In Section 11.2, we apply this analysis to binary features and choice features, which turns out to be not without problems. Section 11.3 is concerned with equivalences between attribute-value predicates that arise as consequences of the logical form of these expressions.

11.1 Regimentation and Formalization

The rationale of this section is to devise a formal language of attributive descriptions on the basis of their natural language origins. The underlying assumption is that formalization of a certain pre-formal way of referring to a domain of discourse should start with the very discourse itself. Our approach thus stands in contrast to those that first stipulate a formal language which is then equipped with a model-theoretic semantics (or a translation into predicate logic). It also differs from the position that (mathematical) models of the entities in question are the primary source for designing a formal language.

11.1.1 *Attribute-Value Predicates*

Attributive descriptions as considered here are monadic predicates of the form ‘someone whose father is a lawyer’. Ascribing them to the members of a universe of discourse, say, a set of students of which one is Mary, yields predications like ‘Mary is someone whose father is a lawyer’, which is just a slightly complicated way of saying that Mary’s father is a lawyer. It is common in the context of feature-based classification to speak of the father relation as an *attribute* or *feature* whereas the father of Mary is said to be the *value* of Mary with respect to this feature; feature values in turn can be of a certain *sort*, here of the sort lawyer. So under this usage, an attribute is always ascribed not to a single entity but to a pair of entities.

Our goal is to reveal the logical structure of attributive descriptions by stepwise regimentation and formalization. The first step is to paraphrase 'someone whose father is a lawyer' by the somewhat clumsy expression 'someone such that the father of her or him is a lawyer'. Its gain shows up when (coreferring) pronouns are replaced by (identical) *variables*:

an x such that the father of x is a lawyer.

The expression 'an x such that' works as an operator of *predicate abstraction*: putting it in front of an open sentence like 'the father of x is a lawyer' produces a monadic predicate. Following Quine (1982, Sect. 21), the construction 'an x such that $\ldots x\ldots$' will be rendered symbolically '$\{x \mid \ldots x\ldots\}$'. Our example predicate then becomes '$\{x \mid$ the father of x is a lawyer$\}$'.

This leaves us with the matrix sentence 'the father of x is a lawyer', which is a predication consisting of the predicate 'lawyer' and the singular description 'the father of x'. Let us briefly indicate how to cope with the latter by eliminating it along Russellian lines.[1] Rephrasing 'the father of x' by 'the y such that y is father of x' and writing 'ι' for the definite article gives '$\iota\{y \mid y$ is father of $x\}$', whose more common notation is '$\iota y(y$ is father of $x)$'. So, with 'F' standing for the dyadic predicate '$\{yx \mid y$ is father of $x\}$' and 'A' for '$\{x \mid x$ is a lawyer$\}$', the matrix 'the father of x is a lawyer' becomes '$A(\iota yFyx)$'. Now we eliminate the singular description à la Russell: take '$A(\iota yFyx)$' to mean '$\exists z(z = \iota yFyx \wedge Az)$' and suppose that $z = \iota yFyx$ just in case Fzx and $\forall yw(Fyx \wedge Fwx \rightarrow y = w)$, in words: z is father of x and it is the only one.[2]

All in all, consecutive regimentation and formalization has lead us to the following scheme for attributive descriptions, where 'F' stands for a (dyadic) attribute predicate and 'A' for a (monadic) sort predicate:

$$\{x \mid \exists y(Fyx \wedge Ay) \;\wedge\; \forall yz(Fyx \wedge Fzx \rightarrow y = z)\}.$$

If the dyadic predicate 'F' is presupposed as *functional*, that is, if

(11.1) $$\forall xyz(Fxz \wedge Fyz \,\rightarrow\, x = y),$$

then the scheme of attributive description can be abbreviated by '$F\!:\!A$', where

[1] See e.g. Quine 1960, §38, Quine 1982, Sect. 44, or Forbes 1994, Chap. 7, §2.

[2] It should be emphasized that we do not claim Russell's analysis to apply to all singular descriptions in natural language. The claim is rather that it is appropriate for the restricted (scientific) discourse of attributive descriptions.

(11.2) $$F\!:\!A \;\equiv\; \{x \mid \exists y(Fyx \wedge Ay)\}.$$

We call '$F\!:\!A$' the *inverse image* or the *Peirce product* of 'A' by 'F'. So '$:$' can be seen as a predicate operator that combines a dyadic predicate 'F' and a monadic predicate 'A' to a monadic one. In the following we refer to predicates of the form '$F\!:\!A$' as *attribute-value predicates*, where 'F' is presupposed as functional. The corresponding schematic formulation in natural language is 'someone (or something) whose F is an A'.

(11.3) Remark The curly brace notation for predicate abstraction conforms to the standard notation of set theory when predicates double as names of their extensions. Correspondingly, the set $F\!:\!A$, which is $\{x \mid \exists y(Fyx \wedge Ay)\}$, is the extension of the monadic predicate '$F\!:\!A$'. So the predicate operator '$:$' doubles as a name for a two-place operation that takes a binary relation F and set A to the set $F\!:\!A$, the Peirce product of A by F.

(11.4) Remark The identification of functional relations with one-many relations probably bewilders many of the readers, since the converse is common nowadays. (A relation F is said to be *one-many* if no two things bear F to the same thing, i.e. if F satisfies (11.1).) The one-many definition can be found, for example, in the work of Peano, Gödel, Tarski, and Quine. We prefer this "old-fashioned" convention to the "modern" one because it is more natural and technically superior. It is worth to cite Quine (1969, p. 24) on these matters:

> My way (Peano's, Gödel's) is natural in that if we are going to identify functions with relations at all, it is natural to identify the square (or father) function with the square (or father) relation; and certainly the square relation is the relation of square to root, as the father relation is the relation of father to child. [...]
> The reverse way might be said to be more natural on one count: a function transforms the argument into the value, thus leads from the argument *to* the value, and thus is naturally identified with the relation of the argument *to* the value. But the reasoning seems lame to me.

Another pressing argument in favor of the one-many convention runs as follows.[3] The composite $F \circ G$ of two binary relations F and G is standardly defined to be $\{xy \mid \exists z(Fxz \wedge Gzy)\}$. This definition should surely not depend on whether F and G are functional or not. Moreover, concerning functional application, $(F \circ G)(x)$ should be the same as $F(G(x))$. These two conditions together inevitably imply the one-many definition of functionality.[4]

[3] See also Quine *ibid*.

[4] Perhaps the many-one convention has its roots in graphical (or tabular) presentations of argument-value pairs. If arguments correspond to points on the x-axis and values to points on the y-axis, then the standard vector representation of points of the plane leads to pairs with argument and value respectively as first and second component.

Furthermore notice that the one-many convention necessarily leads to defining the Peirce product or inverse image as in (11.2), which is converse to the definition typically given in the literature.[5] This is so because if F is functional then clearly the image of B by F must be $\{y \mid \exists x(Fyx \wedge Bx)\}$; consequently, the inverse image of A by F is $\{x \mid \exists y(Fyx \wedge Ay)\}$.

11.1.2 *Attribute Composition*

The attribute predicates occurring in attributive descriptions can be composed of others. An example is the predicate 'someone whose father's car is a convertible'. It can be paraphrased by:

an x such that the father of x is someone whose car is a convertible,

which is obviously an attributive description whose sort predicate 'someone whose car is a convertible' is an attributive description in turn. With 'F' for '$\{yx \mid y$ is father of $x\}$', 'G' for '$\{yx \mid y$ is car of $x\}$', and 'A' for '$\{x \mid x$ is a convertible$\}$', the attribute-value predicate can be condensed to '$F{:}(G{:}A)$', with 'F' and 'G' assumed as functional.

The translation of '$F{:}(G{:}A)$' into quantificational form by (11.2) allows us to derive a different but equivalent attribute-value predicate; for

$$\{x \mid \exists y(Fyx \wedge \exists z(Gzy \wedge Az))\} \ \equiv\ \{x \mid \exists z(\exists y(Fyx \wedge Gzy) \wedge Az)\}.$$

Consequently, if '$\shortmid$' is the predicate operator of *reverse (relational) composition*,[6] that is,

$$F{\shortmid}G \ \equiv\ \{xy \mid \exists z(Fzy \wedge Gxz)\},$$

then '$F{:}(G{:}A)$' is equivalent to the attribute-value predicate '$(F{\shortmid}G){:}A$'. It thus makes sense to speak of a *composed* attribute predicate '$F{\shortmid}G$'.

The scheme '$(F{\shortmid}G){:}A$' can also be derived by paraphrasing the original expression 'someone whose father's car is a convertible' as follows:

an x such that the car of the father of x is a convertible.

Here the complex singular description 'the car of the father of x' needs further analysis. It is not difficult to see that Russell's elimination method applied twice leads to '$(F{\shortmid}G){:}A$', where 'F' and 'G' and hence '$F{\shortmid}G$' are presupposed to be functional.

[5] Cf. e.g. Brink et al. 1994 or de Rijke 1995.

[6] As it happens, our use of the bar (accidentally) is opposite to the one initiated by Russell, who used it for relational composition (also known as *relative product*).

11.1.3 *Internal and External Negation*

It was already perceived by Russell (1905) that his method of eliminating definite descriptions gives rise to two different readings when applied to negated contexts. Consider the following classical example:

(11.5) an x such that the king of x is not bald,

in symbols: $\{x \mid \neg B(\iota y K y x)\}$. The matrix sentence $\neg B(\iota y K y x)$ has two possible readings depending on whether the negation sign takes narrow scope as in '$(\neg B)(\iota y K y x)$' or wide scope as in '$\neg(B(\iota y K y x))$'. This is the distinction between *term* or *internal negation* and *predicate denial* or *external negation*. Application of Russel's elimination method (cf. Section 11.1.1) leads respectively to

$$\exists z(z = \iota y K y x \,\wedge\, \neg B z) \quad \text{and} \quad \neg\exists z(z = \iota y K y x \,\wedge\, B z),$$

and further to

$$\exists y(Kyx \wedge \neg By) \wedge \forall yz(Kyx \wedge Kzx \rightarrow y = z) \quad \text{and}$$
$$\neg\exists y(Kyx \wedge By) \vee \neg\forall yz(Kyx \wedge Kzx \rightarrow y = z).$$

Suppose K is functional. Then the second clauses of the previous two expressions are redundant and, hence, the two readings of the attribute-value predicate (11.5) come down to '$K{:}\neg B$' and '$\neg(K{:}B)$'. Furthermore,

(11.6) $$K{:}\neg B \;\equiv\; K{:}\mathrm{V} \,\wedge\, \neg(K{:}B).$$

Proof. $\neg\exists y(Kyx \wedge By)$ iff $\forall y(Kyx \rightarrow \neg By)$. Moreover, since K is functional, $\exists y(Kyx \wedge \neg By)$ iff $\exists y Kyx \wedge \forall y(Kyx \rightarrow \neg By)$. ⌟

The relation between internal and external negation as expressed by (11.6) nicely illustrates the existential presupposition implicit in internal negation.

11.1.4 *Proper Names*

When attributive descriptions are employed in practice, say, in records of employees, attribute values are denoted by *proper names* more often than not. The attribute 'birthplace' provides a typical example. If predicates like 'someone whose birthplace is London' are to put into the Procrustes bed of attribute-value predicates, then proper names, here 'London', have to be reanalyzed as sort predicates in one way or another. A standard method is to consider names as definite descriptions and to eliminate the latter in the usual way.

Let us work this out for the given example, whose regimented form is 'an x such that the birthplace of x is London'. Notice that the 'is' in 'the birthplace of x is London' is not the copula of predication but that of identity. So, instead of (11.2), we get $\{x \mid \exists y(Fyx \wedge y = \text{London})\}$, with '$F$' for '$\{yx \mid y$ is birthplace of $x\}$'. The trick is now to consider 'London' as a monadic predicate and to replace '$y = \text{London}$' by '$y = \iota\,\text{London}$'.[7] Elimination of the definite description leads to $Ny \wedge \forall z(Nz \rightarrow y = z)$, where '$N$' stands for 'London'. Hence, if we presuppose that 'N' denotes at most one entity, i.e.

(11.7) $$\forall xy(Nx \wedge Ny \rightarrow x = y),$$

the example predicate can be expressed in the standard form '$F{:}N$'. As an immediate consequence of (11.7), it follows that

$$F{:}A \;\wedge\; F{:}N \;\wedge\; G{:}N \;\preceq\; G{:}A$$

is a valid statement for arbitrary dyadic F and G and monadic A.

11.2 Two Problematic Cases

The analysis of attributive descriptions developed so far rests on the assumption that attributes are binary functional relations. This is surely true for attribute terms like 'father'. But our first examples in Chapter 10, namely binary features and choice features, turn out to be rather problematic in this respect.

11.2.1 *Binary Features*

Let us consider binary features first. Taking for example HUMAN:$+$ as an attribute-value predicate in the sense of Section 11.1.1 means to assume that HUMAN and $+$ are respectively dyadic and monadic predicates. Then, by definition, an entity x of the universe of discourse satisfies HUMAN:$+$ just in case there is a (unique) entity y such that $\langle y, x\rangle$ satisfies HUMAN and y satisfies $+$. The reader will presumably find this rather awkward: not only are there $+$'s but they also bear a functional relation named HUMAN to human beings. No less awkward is the straighforward reformulation of HUMAN:$+$ into natural language along the lines of Section 11.1.1: the expression 'something whose human being is a $+$' is hardly intelligible.

It is nevertheless possible to treat terms of this type as attribute-value predicates. The following construction presumes that sentences denote *truth values*, say $\mathbf{t}$ or $\mathbf{f}$, and makes use of *functional abstraction*. Let HUMAN denote the function $\lambda x(x$ is a human being$)$, which takes x to $\mathbf{t}$ if x is a human being and

[7] See Quine 1960, §§37–39 for a more detailed exposition.

to $\mathbf{f}$ otherwise.[8] Assume now that $\mathbf{t}$ is the only entity satisfying the predicate $+$ (whereas $\mathbf{f}$ is the only entity satisfying $-$). Then x satisfies HUMAN:$+$ iff there is a y such that $y = (x \text{ is a human being})$ and y satisfies $+$, i.e. is identical to $\mathbf{t}$. In short, x satisfies HUMAN:$+$ iff x is a human being.

Though viable in principle, the presented solution has two disadvantages, of which one is conceptual and the other is technical in nature. First, one may criticize the reification of truth values since "it is obviously desirable to analyze discourse in such a way as not to impute special ontological presuppositions to portions of discourse which are innocent of them."[9] But the decision to introduce predicates $+$ and $-$ implies that models contain "truth value entities". Second, $+$ and $-$ are implicitly assumed to denote singleton sets. In other words, they are subject to the presupposition (11.7).

There is a simple way to abandon sort predicates for truth values altogether. Simply use two sort predicates *human* and *non-human* in place of HUMAN:$+$ and HUMAN:$-$. In addition, one has to ensure axiomatically that *human* and *non-human* are incompatible. Though incompatibility is needed in the case of $+$ and $-$ too, one might object that stating this explicitly for *human* and *non-human* lacks generality; for it has also to be done for every other such pair. Alternatively, we could make negation syntactically transparent and thereby accessible to inference schemes. Within the realm of observational logic, this strategy has to be qualified insofar as negation is unproblematic in statements, whereas predicate negation in positive assertions has to be "incorporated" into the predicate (see Sections 5.4.1 and 12.1.4).

11.2.2 *Choice Features*

The ambivalence between choice features and choice systems has already been pointed out in Section 10.1.2. There we observed that e.g. CASE is best seen as a choice system, i.e. as a set consisting of the monadic predicates *nom*, *acc*, etc. Suppose we want CASE:*acc* nevertheless to work as an attribute-value predicate. Intuitively, a (linguistic) entity x should satisfy CASE:*acc* if x is accusative, where accusative is a case. On the other hand, by definition (11.2), x satisfies CASE:*acc* if there is a y such that $\langle y, x \rangle$ satisfies CASE and y satisfies *acc*. To bring the latter analysis in line with the former, let nom be the set $\{x \mid x \text{ is nominative}\}$ of all nominative (linguistic) entities, acc the set of all accusative entities, etc; in addition, let *acc* be the predicate '$\{y \mid y = acc\}$' and let CASE be '$\{yx \mid x \in y \in Case\}$', where $Case$ is $\{nom, acc, \ldots\}$. Then CASE is functional – at least if no entity has more than one case, i.e. if the members of $Case$ are pairwise disjoint; and x satisfies CASE:*acc* iff $x \in acc \in Case$, as desired.

---

[8]The λ-notation for functional abstraction is Curry's; the idea to take sentences as names for truth values goes back to Frege; cf. e.g. Quine 1976, Sect. II.

[9]Quine 1953, pp. 115f

So, in order to reconcile the specification of choice features with attributive descriptions, we introduced *sets of entities* as entities into our universe of discourse.[10] This decision is neither good nor bad from the viewpoint of formal theory. But at least, it should be made consciously when formalizing the discourse in question. For ontic commitments are coupled to predicates and thus to extensions, be it in the domain of discourse or in a formal model of the latter. Moreover notice that we again made use of sort predicates that denote singletons.

Let us briefly indicate how the given formalization of choice feature specifications is directly derivable from their natural language counterparts. Consider the attribute-value predicate ‘something whose color is red’. (For simplicity, assume that everything in the universe of discourse is either red or not.) Applying the ‘such that’ paraphrase leads to ‘an x such that the color of x is red’. Observe that the ‘is’ in ‘the color of x is red’ means not predication but identity; for the color of x is not something red but is identical to red. Hence we get $\{x \mid \exists y(Fyx \wedge y = red)\}$ instead of (11.2), where ‘F’ stands for ‘$\{yx \mid y$ is a color of $x\}$’. With colors interpreted extensionally, red is the set of red things and ‘F’ can be replaced by ‘$\{yx \mid x \in y$ and y is a color$\}$’. This is precisely the type of analysis used above for ‘something whose case is accusative’. The presupposition of functionality here means that for every two colors y and z, if $x \in y \cap z$ then $y = z$; in short, everything has at most one color.

(11.8) Remark The difficulty of handling choice features within the realm of attributive descriptions has also been noticed by Marcus Kracht (1995, p. 448):

> “[...] even though it is optically pleasing to regard CASE as intrinsically relational, the initial attraction results from a superficial linguistic analogy of $\langle\text{CASE}\rangle\mathsf{acc}$ with *the case of this item is accusative*, thus suggesting that there is an underlying relation. Yet, the expression *my car is red* does not indicate that the expression *car* is relational, neither does *this item’s case is accusative* lead to the conclusion that *case* denotes a relation. There are only cars and in particular red ones, and there are cases, in particular accusative.”

Though Kracht’s conclusion is right for the most part, his argument is not. The attribute-value predicate in question is ‘someone whose car is red’, or, more formally, ‘an x such that the car of x is red’. This expression fits perfectly into the standard formalization scheme for attribute-value predicates: take the Peirce product of monadic predicate ‘red’ by the dyadic predicate ‘$\{yx \mid y$ is a car of $x\}$’ and assume functionality, i.e. assume people to own at most one car. So, though Kracht is right in saying that ‘car’ is not relational, this is not

[10] Alternatively, one can employ semantic ascent: Instead of using the sets nom, acc, etc, one can use their *names*, i.e. the predicates they are the extensions of. Then $Case$ is a set of predicates and CASE is ‘$\{yx \mid x$ satisfies $y \in Case\}$’.

at issue here. The relational component of ‘is a car of’ is that of *ownership*. Kracht’s example is thus not well suited for pointing out the problems of using choice features in attribute-value predicates.

11.3 The Logic of Attributive Descriptions

In this section we study more closely equivalences between attribute-value predicates and, in particular, how the Peirce product interacts with Λ, disjunction and conjunction. Equivalence of attribute-value predicates here means equivalence of their translations into quantificational form, under the condition that attributes are functional.

Let us start with equivalences that do not presuppose functionality. The following statements are logically valid:

(11.9) $$F{:}\Lambda \;\equiv\; \Lambda.$$

Proof. $\forall x(\exists y(Fyx \wedge \neg(y=y)) \leftrightarrow \neg(x=x))$ iff $\forall x(\neg\exists y(Fyx \wedge \neg(y=y)))$ iff $\forall xy(\neg Fxy \vee y=y)$, which is a tautology. ⌟

(11.10) $$F{:}(A \vee B) \;\equiv\; F{:}A \vee F{:}B.$$

Proof. $\forall x(\exists y(Fyx \wedge (Ay \vee By)) \leftrightarrow \exists y(Fyx \wedge Ay) \vee \exists y(Fyx \wedge By))$. ⌟

(11.11) Proposition The *prefixing rule* (P) is a sound inference scheme.

$$\frac{A \equiv B}{F{:}A \equiv F{:}B} \qquad \text{(P)}$$

Proof. If $\forall y(Ay \leftrightarrow By)$ then $\forall x(\exists y(Fyx \wedge Ay) \leftrightarrow \exists y(Fyx \wedge By))$. ⌟

For the next statement to be valid, it is essential that F is functional; its straightforward proof is left to the reader.

(11.12) $$F{:}(A \wedge B) \;\equiv\; F{:}A \wedge F{:}B.$$

In addition, all axioms and inference rules of observational logic remain valid for attribute-value predicates (cf. Section 6.3). Needless to say that the valid statements (11.9) to (11.12) have not been picked out arbitrarily. In fact, together with the calculus $OC_{\equiv}$ of observational logic they constitute a strongly complete calculus for attribute-value statements (see Section 12.2.4). Notice that the soundness of that calculus has just been shown.

We close this section by stating two straightforward logical equivalences that involve attribute composition:

$$F|(G|H) \;\equiv\; (F|G)|H, \qquad F{:}(G{:}A) \;\equiv\; (F|G){:}A.$$

12

Attribute-Value Theories

Attribute-value statements and theories over a given set of attribute and sort predicates are formally introduced in Section 12.1. Since we take attribute-value statements to stand for statements of first-order predicate logic, the latter provide the former with the standard definition of interpretation and model. So-called *feature systems* of an attribute-value theory are defined as first-order models of that theory, with attributes interpreted by functional relations.

In Section 12.2, we generalize the algebraic approach of Chapter 6 from observational theories to attribute-value theories. The appropriate algebraic concept is that of a *feature algebra*, which is an observational algebra together with operations that preserve 0, $\wedge$, and $\vee$. The fact that each attribute-value theory has a universal model in a feature algebra is employed to prove the completeness of the attribute-value calculus sketched in Section 11.3. Moreover, the universal model gives rise to a canonical feature system of the given theory, which is the analogue of the canonical model of an observational theory.

Section 12.3 describes the canonical feature system in more detail. Its universe can be identified with the set of *term feature trees* of the theory in question. In order to characterize the specialization ordering on the canonical universe, attribute-value theories are expressed as observational theories, which allows us to apply the results of Part I and II. The final Section 12.4 addresses possible extensions and topics for future research.

12.1 Attribute-Value Theories and Feature Systems

12.1.1 *Term Algebra*

Let L be a set of (primitive) attribute predicates and S be a set of sort predicates (terms, symbols, etc), where S does not contain V or Λ. The pair $\langle L, S\rangle$ is also called an *attribute-value* or *feature signature*.

The set $\mathrm{T}[L, S]$ of *attribute-value predicates over* $\langle L, S\rangle$ is inductively defined as follows: V, Λ and all members of S belong to $\mathrm{T}[L, S]$, as well as

$$\varphi \wedge \psi, \qquad \varphi \vee \psi, \qquad \text{and} \qquad l{:}\varphi$$

whenever $\varphi, \psi \in \mathrm{T}[L, S]$ and $l \in L$. We call an attribute-value predicate *primitive* if it belongs to S or is of the form $l{:}\varphi$, where φ is primitive or Λ or V. Let $\Sigma[L, S]$ be the set of primitive attribute-value predicates.[1]

There are different ways to give $\mathrm{T}[L, S]$ an algebraic structure. For example, one can take $:$ as an operation, which leads to a two-sorted algebra. Here we prefer the view that each $l \in L$ determines a one-place operation o_l on $\mathrm{T}[L, S]$ that takes φ to $l{:}\varphi$. In addition, $\mathrm{T}[L, S]$ is equipped with the two-place operations $\wedge$ and $\vee$ and the zero-place operations Λ and V (see Section 5.1.2). A standard argument from universal algebra ensures the existence of *unique homomorphic extensions*:

(12.1) Proposition If f is a function from S to an algebra A of the same type as $\mathrm{T}[L, S]$, then there is a unique algebra homomorphism $\hat{f}$ from $\mathrm{T}[L, S]$ to A such that $\hat{f}(s) = f(s)$ for all $s \in S$.

The lack of an explicit operator for attribute composition can be compensated as follows. Define $p{:}\varphi$ inductively such that

$$\epsilon{:}\varphi \;=\; \varphi \qquad \text{and} \qquad pl{:}\varphi \;=\; p{:}(l{:}\varphi),$$

for all $p \in L^*$, $l \in L$, and $\varphi \in \mathrm{T}[L, S]$.[2] It then follows that $pq{:}\varphi \;=\; p{:}(q{:}\varphi)$ for all $p, q \in L^*$.

(12.2) Remark *(Path Terms)* It is of course also possible to give a definition of attribute-value predicates that explicitly takes attribute composition into account. Define the set $\mathrm{P}[L]$ of *composed attributes* or *(attribute) paths over* L by requiring that $L \subseteq \mathrm{P}[L]$ and that $p{\mid}l$ belongs to $\mathrm{P}[L]$ whenever $p \in \mathrm{P}[L]$ and $l \in L$. The set $\mathrm{T}[L, S]$ of attribute-value predicates over $\langle L, S\rangle$ is then the least superset of $S \cup \{\mathrm{V}, \Lambda\}$ such that $\varphi \wedge \psi$, $\varphi \vee \psi$, and $p{:}\varphi$ belong to $\mathrm{T}[L, S]$ whenever $\varphi, \psi \in \mathrm{T}[L, S]$ and $p \in \mathrm{P}[L]$. Notice that $p{\mid}l{:}\varphi$ and $p{:}(l{:}\varphi)$ are different terms under this construction. Viewed algebraically, $\mathrm{T}[L, S]$ carries the structure of an algebra with operations $\wedge$, $\vee$, Λ, and V. In addition, the operators $:$ and $\mid$ give rise to operations from $\mathrm{P}[L] \times L$ to $\mathrm{P}[L]$ and from $\mathrm{P}[L] \times \mathrm{T}[L, S]$ to $\mathrm{T}[L, S]$, respectively.

[1] We do not count V and Λ as primitive in order to be compatible with the terminology of Chapter 5.

[2] L^* is the free monoid (the set of finite strings) over L with unit (empty string) ϵ.

12.1.2 *Theories and Systems*

Let $\langle L, S\rangle$ be an attribute-value signature. Since attribute predicates are dyadic and sort predicates are monadic, an *interpretation* of L and S in the usual sense of predicate logic consists of a *universe* U and an *interpretation function* M that takes each element of L to a binary relation on U and each element of S to a subset of U. Because attributes are assumed to be functional, we are interested in models of the (first-order) theory F (or F_L) consisting of the statements

$$\textbf{\textit{(12.3)}} \qquad \forall xyz(lxz \wedge lyz \rightarrow x = y),$$

for every $l \in L$.[3] Following Rounds (1997) we call such models *feature systems* over $\langle L, S\rangle$.[4]

Let $\langle M, U\rangle$ be a feature system over $\langle L, S\rangle$. By employing the standard first-order interpretation of logical connectives and quantifiers, the interpretation function M can be extended to all attribute-value terms over $\langle L, S\rangle$. So M takes the Peirce product $l{:}\varphi$, which stands for $\{x \mid \exists y(lyx \wedge \varphi y)\}$, to the Peirce product of sets

$$\textbf{\textit{(12.4)}} \qquad M(l){:}M(\varphi) \;=\; \{x \mid \exists y(\langle y, x\rangle \in M(l) \wedge y \in M(\varphi))\}.$$

The interpretation of Λ, V, $\varphi \wedge \psi$ and $\varphi \vee \psi$ is that of Chapter 5: $M(\Lambda) = \varnothing$, $M(\mathrm{V}) = U$, $M(\varphi \wedge \psi) = M(\varphi) \cap M(\psi)$, and $M(\varphi \vee \psi) = M(\varphi) \cup M(\psi)$.

In addition, it is suitable to extend the interpretation M to all $p \in L^*$ such that $M(\epsilon) = \{xy \mid x = y\}$ and

$$M(pl) \;=\; M(p) \mid M(l) \;=\; \{xy \mid \exists z(\langle z, y\rangle \in M(p) \wedge \langle x, z\rangle \in M(l))\}.$$

By straightforward induction, it follows that $M(p{:}\varphi) = M(p){:}M(\varphi)$, for all $p \in L^*$ and $\varphi \in \mathrm{T}[L, S]$. That is, $x \in M(p{:}\varphi)$ iff there is a (necessarily unique) y such that $\langle y, x\rangle \in M(p)$ and $y \in M(\varphi)$. For the sake of notational parsimony we also write '$p \cdot x$' instead of 'the y such that $\langle y, x\rangle \in M(p)$' (thereby reviving to some extent the definite description eliminated in Section 11.1.1). This convention allows us, for instance, to write:

$$\textbf{\textit{(12.5)}} \qquad M(p{:}\varphi) \;=\; \{x \mid p \cdot x \in M(\varphi)\}.$$

[3] The theory F is a natural starting point for approaches to feature logic within the realm of first-order logic; see e.g. Smolka 1992, Aït-Kaci et al. 1994.

[4] Smolka (1992) speaks of *feature algebras* instead, which differs from our usage of this term in Section 12.2 below.

Furthermore notice that $pq \cdot x = q \cdot (p \cdot x)$, for $p, q \in L^*$.

Every feature system $\langle M, U\rangle$ over $\langle L, S\rangle$ determines a *satisfaction relation* $\models$ from U to $\mathrm{T}[L, S]$ with $x \models \varphi$ iff $x \in M(\varphi)$. As in the observational case (see Section 5.1.4), one can define a *specialization relation* $\sqsubseteq$ on U such that

$$x \sqsubseteq y \qquad \text{iff} \qquad \forall \varphi \in \mathrm{T}[L, S]\,(x \models \varphi \;\rightarrow\; y \models \varphi).$$

Recall that M is said to satisfy *identity of indiscernibles* in case $\sqsubseteq$ is antisymmetric.

An *attribute-value statement* over $\langle L, S\rangle$ is a statement of the form $\varphi \preceq \psi$, where φ and ψ are attribute-value predicates over $\langle L, S\rangle$. An *attribute-value theory* Γ over $\langle L, S\rangle$ is a set of such statements. By a *feature system of* Γ we mean a (first-order) model of $\mathrm{F} \cup \Gamma$.

(12.6) Example Suppose $L = \{\textsc{f}, \textsc{g}\}$ and $S = \{a, b, c\}$. Let Γ be the attribute-value theory over $\langle L, S\rangle$ consisting of the statements

$$a \wedge b \preceq \Lambda, \qquad c \preceq b, \qquad a \preceq \textsc{f}{:}b, \qquad \textsc{f}{:}\mathrm{V} \preceq a.$$

Figure 76 shows several examples of feature systems of Γ depicted in form of directed graphs. The nodes of a graph represent the members of the universe U; if a node is labeled with a subset X of S, the respective element of U satisfies every member of X; if an arc is labeled with an attribute symbol l, then $y = l \cdot x$, where x and y are the members of U represented by the source and the target node of the arc, respectively. Observe that feature systems of Γ need not be finite (d) nor acyclic (e) nor rooted (f) nor connected (g). Moreover notice that the feature systems (b), (d), and (f) do not satisfy identity of indiscernibles. Figure 77 shows three interpretations of $\langle L, S\rangle$ that are *not* feature systems of Γ. Example (a) is not a feature system at all since $\textsc{f}$ is interpreted by a non-functional relation, whereas (b) and (c) are feature systems but not feature systems of Γ.

A feature system over $\langle L, S\rangle$ with universe U is a *feature tree* if there is unique $r \in U$, called the *root*, such that for every $x \in U$ there is a unique $p \in L^*$ with $x = p \cdot r$. In Figure 76, for instance, the feature systems (a), (b), and (d) are feature trees. We shall see below that the feature trees of an attribute-value theory Γ provide a complete set of feature systems of Γ in the sense that an attribute-value statement is deducible from Γ modulo F_L just in case the statement is true with respect to all feature trees of Γ.[5]

[5] An attribute-value statement is said to be *deducible from* Γ *modulo* F if it is deducible from $\mathrm{F} \cup \Gamma$. Two attribute-value predicates φ and ψ are said to be *equivalent modulo* F if F entails $\varphi \equiv \psi$.

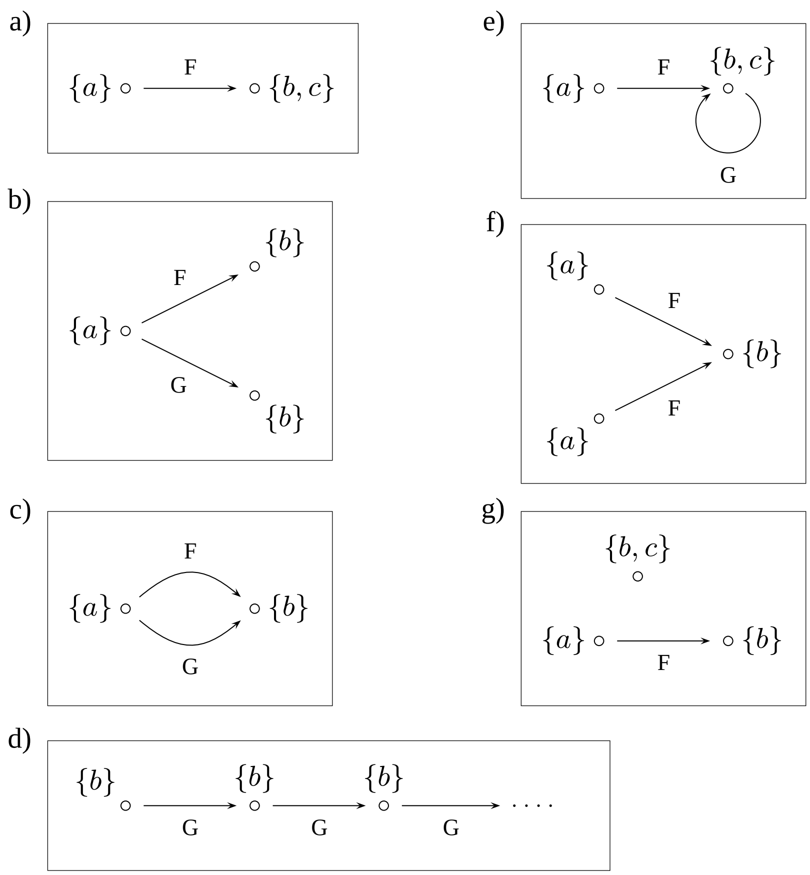

FIGURE 76 Examples of feature systems of Γ

(12.7) Example *(Term feature trees)* Let U be a nonempty, prefix-closed subset of L^*; that is, $\epsilon \in U$ and if $pl \in U$ then $p \in U$, for all $l \in L$ and $p \in L^*$. Suppose $M(l) = \{\langle pl, p\rangle \mid pl \in U\}$ and $M(s) \subseteq U$, for all $l \in L$ and $s \in S$. Clearly M is a feature tree over $\langle L, S\rangle$ with universe U and root ϵ. We refer to feature trees of this form as *term feature trees* over $\langle L, S\rangle$.[6]

A *homomorphism f of feature systems* over $\langle L, S\rangle$ from M to M' is a homomorphism of interpretations; that is, for all $s \in S$, $l \in L$, and members x, y of the universe U of M,

$$\text{if} \quad x \models s \quad \text{then} \quad f(x) \models s, \quad \text{and}$$

$$\text{if} \quad y = l \cdot x \quad \text{then} \quad f(y) = l \cdot f(x).$$[7]

[6] Put differently, a term feature tree over $\langle L, S\rangle$ is a partial function from L^* to $\wp(S)$ with a nonempty, prefix-closed domain of definition.

[7] That is, if $x \in M(s)$ then $f(x) \in M'(s)$, and if $\langle y, x\rangle \in M(l)$ then $\langle f(y), f(x)\rangle \in M'(l)$.

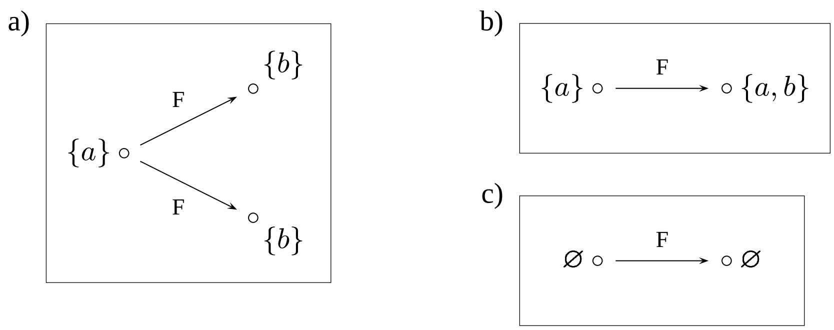

FIGURE 77 Interpretations of $\langle L, S\rangle$ that are not feature systems of Γ

By term induction, it follows that, for every $\varphi \in \mathrm{T}[L, S]$,

(12.8)
$$f(M(\varphi)) \subseteq M'(\varphi).$$

Suppose $\langle M, U\rangle$ is a feature tree over $\langle L, S\rangle$ with root r. Then the set $U' = \{p \mid p \cdot r \in U\}$ is nonempty and prefix-closed. Moreover, $p = q$ whenever $p \cdot r = q \cdot r$. So the function from U' to U that takes p to $p \cdot r$ is one-to-one and onto. With $M'(l) = \{\langle pl, p\rangle \mid pl \in U'\}$ and $M'(s) = \{p \mid p \cdot r \models s\}$, for $l \in L$ and $s \in S$, we get thus a term feature tree $\langle M', U'\rangle$ that is isomorphic to $\langle M, U\rangle$; in short:

(12.9) Proposition Every feature tree over $\langle L, S\rangle$ is isomorphic to a term feature tree over $\langle L, S\rangle$.

For example, the term feature trees isomorphic to the feature trees (b) and (d) of Figure 76 have universe $\{\epsilon, \mathrm{F}, \mathrm{G}\}$ and G^*, respectively.

12.1.3 *Normal Forms*

We say that an attribute-value predicate over $\langle L, S\rangle$ has *observational (normal) form* if it is an observational predicate over the set $\Sigma[L, S]$ of primitive attribute-value predicates, i.e., if it is built of primitive attribute-value predicates plus V and Λ by finite conjunction and disjunction. Inductive application of (11.10) and (11.12) shows:

(12.10) Proposition Every attribute-value predicate over $\langle L, S\rangle$ is logically equivalent modulo F_L to an observational predicate over the set $\Sigma[L, S]$ of primitive attribute-value predicates.

Consequently, by (5.4), every attribute-value predicate is equivalent to one in disjunctive (or conjunctive) normal form.

Instead of restructuring attribute-value terms to the effect that no $\wedge$ or $\vee$ is inside the scope of $:$, one can also proceed the other way around and extract common prefixes as far as possible. To keep things simple, let us consider only attribute-value predicates without disjunction. One easily checks that every disjunction-free attribute-value predicate is equivalent to one of the form $\varphi_1 \wedge \ldots \wedge \varphi_n$, where each φ_i either belongs to $S \cup \{\mathrm{V}, \Lambda\}$ or is of the form $l_i : \psi_i$ such that the l_i's are all different, and the ψ_i's are conjunctive predicates of the described form in turn. The resulting normal form is closely related to the matrix notation of Section 10.1, which also rests on the extraction of common prefixes. Vice versa, reading the matrix "line by line" returns the corresponding conjunction of primitive attribute-value predicates (with the exception that these terms may still contain conjunctions of sort predicates).

(12.11) Example Consider the term $\textsc{f}:(\textsc{g}:a) \wedge \textsc{f}:(\textsc{g}:b) \wedge \textsc{g}:(\textsc{f}:b \wedge \textsc{h}:c)$. It is equivalent (modulo functionality) to the conjunction of primitive terms

$$\textsc{f}:(\textsc{g}:a) \wedge \textsc{f}:(\textsc{g}:b) \wedge \textsc{g}:(\textsc{f}:b) \wedge \textsc{g}:(\textsc{h}:c),$$

whereas extracting common attributes gives $\textsc{f}:(\textsc{g}:(a \wedge b)) \wedge \textsc{g}:(\textsc{f}:b \wedge \textsc{h}:c)$. Compare these terms with the following matrix:

$$\begin{bmatrix} \textsc{f} \mid \textsc{g} & a \wedge b \\ \textsc{g} & \begin{bmatrix} \textsc{f} & b \\ \textsc{h} & c \end{bmatrix} \end{bmatrix}$$

12.1.4 *Negation*

We saw in Section 5.4.1 that negation as a predicate operator does not affect the logical expressivity of observational statements. Here we shall see that the same holds for attribute-value statements. Suppose the predicate operators $\neg$ and $\rightarrow$ are added to the term constructors for attribute-value predicates. Let us refer to the resulting monadic predicates briefly as *Boolean predicates* over $\langle L, S \rangle$ (thereby adapting the terminology of Section 5.4.1). A *universal statement* over $\langle L, S \rangle$ is a statement of the form $\forall \alpha$, where α is a Boolean predicate over $\langle L, S \rangle$. We claim that every such statement is logically equivalent modulo F_L to a finite conjunction of attribute-value statements over $\langle L, S \rangle$.

Every Boolean predicate can be transformed into a conjunctive normal form that consists of primitive attribute-value predicates or their negation. To see this, notice that internal negation can be eliminated by applying (11.6), here repeated as

(12.12) $$F:\neg A \;\equiv\; F:\mathrm{V} \;\wedge\; \neg(F:A),$$

which holds for functional F. The remaining transformation into a finite conjunction of attribute-value statements is completely analogous to that in the proof of (5.50). To summarize:

(12.13) Proposition Every universal statement over $\langle L, S\rangle$ is equivalent modulo F_L to a finite conjunction of attribute-value statements over $\langle L, S\rangle$.

(12.14) Example Consider the statement $\varphi \preceq \textsc{person}{:}\neg \textit{third}$, where φ is an arbitrary attribute-value predicate over a given feature signature. Since $\varphi \preceq \neg\psi$ can be replaced by $\varphi \wedge \psi \preceq \Lambda$, it follows by (12.12) and elimination of conjunction that the statement under consideration is equivalent (modulo F) to the conjunction of the attribute-value statements

(12.15) $\varphi \preceq \textsc{person}{:}\mathrm{V}$ and $\varphi \wedge \textsc{person}{:}\textit{third} \preceq \Lambda$.

Now suppose we are working in an attribute-value theory Γ containing the statements

$$\textsc{person}{:}\mathrm{V} \preceq \textsc{person}{:}\textit{person}, \quad \textit{person} \equiv \textit{first} \vee \textit{second} \vee \textit{third},$$
$$\textit{first} \wedge \textit{second} \equiv \textit{first} \wedge \textit{third} \equiv \textit{second} \wedge \textit{third} \equiv \Lambda,$$

which provides a typical example of a choice system in the guise of feature logic. Then the predicates $\textsc{person}{:}\neg\textit{third}$ and $\textsc{person}{:}(\textit{first} \vee \textit{second})$ are first-order equivalent modulo $\mathrm{F} \cup \Gamma$, as one easily verifies.

The use of negation in the setting of (12.14) is sometimes referred to as “abbreviatory” (Carpenter 1992, p. 111). This is in accordance with the view presented here, as long as negation in statements is the issue. If, in contrast, negated attribute-value predicates are to be used in positive assertions, then some form of term negation in the sense of Section 5.4.2 has to be employed.

12.2 Models in Feature Algebras

Feature algebras are to attribute-value theories what observational algebras are to observational theories. The canonical feature algebra of an attribute-value theory, to be defined in Section 12.2.3, gives us a complete calculus for attribute-value statements via its construction and a generic feature system via its prime spectrum representation. The following two sections provide the necessary background.

12.2.1 *Feature Algebras*

Let A be an observational algebra, i.e. a distributive lattice with zero and unit. We call a one-place operation o on A a 0-*homomorphism* on A, if o preserves 0, $\wedge$, and $\vee$, that is, if

$$o(0) = 0, \qquad o(a \wedge b) = o(a) \wedge o(b), \qquad o(a \vee b) = o(a) \vee o(b).$$

A 0-homomorphism thus differs from a homomorphism of observational algebras in that it does not necessarily preserve the unit.

(12.16) Example *(Peirce operation)* Let R be a functional relation on a set U. Then the operation on $\wp(U)$ that takes $V \subseteq U$ to the Peirce product

$$R{:}V = \{x \mid \exists y(\langle y, x\rangle \in R \wedge y \in V)\}$$

of V by R is a 0-homomorphism on the powerset algebra $\langle \wp(U), \cap, \cup, \varnothing, U\rangle$. This is so because the formulas (11.9), (11.10), and (11.12) are valid. A 0-homomorphism given in this way on an observational algebra of sets is henceforth referred to as an *inverse image* or *Peirce operation*.

By a *feature algebra A over L* we mean an observational algebra together with a family $\langle o_l\rangle_{l \in L}$ of 0-homomorphisms on A.[8] Instead of $o_l(a)$ we also write $l{:}a$, thereby indicating that a feature algebra over L can be seen as an observational algebra A together with a function $:$ from $L \times A$ to A such that, for all $l \in L$ and $a, b \in A$, $l{:}0 = 0$, $l{:}(a \wedge b) = l{:}a \wedge l{:}b$, and $l{:}(a \vee b) = l{:}a \vee l{:}b$. Furthermore, the function $:$ can be inductively extended to one from $L^* \times A$ to A such that, for all $p \in L^*$, $l \in L$, and $a \in A$,

$$\epsilon{:}a = a \qquad \text{and} \qquad pl{:}a = p{:}(l{:}a).$$

One easily verifies by induction that, for all $p \in L^*$,

$$p{:}0 = 0, \qquad p{:}(a \wedge b) = p{:}a \wedge p{:}b, \qquad p{:}(a \vee b) = p{:}a \vee p{:}b.$$

[8] The corresponding notion in modal logic is that of a *multi-modal algebra*; cf. e.g. Kracht 1999 or Blackburn et al. 2001. In standard modal logic, negation is included whereas functionality is usually not presumed. Consequently, Boolean algebras take the place of observational algebras and the (existential modal) operations need not preserve conjunction. (The modal logic enforcing functionality is the multi-modal version of $\mathbf{K.Alt}_1$; see e.g. Kracht 1995.)

In case a feature algebra is an observational algebra of sets over a set U and its 0-homomorphisms are Peirce operations, we speak of a *complex feature algebra* over U. A complex feature algebra over U is called *full*, if its underlying observational algebra is the powerset algebra $\wp(U)$ over U.[9]

(12.17) Example *(Extension algebra)* Complex feature algebras naturally arise from feature systems. Suppose Γ is an attribute-value theory over $\langle L, S\rangle$ and $\langle M, U\rangle$ is a feature system of Γ. The set of extensions

$$\Omega_M = \{M(\varphi) \mid \varphi \in \mathrm{T}[L, S]\}$$

carries the structure of an observational algebra of sets over U; cf. (6.1). For each $l \in L$, let o_l be the Peirce operation given by the functional relation $M(l)$ on U, i.e. $o_l(V) = M(l){:}V$; cf. (12.16). Since $M(l){:}M(\varphi) = M(l{:}\varphi)$, o_l can be restricted to an operation on Ω_M. So the observational algebra Ω_M together with the operations $\langle o_l\rangle_{l\in L}$ is a complex feature algebra over L, and M is a homomorphism from $\mathrm{T}[L, S]$ to Ω_M.

(12.18) Example *(Feature tree algebra)* Let U be a nonempty, prefix-closed subset of L^*; cf. (12.7). Then $R_l = \{\langle pl, p\rangle \mid pl \in U\}$ is a functional relation on U, for every $l \in L$. Let $\mathcal{T}$ be the full complex feature algebra with universe U and Peirce operations determined by the relations R_l (that is, $V \subseteq U$ is taken to $R_l{:}V = \{p \mid pl \in V\}$). Feature algebras of this form will be referred to as *feature tree algebras* over L.

Homomorphisms of feature algebras are defined in accordance with the general principles of universal algebra as functions that preserve all operations. That is, if A and B are feature algebras over L, a *homomorphism of feature algebras over* L from A to B is a function h of carrier sets that preserves 0, 1, $\wedge$, $\vee$, and o_l for every $l \in L$; in particular, $h \circ o_l = o_l \circ h$.

Let A be a feature algebra over L and ε the inclusion of a subset S of A into A. Then A is *generated* by S (as a feature algebra over L) if the homomorphic extension $\hat{\varepsilon}$ from $\mathrm{T}[L, S]$ to A is onto. The feature algebra A (over L) is *freely generated* by S, if every function from S to a feature algebra B over L factors uniquely by ε and a homomorphism from A to B.

Consider finally the case that A_0 is an observational algebra which is a subalgebra of A (as an observational algebra); that is, the inclusion function ε from A_0 into A is a homomorphism of observational algebras. Then A is said to be *freely generated* (as a feature algebra over L) by (the observational

[9] This terminology is in accordance with that of modal logic; cf. e.g. Blackburn et al. 2001. According to Goldblatt 1989, 2001, it goes back to an old usage of ‘complex’ by Frobenius in the 1880’s.

algebra) A_0 if every homomorphism (of observational algebras) from A_0 to a feature algebra B over L factors uniquely by ε and a homomorphism of feature algebras from A to B.

12.2.2 *The Prime Spectrum*

We have seen in Section 6.2.1 that every observational algebra A is isomorphic to an observational algebra of sets. Concretely, the function that takes $a \in A$ to the set $\mathcal{P}(a)$ of prime filters of A with member a is an embedding of A into the powerset algebra $\wp(\mathrm{P}(A))$ over the set $\mathrm{P}(A)$ of prime filters of A. This result will now be generalized to feature algebras: every feature algebra is embeddable into a full complex feature algebra.

It suffices to show that every 0-homomorphism on A can be realized as a Peirce operation on $\{\mathcal{P}(a) \mid a \in A\}$. More precisely, if o is a 0-homomorphism on A, we ask for a functional relation $\mathcal{R}$ on $\mathrm{P}(A)$ such that $\mathcal{P}(o(a)) = \mathcal{R}{:}\mathcal{P}(a)$. One can choose

(12.19) $$\mathcal{R} = \{\langle Q, P\rangle \in \mathrm{P}(A) \times \mathrm{P}(A) \mid Q = o^{-1}(P)\}.$$ [10]

Proof. Suppose $P \in \mathrm{P}(A)$. We need to show that $P \in \mathcal{P}(o(a))$, i.e. $o(a) \in P$, iff there is a prime filter Q such that $Q = o^{-1}(P)$ and $a \in Q$. It suffices to show that $o^{-1}(P)$ is a prime filter if $o(a) \in P$.[11] Suppose $o(a) \in P$. Then $o^{-1}(P)$ is nonempty. In addition, $0 \notin o^{-1}(P)$ since $o(0) = 0 \notin P$. Moreover, $a \in o^{-1}(P)$ and $b \in o^{-1}(P)$ iff $o(a) \in P$ and $o(b) \in P$ iff $o(a \wedge b) \in P$ iff $a \wedge b \in o^{-1}(P)$. Similarly, $a \vee b \in o^{-1}(P)$ iff $a \in o^{-1}(P)$ or $b \in o^{-1}(P)$. ⌟

We have thus proved the following representation theorem:

(12.20) Theorem Every feature algebra is isomorphic to a complex feature algebra.

(12.21) Remark Representation theorems of this type go back to Jónsson and Tarski (1951), who considered Boolean algebras with operations that preserve

---

[10] Recall that a functional relation is a one-many relation – cf. (11.4). Furthermore notice a rather unsatisfying redundancy of notation: $o^{-1}(P)$ is just the same as $o{:}P$. The better way would be to stick solely to the colon operator for inverse images. But then we should consequently introduce an image operator, for which since the days of Russell double quote marks are in use: o“P is $o(P)$. And even better, one could also make function application explicit by an operator, e.g. by Russell's single quote mark; so o‘x is $o(x)$. The technical advantages of such a notational precision should be clear to the reader. Alas, readability, which hinges on conventions too, was the reason not to do so.

[11] The only difference between the rest of the proof and that of (8.23) is that in the case of 0-homomorphisms one cannot use preservation of 1 to show that the inverse image of a prime filter is nonempty.

0 and $\vee$, the so-called *hemimorphisms* of Halmos (1955).[12] An extension to distributive lattices with zero and unit is given by Hansoul (1983). Hemimorphisms naturally arise as Peirce operations, where the underlying relation is not required to be functional. Conversely, it can be shown that every hemimorphism o on an observational algebra A can be represented by the Peirce operation that is given by the relation

$$\mathcal{R} \;=\; \{\langle Q,P\rangle \in \mathrm{P}(A)\times\mathrm{P}(A) \mid o(Q)\subseteq P\}.$$[13]

Notice, however, that this definition differs from (12.19) even if o preserves $\wedge$. In particular, $\mathcal{R}$ is not necessarily functional in this case.

From Chapter 6 onwards we frequently made use of the fact that there is a one-to-one correspondence between the prime filters of an observational algebra and the homomorphisms from that algebra to the algebra $\mathbb{2}$ of two elements – cf. (6.11). This result can be generalized to feature algebras as follows. Suppose P is a prime filter of a feature algebra A over L. Then, for all $a\in A$, if $p{:}a\in P$ then $p{:}1\in P$ (since $a\leqslant 1$ and thus $p{:}a\leqslant p{:}1$). Consequently,

$$U_P \;=\; \{p\in L^* \mid p{:}1\in P\}$$

is nonempty, since $\epsilon\in U_P$, and prefix-closed, that is, $p\in U_P$ whenever $pl\in U_P$. Moreover, the set

$$H_P(a) \;=\; \{p\in L^* \mid p{:}a\in P\}$$

is a subset of U_P for every $a\in A$. Let $\mathcal{T}_P$ be the feature tree algebra given by U_P according to (12.18). Recall that the 0-homomorphisms on $\mathcal{T}_P$ are the Peirce operations determined by the relations $R_l=\{\langle pl,p\rangle \mid pl\in U_P\}$, for $l\in L$.

(12.22) Lemma H_P is a homomorphism of feature algebras from A to $\mathcal{T}_P$.

Proof. By definition, $H_P(1)=U_P$. Since P is prime, $p{:}0=0\notin P$ and thus $H_P(0)=\varnothing$. Moreover, $p{:}(a\wedge b)=p{:}a\wedge p{:}b\in P$ iff $p{:}a\in P$ and $p{:}b\in P$; hence $H_P(a\wedge b)=H_P(a)\cap H_P(b)$. Correspondingly, $H_P(a\vee b)=H_P(a)\cup H_P(b)$, because $p{:}(a\vee b)=p{:}a\vee p{:}b\in P$ iff $p{:}a\in P$ or $p{:}b\in P$. Finally, $H_P(l{:}a)=\{p \mid pl{:}a\in P\}=\{p \mid pl\in H_P(a)\}=R_l{:}(H_P(a))$. $\lrcorner$

Conversely, suppose H is a homomorphism of feature algebras from A to some feature tree algebra over L. Then

[12] See e.g. Blackburn et al. 2001 for the standard application of the Jónsson-Tarski theorem in modal logic.

[13] See e.g. Dunn and Hardegree 2001, Sect. 8.12 for a proof.

(12.23) $$P \;=\; \{a\in A \mid \epsilon\in H(a)\}$$

is clearly a prime filter of A. To see that this construction is inverse to the foregoing one, observe that

$$H(p{:}a) \;=\; \{q \mid qp\in H(a)\}.$$

Hence $p\in H(a)$ iff $\epsilon\in H(p{:}a)$ iff $p{:}a\in P$. We can conclude:

(12.24) Theorem Let A be a feature algebra over L. There is a one-to-one correspondence between the prime spectrum of A and the set of homomorphisms from A to feature tree algebras over L.

Notice that if L is empty then L^* is $\{\epsilon\}$ and hence $\{\varnothing, \{\epsilon\}\}$ is the only feature tree algebra over L. So (6.11) is a special case of (12.24).

12.2.3 *The Canonical (Algebraic) Model*

Algebraic models of observational theories have been introduced in Section 6.1. Here we adapt this concept to attribute-value theories. Suppose Γ is an attribute-value theory over $\langle L, S\rangle$ and A is a feature algebra over L. An *interpretation of S with values in A* is a function m from S to A. The interpretation m is an *A-valued model of* Γ iff its homomorphic extension $\widehat{m}$ to $\mathrm{T}[L, S]$ satisfies $\widehat{m}(\varphi) \leqslant \widehat{m}(\psi)$ for every statement $\varphi \preceq \psi$ of Γ. Notice that if m is an A-valued model of Γ and h is a homomorphism of feature algebras (over L) from A to B then $h \circ m$ is a model of Γ with values in B.

(12.25) Example *(Feature systems as algebraic models)* Let Γ be an attribute-value theory over $\langle L, S\rangle$. Every feature system of Γ can be regarded as a model of Γ in a complex feature algebra, and vice versa. To see this, recall from (12.17) that the extension algebra Ω_M of a feature system $\langle M, U\rangle$ of Γ is a complex feature algebra over L, where the Peirce operations on Ω_M are determined by the functional relations $M(l)$ on U. Let m be the interpretation of S in Ω_M that takes $s\in S$ to $M(s)$. Then, by (12.1), $\widehat{m}(\varphi) = M(\varphi)$, for all $\varphi\in\mathrm{T}[L, S]$. Since M is a feature system of Γ, it follows that m is a model of Γ with values in Ω_M. This correspondence between feature systems and models in complex feature algebras works the other way around as well. For suppose m is a model of Γ in a complex feature algebra $\mathcal{A}$ over L with universe U. Then there are functional relations R_l on U, for every $l\in L$, such that $o_l(V) = R_l{:}V$ for all $V\in\mathcal{A}$. Let M be the feature system over $\langle L, S\rangle$ with $M(s) = m(s)$ and $M(l) = R_l$. It follows that $M(\varphi) = \widehat{m}(\varphi)$ for all $\varphi\in\mathrm{T}[L, S]$. Consequently, M is a feature system of Γ, since m is a model of Γ.

An algebraic model $\langle m, A\rangle$ of Γ is called *universal* if, for every algebraic model $\langle m', A'\rangle$ of Γ, m' factors uniquely through m by a homomorphism h from A to A'. Universal models of Γ, if existent, are unique up to isomorphism. In order to construct such a universal model in a canonical way we can follow the guidelines of universal algebra – in analogy to Section 6.1.1.

The *canonical feature algebra* $\mathrm{L}(\Gamma)$ associated with Γ is the quotient of $\mathrm{T}[L, S]$ modulo the least congruence relation $\cong_\Gamma$ determined by Γ and all defining equations of a feature algebra over L. In particular,

$$l{:}\Lambda \cong_\Gamma \Lambda, \quad l{:}(\varphi \wedge \psi) \cong_\Gamma l{:}\varphi \wedge l{:}\psi, \quad l{:}(\varphi \vee \psi) \cong_\Gamma l{:}\varphi \vee l{:}\psi,$$

and if $\varphi \cong_\Gamma \psi$ then $l{:}\varphi \cong_\Gamma l{:}\psi$. The induced 0-homomorphisms on $\mathrm{L}(\Gamma)$ take $[\varphi]$ to $[l{:}\varphi]$, for every $l \in L$. By definition of $\cong_\Gamma$, the interpretation m_Γ of S in $\mathrm{L}(\Gamma)$ with $\mathrm{m}_\Gamma(s) = [s]_{\cong_\Gamma}$ is a model of Γ in $\mathrm{L}(\Gamma)$, henceforth called the *canonical (algebraic) model* of Γ. The same argument as in the proof of (6.2) leads to:

(12.26) Proposition The canonical algebraic model of an attribute-value theory is universal.

There is thus a one-to-one correspondence between the set $\mathrm{Mod}(\Gamma, A)$ of models of Γ in a feature algebra A and the set $\mathrm{Hom}(\mathrm{L}(\Gamma), A)$ of homomorphisms from $\mathrm{L}(\Gamma)$ to A, where a homomorphism h from $\mathrm{L}(\Gamma)$ to A corresponds to the model $h \circ \mathrm{m}_\Gamma$ of Γ in A; in short,

(12.27) $$\mathrm{Hom}(\mathrm{L}(\Gamma), A) \;\simeq\; \mathrm{Mod}(\Gamma, A).$$

The canonical algebraic model of Γ gives rise to a feature system of Γ as follows. According to (12.20), $\mathrm{L}(\Gamma)$ is isomorphic to a complex feature algebra over the prime spectrum $\mathrm{P}(\mathrm{L}(\Gamma))$ of $\mathrm{L}(\Gamma)$. By (12.27), it follows that there is a model of Γ with values in this complex feature algebra, which in turn, by (12.25), uniquely determines a feature system M_Γ of Γ with universe $\mathrm{P}(\mathrm{L}(\Gamma))$, henceforth called the *canonical feature system* of Γ. Since $\mathrm{M}_\Gamma(\varphi)$ is the set $\mathcal{P}([\varphi])$ of all prime filters of $\mathrm{L}(\Gamma)$ with member $[\varphi]$, and $\mathcal{P}([\varphi]) = \mathcal{P}([\psi])$ iff $[\varphi] = [\psi]$, we have that

(12.28) $$\mathrm{M}_\Gamma(\varphi) = \mathrm{M}_\Gamma(\psi) \qquad \text{iff} \qquad \varphi \cong_\Gamma \psi,$$

for all attribute-value predicates φ and ψ over $\langle L, S\rangle$. In Section 12.2.4 below we will use (12.28) to show that the "logic" underlying the construction of $\cong_\Gamma$ is complete with respect to first-order entailment modulo F. A closer examination of the canonical feature system will follow in Section 12.3.

We conclude this section by generalizing the presentation of algebras by theories from the observational case (see Section 6.1.3) to feature algebras and theories. A feature algebra A over L is said to be *presented* by an attribute-value theory Γ over $\langle L, S\rangle$ if $A \simeq \mathrm{L}(\Gamma)$.

(12.29) Proposition Every feature algebra over L generated by S is presented by an attribute-value theory over $\langle L, S\rangle$.

Proof. Suppose A is a feature algebra over L generated by S, and ε is the inclusion of S into A. Let $\cong$ be the congruence kernel of the homomorphic extension $\hat{\varepsilon}$ of ε. Then $A \simeq \mathrm{T}[L, S]/{\cong}$ because $\hat{\varepsilon}$ is onto by assumption. Now let Γ be the set of all attribute-value statements $\varphi \equiv \psi$ over $\langle L, S\rangle$ such that $\varphi \cong \psi$. Then Γ presents A. ⌟

12.2.4 *Feature Logic*

Suppose Γ is an attribute-value theory and $\varphi \equiv \psi$ is an attribute-value statement over $\langle L, S\rangle$. Since the canonical feature system M_Γ of Γ is a first-order model of $\mathrm{F}\cup\Gamma$, it follows that $\mathrm{M}_\Gamma(\varphi) = \mathrm{M}_\Gamma(\psi)$ whenever $\mathrm{F}\cup\Gamma$ entails $\varphi \equiv \psi$. On the other hand, $\mathrm{F}\cup\Gamma$ entails $\varphi \equiv \psi$ whenever $\varphi \cong_\Gamma \psi$, because the schemes (11.9) to (11.12) that underly the construction of the congruence closure $\cong_\Gamma$ are first-order deducible modulo F. So, by (12.28),

(12.30) $$\mathrm{F} \cup \Gamma \vdash \varphi \equiv \psi \qquad \text{iff} \qquad \varphi \cong_\Gamma \psi.$$

This fact immediately provides us with a sound and complete inference calculus for attribute-value statements. We simply have to mimic the algebraic construction of the congruence relation $\cong_\Gamma$ by axioms and rules. To this end, it suffices to supplement the calculus $OC_\equiv$ of Section 6.3.1 with the following axiom and inference schemes:

$$F{:}\Lambda \;\equiv\; \Lambda$$

$$F{:}(A \wedge B) \;\equiv\; F{:}A \wedge F{:}B$$

$$F{:}(A \vee B) \;\equiv\; F{:}A \vee F{:}B$$

$$\frac{A \;\equiv\; B}{F{:}A \;\equiv\; F{:}B} \quad (\mathrm{P}_\equiv)$$

Let $FC_\equiv$ be the resulting extension of $OC_\equiv$. Then, by definition, $\varphi \cong_\Gamma \psi$ iff Γ entails $\varphi \equiv \psi$ by $FC_\equiv$. Hence, by (12.30):

(12.31) Theorem The calculus $\mathit{FC}_{\equiv}$ is sound and strongly complete with respect to first-order entailment modulo F, that is,

$$\mathrm{F} \cup \Gamma \vdash \varphi \equiv \psi \qquad \text{iff} \qquad \Gamma \vdash_{\mathit{FC}_{\equiv}} \varphi \equiv \psi.$$

To get a sound and strongly complete calculus for conditional statements, it is enough to add the following schemes to the calculus $\mathit{OC}_{\preceq}$ of Section 6.3.2:

$$F{:}\Lambda \;\preceq\; \Lambda$$

$$F{:}A \wedge F{:}B \;\preceq\; F{:}(A \wedge B)$$

$$F{:}(A \vee B) \;\preceq\; F{:}A \vee F{:}B$$

$$\frac{A \;\preceq\; B}{F{:}A \;\preceq\; F{:}B} \quad (\mathrm{P}_{\preceq})$$

For $\Lambda \preceq F{:}\Lambda$ is an instance of (Q), $F{:}(A \wedge B) \;\preceq\; F{:}A \wedge F{:}B$ is deducible from $A \wedge B \preceq A$ and $A \wedge B \preceq B$ via prefixing ($\mathrm{P}_{\preceq}$) and introduction of $\wedge$ ($\mathrm{I}_{\wedge}$), and so on. Since biconditional theories can be transformed into conditional ones and vice versa (see (5.1)), we write *FC* when it does not matter whether to use $\mathit{FC}_{\equiv}$ or $\mathit{FC}_{\preceq}$.

12.3 The Generic Feature System

Let Γ be an attribute-value theory over $\langle L, S\rangle$. Combining (12.28) and (12.30) shows that the canonical feature system M_{Γ} of Γ satisfies *equivalence of coextensives* in the sense that

$$\mathrm{M}_{\Gamma}(\varphi) = \mathrm{M}_{\Gamma}(\psi) \qquad \text{iff} \qquad \mathrm{F} \cup \Gamma \vdash \varphi \equiv \psi.$$

Moreover, M_{Γ} satisfies *identity of indiscernibles* since $\mathrm{M}_{\Gamma}(\varphi)$ is the set of all prime filters of $\mathrm{L}(\Gamma)$ with member $[\varphi]$. Employing the terminology introduced for observational theories, we call a feature system of Γ *generic* in case it is isomorphic to M_{Γ}; the universe of such a feature system will be called “the” *generic universe* of Γ, its members the *generic entities* determined by Γ.

Let us explicate the interpretation of attribute predicates under M_{Γ}. By construction of M_{Γ}, the functional relation $\mathrm{M}_{\Gamma}(l)$ on the prime spectrum of $\mathrm{L}(\Gamma)$ is defined along (12.19), that is,

(12.32) $$Q = l \cdot P \quad \text{iff} \quad Q = o_l^{-1}(P) = \{[\varphi] \mid [l{:}\varphi] \in P\},$$

for all prime filters P and Q of $\mathrm{L}(\Gamma)$. (Recall that, by definition, $Q = l \cdot P$ iff $\langle Q, P\rangle \in \mathrm{M}_{\Gamma}(l)$.)

12.3.1 *Representation by Term Feature Trees*

The one-to-one correspondence between prime filters of $\mathrm{L}(\Gamma)$ and homomorphisms from $\mathrm{L}(\Gamma)$ to feature tree algebras gives rise to a representation of generic entities by term feature trees in the following way. According to (12.24), every prime filter P of $\mathrm{L}(\Gamma)$ uniquely corresponds to a homomorphism H_P from $\mathrm{L}(\Gamma)$ to a feature tree algebra $\mathcal{T}_P$ over L with carrier set

$$U_P \;=\; \{p \in L^* \mid [p{:}\mathrm{V}] \in P\}.$$

Application of (12.27) gives us an algebraic model of Γ with values in $\mathcal{T}_P$, which, by (12.25), uniquely determines a term feature tree M_P of Γ with universe U_P such that

$$M_P(l) = \{\langle pl, p\rangle \mid pl \in U_P\} \quad \text{and} \quad M_P(\varphi) = \{p \in U_P \mid [p{:}\varphi] \in P\}.$$

Conversely, if M is a term feature tree of Γ then $P = \{[\varphi] \mid \epsilon \in M(\varphi)\}$ is a prime filter of $\mathrm{L}(\Gamma)$ such that $M = M_P$; cf. (12.23).

So we can take the term feature trees of Γ to represent the generic entities of Γ, where a term feature tree M of Γ satisfies an attribute-value predicate φ iff $\epsilon \in M(\varphi)$. For every two term feature trees M and M' of Γ, it follows by (12.32) that $M' = l \cdot M$ iff $M'(\varphi) = M(l{:}\varphi)$ for all $\varphi \in \mathrm{T}[L, S]$. In particular, the universe $M'(\mathrm{V})$ of M' is $M(l{:}\mathrm{V}) = \{p \mid lp \in M(\mathrm{V})\}$ in this case.

As for specialization, which is set inclusion on $\mathrm{P}(\mathrm{L}(\Gamma))$, we have

$$M \sqsubseteq M' \qquad \text{iff} \qquad \forall\varphi(M(\varphi) \subseteq M'(\varphi)),$$

which is equivalent to: $\forall\varphi(\epsilon \in M(\varphi) \rightarrow \epsilon \in M'(\varphi))$. Put differently, M is specialized by M' iff the universe of M is a subset of the universe of M' and the inclusion function is a homomorphism of feature trees from M to M'. To sum up:

(12.33) Proposition The generic feature system of Γ is isomorphic to the feature system of term feature trees of Γ.

Given a generic feature system $\langle \tilde{M}, \tilde{U}\rangle$ of Γ, one can recover the term feature tree $\langle M_x, U_x\rangle$ of Γ corresponding to $x \in \tilde{U}$ as follows. Since $p \in M_x(\varphi)$ iff $\epsilon \in M_x(p{:}\varphi)$ iff $x \in \tilde{M}(p{:}\varphi)$ iff $p \cdot x \in \tilde{M}(\varphi)$, we have:

$$M_x(\varphi) \;=\; \{p \mid p \cdot x \models \varphi\} \qquad \text{and} \qquad U_x = \{p \mid p \cdot x \models \mathrm{V}\}.$$

So we get M_x by “unraveling” $\tilde{M}$ at x. Notice that the function f from U_x to $\tilde{U}$ that takes p to $p \cdot x$ is a homomorphism of feature systems with $p \models \varphi$

iff $f(p) \models \varphi$. In general, f is not an embedding; for in contrast to the generic feature system of Γ, the feature trees of Γ do not necessarily satisfy identity of indiscernibles.

(12.34) Example Suppose $L = \{\textsc{next}\}$ and $S = \{\mathit{cont}, \mathit{stop}\}$. Let Γ be the attribute-value theory over $\langle L, S\rangle$ consisting of the statements

$$\mathit{cont} \wedge \mathit{stop} \;\preceq\; \Lambda \qquad \text{and} \qquad \textsc{next}{:}\mathrm{V} \;\preceq\; \mathit{cont}.$$

It is not difficult to list the term feature trees of Γ. Let U_k be $\{\textsc{next}^i \mid i \leqslant k\}$, with $0 \leqslant k < \omega$, and U_ω be $\{\textsc{next}^i \mid i < \omega\}$. For each $k < \omega$ there are three term feature trees of Γ with universe U_k. They differ in that $\textsc{next}^k$ satisfies *cont* or *stop* or neither *cont* nor *stop*, respectively; in all three cases $\textsc{next}^i$ satisfies *cont* for $i < k$. In addition, there is one term feature tree with universe U_ω such that $\textsc{next}^i$ satisfies *cont* for all $i < \omega$. Figure 78 shows the generic feature system of Γ represented by the term feature trees of Γ, as well as the specialization relation between these trees. Figure 79, in contrast, abstracts away from a specific representation of the generic entities. The sole infinite term feature tree of Γ, for example, can be recovered from this abstract representation by unraveling the generic feature system at its loop node. Figure 80 shows all term feature trees of Γ that are maximal with respect to specialization. They constitute the generic universe of the attribute-value theory over $\langle L, S\rangle$ that consists of the statements

$$\mathit{cont} \wedge \mathit{stop} \;\equiv\; \Lambda, \quad \mathit{cont} \vee \mathit{stop} \;\equiv\; \mathrm{V}, \quad \textsc{next}{:}\mathrm{V} \;\equiv\; \mathit{cont},$$

which is essentially a rule completion of Γ in the sense of Section 8.4.2.

12.3.2 *Feature Logic as Observational Theory*

According to (12.10), every attribute-value theory over $\langle L, S\rangle$ is equivalent (modulo F) to one in observational form, i.e. to an observational theory over the set $\Sigma[L, S]$ of primitive attribute-value predicates. Clearly we cannot expect the calculus *OC* of observational logic to be complete for attribute-value statements in observational form, because *OC* is insensitive to the inner logical structure of attribute-value predicates.

Prefixing, for instance, is not covered by *OC*. More precisely, since the prefixing rule (P) does not respect observational forms, the prefixing scheme at issue here is

$$\frac{\varphi \;\preceq\; \psi}{(l{:}\varphi)^{\circ} \;\preceq\; (l{:}\psi)^{\circ}} \quad (\pi_l)$$

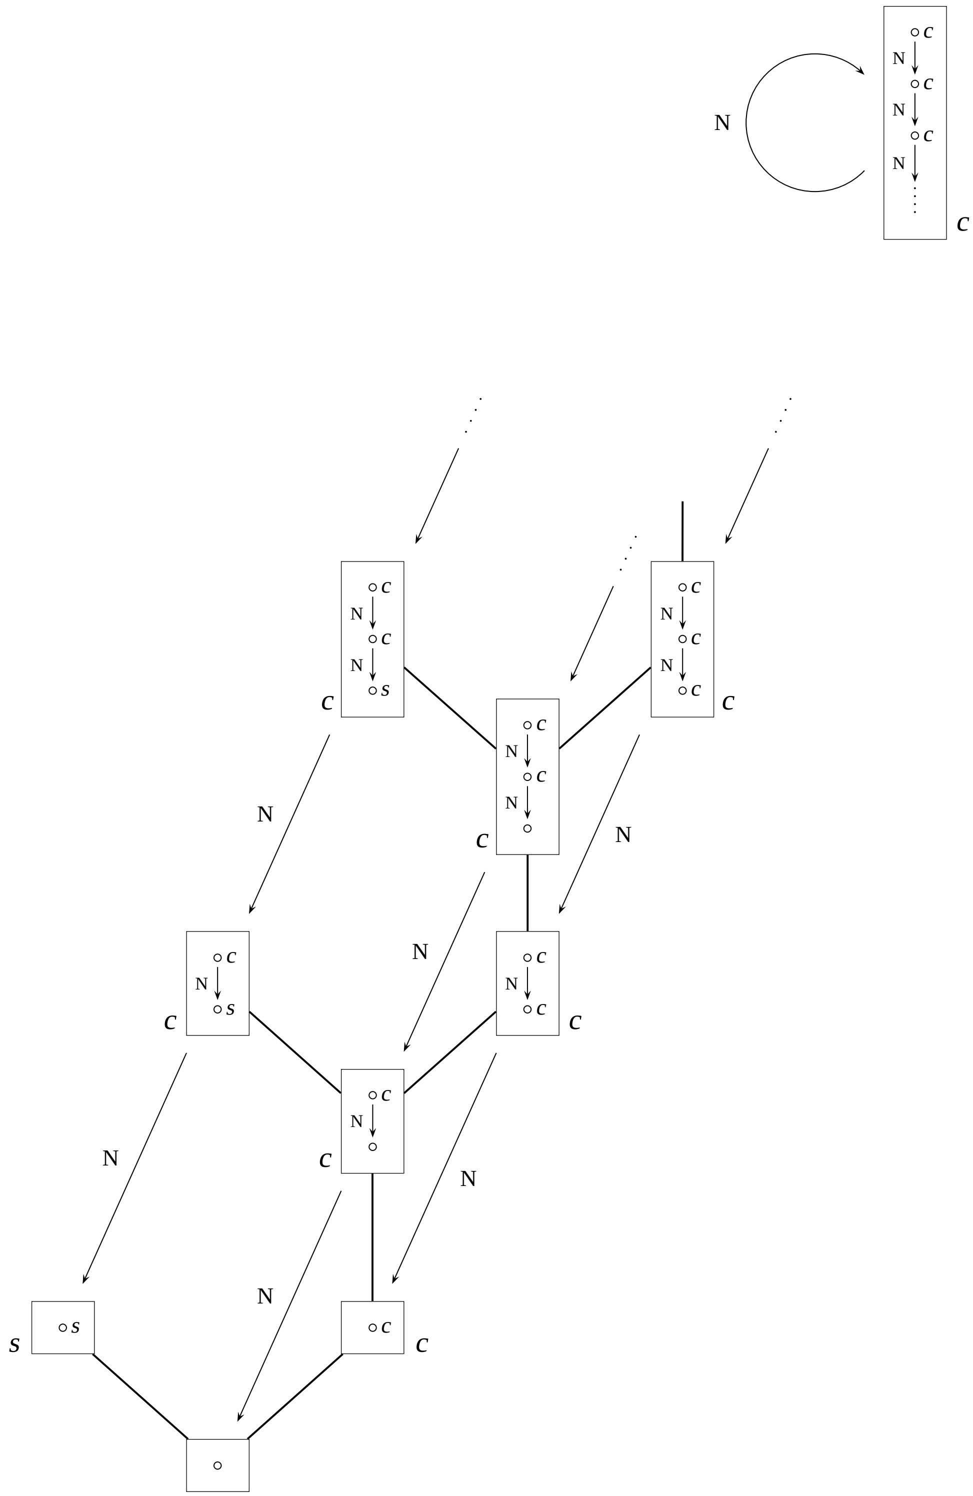

FIGURE 78 The feature system of term feature trees of Γ

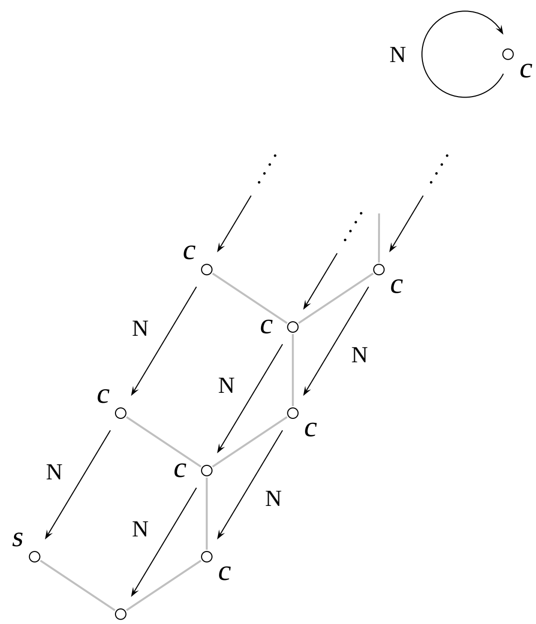

FIGURE 79 Generic feature system of Γ

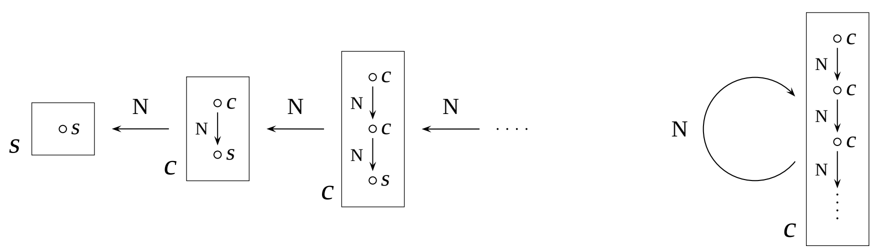

FIGURE 80 Generic feature system of rule completion of Γ

where φ and ψ are observational predicates over $\Sigma[L,S]$ and $(l{:}\varphi)^{\mathrm{o}}$ and $(l{:}\psi)^{\mathrm{o}}$ are observational forms of $l{:}\varphi$ and $l{:}\psi$, respectively, which we define as follows. For $p \in L^*$ let $(p{:}\varphi)^{\mathrm{o}}$ be $p{:}\varphi$ in case φ belongs to $\Sigma[L,S] \cup \{\Lambda, \mathrm{V}\}$, and extend this definition inductively to all observational predicates over $\Sigma[L,S]$ such that $(p{:}(\varphi \wedge \psi))^{\mathrm{o}} = (p{:}\varphi)^{\mathrm{o}} \wedge (p{:}\psi)^{\mathrm{o}}$ and $(p{:}(\varphi \vee \psi))^{\mathrm{o}} = (p{:}\varphi)^{\mathrm{o}} \vee (p{:}\psi)^{\mathrm{o}}$. Then $(p{:}\varphi)^{\mathrm{o}}$ is logically equivalent to $p{:}\varphi$ and has observational form. Moreover, $(\epsilon{:}\varphi)^{\mathrm{o}} = \varphi$ and $(p{:}(q{:}\varphi)^{\mathrm{o}})^{\mathrm{o}} = (pq{:}\varphi)^{\mathrm{o}}$. Of course, ($\pi_l$) is provable in $FC_{\preceq}$, as one easily checks.

Given an observational theory Γ over $\Sigma[L,S]$, let $\Gamma^{(\pi)}$ be its prefix closure in observational form, i.e. the closure of Γ with respect to (π_l), for all $l \in L$. Then, in general, $\Gamma^{(\pi)}$ still does not *OC*-entail all observational attribute-value statements entailed by *FC*; witness the following two statement schemes:

(12.35) $$p{:}\Lambda \preceq \Lambda \qquad \text{and} \qquad pq{:}s \preceq p{:}\mathrm{V},$$

with $p, q \in L^*$ and $s \in S$. The first scheme is a consequence of the axiom scheme $F{:}\Lambda \preceq \Lambda$ of *FC* whereas the second can be deduced by prefixing from the scheme (U) of *OC* (see Figure 51).

Let T be the observational theory over $\Sigma[L, S]$ given by the two statement schemes of (12.35). We claim that for every observational statement $\varphi \preceq \psi$ over $\Sigma[L, S]$,

(12.36) $$\mathrm{T} \cup \Gamma^{(\pi)} \vdash_{OC} \varphi \preceq \psi \qquad \text{iff} \qquad \Gamma \vdash_{FC} \varphi \preceq \psi.$$

One way to verify (12.36) is to show that every proof of $\varphi \preceq \psi$ from Γ by $FC_{\preceq}$ can be transformed into one from $\mathrm{T} \cup \Gamma^{(\pi)}$ by $OC_{\preceq}$. The following argument, in contrast, makes use of the canonical model of $\mathrm{T} \cup \Gamma^{(\pi)}$ as an observational theory (see Section 5.2). We show that every member X of the canonical model $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ of $\mathrm{T} \cup \Gamma^{(\pi)}$ gives rise to a term feature tree M_X of Γ such that $X \models \varphi$ iff $\epsilon \in M_X(\varphi)$, for all observational predicates φ over $\Sigma[L, S]$. In case Γ entails $\varphi \preceq \psi$ by *FC* it then follows that $M_X(\varphi) \subseteq M_X(\psi)$, for all X, hence $[\![\varphi]\!] \subseteq [\![\psi]\!]$, and thus $\mathrm{T} \cup \Gamma^{(\pi)} \vdash_{OC} \varphi \preceq \psi$. (Here $[\![\varphi]\!]$ is the extension of φ in $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$, i.e. the set of all $X \in \mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ with $X \models \varphi$.)

In order to associate with every member X of $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ a term feature tree M_X, observe that the set

$$U_X = \{p \in L^* \mid X \models p{:}\mathrm{V}\}$$

is nonempty, because $X \models \mathrm{V}$, and prefix-closed, since $pq{:}\mathrm{V} \preceq p{:}\mathrm{V}$ belongs to T, for all $q \in L^*$. Let M_X be the term feature tree over $\langle L, S\rangle$ with universe U_X and

$$M_X(l) = \{\langle pl, p\rangle \mid pl \in U_X\}, \qquad M_X(s) = \{p \in L^* \mid X \models p{:}s\},$$

for $l \in L$ and $s \in S$. Notice that $M_X(s) \subseteq U_X$ because $p{:}s \preceq p{:}\mathrm{V}$ is in T. Moreover, for every observational predicate φ over $\Sigma[L, S]$,

(12.37) $$M_X(\varphi) = \{p \in L^* \mid X \models (p{:}\varphi)^{\mathrm{o}}\}.$$

Proof. We have $p \in M_X(q{:}s)$ iff $pq \in M_X(s)$ iff $X \models pq{:}s$, which proves the primitive case. By definition, $p \in M_X(\mathrm{V}) = U_X$ iff $X \models p{:}\mathrm{V}$. Again by definition, $M_X(\Lambda) = \varnothing$; moreover, $X \not\models p{:}\Lambda$, for all $p \in L^*$, because $p{:}\Lambda \preceq \Lambda$ belongs to T. Now term induction: $p \in M_X(\varphi \wedge \psi)$ iff $p \in M_X(\varphi) \cap M_X(\psi)$ iff $X \models (p{:}\varphi)^{\mathrm{o}} \wedge (p{:}\varphi)^{\mathrm{o}} = (p{:}(\varphi \wedge \psi))^{\mathrm{o}}$. Analogously for disjunction. ⌟

(12.38) Lemma If $X \in \mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ then M_X is a term feature tree of Γ.

Proof. Suppose $(\varphi \preceq \psi) \in \Gamma$. If $X \models (p{:}\varphi)^{\mathrm{o}}$ then $X \models (p{:}\psi)^{\mathrm{o}}$, because X is consistently closed with respect to $\Gamma^{(\pi)}$. Hence $M_X(\varphi) \subseteq M_X(\psi)$, by (12.37). ⌟

As mentioned before, this is enough to prove (12.36). To repeat the argument, if Γ *FC*-entails an observational statement $\varphi \preceq \psi$ over $\Sigma[L, S]$ then, by (12.38), $M_X(\varphi) \subseteq M_X(\psi)$, for all $X \in \mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$, hence $[\![\varphi]\!] \subseteq [\![\psi]\!]$, by (12.37), and thus $\mathrm{T} \cup \Gamma^{(\pi)} \vdash_{OC} \varphi \preceq \psi$, as shown in Section 6.3.

Consequently, the canonical feature algebra $\mathrm{L}(\Gamma)$ of Γ is isomorphic (as an observational algebra) to the Lindenbaum algebra of $\mathrm{T} \cup \Gamma^{(\pi)}$ (in the sense of Section 6.2.2), where $[\varphi] \in \mathrm{L}(\Gamma)$ is taken to $[\varphi']$ with φ' an observational form of φ. So the Lindenbaum algebra of $\mathrm{T} \cup \Gamma^{(\pi)}$ can be equipped with the structure of a feature algebra over L where the 0-homomorphisms take $[\varphi]$ to $[(l{:}\varphi)^{\mathrm{o}}]$, for each $l \in L$. In particular, the Lindenbaum algebra of T is (isomorphic to) the feature algebra over L that is freely generated by S.

Furthermore, it follows that $\mathrm{L}(\Gamma)$ and the Lindenbaum algebra of $\mathrm{T} \cup \Gamma^{(\pi)}$ have order-isomorphic prime spectra; hence:

(12.39) Proposition Let Γ be an attribute-value theory over $\langle L, S\rangle$ in observational form. Then the generic universe of Γ is order-isomorphic to the generic universe of the observational theory $\mathrm{T} \cup \Gamma^{(\pi)}$ over $\Sigma[L, S]$.

By (12.33), we have thus a one-to-one correspondence between $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ and the set of term feature trees of Γ, where $X \in \mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ is taken to M_X. Its inverse takes every term feature tree M of Γ to the set of all $\varphi \in \Sigma[L, S]$ such that $M \models \varphi$, i.e. $\epsilon \in M(\varphi)$. As for the interpretation of $l \in L$ by a functional relation on $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$, we have $Y = l \cdot X$ iff $Y = \{\varphi \mid l{:}\varphi \in X\}$.

Since the observational theory T is a simple inheritance theory with unary exclusions, it follows by the results of Chapter 4 that the ordered generic universe of T is a complete completely distributive algebraic lattice:

(12.40) Proposition The set of term feature trees over $\langle L, S\rangle$ ordered by specialization is a complete completely distributive algebraic lattice, whose compact elements are the finite trees.

Proof. Every compact element X of $\mathrm{C}(\mathrm{T})$ is the least satisfier of a conjunctive attribute-value predicate $p_1{:}s_1 \wedge \ldots \wedge p_n{:}s_n$ over $\langle L, S\rangle$, and vice versa. The definition (12.35) of T implies that the universe U_X of the corresponding term feature tree is the set of all prefixes of the p_i's, which is finite. ⌟

If L and S are both countable, 'algebraic' can be replaced by 'ω-algebraic'.

More generally, let Γ be an attribute-value theory over $\langle L, S \rangle$ in observational form, i.e. an observational theory over $\Sigma[L, S]$. If Γ satisfies any of the properties listed in Table 1 except the last two ones, then so does $\Gamma^{(\pi)}$ and therefore $\mathrm{T} \cup \Gamma^{(\pi)}$. For instance, if Γ is a simple inheritance theory with exclusions then $\mathrm{C}(\mathrm{T} \cup \Gamma^{(\pi)})$ and hence the ordered generic universe of Γ is a completely distributive Scott domain. An example is provided by (12.34).

12.4 Prospects

In this last section, we address possible extensions to the foregoing presentation of attribute-value theories and list some topics for future research.

A first topic of research is to properly define morphisms between attribute-value theories over *different* feature signatures and to study their effect on the respective generic feature systems. The category of attribute-value theories thus defined would then allow to apply the standard categorical constructions like forming coproducts and inductive limits. In particular, attribute-value theories could be systematically composed of or decomposed into simpler ones.

12.4.1 *Attribute-Value Identity and Feature Structures*

As mentioned at the close of Section 10.1.4, the classification of syntactic phenomena by feature-based grammatical theories typically employs descriptive means for expressing that attributes have identical values. It is thus a natural task to integrate this type of description into the framework developed so far, which then leads to the standard version of feature logic as presented for example in Kasper and Rounds 1990 and Carpenter 1992.

It is not difficult to extend the method of regimentation and formalization of Chapter 11 to this case. Consider the predicate 'someone whose father is her employer', which is true of everybody whose father (value) is identical to her employer (value). Regimentation and formalization leads to predicates of the form '$\{x \mid \exists y(Fyx \wedge Gyx)\}$', usually abbreviated by '$F \doteq G$', with F and G presupposed as functional. In similar vein, one can introduce descriptive means for expressing arbitrary relations between attribute values (see Osswald 1999b). As in Section 11.3, the logical form of these attribute-value predicates allows to deduce logical equivalences between them.

Completeness proofs for feature logic with identity have been given by Moss (1992) and Moshier (1993). An outline how to proof the strong completeness of feature logic with identity and relations by reducing it to observational logic can be found in Osswald 1999b. The generic entities of an attribute-value theory with identity, when viewed as an observational theory, can be represented by the *rooted feature systems* of that theory, which are also known as *feature structures*. A satisfying algebraic approach to feature logic with identity, however, is still missing.

12.4.2 *General Observational Logic*

Finally, it is worth mentioning that there is also a first-order version of observational logic that allows primitive predicates of arbitrary adicity (see Vickers 1993 or Mac Lane and Moerdijk 1992, Chap. X). Observational predicates over these primitives are then built by finite conjunction, disjunction, and existential quantification of primitive predicates plus V, Λ, and $=$.[14] Observational statements are universally quantified conditionals, with observational predicates of the same adicity as antecedent and consequent. Clearly, attribute-value theories with identity and relations are subsumed by this more general approach.

[14] Vickers allows disjunctions to be infinite.

Appendix

Lattices and Order

Ordered Sets

An *ordered set* (also *partially ordered set* or *partial order*) consists of a set P and a reflexive, transitive, and anti-symmetric relation $\sqsubseteq$ on P. The *dual* of an ordered set $\langle P, \sqsubseteq \rangle$ is the ordered set $\langle P, \sqsupseteq \rangle$, where $\sqsupseteq$ is the converse of the relation $\sqsubseteq$. An ordered set is a *chain* if $x \sqsubseteq y$ or $y \sqsubseteq x$, for all $x, y \in P$; it is an *antichain* in case $x = y$ whenever $x \sqsubseteq y$.

Suppose P and Q are ordered sets. A function f from P to Q is called *order-preserving* if $f(x) \sqsubseteq f(y)$ whenever $x \sqsubseteq y$, for all $x, y \in P$. If in addition $x \sqsubseteq y$ whenever $f(x) \sqsubseteq f(y)$, then f is called an *order embedding*. An *order isomorphism* is a bijective order embedding.

A subset S of an ordered set P is called *downwards closed* (or *down-set* or *decreasing set* or *order ideal*) if $\downarrow S \subseteq S$, where $\downarrow S = \{x \mid \exists y \in S(x \sqsubseteq y)\}$. Instead of '$\downarrow\{x\}$' we also write '$\downarrow x$'. An *upwards closed* subset S is one with $\uparrow S \subseteq S$, where $\uparrow$ is defined dually to $\downarrow$. A subset S of P is *upwards directed* if it is nonempty and every two members of S have a common upper bound in S; it is *downwards directed* if it is upwards directed with respect to the dual order.

The least and the greatest elements of P, if existent, are respectively called *bottom* and *top* (notation: $\bot$ and $\top$). A subset S of P is *bounded above* or *below* if there is respectively an upper or lower bound of S in P. A *least upper bound* of a subset S of P, if existent, is unique and is called a *supremum* of S; notation: $\bigsqcup S$. *Greatest lower bounds* are referred to as *infima* ($\bigsqcap S$). Notice that $\bigsqcup \varnothing = \bot$ and $\bigsqcap \varnothing = \top$. The supremum of two elements x and y of S is also called the *join* of x and y, in symbols: $x \sqcup y$; the infimum of x and y is referred to as their *meet*, written $x \sqcap y$.

An ordered set P is said to be *directed complete* (or a *dcpo*) if each of its (upwards) directed subsets has a supremum in P. An order-preserving function of dcpos that preserves suprema of directed sets is called *Scott-continuous*.

Lattices

Lattices are nonempty ordered sets where every two elements have a join and a meet. By taking $\sqcup$ and $\sqcap$ as binary operations, lattices can be considered as algebras of type $\langle 2, 2\rangle$. The ordering $\sqsubseteq$ is related to the operations $\sqcup$ and $\sqcap$ according to:

$$x = x \sqcap y \quad \text{iff} \quad x \sqsubseteq y \quad \text{iff} \quad x \sqcup y = y.$$

If L is a lattice then, for all $x, y, z \in L$, the operations $\sqcup$ and $\sqcap$ satisfy the laws:

$$\begin{aligned}
&(x \sqcup y) \sqcup z = x \sqcup (y \sqcup z) && (x \sqcap y) \sqcap z = x \sqcap (y \sqcap z)\\
&x \sqcup y = y \sqcup x && x \sqcap y = y \sqcap x\\
&x \sqcup x = x && x \sqcap x = x\\
&x \sqcup (x \sqcap y) = x && x \sqcap (x \sqcup y) = x
\end{aligned}$$

which are respectively known as *associativity*, *commutativity*, *idempotency*, and *absorption*. Conversely, if two binary operations $\sqcup$ and $\sqcap$ on a set L satisfy these laws, then the ordering $\sqsubseteq$ on L defined by $x \sqsubseteq y$ iff $x \sqcup y = y$ (or $x = x \sqcap y$) turns L into a lattice. So lattices can be characterized order-theoretically or algebraically.

Top and bottom of a lattice L viewed as an algebra are called *unit* and *zero*, respectively; in symbols, 1 and 0. Unit and zero are characterized by the equations $x \sqcap 1 = x$ and $x \sqcup 0 = x$, or equivalently, $x \sqcup 1 = 1$ and $x \sqcap 0 = 0$, for all $x \in L$. Given a lattice L with unit and zero, an element $b \in L$ is said to be a *complement* of an element $a \in L$ if $a \sqcap b = 0$ and $a \sqcup b = 1$.

A lattice L is called *distributive* if the following distributivity law holds for all $x, y, z \in L$:

$$x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z).$$

In a distributive lattice every element has at most one complement. A *Boolean lattice* is a distributive lattice with unit and zero whose every element has a complement.

A lattice is *complete* if it has suprema and infima for all subsets. A complete lattice L is said to be *completely distributive* if, for every nonempty set $\mathcal{S}$ of nonempty subsets of L,

$$\bigsqcap\{\bigsqcup X \mid X \in \mathcal{S}\} = \bigsqcup\{\bigsqcap f(\mathcal{S}) \mid f \text{ is a selector of } \mathcal{S}\}.$$[15]

[15] A *selector* or *choice function* of a system $\mathcal{S}$ of sets is a function from $\mathcal{S}$ to $\bigcup \mathcal{S}$ such that $f(X) \in X$ for every nonempty member X of $\mathcal{S}$.

Notice that $\mathcal{S} = \{\{x\}, \{y, z\}\}$ gives the the finite distributivity law introduced above.

Filters and Ideals

A *filter* of a lattice L is a nonempty subset F of L such that if $a, b \in F$ then $a \sqcap b \in F$, and if $a \in F$ and $a \sqsubseteq b \in L$ then $b \in F$. A filter F is *prime* if $F \neq L$, and $a \sqcup b \in F$ implies that $a \in F$ or $b \in F$. It is easy to see that if L has zero and unit, $F \subseteq L$ is a prime filter of L just in case $0 \notin F$, $1 \in F$,

$$a \sqcap b \in F \quad \text{iff} \quad a \in F \quad \text{and} \quad b \in F, \qquad a \sqcup b \in F \quad \text{iff} \quad a \in F \quad \text{or} \quad b \in F.$$

The set $\uparrow a = \{b \in L \mid a \sqsubseteq b\}$ is a filter of L, for every $a \in L$. Filters of this form are called *principal*. The filter $\uparrow a$ is prime iff $a \neq 0$ and a is *join-irreducible*, that is, if $a = b \sqcup c$, with $b, c \in L$, then $a = b$ or $a = c$.

An *ideal* of L is a nonempty subset I of L such that if $a, b \in I$ then $a \sqcup b \in I$, and if $a \in I$ and $a \sqsupseteq b \in L$ then $b \in I$. An ideal I of L is *prime* if $I \neq L$ and $a \sqcap b \in I$ implies that $a \in I$ or $b \in I$. Clearly I is a prime ideal of L iff $L \setminus I$ is a prime filter.

Prime Ideal Theorem Let L be a distributive lattice. If I is an ideal of L and F is a filter of L with $I \cap F = \varnothing$ then there exists a prime ideal J of L such that $I \subseteq J$ and $J \cap F = \varnothing$.[16]

Since complements of prime ideals are prime filters, the Prime Ideal Theorem ensures the existence of prime filters too. For example, given two elements a, b of a distributive lattice L with $a \not\sqsubseteq b$, the Prime Ideal Theorem applied to the ideal $\downarrow b$ and the filter $\uparrow a$ yields a prime filter F of L such that $a \in F$ and $b \notin F$.

Some Concepts from Category Theory

Categories

A *category* $\mathbf{C}$ consists of a class of *objects* and, for each pair $\langle A, B\rangle$ of objects, a set $\mathrm{Hom}_{\mathbf{C}}(A, B)$ of *morphisms* from A (the *source*) to B (the *target*), where the morphism sets are pairwise disjoint. In addition, for every two morphisms f from A to B and g from B to C there is a morphism $g \circ f$ from A to C, the *composite* of g and f, subject to the condition that $\circ$ is associative. Moreover, for every object B there is an identity morphism $\imath_B$ from B to itself such that $\imath_B \circ f = f$ and $g \circ \imath_B = g$.

An *isomorphism* is a morphism that has a two-sided inverse with respect to composition. A morphism f from A to B is a *monomorphism* if, for every two

[16] See e.g. Davey and Priestley 1990, Chap. 9 for a proof.

morphisms g and h from D to A, $g = h$ whenever $f \circ g = f \circ h$. The morphism f is an *epimorphism* if, for every two morphisms g and h from B to C, $g = h$ whenever $g \circ f = h \circ f$.

A *subcategory* $\mathbf{C}$ of a category $\mathbf{D}$ is a category that consists of some of the objects and morphisms of $\mathbf{D}$, where the identity morphism of each object is the identity morphism of the object in $\mathbf{D}$, and source and target of each morphism are those of the morphism in $\mathbf{D}$. The subcategory $\mathbf{C}$ is called *full*, if the $\mathbf{C}$-morphisms between two objects of $\mathbf{C}$ coincide with the $\mathbf{D}$-morphisms between these two objects.

Suppose $\mathbf{C}$ is a category and $\sim$ is a *congruence relation* on the $\mathbf{C}$-morphisms with respect to $\circ$, that is, $\sim$ is reflexive, transitive, and symmetric, and if $f \sim g$ and $h \sim k$ then $h \circ f \sim k \circ g$, for all morphisms f, g from A to B and h, k from B to C. The *quotient category* $\mathbf{C}/\!\sim$ of $\mathbf{C}$ by $\sim$ has the same objects as $\mathbf{C}$, whereas its morphisms are the equivalence classes of $\sim$.

The *dual* or *opposite* of a category is the category that consists of the same objects and morphisms but with source and target interchanged. So a morphism from A to B in the category $\mathbf{C}$ is a morphism from B to A in the dual category of $\mathbf{C}$.

Functors

A *covariant functor* F from a category $\mathbf{C}$ to a category $\mathbf{D}$ is a function that takes each $\mathbf{C}$-object A to a $\mathbf{D}$-object $F(A)$ and each $\mathbf{C}$-morphism f from A to B to a $\mathbf{D}$-morphism $F(f)$ from $F(A)$ to $F(B)$ such that

$$F(f \circ g) \;=\; F(f) \circ F(g) \qquad \text{and} \qquad F(\imath_A) \;=\; \imath_{F(A)}.$$

A *contravariant functor* from $\mathbf{C}$ to $\mathbf{D}$ is a covariant functor from $\mathbf{C}$ to the dual category of $\mathbf{D}$.

For each category $\mathbf{C}$ there is an *identity functor* $I_{\mathbf{C}}$ that takes each object and each morphism to itself. The *composite* of two functors F and G (as functions on objects and morphisms) from $\mathbf{C}$ to $\mathbf{D}$ and from $\mathbf{D}$ to $\mathbf{E}$, respectively, is the functor $G \circ F$ from $\mathbf{C}$ to $\mathbf{E}$.

A functor F from $\mathbf{C}$ to $\mathbf{D}$ is *full* if for every two $\mathbf{C}$-objects A and B and every $\mathbf{D}$-morphism g from $F(A)$ to $F(B)$ there is a $\mathbf{C}$-morphism f from A to B such that $g = F(f)$. The functor F is *faithful* when for every two $\mathbf{C}$-objects A and B and every two $\mathbf{C}$-morphisms f and g from A to B with $F(f) = F(g)$ it follows that $f = g$.

If $\mathbf{C}$ is a subcategory of $\mathbf{D}$ then the associated inclusion function is a faithful functor from $\mathbf{C}$ to $\mathbf{D}$, called the *inclusion functor*. If $\mathbf{C}/\!\sim$ is the quotient category of $\mathbf{C}$ by a congruence $\sim$, then the *quotient functor* from $\mathbf{C}$ to $\mathbf{C}/\!\sim$ that takes each object to itself and each morphism to its equivalence class modulo $\sim$ is a full functor by construction.

Every object A of a category $\mathbf{C}$ determines a *covariant Hom-functor* from $\mathbf{C}$ to $\mathbf{Set}$, which takes every $\mathbf{C}$-object B to the set $\mathrm{Hom}_{\mathbf{C}}(A, B)$ and every $\mathbf{C}$-morphism f from B to C to the function from $\mathrm{Hom}_{\mathbf{C}}(A, B)$ to $\mathrm{Hom}_{\mathbf{C}}(A, C)$ that takes g to $f \circ g$. The *contravariant Hom-functor*, in contrast, takes B to $\mathrm{Hom}_{\mathbf{C}}(B, A)$ and f to the function from $\mathrm{Hom}_{\mathbf{C}}(C, A)$ to $\mathrm{Hom}_{\mathbf{C}}(B, A)$ that takes g to $g \circ f$

A *forgetful functor* is a functor whose only effect is to "forget" part of the mathematical structure of the objects. For example, the forgetful functor from the category of groups to $\mathbf{Set}$ takes each group to its underlying set and each group homomorphism to itself (as a function).

Let F and G be functors from $\mathbf{C}$ to $\mathbf{D}$. A *natural transformation* τ from F to G assigns to each $\mathbf{C}$-object A a $\mathbf{D}$-morphism τ_A from $F(A)$ to $G(A)$ such that, for every $\mathbf{C}$-morphism f from A to B,

$$\tau_B \circ F(f) = G(f) \circ \tau_A.$$

If each of the morphisms τ_A is an isomorphism then τ is said to be a *natural isomorphism* from F to G.

Equivalences and Adjoints

A functor F from $\mathbf{C}$ to $\mathbf{D}$ is called an *equivalence of categories*, if there is a functor G from $\mathbf{D}$ to $\mathbf{C}$ such that $G \circ F$ and $F \circ G$ are naturally isomorphic to $I_{\mathbf{C}}$ and $I_{\mathbf{D}}$, respectively. If $\mathbf{C}$ is equivalent to the dual of $\mathbf{D}$, one speaks of a *dual equivalence*.

Fact A functor F from a category $\mathbf{C}$ to a category $\mathbf{D}$ is an equivalence of categories iff F is full and faithful and every object of $\mathbf{D}$ is isomorphic to $F(A)$ for some object A of $\mathbf{C}$.[17]

Given two functors F from $\mathbf{C}$ to $\mathbf{D}$ and G from $\mathbf{D}$ to $\mathbf{C}$, the functor F is said to be *left adjoint* to G (and G is said to be *right adjoint* to F) if there is a natural transformation ε from $F \circ G$ to $I_{\mathbf{D}}$ such that the function from $\mathrm{Hom}_{\mathbf{C}}(A, G(B))$ to $\mathrm{Hom}_{\mathbf{D}}(F(A), B)$ that takes f to $\varepsilon_B \circ F(f)$ is a bijection, for all $\mathbf{C}$-objects A and $\mathbf{D}$-objects B. This condition is equivalent to the existence of a natural transformation η from $I_{\mathbf{C}}$ to $G \circ F$ such that the function from $\mathrm{Hom}_{\mathbf{D}}(F(A), B)$ to $\mathrm{Hom}_{\mathbf{C}}(A, G(B))$ that takes g to $G(g) \circ \eta_A$ is a bijection.

[17] See e.g. Mac Lane 1971, Sect. IV.4.

Bibliography

Abramsky, Samson. 1991. Domain Theory in Logical Form. *Journal of Pure and Applied Logic* 51:1–77.

Abramsky, Samson, Dov M. Gabbay, and T. S. E. Maibaum (ed.). 1992. *Handbook of Logic in Computer Science, Vol. 1: Background: Mathematical Structures*. Oxford: Oxford University Press.

Abramsky, Samson, Dov M. Gabbay, and T. S. E. Maibaum (ed.). 1994. *Handbook of Logic in Computer Science, Vol. 3: Semantic Structures*. Oxford: Oxford University Press.

Abramsky, Samson, Dov M. Gabbay, and T. S. E. Maibaum (ed.). 1995. *Handbook of Logic in Computer Science, Vol. 4: Semantic Modelling*. Oxford: Oxford University Press.

Abramsky, Samson, Dov M. Gabbay, and T. S. E. Maibaum (ed.). 2000. *Handbook of Logic in Computer Science, Vol. 5: Logical and Algebraic Methods*. Oxford: Oxford University Press.

Abramsky, Samson, and Achim Jung. 1994. Domain Theory. In Abramsky et al. 1994, 1–168.

Abramsky, Samson, and Steven Vickers. 1993. Quantales, Observational Logic and Process Semantics. *Mathematical Structures in Computer Science* 3:161–227.

Aho, Alfred V., John E. Hopcroft, and Jeffrey D. Ullman. 1983. *Data Structures and Algorithms*. Reading, MA: Addison-Wesley.

Aït-Kaci, Hassan. 1984. *A Lattice-Theoretic Approach to Computation Based on a Calculus of Partially Ordered Typed Structures*. Doctoral dissertation, University of Pennsylvania.

Aït-Kaci, Hassan, Robert Boyer, Patrick Lincoln, and Roger Nasr. 1989. Efficient Implementation of Lattice Operations. *ACM Transactions on Programming Languages and Systems* 11(1):115–146.

Aït-Kaci, Hassan, Andreas Podelski, and Gert Smolka. 1994. A Feature Logic Constraint System for Logic Programming with Entailment. *Theoretical Computer Science* 122:263–283.

Backofen, Rolf, and Gert Smolka. 1995. A Complete and Recursive Feature Theory. *Theoretical Computer Science* 146:243–268.

Balbes, Raymond. 1971. On the Partially Ordered Set of Prime Ideals of a Distributive Lattice. *Canadian Journal of Mathematics* 23:866–874.

Balbes, Raymond. 1978. On the Partially Ordered Set of Prime Ideals of a Distributive Lattice II. *Bulletin de l'Académie Polonaise des Sciences. Série des Sciences Mathématiques, Astronomiques et Physiques* 26(9–10):771–773.

Balbes, Raymond, and Philip Dwinger. 1974. *Distributive Lattices*. Columbia, Missouri: University of Missouri Press.

Barr, Avron, and Edward E. Feigenbaum (ed.). 1981. *The Handbook of Artificial Intelligence*. Los Altos, CA: William Kaufmann.

Barr, Michael, and Charles Wells. 1995. *Category Theory for Computing Science*. Englewood Cliffs, NJ: Prentice Hall. 2nd edition.

Barton, G. Edward, Robert C. Berwick, and Eric Sven Ristad. 1987. *Computational Complexity and Natural Language*. Cambridge, MA: MIT Press.

Barwise, Jon. 1992. Information Links in Domain Theory. In *Proceedings of the 7th International Conference on the Mathematical Foundations of Programming Semantics*, ed. S. Brookes, M. Main, A. Melton, M. Mislove, and D. Schmidt, 168–192. Lecture Notes in Computer Science 598. Berlin. Springer.

Barwise, Jon. 1993. Constraints, Channels, and the Flow of Information. In *Situation Theory and Its Applications, Vol. 3*, ed. P. Aczel, D. Israel, Y. Katagiri, and S. Peters. 3–27. CSLI Lecture Notes, No. 37. Stanford, CA: CSLI Publications.

Barwise, Jon, and John Etchemendy. 1990. Information, Infons, and Inference. In *Situation Theory and Its Applications, Vol. 1*, ed. Robin Cooper, Kuniaki Mukai, and John Perry. 33–78. CSLI Lecture Notes, No. 22. Stanford, CA: CSLI Publications.

Barwise, Jon, and Jerry Seligman. 1997. *Information Flow. The Logic of Distributed Systems*. Cambridge Tracts in Theoretical Computer Science 44. Cambridge: Cambridge University Press.

Batarekh, Aida, and V. S. Subrahmanian. 1989. The Query Topology in Logic Programming. In *Proceedings of the 6th Annual Symposium on Theoretical Aspects of Computer Science*, 375–387. Lecture Notes in Computer Science 349. Berlin. Springer.

Birkhoff, Garrett. 1967. *Lattice Theory*. AMS Colloquium Publications, Vol. 25. Providence, RI: AMS. 3rd edition.

Blackburn, Patrick. 1993. Modal Logic and Attribute Value Structures. In *Diamonds and Defaults*, ed. Maarten de Rijke. Dordrecht: Kluwer.

Blackburn, Patrick. 1994. Structures, Languages and Translations: the Structural Approach to Feature Logic. In Rupp et al. 1994, 1–27.

Blackburn, Patrick, and Maarten de Rijke (ed.). 1997. *Specifying Syntactic Structures*. Studies in Logic, Language, and Computation. Stanford, CA: CSLI Publications.

Blackburn, Patrick, Maarten de Rijke, and Yde Venema. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science 53. Cambridge: Cambridge University Press.

Blackburn, Patrick, and Edith Spaan. 1993. A Modal Perspective on the Computational Complexity of Attribute Value Grammar. *Journal of Logic, Language, and Information* 2:129–169.

Borsley, Robert D. 1996. *Modern Phrase Structure Grammar*. Oxford: Blackwell.

Brachman, Ronald J. 1979. On the Epistemological Status of Semantic Networks. In *Associative Networks: Representation und Use of Knowledge by Computers*, ed. N. Findler. 3–50. New York: Academic Press.

Brachman, Ronald J., D. McGuinness, P. Patel-Schneider, L. Resnick, and A. Borgida. 1991. Living with CLASSIC: When and how to use a KL-ONE-like language. In *Principles of Semantic Networks*, ed. John F. Sowa. 401–456. San Mateo, CA: Morgan Kaufmann.

Brachman, Ronald J., and James G. Schmolze. 1985. An Overview of the KL-ONE Knowledge Representation System. *Cognitive Science* 9(2):171–216.

Brew, Chris. 1991. Systemic Classification and its Efficiency. *Computational Linguistics* 17(4):375–408.

Brink, Chris, Katarina Britz, and Renate A. Schmidt. 1994. Peirce Algebras. *Formal Aspects of Computing* 6:339–358.

Brink, Chris, and Ingrid Rewitzky. 2001. *A Paradigm for Program Semantics*. Studies in Logic, Language, and Computation. Stanford, CA: CSLI Publications.

Briscoe, Ted, Ann Copestake, and Valeria de Paiva (ed.). 1993. *Inheritance, Defaults, and the Lexicon*. Cambridge: Cambridge University Press.

Calder, Jo. 1999. The Horn Subset of Systemic Networks. In *Proceedings of the 6th Meeting on Mathematics of Language*, 203–212. Orlando, FL. University of Central Florida.

Carpenter, Bob. 1992. *The Logic of Typed Feature Structures*. Cambridge Tracts in Theoretical Computer Science 32. Cambridge: Cambridge University Press.

Carpenter, Bob. 1993. Skeptical and Credulous Default Unification with Applications to Templates and Inheritance. In Briscoe et al. 1993, 13–37.

Carpenter, Bob, and Paul John King. 1995. The Complexity of Closed World Reasoning in Constraint-Based Grammar Theories. In *Proceedings of the Fourth Meeting on Mathematics of Language*. University of Pennsylvania.

Carpenter, Bob, and Carl J. Pollard. 1991. Inclusion, Disjointness and Choice: The Logic of Linguistic Classification. In *Proceedings of the 29th Annual Meeting of the ACL*, 9–16.

Carpenter, Bob, and Richmond Thomason. 1990. Inheritance Theory and Path-Based Reasoning: An Introduction. In Kyburg et al. 1990, 309–343.

Caseau, Yves, Michel Habib, Lhouari Nourine, and Olivier Raynoud. 1999. Encoding of Multiple Inheritance Hierarchies and Partial Orders. *Computational Intelligence* 15(1):50–62.

Cederquist, Jan, and Thierry Coquand. 2000. Entailment Relations and Distributive Lattices. In *Lecture Notes in Logic, 13 – Logic Colloquium '98*, 110–123. A. K. Peters.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam, and Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.

Cignoli, R., L. Lafalce, and A. Petrovich. 1991. Remarks on Priestley Duality for Distributive Lattices. *Order* 8:299–315.

Cohn, Anthony G. 1992. Completing Sort Hierarchies. *Computers and Mathematics with Applications* 23(6–9):477–491.

Cohn, Paul M. 1982. *Universal Algebra*. Dordrecht: D. Reidel.

Copestake, Ann. 1993. Defaults in Lexical Representation. In Briscoe et al. 1993, 223–245.

Copestake, Ann. 2002. *Implementing Typed Feature Structure Grammars*. CSLI Lecture Notes, No. 110. Stanford, CA: CSLI Publications.

Coquand, Thierry, and Guo-Qiang Zhang. 2000. Sequents, Frames, and Completeness. In *Computer Science Logic, 14th Annual Conference of the EACSL*, 277–291. Lecture Notes in Computer Science 1862. Berlin. Springer.

Crawley, Peter, and Robert P. Dilworth. 1973. *Algebraic Theory of Lattices*. Englewood Cliffs, NJ: Prentice Hall.

Curry, Haskell B. 1963. *Foundations of Mathematical Logic*. New York: McGraw-Hill.

Dabrowski, Andrew, Lawrence S. Moss, and Rohit Parikh. 1996. Topological Reasoning and the Logic of Knowledge. *Annals of Pure and Applied Logic* 78:73–110.

Davey, Brian A. 1973. A Note on Representable Posets. *Algebra Universalis* 3:345–347.

Davey, Brian A., and Hilary A. Priestley. 1990. *Introduction to Lattices and Order*. Cambridge: Cambridge University Press.

Davey, Brian A., and Hilary A. Priestley. 1998. Distributive Lattices and Duality. In Grätzer 1998, 499–517.

Dawar, Anuj, and K. Vijay-Shanker. 1990. An Interpretation of Negation in Feature Structure Descriptions. *Computational Linguistics* 16(1):11–21.

de Paiva, Valeria. 1993. Types and Constraints in the LKB. In Briscoe et al. 1993, 164–189.

de Rijke, Maarten. 1995. The Logic of Peirce Algebras. *Journal of Logic, Language, and Information* 4:227–250.

Dimov, Georgi, and Dimiter Vakarelov. 1998. On Scott consequence systems. *Fundamenta Informaticae* 33(1):43–70.

Dörre, Jochen, Esther König, and Dov Gabbay. 1996. Fibred Semantics for Feature-Based Grammar Logic. *Journal of Logic, Language, and Information* 5:387–422.

Dörre, Jochen, and William C. Rounds. 1992. On Subsumption and Semi-Unification in Feature Algebras. *Journal of Symbolic Computation* 13(6):441–461.

Droste, Manfred. 1989. Event Structures and Domains. *Theoretical Computer Science* 68:37–47.

Droste, Manfred, and Rüdiger Göbel. 1990. Non-Deterministic Information Systems and their Domains. *Theoretical Computer Science* 75:289–309.

Dunn, J. Michael. 1986. Relevance Logic and Entailment. In Gabbay and Guenthner 1986, 117–224.

Dunn, J. Michael. 1995. Positive Modal Logic. *Studia Logica* 55:301–317.

Dunn, J. Michael, and Gary Hardegree. 2001. *Algebraic Methods in Philosophical Logic*. Oxford Logic Guides No. 41. Oxford: Oxford University Press.

Düntsch, Ivo. 1982. On a Problem of Chen and Grätzer. *Algebra Universalis* 14:401.

Eco, Umberto. 1984. *Semiotics and the Philosophy of Language*. London: MacMillan.

Eisenberg, Peter. 1999. *Grundriß der Deutschen Grammatik: Der Satz*. Stuttgart: J. B. Metzler.

Evans, Roger, and Gerald Gazdar. 1996. DATR: A Language for Lexical Knowledge Representation. *Computational Linguistics* 22(2):167–216.

Fall, Andrew. 1996. *Reasoning with Taxonomies*. Doctoral dissertation, School of Computing Science, Simon Fraser University.

Fiech, Adrian. 1996. Colimits in the category DCPO. *Mathematical Structures in Computer Science* 6(5):455–468.

Forbes, Graeme. 1994. *Modern Logic*. Oxford: Oxford University Press.

Gabbay, Dov M., and Franz Guenthner (ed.). 1986. *Handbook of Philosophical Logic, Vol. III: Alternatives to Classical Logic*. Dordrecht: D. Reidel.

Gabbay, Dov M., C. J. Hogger, and J. A. Robinson (ed.). 1994. *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford: Oxford University Press.

Gallier, Jean H. 1986. *Logic for Computer Science*. New York: Harper & Row.

Ganter, Bernhard. 1999. Attribute Exploration with Background Knowledge. *Theoretical Computer Science* 217:215–233.

Ganter, Bernhard, and Rudolf Wille. 1999. *Formal Concept Analysis*. Berlin: Springer.

Gärdenfors, Peter. 1994. How Logic Emerges from the Dynamics of Information. In van Eijck and Visser 1994, 49–77.

Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum, and Ivan Sag. 1985. *Generalized Phrase Structure Grammar*. Oxford: Blackwell.

Gehrke, Mai, and Bjarni Jónsson. 1994. Bounded Distributive Lattices with Operators. *Mathematica Japonica* 40(2):207–215.

Gierz, Gerhard, Karl Heinrich Hofmann, Klaus Keimel, Jimmie D. Lawson, Michael W. Mislove, and Dana S. Scott. 1980. *A Compendium of Continuous Lattices*. Berlin: Springer.

Girard, Jean-Yves, Yves Lafont, and Paul Taylor. 1989. *Proofs and Types*. Cambridge Tracts in Theoretical Computer Science 7. Cambridge: Cambridge University Press.

Godin, Robert, and Rokia Missaoui. 1994. An Incremental Concept Formation Approach for Learning from Databases. *Theoretical Computer Science* 133(2):387–419.

Goldblatt, Robert. 1976. Metamathematics of Modal Logic. In Goldblatt 1993, 9–79.

Goldblatt, Robert. 1989. Varieties of Complex Algebras. *Journal of Pure and Applied Logic* 44:173–242.

Goldblatt, Robert. 1993. *Mathematics of Modality*. CSLI Lecture Notes, No. 43. Stanford, CA: CSLI Publications.

Goldblatt, Robert. 2001. Mathematical Modal Logic: A View of Its Evolution. To appear in: *A History of Mathematical Logic*, ed. D. van Dalen, J. Dawson and A. Kanamori.

Graham, Ronald, Martin Grötschel, and László Lovász (ed.). 1995. *Handbook of Combinatorics*. Amsterdam: North-Holland.

Grätzer, George. 1979. *Universal Algebra*. New York: Springer. 2nd edition.

Grätzer, George. 1998. *General Lattice Theory*. Basel: Birkhäuser. 2nd edition.

Guigues, J.-L., and Vincent Duquenne. 1986. Familles minimales d'implications informatives resultant d'un tableau de données binaires. *Mathématiques et Sciences Humaines* 95:5–18.

Gunter, Carl A. 1985. Profinite Solutions for Recursive Domain Equations. Technical Report CMU-CS-85-107. Department of Computer Science, Carnegie Mellon University.

Gunter, Carl A., and Dana S. Scott. 1990. Semantic Domains. In van Leeuwen 1990b, 635–674.

Habib, Michel, and Lhouari Nourine. 1994. Bit-Vector Encoding for Partially Ordered Sets. In *Proceedings of the International Workshop on Orders, Algorithms, and Applications (ORDAL) '94*, ed. Vincent Bouchitteé and Michel Morvan, 1–12. Lecture Notes in Computer Science 831. Berlin. Springer.

Halliday, Michael A. K. 1994. *An Introduction to Functional Grammar*. London: Edward Arnold. 2nd edition.

Halmos, Paul R. 1955. Algebraic Logic, I. Monadic Boolean Algebras. In Halmos 1962, 37–72.

Halmos, Paul R. 1962. *Algebraic Logic*. New York: Chelsea.

Hansoul, Georges. 1983. A Duality for Boolean Algebras with Operators. *Algebra Universalis* 17:34–49.

Hartrumpf, Sven. 2002. *Hybrid Disambiguation in Natural Language Analysis*. Doctoral dissertation, FernUniversität Hagen.

Hartrumpf, Sven, and Hermann Helbig. 2002. The Generation and Use of Layer Information in Multilayered Extended Semantic Networks. In *Proceedings of the 5th International Conference on Text, Speech, and Dialogue (TSD-2002)*. Berlin: Springer.

Hartrumpf, Sven, Hermann Helbig, and Rainer Osswald. 2003. The Semantically Based Computer Lexicon HaGenLex – Structure and Technological Environment. *Traitement automatique des langues* 44(2). To appear.

Hayes, Patrick J. 1979. The Logic of Frames. In *Frame Conceptions and Text Understanding*, ed. Dieter Metzing. Berlin: de Gruyter.

Hegner, Stephen J. 1994. Properties of Horn Clauses in Feature-Structure Logic. In Rupp et al. 1994, 111–147.

Helbig, Hermann. 2001. *Die semantische Struktur natürlicher Sprache: Wissensrepräsentation mit MultiNet*. Berlin: Springer.

Helbig, Hermann, and Carsten Gnörlich. 2002. Multilayered Extended Semantic Networks as a Language for Meaning Representation in NLP Systems. In *Computational Linguistics and Intelligent Text Processing (CICLing 2002)*, ed. Alexander Gelbukh, 69–85. Lecture Notes in Computer Science 2276. Berlin. Springer.

Helbig, Hermann, and Sven Hartrumpf. 1997. Word Class Functions for Syntactic-Semantic Analysis. In *Proceedings of the 2nd International Conference on Recent Advances in Natural Language Processing (RANLP-97)*, 312–317.

Helbig, Hermann, and Marion Schulz. 1997. Knowledge Representation with MESNET – A Multilayered Extended Semantic Network. In *Proceedings of the AAAI Spring Symposium on Ontological Engineering*.

Hopcroft, John E., and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory, Languages, and Computation*. Reading, MA: Addison-Wesley.

Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press.

Horty, John F. 1994. Some Direct Theories of Nonmonotonic Inheritance. In Gabbay et al. 1994, 111–187.

Jakobson, Roman, and Morris Halle. 1956. *Fundamentals of Language*. Den Haag: Mouton.

Johnson, David E., and Lawrence S. Moss. 1994. An Overview of Stratified Feature Grammar. In *Current Issues in Mathematical Linguistics*, ed. Carlos Martín-Vide. 103–119. Amsterdam: North-Holland.

Johnson, Mark. 1991. Features and Formulae. *Computational Linguistics* 17(2):131–151.

Johnson, Mark. 1994a. Computing with Features as Formulae. *Computational Linguistics* 20(1):1–25.

Johnson, Mark. 1994b. Two Ways of Formalizing Grammars. *Linguistics and Philosophy* 17:221–248.

Johnstone, Peter T. 1982. *Stone Spaces*. Cambridge Studies in Advanced Mathematics 3. Cambridge: Cambridge University Press.

Jónsson, Bjarni, and Alfred Tarski. 1951. Boolean Algebras with Operators, Part I. *American Journal of Mathematics* 73:891–939.

Jung, Achim, Mathias Kegelmann, and M. Andrew Moshier. 1999. Multi Lingual Sequent Calculus and Coherent Spaces. *Fundamenta Informaticae* 37:369–412.

Kasper, Robert, and William C. Rounds. 1990. The Logic of Unification in Grammar. *Linguistics and Philosophy* 13:35–58.

Katz, Jerrold J., and Jerry A. Fodor. 1963. The Structure of a Semantic Theory. *Language* 39:170–210.

Keller, Bill. 1993. *Feature Logics, Infinitary Descriptions and Grammar*. CSLI Lecture Notes, No. 44. Stanford, CA: CSLI Publications.

King, Paul John. 1999. Towards Truth in Head-Driven Phrase Structure Grammar. In *Tübingen Studies in Head-Driven Phrase Structure Grammar*, 301–352. Arbeitspapiere des SFB 340, No. 132. University of Tübingen.

Koenig, Jean-Pierre. 1999. *Lexical Relations*. Stanford, CA: CSLI Publications.

Kornai, András. 1993. The Generative Power of Feature Geometry. *Annals of Mathematics and Artificial Intelligence* 8:37–46.

Kracht, Marcus. 1989. On the Logic of Category Definitions. *Computational Linguistics* 15(2):111–113.

Kracht, Marcus. 1995. Is There a Genuine Modal Perspective on Feature Structures? *Linguistics and Philosophy* 18:401–458.

Kracht, Marcus. 1999. *Tools and Techniques in Modal Logic*. Studies in Logic and the Foundations of Mathematics 142. Amsterdam: North-Holland.

Kress, Gunther R. (ed.). 1976. *Halliday: System and Function in Language*. London: Oxford University Press.

Krieger, Hans-Ulrich. 1994. Derivation without Lexical Rules. In Rupp et al. 1994, 277–313.

Krieger, Hans-Ulrich. 2001. Greatest Model Semantics for Typed Feature Structures. *Grammars* 4(2):139–165.

Krieger, Hans-Ulrich, and John Nerbonne. 1993. Feature-Based Inheritance Networks for Computational Lexicons. In Briscoe et al. 1993, 90–136.

Kyburg, Henry E., Ronald P. Loui, and Greg N. Carlson (ed.). 1990. *Knowledge Representation and Defeasible Reasoning*. Dordrecht: Kluwer.

Landman, Fred. 1991. *Structures for Semantics*. Dordrecht: Kluwer.

Lawvere, F. William, and Stephen Hoel Schanuel. 1997. *Conceptual Mathematics*. Cambridge: Cambridge University Press.

Link, Godehard. 1998. *Algebraic Semantics in Language and Philosophy*. CSLI Lecture Notes, No. 74. Stanford, CA: CSLI Publications.

Lyons, John. 1977. *Semantics, Vol. I*. Cambridge: Cambridge University Press.

Lyons, John. 1995. *Linguistic Semantics*. Cambridge: Cambridge University Press.

Mac Lane, Saunders. 1971. *Categories for the Working Mathematician*. New York: Springer.

Mac Lane, Saunders, and Ieke Moerdijk. 1992. *Sheaves in Geometry and Logic*. New York: Springer.

Magnan, François, and Gonzalo E. Reyes. 1994. Category Theory as a Conceptual Tool in the Study of Cognition. In *The Logical Foundations of Cognition*, ed. John Macnamara and Gonzalo E. Reyes. 57–90. Vancouver Studies in Cognitive Science, Vol. 4. Oxford: Oxford University Press.

Malouf, Robert, John Caroll, and Ann Copestake. 2000. Efficient Feature Structure Operations without Compilation. *Natural Language Engineering* 6(1):29–46.

Manes, Ernest G. 1976. *Algebraic Theories*. New York: Springer.

Mellish, Chris. 1988. Implementing Systemic Classification by Unification. *Computational Linguistics* 14(1):40–51.

Miller, Philip H. 1999. *Strong Generative Capacity: The Semantics of Linguistic Formalism*. Stanford, CA: CSLI Publications.

Mitchell, Tom M. 1997. *Machine Learning*. New York: McGraw-Hill.

Mönnich, Uwe. 1997. Zur bereichstheoretischen Semantik von Grammatikformalismen. Technical Report 111. Tübingen: University of Tübingen, SFB 340.

Möring, Rolf H. 1989. Computationally Tractable Classes of Ordered Sets. In *Algorithms and Order*, ed. I. Rival. 105–193. Dordrecht: Kluwer.

Moshier, M. Andrew. 1993. On Completeness Theorems for Feature Logics. *Annals of Mathematics and Artificial Intelligence* 8:175–213.

Moshier, M. Andrew. 1995. A Rational Construction of the Domain of Feature Structures. *Journal of Logic, Language, and Information* 4:111–143.

Moshier, M. Andrew. 1997. Is HPSG Featureless or Unprincipled? *Linguistics and Philosophy* 20:669–695.

Moshier, M. Andrew, and Carl J. Pollard. 1994. The Domain of Set-Valued Feature Structures. *Linguistics and Philosophy* 17:607–631.

Moss, Lawrence S. 1992. Completeness Theorems for Logics of Feature Structures. In *Logic from Computer Science*, ed. Yiannis N. Moschovakis. 387–403. New York: Springer.

Moss, Lawrence S., and Jerry Seligman. 1994. Classification Domains and Information Links: A Brief Survey. In van Eijck and Visser 1994, 112–124.

Nebel, Bernhard, and Gert Smolka. 1990. Representation and Reasoning with Attributive Descriptions. In *Sorts and Types in Artificial Intelligence*, ed. K. H. Bläsius, U. Hedtstück, and C.-R. Rollinger. 112–139. LNAI 418. Berlin: Springer.

Nebel, Bernhard, and Gert Smolka. 1991. Attributive Description Formalisms . . . and the Rest of the World. In *Text Understanding in LILOG*, ed. O. Herzog and C.-R. Rollinger. 439–452. LNAI 546. Berlin: Springer.

Negri, Sara, and Jan von Plato. 2001. *Structural Proof Theory*. Cambridge: Cambridge University Press.

Nerode, Anil, and Richard A. Shore. 1993. *Logic for Applications*. New York: Springer.

Oldager, Nikolaj. 2000. Taxonomies with Lattice Algebras. In *Proceedings of the 1st International Workshop on Ontologies and Lexical Knowledge Bases (OntoLex)*, ed. K. I. Simov and A. Kiryakov, 106–117.

Oles, Frank J. 2000. An Application of Lattice Theory to Knowledge Representation. *Theoretical Computer Science* 249:163–196.

Osswald, Rainer. 1999a. On Conditional Information in Feature-Based Theories. In *Proceedings of the 6th Meeting on Mathematics of Language*, 89–104. Orlando, FL. University of Central Florida.

Osswald, Rainer. 1999b. Semantics for Attribute-Value Theories. In *Proceedings of the Twelfth Amsterdam Colloquium*, ed. Paul Dekker, 199–204. Amsterdam. University of Amsterdam, ILLC.

Osswald, Rainer. 2001. Classifying Classification. In *Proceedings of the Joint Conference on Formal Grammar and Mathematics of Language (FG/MOL-01)*, ed. Geert-Jan Kruijff, Larry Moss, and Dick Oehrle. Electronic Notes in Theoretical Computer Science 53.

Osswald, Rainer. 2003a. Assertions, Conditionals, and Defaults. In *Proceedings of the 1st Workshop on Conditionals, Information, and Inference*, ed. Gabriele Kern-Isberner and Wilhelm Rödder. Lecture Notes in Artificial Intelligence. Berlin. Springer. To appear.

Osswald, Rainer. 2003b. Generic Ontology of Linguistic Classification. In *Foundations of the Formal Sciences II – Applications of Mathematical Logic in Philosophy and Linguistics*, ed. Benedikt Löwe, Wolfgang Malzkorn, and Thoralf Räsch. 203–212. Trends in Logic 17. Dordrecht: Kluwer.

Osswald, Rainer, and Wiebke Petersen. 2002. Induction of Classifications from Linguistic Data. In *Proceedings of the ECAI-Workshop on Advances in Formal Concept Analysis for Knowledge Discovery in Databases*. Lyon, July.

Osswald, Rainer, and Wiebke Petersen. 2003. A Logical Approach to Data Driven Classification. In *Proceedings of the German Conference on Artificial Intelligence (KI-2003)*. Lecture Notes in Artificial Intelligence. Berlin. Springer. To appear.

Palmgren, Erik, and Viggo Stoltenberg-Hansen. 1995. Logically Presented Domains. In *10th Annual IEEE Symposium on Logic in Computer Science*, 455–463. IEEE Computer Society Press.

Patten, Terry, and Graeme Ritchie. 1987. A Formal Model of Systemic Grammar. In *Natural Language Generation*, ed. G. Kempen. 279–299. Dordrecht: Kluwer.

Penn, Gerald. 2000. *The Algebraic Structure of Attributed Type Signatures*. Doctoral dissertation, Carnegie Mellon University, School of Computer Science.

Pereira, Fernando C. N., and Stuart M. Shieber. 1984. The Semantics of Grammar Formalisms Seen as Computer Languages. In *Proceedings of the 10th International Conference on Computational Linguistics*, 123–129.

Petersen, Wiebke. 2001. A Set-Theoretic Approach for the Induction of Inheritance-Hierarchies. In *Proceedings of the Joint Conference on Formal Grammar and Mathematics of Language (FG/MOL-01)*, ed. Geert-Jan Kruijff, Larry Moss, and Dick Oehrle. Electronic Notes in Theoretical Computer Science 53.

Petrovich, Alejandro. 1996. Distributive Lattices with an Operator. *Studia Logica* 56:205–224.

Poigné, Axel. 1992. Basic Category Theory. In Abramsky et al. 1992, 413–640.

Pollard, Carl J. 1989. Sorts in Unification-Based Grammar and What They Mean. Manuscript.

Pollard, Carl J. 1999. Strong Generative Capacity in HPSG. In *Lexical and Constructional Aspects of Linguistic Explanation*, ed. Gert Webelhuth, Jean-Pierre Koenig, and Andreas Kathol. 281–297. Stanford, CA: CSLI Publications.

Pollard, Carl J., and M. Andrew Moshier. 1990. Unifying Partial Descriptions of Sets. In *Information, Language, and Cognition*, ed. P. H. Hanson. 285–322. Vancouver Studies in Cognitive Science, Vol. 1. Oxford: Oxford University Press.

Pollard, Carl J., and Ivan A. Sag. 1987. *Information-Based Syntax and Semantics, Vol. 1*. CSLI Lecture Notes, No. 13. Stanford, CA: CSLI Publications.

Pollard, Carl J., and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Chicago: University of Chicago Press.

Pulman, Stephen G. 1996. Unification Encodings of Grammatical Notations. *Computational Linguistics* 22(3):295–327.

Quine, W. V. 1951. *Mathematical Logic*. Cambridge, MA: Harvard University Press. 2nd edition.

Quine, W. V. 1953. Logic and the Reification of Universals. In Quine 1961, 102–129.

Quine, W. V. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Quine, W. V. 1961. *From a Logical Point of View*. Cambridge, MA: Harvard University Press. 2nd edition.

Quine, W. V. 1969. *Set Theory and its Logic*. Cambridge, MA: Harvard University Press. 2nd edition.

Quine, W. V. 1976. Algebraic Logic and Predicate Functors. In *The Ways of Paradox and Other Essays*. 284–307. Cambridge, MA: Harvard University Press, 2nd edition.

Quine, W. V. 1982. *Methods of Logic*. Cambridge, MA: Harvard University Press. 4th edition.

Ramsay, Allan. 1990. Disjunction without Tears. *Computational Linguistics* 16(3):171–174.

Rasiowa, Helena. 1974. *An Algebraic Approach to Non-Classical Logics*. Amsterdam: North-Holland.

Rautenberg, Wolfgang. 1979. *Klassische und nichtklassische Aussagenlogik*. Braunschweig: Vieweg.

Reape, Mike. 1994. A Feature Value Logic with Intensionality, Nonwellfoundedness and Functional and Relational Dependencies. In Rupp et al. 1994, 77–110.

Rounds, William C. 1997. Feature Logics. In van Benthem and ter Meulen 1997, 475–533.

Rounds, William C., and Guo-Qiang Zhang. 1995. Suggestions for a Non-monotonic Feature Logic. Technical Report CS-R0551. Amsterdam: Centrum voor Wiskunde en Informatica.

Rounds, William C., and Guo-Qiang Zhang. 2001. Clausal Logic and Logic Programming in Algebraic Domains. *Information and Computation* 171(2):183–200.

Rupp, C. J., Michael A. Rosner, and Rod L. Johnson (ed.). 1994. *Constraints, Language and Computation*. San Diego, CA: Academic Press.

Russell, Bertrand. 1905. On Denoting. *Mind* 14:479–493.

Sag, Ivan A., and Thomas Wasow. 1999. *Syntactic Theory: A Formal Introduction*. Stanford, CA: CSLI Publications.

Schnelle, Helmut. 1973. *Sprachphilosophie und Linguistik*. Reinbek bei Hamburg: Rowohlt Taschenbuch Verlag.

Schubert, Horst. 1972. *Categories*. Berlin: Springer.

Schulz, Marion. 1999. *Eine Werkbank zur interaktiven Erstellung semantikbasierter Computerlexika*. Doctoral dissertation, FernUniversität Hagen.

Scott, Dana S. 1974. Completeness and Axiomatizability in Many-Valued Logic. In *Proceedings of the Tarski Symposium*, ed. Leon Henkin, John Addision, C. C. Chang, William Craig, Dana Scott, and Robert Vaught. 411–435. Providence, RI: AMS.

Scott, Dana S. 1982. Domains for Denotational Semantics. In *Proceedings of the 9th Colloquium on Automata, Languages and Programming*, ed. M. Nielsen and E. M. Schmidt, 577–613. Lecture Notes in Computer Science 140. Berlin. Springer.

Seda, Anthony Karel. 1995. Topology and the Semantics of Logic Programs. *Fundamenta Informaticae* 24(4):359–386.

Sells, Peter. 1985. *Lectures on Contemporary Syntactic Theories*. CSLI Lecture Notes, No. 3. Stanford, CA: CSLI Publications.

Shieber, Stuart M. 1992. *Constraint-Based Grammar Formalisms*. Cambridge, MA: MIT Press.

Sikkel, Klaas. 1997. *Parsing Schemata*. Berlin: Springer.

Smolka, Gert. 1992. Feature-Constraint Logics for Unification Grammars. *Journal of Logic Programming* 12:51–87.

Smolka, Gert, and Hassan Aït-Kaci. 1989. Inheritance Hierarchies: Semantics and Unification. *Journal of Symbolic Computation* 7:343–370.

Smyth, Michael B. 1983. Power Domains and Predicate Transformers: A Topological View. In *Proceedings of ICALP'83*, ed. Josep Díaz, 662–675. Lecture Notes in Computer Science 154. Berlin. Springer.

Smyth, Michael B. 1992. Topology. In Abramsky et al. 1992, 641–761.

Sofronie-Stokkermans, Viorica. 1999. On the Universal Theory of Varieties of Distributive Lattices with Operators: Some Decidability and Complexity Results. In *Proceedings of the 16th Conference on Automated Deduction*, 157–171. Lecture Notes in Artificial Intelligence 1632. Berlin. Springer.

Sofronie-Stokkermans, Viorica. 2000. On unification for bounded distributive lattices. In *Proceedings of the 17th Conference on Automated Deduction,* 465–481. Lecture Notes in Artificial Intelligence 1831. Berlin. Springer.

Sowa, John F. 2000. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Pacific Grove, CA: Brooks/Cole.

Speed, Terence P. 1972a. On the Order of Prime Ideals. *Algebra Universalis* 2:85–87.

Speed, Terence P. 1972b. Profinite Posets. *Bulletin of the Australian Mathematical Society* 6:177–183.

Stolcke, Andreas. 1989. Unification as Constraint Satisfaction in Structured Connectionist Networks. *Neural Computation* 1:559–567.

Stoltenberg-Hansen, Viggo, Ingrid Lindström, and Edward R. Griffor. 1994. *Mathematical Theory of Domains*. Cambridge Tracts in Theoretical Computer Science 22. Cambridge: Cambridge University Press.

Stoltenberg-Hansen, Viggo, and J. V. Tucker. 1995. Effective Algebras. In Abramsky et al. 1995, 357–526.

Stumme, Gerd. 1998. Distributive Concept Exploration – A Knowledge Acquisition Tool in Formal Concept Analysis. In *KI-98: Advances in Artificial Intelligence*, ed. Christian Retorè, 117–128. Lecture Notes in Artificial Intelligence 1504. Berlin. Springer.

Tarski, Alfred. 1946. *Introduction to Logic and to the Methodology of Deductive Sciences*. Oxford: Oxford University Press.

Taylor, John R. 1989. *Linguistic Categorization*. Oxford: Clarendon Press.

Thomason, Richmond H. 1997. Nonmonotonicity in Linguistics. In van Benthem and ter Meulen 1997, 777–831.

Thomason, Richmond H., John F. Horty, and David S. Touretzky. 1987. A Calculus for Inheritance in Monotonic Semantic Nets. In *Methodologies for Intelligent Systems*, ed. Z. W. Ras and M. Zemankova, 280–287. Amsterdam. Elsevier.

Tokizawa, Masamichi, and Tatsuro Kanki. 1985. Quotient Categories and Related Theorems of the Homomorphism Type. *Chuo University. Faculty of Science and Engineering. Bulletin* 28:83–100.

Treinen, Ralf. 1997. Feature Trees over Arbitrary Structures. In Blackburn and de Rijke 1997, 185–211.

Troelstra, A. S., and H. Schwichtenberg. 2000. *Basic Proof Theory*. Cambridge Tracts in Theoretical Computer Science 43. Cambridge: Cambridge University Press. 2nd edition.

Trotter, William T. 1995. Partially Ordered Sets. In Graham et al. 1995, 433–480.

van Benthem, Johan, and Alice ter Meulen (ed.). 1997. *Handbook of Logic and Language*. Amsterdam: North-Holland.

van Dalen, Dirk. 1986. Intuitionistic Logic. In Gabbay and Guenthner 1986, 225–339.

van Dalen, Dirk. 1994. *Logic and Structure*. Berlin: Springer. 3rd edition.

van Eijck, Jan, and Albert Visser (ed.). 1994. *Logic and Information Flow*. Cambridge, MA: MIT Press.

van Leeuwen, Jan (ed.). 1990a. *Handbook of Theoretical Computer Science, Vol. A*. Amsterdam: Elsevier.

van Leeuwen, Jan (ed.). 1990b. *Handbook of Theoretical Computer Science, Vol. B*. Amsterdam: Elsevier.

Vickers, Steven. 1989. *Topology via Logic*. Cambridge Tracts in Theoretical Computer Science 5. Cambridge: Cambridge University Press.

Vickers, Steven. 1992. Geometric Theories and Databases. In *Applications of Categories in Computer Science*, ed. M. P. Fourman, P. T. Johnstone, and A. M. Pitts, 288–314. London Mathematical Society Lecture Note Series 177. Cambridge. Cambridge University Press.

Vickers, Steven. 1993. Geometric Logic in Computer Science. In *Proceedings of the First Imperial College Department of Computing Workshop on Theory and Formal Methods*, ed. G. L. Burn, S. J. Gay, and M. D. Ryan, 37–54.

Vickers, Steven. 1999. Topology via Constructive Logic. In *Logic, Language, and Computation, Vol. 2*, ed. Jonathan Ginzburg, Lawrence S. Moss, and Maarten de Rijke. Stanford, CA: CSLI Publications.

Wechler, Wolfgang. 1992. *Universal Algebra for Computer Scientists*. Berlin: Springer.

Wedekind, Jürgen. 1994. Some Remarks on the Logic of Unification Grammars. In Rupp et al. 1994, 29–76.

Wille, Rudolf. 1992. Concept Lattices and Conceptual Knowledge Systems. *Computers and Mathematics with Applications* 23(6–9):493–515.

Winograd, Terry. 1983. *Language as a Cognitive Process: Syntax*. Reading, MA: Addison-Wesley.

Winskel, Glynn. 1988. An Introduction to Event Structures. In *Linear Time, Branching Time and Partial Order in Logics and Models of Concurrency*, ed. J. W. de Bakker, W.-P. Roever, and G. Rozenberg, 364–397. Lecture Notes in Computer Science 354. Berlin. Springer.

Winskel, Glynn. 1993. *The Formal Semantics of Programming Languages*. Cambridge, MA: MIT Press.

Winskel, Glynn, and Mogens Nielsen. 1995. Models for Concurrency. In Abramsky et al. 1995, 1–148.

Wintner, Shuly, and Nissim Francez. 1999. Off-Line Parsability and the Well-Foundedness of Subsumption. *Journal of Logic, Language, and Information* 8:1–16.

Wirsing, Martin. 1990. Algebraic Specification. In van Leeuwen 1990b, 675–788.

Woods, William W., and James G. Schmolze. 1992. The KL-ONE Family. *Computers and Mathematics with Applications* 23(2–5):133–177.

Young, Mark A., and William C. Rounds. 1993. A Logical Semantics for Nonmonotonic Sorts. In *Proceedings of the 31st Annual Meeting of the ACL*, 209–215.

Zajac, Rémi. 1992. Inheritance and Constraint-Based Grammar Formalisms. *Computational Linguistics* 18(2):159–182.

Zhang, Guo-Qiang. 1991. *Logic of Domains*. Basel: Birkhäuser.

Zhang, Guo-Qiang. 1994. A Representation of SFP. *Information and Computation* 110(1):233–263.

Index