Chapter 1

Introduction

1.1 Introduction

The field of *algebraic logic* assumed its modern systematic form, known as *abstract algebraic logic*, with the appearance of the pioneering "Memoirs" monograph of Blok and Pigozzi [35]. In this celebrated monograph one can find clearly discernible the seeds and the foundations of almost all subsequent developments in the field and, consequently, also, the foundations on which most parts of the work and of the developments detailed in the present monograph are based.

Related to the term "abstract algebraic logic", another of the pioneers of the field, Josep Maria Font, in a more recent textbook, titled "Abstract Algebraic Logic An Introductory Textbook" [86], advocates that the name should continue to be simply *algebraic logic* and that, as is the case with most other fields of Mathematics, Logic and Science, the abstraction, to which the term "abstract" refers, is part of the natural evolution of the same field, and should not be construed as constituting a special subfield justifying a special naming or rebranding.

In a similar sense, one may share the same belief for *categorical abstract algebraic logic*, which is also another natural evolution of algebraic logic and, therefore, according to this point of view, should also be referred to, simply, as *algebraic logic*. It may, in fact, be preferable to refer to the underlying formalizations of the logical systems treated in each particular context than to rebrand the entire field. So instead of referring to "abstract algebraic logic", we may say "algebraic logic as applied to sentential logics" (or "to deductive systems") and, similarly, "algebraic logic as applied to logics formalized as institutions or π-institutions", instead of using "categorical abstract algebraic logic" for the latter. For now, however, the traditional names have stuck and have been used widely, with well-discernible meanings, and we use them freely, as is also done in [86].

In "traditional" algebraic logic, which may be viewed to have started with the work of Tarski [5], the underlying formalism consists of *sentential logics* or *deductive systems*. These are pairs $\mathcal{S} = \langle \mathcal{L}, \vdash_{\mathcal{S}} \rangle$, where $\mathcal{L}$ is an algebraic language (a set of operation symbols with specified finite arities) and $\vdash_{\mathcal{S}}$ is a *consequence relation* on the absolutely free algebra $\mathbf{Fm}_{\mathcal{L}}(V)$ generated by a countable set V of variables. That is, $\vdash_{\mathcal{S}} \subseteq \mathcal{P}(\mathrm{Fm}_{\mathcal{L}}(V)) \times \mathrm{Fm}_{\mathcal{L}}(V)$, satisfies the following, for all $\Gamma \cup \Delta \cup \{\varphi\} \subseteq \mathrm{Fm}_{\mathcal{L}}(V)$,

Inflation: $\Gamma \vdash_{\mathcal{S}} \varphi$, for all $\varphi \in \Gamma$;

Monotonicity: $\Gamma \vdash_{\mathcal{S}} \varphi$ and $\Gamma \subseteq \Delta$ imply $\Delta \vdash_{\mathcal{S}} \varphi$;

Idempotency: $\Gamma \vdash_{\mathcal{S}} \varphi$ and $\Delta \vdash_{\mathcal{S}} \gamma$, for all $\gamma \in \Gamma$, imply, $\Delta \vdash_{\mathcal{S}} \varphi$;

Structurality: $\Gamma \vdash_{\mathcal{S}} \varphi$ implies $\sigma(\Gamma) \vdash_{\mathcal{S}} \sigma(\varphi)$, for all endomorphisms $\sigma : \mathbf{Fm}_{\mathcal{L}}(V) \to \mathbf{Fm}_{\mathcal{L}}(V)$.

Equivalently, $\mathcal{S}$ may be expressed in terms of a *structural closure operator* $C_{\mathcal{S}}$, i.e., a function $C_{\mathcal{S}} : \mathcal{P}(\mathrm{Fm}_{\mathcal{L}}(V)) \to \mathcal{P}(\mathrm{Fm}_{\mathcal{L}}(V))$, satisfying, for all $\Gamma \cup \Delta \subseteq \mathrm{Fm}_{\mathcal{L}}(V)$:

Inflation: $\Gamma \subseteq C_{\mathcal{S}}(\Gamma)$;

Monotonicity: $C_{\mathcal{S}}(\Gamma) \subseteq C_{\mathcal{S}}(\Delta)$, for all $\Gamma \subseteq \Delta$;

Idempotency: $C_{\mathcal{S}}(C_{\mathcal{S}}(\Gamma)) \subseteq C_{\mathcal{S}}(\Gamma)$;

Structurality: $\sigma(C_{\mathcal{S}}(\Gamma)) \subseteq C_{\mathcal{S}}(\sigma(\Gamma))$, for all endomorphisms $\sigma : \mathbf{Fm}_{\mathcal{L}}(V) \to \mathbf{Fm}_{\mathcal{L}}(V)$.

The equivalence is established by setting, on the one hand, for all $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}(V)$,

$$C_{\mathcal{S}}(\Gamma) = \{\varphi \in \mathrm{Fm}_{\mathcal{L}}(V) : \Gamma \vdash_{\mathcal{S}} \varphi\},$$

and, on the other, for all $\Gamma \cup \{\varphi\} \subseteq \mathrm{Fm}_{\mathcal{L}}(V)$,

$$\Gamma \vdash_{\mathcal{S}} \varphi \quad \text{iff} \quad \phi \in C_{\mathcal{S}}(\Gamma).$$

The reliance on sentential logics as the underlying formalism of the theory persists when passing to abstract algebraic logic. The reader is referred to the aforementioned [35, 86], as well as to the standard reference [64] by Janusz Czelakowski, another pioneer in the field, all clearly showcasing the primary role of this framework in all related developments and investigations.

By contrast, in this monograph the underlying logical formalism consists of π-institutions [33]. This formalism encompasses systems with varying signatures and quantifiers in a more direct way than allowed by the formalism of sentential logics (see Appendix C of [35], as well as the work on cylindric [15, 27] and polyadic algebras [9] and related work at the institutional level [100, 101, 102, 103] based and/or closely related to these). The structure of a π-institution forms a modification of the structure of an institution [25, 41], which was introduced in computer science to formalize logical systems for specification and programming, based on semantics. Diaconescu's monograph [79] offers a comprehensive advanced study of institutions and presents a multitude of model theoretic results that can be abstracted from first-order, and other specific logical systems, to the institutional level. On the other hand, in π-institutions, the framework is stripped of the semantic, or model theoretic, aspects and the focus is on the syntax, thus recovering the essential features of the sentential logic framework, without, however, shedding the versatility afforded, and the advantage gained, by incorporating in the object language multiple signatures and signature-changing morphisms. In fact, this inclusion is what gives the area its distinctive and unique character inside (abstract) algebraic logic. This is apparent in all aspects of our studies.

To make clearer the exact relationship between sentential logics and π-institutions, and showcase the fact that the former constitute very narrow special cases of the latter, let us recall the definition of a π-institution. A *π-institution*, as originally defined in [33], is a triple $\mathcal{I} = \langle \mathbf{Sign}, \mathrm{SEN}, C\rangle$, where

- $\mathbf{Sign}$ is an arbitrary category, whose objects are called *signatures* and its morphisms *signature morphisms*;
- $\mathrm{SEN} : \mathbf{Sign} \to \mathbf{Set}$ is a functor giving, for each signature $\Sigma \in |\mathbf{Sign}|$, the set $\mathrm{SEN}(\Sigma)$ of *Σ-sentences*;
- For every $\Sigma \in |\mathbf{Sign}|$, $C_\Sigma : \mathcal{P}(\mathrm{SEN}(\Sigma)) \to \mathcal{P}(\mathrm{SEN}(\Sigma))$ is a closure operator, such that the collection $C = \{C_\Sigma\}_{\Sigma\in|\mathbf{Sign}|}$ satisfies the property of *structurality*, i.e., for all $\Sigma, \Sigma' \in |\mathbf{Sign}|$, all $f \in \mathbf{Sign}(\Sigma, \Sigma')$ and all $\Phi \subseteq \mathrm{SEN}(\Sigma)$,

$$\mathrm{SEN}(f)(C_\Sigma(\Phi)) \subseteq C_{\Sigma'}(\mathrm{SEN}(f)(\Phi)).$$

In the modified (enriched) form that is used in the present monograph, and which was (essentially) introduced in [106], there is an additional component N, which represents a *category of natural transformations* on the sentence functor SEN. Roughly speaking, this category corresponds to clones of algebraic operations on $\{\mathrm{SEN}(\Sigma) : \Sigma \in |\mathbf{Sign}|\}$, under the assumption that all operations are defined uniformly and naturally over all $\mathrm{SEN}(\Sigma)$, for $\Sigma \in |\mathbf{Sign}|$. This accords in style with the algebraic theories of Lawvere [10], which are closely related to the Eilenberg-Moore [11] and the Kleisli [12] constructions. For more details on these, one may consult the classic texts by Mac Lane [16], Pareigis [14], Borceux [45] and Barr and Wells [57]. Thus, we are studying logical systems formalized as quadruples $\mathcal{I} = \langle \mathbf{Sign}, \mathrm{SEN}, N, C\rangle$, which are further recast as pairs

$$\mathcal{I} = \langle \mathbf{F}, C\rangle,$$

where

- $\mathbf{F} = \langle \mathbf{Sign}, \mathrm{SEN}, N\rangle$ expresses the algebraic structure, corresponding to the absolutely free algebra in the case of deductive systems, and
- $C = \{C_\Sigma\}_{\Sigma\in|\mathbf{Sign}|}$ is a family of closure operators, satisfying structurality, which is referred to as a *closure system*, and corresponds to the closure $C_\mathcal{S}$ in the case of sentential logics.

Suppose now that $\mathcal{S} = \langle \mathcal{L}, C_\mathcal{S}\rangle$ is a sentential logic. The standard rendering of it as a π-institution

$$\mathcal{I}^\mathcal{S} = \langle \mathbf{F}^\mathcal{L}, C^\mathcal{S}\rangle,$$

with $\mathbf{F}^\mathcal{L} = \langle \mathbf{Sign}^\mathcal{L}, \mathrm{SEN}^\mathcal{L}, N^\mathcal{L}\rangle$, is given by defining the four components as follows:

- $\mathbf{Sign}^\mathcal{L}$ is a trivial category, with object, say, V;

- $\mathrm{SEN}^{\mathcal{L}} : \mathbf{Sign}^{\mathcal{L}} \to \mathbf{Set}$ is given by $\mathrm{SEN}^{\mathcal{L}}(V) = \mathrm{Fm}_{\mathcal{L}}(V)$;
- $N^{\mathcal{L}}$ is the clone of all $\mathcal{L}$-operations on $\mathrm{Fm}_{\mathcal{L}}(V)$;
- $C^{\mathcal{S}}_V = C_{\mathcal{S}} : \mathcal{P}(\mathrm{Fm}_{\mathcal{L}}(V)) \to \mathcal{P}(\mathrm{Fm}_{\mathcal{L}}(V))$.

It is worth noting that $\mathbf{F}^{\mathcal{L}}$ only depends on $\mathcal{L}$ and V, as was to be expected (since it was deemed to correspond to the algebraic structure), and the deductive apparatus is reflected entirely in the definition of $C^{\mathcal{S}}$. Moreover, the formalism on the logical side does not incorporate substitutions in the object language, even though, since $C^{\mathcal{S}}_V = C_{\mathcal{S}}$ and the latter is structural, we have, for every endomorphism $\sigma : \mathbf{Fm}_{\mathcal{L}}(V) \to \mathbf{Fm}_{\mathcal{L}}(V)$,

$$\sigma(C^{\mathcal{S}}_V(\Phi)) \subseteq C^{\mathcal{S}}_V(\sigma(\Phi)),$$

for all $\Phi \subseteq \mathrm{Fm}_{\mathcal{L}}(V)$. On the algebraic side, on the other hand, e.g., when congruences are to be determined, the inclusion of the clone $N^{\mathcal{L}}$, reflecting the algebraic $\mathcal{L}$-structure, forces congruences at the institutional level to exactly correspond to the familiar $\mathcal{L}$-congruences on the formula algebra in the universal algebraic sense.

The reasons why one might want to develop a theory of algebraization for logical systems formalized as institutions or π-institutions parallel the motivations provided by Blok and Pigozzi [35] for developing a theory of algebraizability for sentential logics.

One of the main motivations is providing a classification of logical systems based on the strength of the ties of their deductive apparatuses with those corresponding to algebraic deductive systems, i.e., deductive systems whose closure systems are induced by algebraic structures. Preferably, when the definitions applicable in the context of logical systems formalized as π-institutions specialize in the way outlined above to π-institutions associated with deductive systems, one would be able to recover the well-known algebraic (or Leibniz) hierarchy of abstract algebraic logic [64, 86]. The finitary and finitely algebraizable sentential logics of [35] form a special class in this hierarchy. In [86], this property is termed *Blok-Pigozzi algebraizability* (see Definition 3.39 of [86]).

Another desideratum is that the definitions should be as general as possible so that, given virtually any π-institution, one would be able, at least in principle, to classify it in one or more of the classes of the hierarchy, based on the strength of its algebraic properties.

Further, an additional reassurance would be provided if the definitions supplied turned out to be robust in the sense that one would be able to obtain, at least for several, if not for most, of them, different characterizations depending on the various viewpoints taken. This was clearly and successfully undertaken in [35] for the class of algebraizable deductive systems. In fact, Blok and Pigozzi obtained several different characterizations whose variety and strength played a major role in convincing other researchers that their definitions were chosen wisely and, as a result, in establishing firmly the new trends in the field and, thus, contributing, in large part, to virtually all subsequent developments. It is hoped that pursuits along the same lines here will prove, at least moderately, successful with respect to similar criteria. In particular, it is hoped that the characterizations of many of the classes presented in this monograph in a variety of ways will prove to many of the readers and to, present and future, researchers in the field satisfactory and motivating, as was the case with the work of Blok and Pigozzi [35].

One last motivation, equally important, however, in significance, comes by taking an adversarial point of view. As Blok and Pigozzi realized when studying sentential logics, and is certainly true also for logics formalized as π-institutions, since they encompass sentential logics, is the fact that many logical systems of historical and/or practical significance failed to be amenable to classical methods of algebraization, such as, e.g., the Lindenbaum-Tarski process. Naturally one is inclined to ask whether those systems can be algebraized in some alternative way, using different techniques, or whether the failure in their algebraization is due to intrinsic reasons. That is, one would like to investigate whether those systems have some innate characteristics, e.g., pertaining to their structural properties, that many, if not all, of them share and that decide their algebraizability status. This is reminiscent of the extensive and intensive research in computational complexity theory in separating various complexity classes [94, 92, 95, 93, 96, 91], where common features and rigorous criteria are sought for classification of problems in hierarchies of complexity classes. As is the case there, such an analysis and rigorous classification presupposes the existence of a formal definition of algebraizability (and of other related properties) so as to delineate formal boundaries and establish criteria that could potentially be used to falsify claims of algebraizability for some logical systems. Such criteria would point to shortcomings and defects of some logical systems as related to qualitative requirements that a logic should satisfy in order to qualify for membership in a corresponding class. It is believed that the definitions adopted here are helpful in establishing such criteria and in setting up boundaries. The examples that are scattered throughout seem to support this assertion, but, of course, the jury is out as far as gathering further evidence in support of, or in criticism and opposition to, this claim.

The notion of algebraizability adopted in this monograph is inspired by the one established for deductive systems by Blok and Pigozzi in [35]. Apart from the technical complications inherent in passing from the sentential to the institutional framework, one substantial difference is that we distinguish between a treatment based on the Leibniz operator, referred to as **semantic**, as contrasted to the one based on interpretations from logic to algebra and vice-versa, which is termed **syntactic**, since it is based on natural transformations corresponding to term operations on the free algebra of terms. In the sentential logic framework, such a distinction is only apparent, since, as it turns out, the two approaches are equivalent and, hence, interchangeable. On the other hand, in π-institutions, the added flexibility afforded in the relation between morphisms (which are treated in the object language in the category of signatures) and clone operations (also part of the framework, but added a posteriori to enhance the algebraic character of the intended studies) means that the syntactic concepts dominate (i.e., are, in general, stronger) than their corresponding semantic counterparts.

The role that theories play in sentential logics is subsumed here by theory families, which consist of deductively closed sets of sentences, one for each signature. They form a complete lattice $\mathbf{ThFam}(\mathcal{I}) = \langle \mathrm{ThFam}(\mathcal{I}), \leq \rangle$, when ordered by signature-wise inclusion $\leq$. To each theory family is associated a congruence system, a collection of equivalence relations on formulas, one for each signature, that satisfy both the congruence property (or substitution property) and invariance under signature morphisms. These also form a complete lattice under signature-wise inclusions, which is denoted by $\mathbf{ConSys}(\mathcal{I}) = \langle \mathrm{ConSys}(\mathcal{I}), \leq \rangle$. The congruence system selected is the largest one compatible with the given theory family and is termed, by analogy with the sentential logic framework, the *Leibniz congruence system* associated with the theory family.

Starting from **semantics**, we say that a π-institution is *algebraizable* if it satisfies two conditions that impose very intimate ties between the lattice of theory families of the π-institution and that of the congruence systems determined by a class of algebraic systems. The first condition is that the Leibniz operator is monotone on theory families. The second is that it is order-reflecting.

On the **syntactic** side, a π-institution is *algebraizable* if, on the one hand, the Leibniz congruence systems are definable via a collection of natural transformations in two arguments and, on the other, if the theory families are definable via a collection of natural transformations in a single argument. In general, parametric arguments are allowed and, by restricting those, we obtain potentially narrower classes.

One of the main theorems established by Blok and Pigozzi in [35] is the characterization of algebraizable sentential logics via the existence of an isomorphism between the theory lattice of the deductive system and the equational theory lattice associated to a class of algebras, which also commutes with substitutions. A characterization along similar lines is established here for logical systems formalized as π-institutions (see, e.g., Section 4.3 or Section 12.3, even though other related forms appear in other places in the monograph, as will be discussed in the overview). In the literature several forms of this theorem and a host of generalizations of increasing power (or generality) have been discussed. A sample list includes [73, 35, 40, 99, 75, 81, 88]. The majority of these deal with deductive equivalence of logical systems, and related lattice-theoretic algebraic structures. They encompass the characterization of algebraizability mentioned above and deal with the case in which mutual interpretations between logical structures induce isomorphisms between lattices of theories and vice versa, under some constraints and special hypotheses, depending on the context under consideration.

Another major characterization theorem provided in [35] for the notion of algebraizability asserts that, roughly speaking, in the context of sentential logics, the aforementioned analogs of the semantic and of the syntactic notions are equivalent. That is, the algebraization attained via the definability of theories and congruences via sets of equations and formulas, respectively, coincides with that ensured by the Leibniz operator being monotone and order reflecting on the lattice of the theories of the logic. This characterization, when abstracted to logics formalized as π-institutions, continues to hold under special provisos, namely, under the hypotheses that the π-institution under consideration has a rich enough supply of natural transformations or, more formally, as will be studied in detail in the monograph, that it has a Leibniz binary reflexive core and an adequate Suszko core.

In [35] as well as in many other works in the field, a considerable amount of emphasis has been placed on, and a substantial amount of effort expended in, studying specific logical systems of historical and/or practical interest from the point of view of algebraizability. This was only natural, given, on the one hand, the desire to showcase the applicability of the theory on logics of particular interest in traditional studies, and, on the other, the urge to investigate the power of falsifiability that the theory provides for those concrete logical systems that had resisted previous attempts at algebraization.

Our point of view, however, is slightly different and, as a consequence, we do not deal with or present such examples. Firstly, the majority of logical systems of historical and/or practical interest have already been dealt with in existing literature. Secondly, since our treatment abstracts and subsumes that of sentential logics and, considerably generalizes it, as was shown above, our goal is not to look at the more concrete, already encompassed by the study of the algebraization of deductive systems, but, rather, to look into the more abstract and discern what can be carried over to that level and how its validity and its applicability compares when applied to new systems and new examples which do not fit exactly, or do not conform at all to, the sentential logic framework. However, these aims and the mode of treatment they motivate should in no way be construed as underestimating the significance, or underplaying the beauty and elegance, of the studies concerned with the concrete and the more specific. After all, it is on those studies that the abstract is based, to those studies' insights, ideas and methodology that an enormous scientific debt is due, and from those studies' successes, and widespread recognition and appreciation, that a relative confidence is drawn regarding the potential usefulness and applicability of the more general framework presented, and elaborated on, in this monograph.

1.2 Fin de siècle: The Golden Age

We give an account of some of the major developments in abstract algebraic logic that occurred mostly, but not exclusively, around the last two decades of the 20th century. This period may be thought of as constituting the golden age of algebraic logic, in the sense that, during this time, there is clearly discernible a passage from an ad-hoc, case-by-case algebraic treatment of logical systems to a well-organized field, with a powerful arsenal of universally applicable concepts, methods and techniques, culminating to the classification of logics in an algebraic hierarchy, known as the Leibniz hierarchy. Needless to say, the foundations for this success were laid much earlier. Likewise, the development continued, and many important results around, and complementing, the main theory were obtained later, into the new millennium, and the area continues to be active. In order to avoid, in our short exposition, reinventing the wheel, we base this account on preexisting sources. We draw the material primarily from the, perhaps best-known, survey of the field by Font, Jansana and Pigozzi [69] and, when needed and/or convenient, the two existing specialized books on the subject by Czelakowski [64] and by Font [86].

Algebraic logic has its origins in the work of George Boole [1, 2], who formalized the "laws of thought" in an algebraic way. The intuition governing this process was made mathematically precise by Tarski [5, 6, 8]. Tarski used the key idea of Lindenbaum of identifying formulas of a logical language with the terms of the absolutely free algebra formed using the logical connectives as operation symbols [3] to give a precise connection between classical propositional calculus and Boolean algebras. This formed the paradigmatic example from which significant inspiration was drawn and on which subsequent developments were based. Furthermore, it served as a kind of testbed for comparing, trying, modifying and calibrating new ideas, methods and techniques. The way Boolean algebras arose as the algebraic counterparts of classical propositional calculus has become known as the Lindenbaum-Tarski method. It has subsequently been used to "algebraize" a variety of propositional systems.

A conceptual shift occurred around 1950 when Rasiowa and Sikorski [7, 20] (see, also, the historical surveys [59, 74]), among others, realized that the Lindenbaum-Tarski method could be applied not only to isolated logics but, rather, to classes of logical systems with an implication connective satisfying certain properties. In passing from a "per logic" or "a la carte" treatment to one addressing classes specified by some abstract properties, one discerns clearly for the first time the seeds of what, later, became known as "abstract algebraic logic". Papers that may be thought of as protoabstract, in the sense that they advance further the main ideas of Rasiowa and Sikorski towards the modern truly abstract era, were the one by Prucnal and Wronski [19] introducing equivalential logics, the ones by Czelakowski introducing protoalgebraic logics [26, 29] and further studying equivalential logics [23, 24] and the one by Blok and Pigozzi [28] studying protoalgebraic logics.

The seismic shift, one might say, in firmly founding and establishing the modern era came in the 1980s with the work of Blok and Pigozzi, which led eventually to the publication of their famous, seminal "Memoirs" monograph [35]. In a way analogous to the preceding three passages, from classical logic and Boolean algebra to the Lindenbaum-Tarski method, from the Lindenbaum-Tarski method to dealing with implicative logics and from implicative logics to abstract properties of deduction, Blok and Pigozzi were able to distil the essential spirit of the association between logic and algebra and, thus, extract and formalize the concept of an algebraizable logic in modern abstract terms and provide landmark characterizations. On the way, they established a very general process of algebraization, applicable to arbitrary logical systems, which has been, since, further refined and used to create the Leibniz hierarchy, often considered the pinnacle - certainly a milestone and a gem - of algebraic logic in general.

Before returning to provide a more detailed account, we take a small break to recount those features of the theory that distinguish the abstract approach from the more traditional treatments and give it its special character. First, as alluded to previously, instead of applying the Lindenbaum-Tarski process in an ad-hoc way, on a case-by-case basis, or, as in Rasiowa's work, to a class of logics sharing a specific connective satisfying certain properties, it applies the abstract process to arbitrary sentential logics and, according to the outcome, classifies them into classes reflecting the closeness of the ties between them and the corresponding algebraic counterparts. In establishing this association and performing the resulting classification, it opens, in parallel, two distinct but closely interrelated directions. On the one hand, it motivates the study of classes of algebras arising as algebraic counterparts of either single or groups of logical systems. On the other, it allows investigating the exact correspondence between metalogical properties of the logical systems at hand and algebraic properties of the classes of their algebraic counterparts.

By now a plethora of works falling distinctly in each of these three directions exist and many will appear as references in the more detailed account that will follow. But to give some indication and pointers, we mention a few of the earliest ones that may be viewed as ground breaking. Concerning the process of algebraization itself and the classification, one should mention Blok and Pigozzi's [28, 35], Czelakowski's [23, 24], Herrmann's [43, 53, 54] and Font and Jansana's [52]. Concerning the study of classes of algebraic counterparts arising from the abstract algebraization process, one should mention [38, 39] dealing with the conjunction-disjunction fragment of classical propositional calculus, as well as Jansana's study of selfextensional logics in [71, 76], with clear precedents in Font and Jansana's [52]. Finally, paradigmatic examples of the study of metalogical and corresponding algebraic properties constitute several works addressing forms of the deduction-detachment theorem, e.g., Czelakowski's [26, 29] and Blok and Pigozzi's [32, 37, 63], the work of Blok and Hoogland on the Beth poperty [72], as well as the work of Czelakowski and Pigozzi concerning interpolation and amalgamation properties [58].

1.3 Outline of Contents by Chapter

We give an outline of the contents of the monograph focusing on the main points of each chapter and describing them by section, using some formal notation, but without providing formal definitions, which will be presented in the main body of the text. This section is very closely related to other sections. First, in Section 1.4, we give a very concise summary, only mentioning the main overarching topics discussed in each chapter. Second, at the beginning of each chapter, a similar overview is provided focusing only on the specific chapter, with the exception that, in those introductions, being closer to the formal treatment, an even more informal narrative is adopted and a concerted effort is made to keep notation at a minimum.

1.3.1 Chapter 2

Chapter 2 presents the basic elements of the theory of algebraic systems, of π-institutions and of the interaction between logical and algebraic structures. These constitute the foundations and form the backbone of our theory throughout the monograph.

Section 2.1 gives an informal introduction to the chapter, akin to the introduction presented here, only containing a little less of formal notation and being more on the narrative, informal, side.

Section 2.2 is the first main section of the chapter. Here, we start by introducing the notion of a *sentence functor* $\mathrm{SEN} : \mathbf{Sign} \to \mathbf{Set}$, which is simply a set-valued functor on an arbitrary category of signatures. It formalizes the carriers on which both algebras and logical systems are based, akin to the underlying universe of a universal algebra. Then we consider *sentence families* of sentence functors, which are families $T = \{T_\Sigma\}_{\Sigma \in |\mathbf{Sign}|}$ of subsets of sentences, one for each signature. These formalize distinguished sets of sentences when one considers logical structures, much like the distinguished sets in logical matrices. A *sentence system* is a sentence family T which is invariant under the action of signature morphisms. Two canonical ways of obtaining from a given sentence family T a sentence system consist of taking the largest sentence system $\overleftarrow{T}$ included in the family T and taking the smallest sentence system $\overrightarrow{T}$ that includes the sentence family T. Sentence functors are related via *sentence morphisms*, which are pairs $\langle F, \alpha \rangle$, F being a functor between the categories of signatures and α a natural transformation mapping sentences to sentences, taking into account the effect of

F. *Special morphisms* are those with surjective and full signature components and *surjective* ones are special ones whose sentence components are also surjective.

We then turn to *relation families* $R = \{R_\Sigma\}_{\Sigma \in |\mathbf{Sign}|}$ over sentence functors. Those assume the place of binary relations. Of the highest interest and importance are *equivalence families* and *equivalence systems*, i.e., equivalence families invariant under the action of signature morphisms. They induce partitions on the components of sentence functors. Equivalence families and systems have important interactions and connections with both sentence families and with morphisms. The notion that relates an equivalence family with a sentence family is that of *compatibility*. An equivalence family R is *compatible* with a sentence family T if each component of the sentence family is a union of blocks of the equivalence family on the same component. The connection between equivalence systems and morphisms goes through the notion of kernels. Namely, the *kernel* $\mathrm{Ker}(\langle F, \alpha\rangle)$ of a morphism $\langle F, \alpha\rangle$ between two sentence functors forms an equivalence system on the domain.

If a set is equipped with operations, we get an algebraic structure. On this algebraic structure, one may reason in an algebraic way about any of the operations that are in its clone, i.e., that can be generated by applying the fundamental operations and the projections and composing them in arbitrary ways. In an analogous fashion, if a sentence functor $\mathrm{SEN} : \mathbf{Sign} \to \mathbf{Set}$ is equipped with a category of natural transformations N, which corresponds to the clone of algebraic operations on an algebra, one obtains an *algebraic system* $\mathbf{A} = \langle \mathbf{Sign}, \mathrm{SEN}, N\rangle$. As algebras play a fundamental role in both logical and algebraic aspects of the traditional theory, so do algebraic systems in the theory developed in the monograph. The role of free algebra is played in this context by that of a *base algebraic system* $\mathbf{F} = \langle \mathbf{Sign}^\flat, \mathrm{SEN}^\flat, N^\flat\rangle$. Moreover, the notion of morphism extends from the context of sentence functors to the context of algebraic systems. The additional stipulation is that they also preserve the algebraic structure that turns the sentence functor into an algebraic system, i.e., that they satisfy the well-known *replacement* or *congruence condition*.

In traditional treatments, in specific contexts, all algebras are considered to be over the same algebraic signature, which is fully captured by the absolutely free algebra over that signature. In the present context, this similarity is captured by fixing a base algebraic system $\mathbf{F}$, as above, and considering only $\mathbf{F}$-*algebraic systems*, which are algebraic systems that, roughly speaking, have similar clones of operations with $\mathbf{F}$ and whose sentences are all images of sentences of $\mathbf{F}$ under a fixed algebraic system morphism $\langle F, \alpha\rangle$. Formally, these are expressed as pairs $\mathcal{A} = \langle \mathbf{A}, \langle F, \alpha\rangle\rangle$, where $\langle F, \alpha\rangle : \mathbf{F} \to \mathbf{A}$ is a surjective algebraic system morphism. The notion of morphism extends further to morphisms between $\mathbf{F}$-algebraic systems.

In Section 2.3, the limelight falls on *congruence systems*, which play in this context the same role that congruences play in the context of universal algebras. The least congruence system on an algebraic system $\mathbf{A}$ is the identity congruence system $\Delta^{\mathbf{A}}$ and the largest one is the full relation system, written $\nabla^{\mathbf{A}}$. These form the min and max elements, respectively, of the complete lattice of congruence systems $\mathbf{ConSys(A)}$ on $\mathbf{A}$. The kernel $\mathrm{Ker}(\langle F, \alpha\rangle)$ of a morphism $\langle F, \alpha\rangle : \mathbf{A} \to \mathbf{B}$ between two algebraic systems forms a congruence system on $\mathbf{A}$. Moreover, congruence systems allow the definition of *quotient algebraic systems*. And, for every algebraic system $\mathbf{A}$ and every one of its quotient systems $\mathbf{A}^\theta := \mathbf{A}/\theta$, there is a canonical morphism $\langle I, \pi^\theta\rangle : \mathbf{A} \to \mathbf{A}^\theta$ onto the quotient algebraic system, whose kernel is exactly the congruence system θ that gave rise to the quotient. All these properties reflect well known properties from the context of congruences and quotients of universal algebras.

Congruence systems inherit from equivalence families the relation of compatibility with given sentence families. The critical property to be established is that for a given sentence family T on an algebraic system $\mathbf{A}$, there exists a largest congruence system on $\mathbf{A}$ that is compatible with T. This is called the *Leibniz congruence system of* T *on* $\mathbf{A}$, is denoted by $\Omega^{\mathbf{A}}(T)$ and plays the role that Leibniz congruences play in the context of traditional abstract algebraic logic. As such, its role in characterizing many of the classes in the algebraic hierarchy studied in the monograph is ubiquitous and, as a consequence, the whole hierarchy is known as the *Leibniz hierarchy*. After introducing the *Leibniz operator* on an algebraic system, we establish two important results concerning it. The first, inspired by a result from the traditional treatment, provides a characterization of the Leibniz operator in terms of the category of natural transformations (i.e., clone operations) of the algebraic system and the sentence family. Roughly speaking it asserts that a pair of sentences are Leibniz related if and only if they are indistinguishable modulo T with respect to the available algebraic apparatus. The second addresses specifically the categorical framework and asserts that the Leibniz congruence system of a sentence family T is dominated by the Leibniz congruence system of the largest sentence system $\overleftarrow{T}$ contained in the sentence family, i.e., that $\Omega^{\mathbf{A}}(T) \leq \Omega^{\mathbf{A}}(\overleftarrow{T})$. The value of this observation in establishing refinements of the traditional hierarchy, as reflected in the present context, is critical and hard to overestimate. Also of importance is the fact that the surjective morphisms between algebraic systems, which form the focus of our work, respect Leibniz congruence systems, in the sense that, if $\langle F, \alpha\rangle : \mathbf{A} \to \mathbf{B}$ is a surjective morphism and T is a sentence family on $\mathbf{B}$, then $\Omega^{\mathbf{A}}(\alpha^{-1}(T)) = \alpha^{-1}(\Omega^{\mathbf{B}}(T))$. Finally, it is worth noting that, in general, the intersection of the Leibniz congruence systems of a collection of sentence families is contained in the Leibniz congruence of the intersection of those sentence families. Significantly, though, the reverse inclusion holds universally on sentence families if and only if the Leibniz operator is monotone on sentence families, a property that does not always hold. In fact, the latter property is used in a critical way, when restricted to special kinds of sentence families, to determine some of the most important classes of logical systems located close to the bottom of the algebraic hierarchy. In addition, it is of great historical significance in many of the most important classical developments in the field.

In Section 2.4, we focus on *congruence systems relative to given classes of algebraic systems*. Given a class K of algebraic systems, all over the same base algebraic system, that is, possessing, in some sense, the same algebraic signature, a congruence system θ on an algebraic system $\mathbf{A}$, not necessarily belonging to K, is called a *congruence system relative to* K or a K-*congruence system* if the quotient $\mathbf{A}^\theta$ belongs to the class K. Naturally, if $\mathbf{A} \in \mathsf{K}$ and the class K happens to be closed under morphic images, then congruence systems relative to K coincide with arbitrary congruence systems. The section introduces another important notion in this context. That of an algebraic system $\mathbf{A}$ being a *subdirect intersection* of a collection of algebraic systems. This means that there exists surjective morphisms $\langle H^i, \gamma^i \rangle : \mathbf{A} \to \mathbf{A}^i$ from the algebraic system to each of the algebraic systems in the given collection and, moreover, the intersection of the kernels of those morphisms is the identity congruence on $\mathbf{A}$. Closure of a class of algebraic systems under subdirect intersections ensures that the collection of congruence systems relative to the class is closed under intersections. Additionally, if the class K contains a trivial algebraic system, then the nabla congruence system happens to be a relative congruence system. Therefore, possession of a trivial algebraic system together with closure under subdirect intersections ensures that the collection of all congruence systems relative to the class forms a complete lattice under signature-wise inclusion.

Suppose that the class K contains a trivial algebraic system and is closed under subdirect intersections so that it makes sense to associate with a given relation family X on its base algebraic system the least congruence system $\Theta^{\mathsf{K}}(X)$ relative to K containing X. An alternative, equally natural, way to associate a congruence system with X is to consider the closure $D^{\mathsf{K}}(X)$ under equational consequence relative to the algebraic systems in the class K. It is proven in this section that the two closures give rise to the same congruence system on the base algebraic system $\mathbf{F}$.

In Section 2.5, we study *varieties of algebraic systems*. There are two possibilities in adopting a choice for the entities that would play the role of equations in this context. The first is to view pairs of sentences as equations. The second is to adopt pairs of natural transformations in the clone as equations. The ones of the latter type are called *natural equations* to differentiate them from those of the former kind which are simply referred to as *equations*. We define formally the notion of *satisfiability* of a given equation and of a given natural equation in an algebraic system and that of validity of a natural equation in an algebraic system. Depending on whether we use equations or natural equations to determine a class of algebraic systems through satisfiability, we obtain two different kinds of varieties. Varieties determined by families of equations are called *semantic varieties*. Those determined by collections of natural equations are called *syntactic varieties*. It turns out that, in general, every syntactic variety is also a semantic variety. The opposite implication does not hold in general. The section concludes by presenting a sufficient condition on the structure of a base algebraic system that ensures that the classes of semantic and syntactic varieties over it coincide.

Much of the work in the first sections of Chapter 1 focuses on the algebraic framework that underlies both the logical and the algebraic aspects of the theory in the monograph. In Section 2.6, we turn to the study of *π-institutions*, the underlying structure of the logical aspects of our theory. The entire monograph assumes that all logical systems are formalized as π-institutions and its main goal is to study the process of their algebraization and to detail the various classes in the hierarchy that is formed by examining their algebraic character. It is needless, thus, to point out the importance of Section 2.6, as it presents the foundational aspects of the logical side of the theory.

We start, here, by defining the notion of *π-institution*. It is a pair $\mathcal{I} = \langle \mathbf{F}, C \rangle$ consisting of a base algebraic system $\mathbf{F}$ and a closure system C on the sentence functor of $\mathbf{F}$. It generalizes the Tarskian concept of a deductive system in that it allows multiple signatures and accommodates morphisms between signatures. To take into account the logical structure imposed on top of the underlying algebraic structure in this context, sentence families and systems are subsumed by *theory families* and *theory systems*. These are sentence families (systems, respectively) each of whose components is closed under logical deduction. The least among these is called the *theorem system* of $\mathcal{I}$. It turns out that, due to the property of structurality, which is key in the study of π-institutions, given a theory family T, $\overleftarrow{T}$ is also closed under deduction, whence it forms that largest theory system included in T. On the other hand, $\overrightarrow{T}$ fails to be closed under deduction in general. That is the reason why the smallest theory system including T is not simply $\overrightarrow{T}$ but, rather, $C(\overrightarrow{T})$.

An important derived concept is that of the π-institution that has as its theory families those theory families of $\mathcal{I} = \langle \mathbf{F}, C \rangle$ which include a given theory system T of $\mathcal{I}$. This is denoted by $\mathcal{I}^T = \langle \mathbf{F}, C^T \rangle$. The construction results in a π-institution whose theorem system is identical with the theory system T of $\mathcal{I}$.

As is the case in most mathematical contexts, objects are accompanied by morphisms between them that preserve the structure of interest in each particular context. *Morphisms between π-institutions* are algebraic morphisms between the underlying algebraic systems that, in addition, preserve the logical structure in the sense that the forward image of the logical closure of a set of sentences is included in the closure of the image of the same set of sentences. Among the most useful characterizations is that a given algebraic morphism is logical if and only if preimages of theory families of the target institution under the morphism constitute theory families of the domain π-institution.

In Section 2.7, we turn to those structures that are intermediate between logic and algebra and facilitate the interplay and the establishment of meaningful ties between the two domains. These are *matrix families*, which correspond to the ordinary logical matrices in the traditional treatment. Roughly speaking a *matrix family* $\mathfrak{A} = \langle \mathcal{A}, T \rangle$ consists of an algebraic system $\mathcal{A}$ together with a sentence family T of the algebraic system. If the sentence family is a system, i.e., invariant under signature morphisms, then the matrix family is called a *matrix system*. Their role is twofold. On the one hand, a given collection of matrix families M, over a base algebraic system $\mathbf{F}$, may be used to define a closure system C^{M}, and hence a π-institution structure $\mathcal{I}^{\mathsf{M}} = \langle \mathbf{F}, C^{\mathsf{M}} \rangle$, on $\mathbf{F}$. On the other, given a π-institution structure $\mathcal{I}$ on $\mathbf{F}$, we may define the class $\mathrm{MatFam}(\mathcal{I})$ of all matrix families whose sentence families are closed under the deductive apparatus of the π-institution. Such sentence families are termed *$\mathcal{I}$-filter families* and, if they happen to be systems, then they are called *$\mathcal{I}$-filter systems*. The collection $\mathrm{FiFam}^{\mathcal{I}}(\mathcal{A})$ of all filter families over the same underlying algebraic system $\mathcal{A}$, ordered by component-wise inclusion, forms a complete lattice and the collection of all filter systems on that same algebraic system forms a complete sublattice of the complete lattice of all filter families.

Among the main results presented in this section are the ones relating morphisms between algebraic systems with preservation of filter families. More precisely, the inverse image of a filter family is a filter family. The situation is more complicated when it comes to direct images. First of all, it only makes sense to define the direct image of a filter family in case the signature functor is an isomorphism. Second, it turns out that, in that case, for the image to also be a filter family on the target algebraic system, we must require additional restrictions. A sufficient condition is that the kernel system of the algebraic morphism be compatible with the filter family in the domain.

This result has particular consequences for the most important type of morphisms considered in the monograph, the canonical quotient morphisms associated with congruence systems on an algebraic system. It asserts that, given a filter family T on the quotient $\mathcal{A}^\theta$, the inverse image $\pi^{\theta^{-1}}(T)$ under the quotient morphism $\langle I, \pi^\theta \rangle : \mathcal{A} \to \mathcal{A}^\theta$ is a filter family on $\mathcal{A}$ and that, moreover, if the congruence system θ is compatible with a filter family T on $\mathcal{A}$, then the quotient T/θ is a filter family on $\mathcal{A}^\theta$.

We consider, by particularizing even further, the Leibniz quotient morphisms, which are those morphisms defined using the Leibniz congruence system that is compatible with a given filter family on the domain. Since, by definition, the Leibniz congruence system $\Omega^{\mathcal{A}}(T)$ associated with a given sentence family T is compatible with that sentence family, it follows that a filter family T on $\mathcal{A}$ gives rise, by passing to the Leibniz quotient $\mathcal{A}/\Omega^{\mathcal{A}}(T)$, to a filter family in the quotient. The corresponding matrix family $\mathfrak{A}/\theta = \langle \mathcal{A}/\Omega^{\mathcal{A}}(T), T/\Omega^{\mathcal{A}}(T)\rangle$ is called a (*Leibniz*) *reduced matrix family*.

The section closes by defining two classes of matrix families and two classes of algebraic systems that play a key role when investigating the algebraic nature of a given π-institution $\mathcal{I}$. The first is the class $\mathrm{MatFam}^*(\mathcal{I})$ of all *Leibniz reduced matrix families* associated with the given π-institution. The second is the class $\mathrm{MatSys}^*(\mathcal{I})$ of all *Leibniz reduced matrix systems*. Finally, on the algebraic side, by considering all algebraic system reducts of the reduced matrix families, we get the class $\mathrm{AlgSys}^*(\mathcal{I})$ of all *family reduced algebraic systems* and, by considering all algebraic system reducts of the reduced matrix systems, we get the class $\mathrm{AlgSys}^\bullet(\mathcal{I})$ of all *system reduced algebraic systems*.

Section 2.8 studies the two related concepts of *axiomatic extensions* and of *filter extensions*. An *axiomatic extension* $\mathcal{I}'$ of a given π-institution $\mathcal{I}$ is a π-institution over the same base algebraic system whose closure system is obtained by that of $\mathcal{I}$ by adding more axioms. More precisely, the consequences $C'(X)$ of a family of sentences X under $\mathcal{I}'$ are the consequences under $\mathcal{I}$ of the same family of sentences, augmented by some fixed family of sentences T, i.e., $C'(X) = C(X \cup T)$. The sentences in T are viewed as axioms in $\mathcal{I}'$. A *filter extension* arises in a similar way. The difference is that one considers filter families over arbitrary algebraic systems and not just theory families over the base algebraic system.

One of the first results in this section provides a characterization of axiomatic extensions. It asserts that axiomatic extensions are characterized by preservation of all those theories that include the theorem system of the extension. An alternative, lifting the condition to arbitrary algebraic systems, asserts that being an axiomatic extension is tantamount to the preservation of filterhood over all algebraic systems, for all those filters that include the least filter over the extension.

The last part of the section deals with *filter generation* over a given matrix family modulo a given π-institution $\mathcal{I}$. It defines the concept and formalizes, in a rather technical proposition, how generation of filters and surjective matrix family morphisms interact.

Section 2.9 turns the focus back to those structures that, like matrix families, play a critical role as intermediate structures in connecting the logical with the algebraic aspects of the theory. *Generalized matrix families* correspond to the generalized matrices of classical algebraic logic and, like generalized matrices, play a critical role in identifying classes of algebraic systems that may be naturally associated with given π-institutions (or classes of π-institutions). The way this association is established sheds light on the strength of ties between the two and on the nature of their interaction, e.g., by revealing which properties may be expected to be shared by the two or transferred from one to the other.

A *generalized matrix family* $\mathbb{A} = \langle \mathcal{A}, \mathcal{T} \rangle$ consists of an underlying algebraic system $\mathcal{A}$ and a collection of sentence families $\mathcal{T}$ of the algebraic system. Such structures may also be used in two ways. They may serve in defining a closure system on a base algebraic system and, therefore, a π-institution structure. On the other hand, given a π-institution $\mathcal{I}$, we may associate with it the collection $\mathrm{GMatFam}(\mathcal{I})$ of those generalized matrices all of whose sentence families are filter families of the π-institution. With any generalized matrix family $\mathbb{A} = \langle \mathcal{A}, \mathcal{T} \rangle$, one may associate its *Tarski congruence system* $\widetilde{\Omega}(\mathbb{A})$ or $\widetilde{\Omega}^{\mathcal{A}}(\mathcal{T})$, an abstraction of the Tarski congruence systems associated with generalized matrices in classical abstract algebraic logic. *Tarski congruence systems* constitute the largest congruence systems on the base algebraic system compatible with all sentence families of the generalized matrix family. Taking the quotient $\mathbb{A}/\widetilde{\Omega}(\mathbb{A})$ of the generalized matrix family by its Tarski congruence system gives a new generalized matrix family $\mathbb{A}^*$, which is called the *Tarski reduction* of $\mathbb{A}$. A *Tarski reduced matrix family* is one that is isomorphic to its reduction, i.e., one whose Tarski congruence system is the identity congruence system on the underlying algebraic system.

There is a close connection between Tarski congruence systems and Leibniz congruence systems. Each generalized matrix system $\mathbb{A} = \langle \mathcal{A}, \mathcal{T} \rangle$ may be viewed as a bundle of matrix families $\{\langle \mathcal{A}, T \rangle : T \in \mathcal{T}\}$, i.e., of those matrix families whose sentence families belong to the collection of sentence families of the generalized matrix family. In that case, the Tarksi congruence system of the generalized matrix family is the intersection (in the component-wise sense) of the Leibniz congruence systems of all matrix families in the corresponding bundle, i.e., $\widetilde{\Omega}^{\mathcal{A}}(\mathcal{T}) = \bigcap_{T \in \mathcal{T}} \Omega^{\mathcal{A}}(T)$.

In a similar way to Tarski congruence systems, one may also consider *Suszko congruence systems* $\widetilde{\Omega}^{\mathcal{A},\mathcal{T}}(T)$ associated with ordinary matrix families $\mathfrak{A} = \langle \mathcal{A}, T \rangle$, and these are also introduced in Section 2.9. Suszko congruence systems of matrix families are defined only in a relative way, by viewing the matrix family $\mathfrak{A} = \langle \mathcal{A}, T \rangle$ as being part of a bundle expressed as a generalized matrix family $\mathbb{A} = \langle \mathcal{A}, \mathcal{T} \rangle$. Then the *Suszko congruence system* of the matrix family is identical to the Tarski congruence system $\widetilde{\Omega}^{\mathcal{A}}(\mathcal{T}^T)$ of the bundle $\langle \mathcal{A}, \mathcal{T}^T \rangle$ consisting of only those sentence families that include (in the component-wise ordering) the sentence family T of the matrix family. Of course, expressed in terms of Lebniz congruence systems, the Suszko congruence system is the intersection of the Leibniz congruence systems of all matrix families determined by the sentence families in the given bundle that include that of the matrix family under consideration, i.e., $\widetilde{\Omega}^{\mathcal{A},\mathcal{T}}(T) = \bigcap_{T \leq T' \in \mathcal{T}} \Omega^{\mathcal{A}}(T')$. As was the case with Tarski congruence systems, we may consider the *Suszko reduction* $\mathfrak{A}^{\mathsf{Su}}$ of a given matrix family $\mathfrak{A}$, obtained by dividing out by the Suszko congruence system $\widetilde{\Omega}^{\mathcal{A},\mathcal{T}}(T)$. And, likewise, we call a matrix family *Suszko reduced*, when its Suszko congruence system is the identity congruence system on the underlying algebraic system.

Part of the significance of the Tarski and of the Suszko operators in algebraic logic is that they form one of the main mechanisms of selecting the "natural" class of algebraic systems to be associated with a given π-institution. Briefly and sketchily, starting from a π-institution $\mathcal{I}$, we obtain the collection $\mathrm{GMatFam}(\mathcal{I})$ of all generalized matrix families $\mathbb{A} = \langle \mathcal{A}, \mathcal{T} \rangle$ whose sentence families $T \in \mathcal{T}$ are filter families of the π-institution. We then compute the Tarski reductions $\mathbb{A}^*$ by dividing out by the corresponding Tarski congruences $\widetilde{\Omega}^{\mathcal{A}}(\mathcal{T})$. This process gives rise to the class $\mathrm{GMatFam}^*(\mathcal{I})$ of all Tarski reduced generalized matrix families and to the class $\mathrm{AlgSys}(\mathcal{I})$ of all their underlying algebraic systems. This class subsumes, in the π-institution framework, the class of algebras which has long been viewed, in the traditional framework, as the most appropriate one to be associated with a given logic and, hence, as constituting the "natural" choice for the algebraic counterpart of the sentential logic. As it turns out, using a similar path, but relying on the Suszko operator, rather than on the Tarski operator, gives rise to exactly the same class of algebraic systems. Tracing the analogous process, one starts from a given π-institution $\mathcal{I}$ and considers all matrix families $\mathfrak{A} = \langle \mathcal{A}, T \rangle$, viewed as part of the bundle $\mathbb{A} = \langle \mathcal{A}, \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \rangle$ of all matrix families associated with the π-institution. Then, one considers the Suszko reductions $\mathfrak{A}^{\mathsf{Su}}$ by dividing out by the corresponding Suszko congruence systems $\widetilde{\Omega}^{\mathcal{I},\mathcal{A}}(T)$. The class of Suszko reduced matrix families obtained in this way is denoted by $\mathrm{MatFam}^{\mathsf{Su}}(\mathcal{I})$. It can then be shown that the class of all algebraic reducts of the matrix families in $\mathrm{MatFam}^{\mathsf{Su}}(\mathcal{I})$ coincides with the class $\mathrm{AlgSys}(\mathcal{I})$.

In Sections 2.7 and 2.9, using the classes of Leibniz reduced matrix families and of Tarski reduced generalized matrix families associated with a given π-institution $\mathcal{I}$, we are able to define the two classes $\mathrm{AlgSys}^*(\mathcal{I})$ and $\mathrm{AlgSys}(\mathcal{I})$ of algebraic systems associated with the π-institution. In Section 2.10, we take up the study of two additional classes of algebraic systems that may be perceived as counterparts of a given π-institution and compare them with those already defined.

Both new classes are based on a single algebraic system, namely the algebraic system $\mathcal{F}/\widetilde{\Omega}(\mathcal{I})$ resulting by considering the quotient of the base algebraic system $\mathcal{F}$ by the Tarski congruence of the collection of all theory families of $\mathcal{I}$. Using this quotient algebraic system, the two classes are formed as the two kinds of varieties that may be generated by it. The first type, called the *semantic variety*, denoted by $\mathbb{V}^{\mathsf{Sem}}(\mathcal{I}) = \mathbb{V}^{\mathsf{Sem}}(\mathcal{F}/\widetilde{\Omega}(\mathcal{I}))$, is the class of all algebraic systems that satisfy all equations that are satisfied by $\mathcal{F}/\widetilde{\Omega}(\mathcal{I})$, i.e., all equations included in $\widetilde{\Omega}(\mathcal{I})$. The second type, called the *syntactic variety*, denoted by $\mathbb{V}^{\mathsf{Syn}}(\mathcal{I}) = \mathbb{V}^{\mathsf{Syn}}(\mathcal{F}/\widetilde{\Omega}(\mathcal{I}))$, is the class of all algebraic systems that satisfy all natural equations that are satisfied by $\mathcal{F}/\widetilde{\Omega}(\mathcal{I})$.

Some results relating the four classes are presented. There is a linear hierarchy of inclusions that is not very difficult to establish. The class $\mathrm{AlgSys}^*(\mathcal{I})$ is the smallest class, followed by $\mathrm{AlgSys}(\mathcal{I})$, which is, in turn, included in

$\mathbb{V}^{\mathsf{Sem}}(\mathcal{I})$, whereas $\mathbb{V}^{\mathsf{Syn}}(\mathcal{I})$ is the largest of the four classes considered. It turns out that all four classes generate the same syntactic variety of algebraic systems, which is identical to $\mathbb{V}^{\mathsf{Syn}}(\mathcal{I})$, since it constitutes already a syntactic variety by definition. The section concludes with an important result showing that the class $\mathrm{AlgSys}(\mathcal{I})$ - perhaps the most important class associated with $\mathcal{I}$ - is closed under subdirect intersections and contains a trivial algebraic system. The usefulness of this fact is that it enables consideration, on any given algebraic system, of the least congruence system relative to $\mathrm{AlgSys}(\mathcal{I})$ generated by a prespecified relation family.

In Section 2.11, we study *equivalence families* and *systems* that are induced by sentence families or collections of sentence families of an algebraic system. The most fundamental among these is the *Frege equivalence family* $\lambda^{\mathbf{A}}(T)$ associated with a sentence family T of an algebraic system $\mathbf{A}$. It identifies two sentences if they are both inside or both outside the sentence family. Its companion *Frege equivalence system* $\Lambda^{\mathbf{A}}(T)$ is the largest equivalence system included in $\lambda^{\mathbf{A}}(T)$. The two Frege equivalences are intimately connected with the Leibniz congruence system $\Omega^{\mathbf{A}}(T)$, the latter being the largest congruence system contained in either of $\lambda^{\mathbf{A}}(T)$ or $\Lambda^{\mathbf{A}}(T)$.

In a way analogous to the extensions of the Leibniz congruence system that give rise to the Tarski and Suszko congruence systems, the Frege relations give rise to two more equivalences with similar roles. Given a collection $\mathcal{T}$ of sentence families of $\mathbf{A}$, the *Carnap equivalence family* $\widetilde{\lambda}^{\mathbf{A}}(\mathcal{T})$ identifies two sentences if they are equivalent modulo T (in the Frege sense) for all $T \in \mathcal{T}$, i.e., $\widetilde{\lambda}^{\mathbf{A}}(\mathcal{T}) = \bigcap_{T\in\mathcal{T}} \lambda^{\mathbf{A}}(T)$. The *Carnap equivalence system* $\widetilde{\Lambda}^{\mathbf{A}}(\mathcal{T})$ is the largest equivalence system included in $\widetilde{\lambda}^{\mathbf{A}}(\mathcal{T})$. The relation connecting Leibniz congruence systems with the Frege equivalences persists here as well, but with the Suszko congruence system in place of the Leibniz one. That is, the Suszko congruence system $\widetilde{\Omega}^{\mathbf{A}}(\mathcal{T})$ is the largest congruence system contained in either of $\widetilde{\lambda}^{\mathbf{A}}(\mathcal{T})$ or $\widetilde{\Lambda}^{\mathbf{A}}(\mathcal{T})$.

Finally, reminiscent of the passage from the Tarski to the Suszko congruence system, given a collection of sentence families $\mathcal{T}$ and $T \in \mathcal{T}$, the *Lindenbaum equivalence family* $\widetilde{\lambda}^{\mathbf{A},\mathcal{T}}(T)$ is the relation family identifying two sentences if they are equivalent modulo every $T' \in \mathcal{T}$, such that $T \leq T'$. The *Lindenbaum equivalence system* $\widetilde{\Lambda}^{\mathbf{A},\mathcal{T}}(T)$ is the largest equivalence system contained in $\widetilde{\lambda}^{\mathbf{A},\mathcal{T}}(T)$, and the Suszko congruence system $\widetilde{\Omega}^{\mathbf{A},\mathcal{T}}(T)$ turns out to be the largest congruence system included in either of $\widetilde{\lambda}^{\mathbf{A},\mathcal{T}}(T)$ or $\widetilde{\Lambda}^{\mathbf{A},\mathcal{T}}(T)$.

The Carnap operators, viewed as operators on collections of sentence families on the same algebraic system, are monotone. The same applies to the Lindenbaum operators, viewed as operators on sentence families relative to the same collection of sentence families. However, the Frege operators do not satisfy a monotonicity property.

In Section 2.12, we are discussing *algebraic subsystems* and *π-subinstitutions*. The starting point is the observation that an algebraic system $\mathbf{A} = \langle \mathbf{Sign}, \mathrm{SEN}, N\rangle$ may contain a *universe*, i.e., a functor $\mathrm{SEN}' : \mathbf{Sign} \to \mathbf{Set}$, such that $\mathrm{SEN}' \leq \mathrm{SEN}$ and closed under the action of natural transformations in N. Then, it is clear that this universe may be used to define an *algebraic subsystem* $\mathbf{A}'$ of $\mathbf{A}$ and, as it turns out, there exists a *canonical injection morphism* $\langle I, j\rangle : \mathbf{A}' \to \mathbf{A}$. Apart from detecting the existence of universes, there is a natural way to generate a universe starting from a given sentence family T of $\mathbf{A}$. This consists of passing, first, to the least sentence system $\overrightarrow{T}$ containing T and, then, closing $\overrightarrow{T}$ under the clone operations in N. This two-step process gives rise to a universe $\nu^{\mathbf{A}}(\overrightarrow{T})$. In case the algebraic system $\mathbf{A}$ supports a π-institution $\mathcal{I} = \langle \mathbf{A}, C\rangle$, then one obtains, for each algebraic subsystem $\mathbf{A}'$ of $\mathbf{A}$, a *π-subsinstitution* $\mathcal{I}' = \langle \mathbf{A}', C'\rangle$ by restricting the action of C on elements of $\mathbf{A}'$. It can be shown that the theory families of $\mathcal{I}'$ are exactly the restrictions of those of $\mathcal{I}$ on the universe giving rise to $\mathbf{A}'$. The section ends with some results relating Leibniz congruence systems of theory families of $\mathcal{I}$ with those of the corresponding theory families of $\mathcal{I}'$. A similar result also holds for Leibniz congruence systems of corresponding filter families of the two π-institutions.

Sections 2.13-2.15 deal with aspects of the “syntactic” apparatus of an algebraic system, i.e., with properties of the natural transformations viewed as term functions. Section 2.13 introduces the framework and studies some connections with the definability of the Leibniz congruence systems. Section 2.14 explores various modes of definability and details their relative power. Section 2.15 studies the effect of parameters and shows that two different possible ways of obtain a parameterless collection of natural transformations out of a given parametric one are essentially equivalent. We provide, next, some more details by section.

Section 2.13 introduces the concepts of *distinguished arguments* and of *parametric arguments* of a collection E of natural transformations. This is a conceptual distinction which becomes important in practice when one differentiates the role they each play when the collection of natural transformations is used to transform sentences, i.e., to produce new sets of sentences from tuples of given ones. The new family of sentences produced from a tuple of sentences $\vec{\phi}$ (possibly with the aid of parameters) is denoted by $E_\Sigma[\vec{\phi}]$, where Σ is the signature of $\vec{\phi}$. Another mode of transformation uses a dual or inverse construction. Namely, given a sentence family T, we consider the set $\overleftarrow{E}(T)$ consisting of all tuples $\vec{\phi}$, such that $E_\Sigma[\vec{\phi}] \leq T$. These tuples all share the same length, which equals the number of distinguished arguments of the transformations in E. The construction has some important properties, e.g., $\overleftarrow{E}$, viewed as an operator on sentence families is monotone and, moreover, commutes with inverse surjective morphisms. But, perhaps, its most important property is that, if E has two distinguished arguments and T is such that $\overleftarrow{E}(T)$ is reflexive, then $\overleftarrow{E}(T)$ includes the Leibniz congruence system $\Omega^{\mathbf{A}}(T)$ of T. Consequently, if $\overleftarrow{E}(T)$ is itself a congruence system compatible with T, then it coincides with $\Omega^{\mathbf{A}}(T)$. Thus, in this case, we may say that, in a specific sense, the Leibniz operator of T is *definable* using the natural transformations in E. We view this as a syntactic definability condition, which plays an important role in establishing the algebraic classification of π-institutions "by syntactic means" in subsequent chapters.

In Section 2.14 we continue the study of natural transformations as means of transforming tuples of sentences to sentences. We look as four possible ways of relating, via a fixed collection E of natural transformations with k distinguished arguments, a k-tuple of sentences $\vec{\phi}$ to a sentence family T. The simplest, *E-local membership*, asserts that $E_\Sigma(\vec{\phi},\vec{\chi}) \subseteq T_\Sigma$, for all values $\vec{\chi}$ of the parametric arguments. The second, *E-global membership*, asserts that $E_{\Sigma'}(\mathrm{SEN}(f)(\vec{\phi}),\vec{\chi}) \subseteq T_{\Sigma'}$ holds for all signatures Σ', all morphisms $f : \Sigma \to \Sigma'$ and all appropriate values of the parameters $\vec{\chi}$. The remaining two, *left E-local membership* and *left E-global membership* mimic the preceding ones except that they use membership in $\overleftarrow{T}$ instead of membership in T. Closer scrutiny of the four modes reveals that the two global memberships are equivalent, followed in strength by left local membership, which, in turn, implies local membership. When a membership property holds for all $\vec{\phi}$, then we attribute it to the collection E itself. In this sense, it turns out that global, local, left global and left local memberships of E in T all coincide.

In Section 2.15, starting from a given collection S of natural transformations, possibly including parametric arguments, we study ways of obtaining a collection that is parameter-free. Here, two of the most natural, for our purposes, ways of doing this turn out to be equivalent, and, hence, release us from the obligation to distinguish between which one is applied in any specific context. Let us assume that S is taken to have k distinguished arguments. Then one way of obtaining from S a parameter-free collection is to replace all parametric arguments with k-ary natural transformations. This results in a collection $\dot{S}$ of k-ary, i.e., parameter-free, natural transformations. The second method builds on the notion of an *anti-monotone property* of natural transformations. These are properties P that a natural transformation either does or does not satisfy and for which an anti-monotonicity property holds, namely, if for all tuples of sentences $\vec{\phi}$, the family of transforms of $\vec{\phi}$ under σ is included in the family of transforms of $\vec{\phi}$ under τ, then τ satisfying P implies that σ also satisfies P. If P also denotes the class of all natural transformations satisfying property P, then we let $\widehat{P}$ be the subclass of P consisting of the parameter-free members of P. The section concludes with the assertion that, for anti-monotone properties P, both constructions $\dot{P}$ and $\widehat{P}$ give the same class of parameter-free natural transformations associated with P.

In Section 2.16, we study *finitarity*. This property holds for a π-institution $\mathcal{I}$ if every sentence ϕ that is derivable from a set Φ of sentences can be derived from some finite subset Φ' of Φ. Finitarity holds for the overwhelming majority of the logics considered in the literature. So it has played a central role in algebraic logic, even though much of the more abstract body of the theory is formalized and developed in a way that encompasses arbitrary, that is, not necessarily finitary, logical systems. A characterization of finitarity using the property of *continuity* is provided. We say that a collection of theory families is *directed* if every finite subcollection is included in some theory family in the collection. A π-institution is *continuous* if the union of a directed collection of theory families is also a theory family. Finitarity and continuity, as it turns out, are equivalent properties.

In the second part of the section, given a finitary π-institution $\mathcal{I}$, we provide a construction of the filter family $C^{\mathcal{I},\mathcal{A}}(X)$ generated by a sentence family X of $\mathcal{A}$. Taking advantage of the finitarity of $\mathcal{I}$, the filter family may be obtained by an incremental process, each step of which adds in the filter family sentences of $\mathcal{A}$ which are derivable, in a certain sense, by finite subsets of sentences that have already been included in the filter family at previous stages of the construction. In this way, the family $\Xi^{\mathcal{I},\mathcal{A}}(X)$ is obtained as the union of the families obtained at all stages and it can be shown that $C^{\mathcal{I},\mathcal{A}}(X) = \Xi^{\mathcal{I},\mathcal{A}}(X)$.

In the last two sections, Sections 2.17 and 2.18, we study *equational consequences* and provide analogs of some well-known fundamental results of universal algebra for classes of algebraic systems.

In Section 2.17, we look at closure families on pairs of sentences, i.e., equations, over a base algebraic system $\mathbf{F}$ that are induced by classes of $\mathbf{F}$-algebraic systems. Given a class K of $\mathbf{F}$-algebraic systems, we say that an equation $\phi \approx \psi$ is a *consequence* of a set E of equations *relative to* K if every algebraic system in K satisfying E also satisfies $\phi \approx \psi$. The resulting consequence family is denoted by D^{K}. It is not necessarily a closure system since it may fail to be structural. It is shown, however, that its theory families are exactly the congruence systems on $\mathbf{F}$ relative to the class K.

The second part of the section deals with a process of generating the closure of a family of equations E relative to an equational axiomatic system Q in an incremental way. Roughly speaking, it formalizes the process of closing under reflexivity, symmetry and transitivity, as well as under replacement and the action of signature morphisms. The family of equations obtained under this step-wise process from axioms Q and hypotheses E is denoted by $\Xi^Q(E)$. In the final result of the section, it is shown that the operator Ξ^Q coincides with D^{K} when Q is taken to be the collection of all equations satisfied by all algebraic systems in K.

Section 2.18, the closing section of the chapter, is inspired by universal algebra. It provides characterizations, in the spirit of Birkhoff's variety and Mal'cev's quasivariety theorems, of classes of algebraic systems defined by equations, quasiequations and generalized quasiequations, also referred to as guasiequations. The section begins by formally defining *equations*, *quasiequations* and *guasiequations* in the context of π-institutions. The relation of *satisfaction* of a syntactic entity of either of the above types in an algebraic system is also formally defined. In the usual way, these satisfaction relations establish Galois connections. The closed sets on the syntactic side form *equational*, *quasiequational* and *guasiequational theories*, whereas, on the semantic side, one obtains *equational*, *quasiequational* and *guasiequational classes of algebraic systems*, respectively. These are, respectively, the classes closed under the semantic variety $\mathbb{V}^{\mathsf{Sem}}$, semantic quasivariety $\mathbb{Q}^{\mathsf{Sem}}$ and semantic guasivariety $\mathbb{G}^{\mathsf{Sem}}$ operators.

To formulate characterizations of these classes, we introduce and study four operators on classes of algebraic systems. Let K be a class of algebraic systems. First, we say that an algebraic system $\mathcal{A}$ is K-*certified* if, for each signature Σ, there exists an algebraic system $\mathcal{A}^{\Sigma}$ in the class K that satisfies exactly the same equations of signature Σ as $\mathcal{A}$. The class K is said to be *abstract* or *closed under* K-*certifications* if every K-certified algebraic system is in K. The operator $\mathbb{C}$ is a closure operator and, if $\mathcal{A} \in \mathbb{C}(\mathsf{K})$, then $\mathcal{A}$ satisfies all guasiequations satisfied by K. Moreover, if K is guasiequational, then it is an abstract class. Next, we say that an algebraic system $\mathcal{A}$ is *directedly* K-*certified* if, for each signature Σ, there exists a collection of algebraic systems $\{\mathcal{A}^{\Sigma,i} : i \in I\}$ in the class K that satisfy two conditions: On the one hand, the collection of all finite sets of equations satisfied by some $\mathcal{A}^{\Sigma,i}$, $i \in I$, is directed and, on the other, the union of all those sets is exactly the set of equations of signature Σ satisfied by $\mathcal{A}$. The class K is said to be *directedly abstract* or *closed under directed* K-*certifications* if every directedly K-certified algebraic system is in K. The operator $\mathbb{C}^*$ is a closure operator. It is shown that, if $\mathcal{A}$ is directedly K-certified, then it satisfies all quasiequations satisfied by K and, furthermore, that directed abstraction is a necessary condition for a class of algebraic systems to be a quasiequational class.

The third operator on classes of algebraic systems is that of taking *subdirect intersections* $\overset{\triangleleft}{\mathbb{I}\mathbb{I}}$. *Subdirect intersections* are collections of morphisms $\langle H^i, \gamma^i \rangle : \mathcal{A} \to \mathcal{A}^i$, $i \in I$, with the same domain, the intersection of whose kernels is the identity system on $\mathcal{A}$. In that case, we also say that $\mathcal{A}$ is a *subdirect intersection* of the $\mathcal{A}^i$'s. This also turns out to be a closure operator on classes of algebraic systems and, in fact, closure under $\overset{\triangleleft}{\mathbb{I}\mathbb{I}}$ is necessary for a class to be guasiequational. The last operator considered is that of taking *morphic images*, denoted by $\mathbb{H}$. It also forms a closure operator on classes of algebraic systems and closure under $\mathbb{H}$ is necessary for a class to be an equational class.

The four operators serve in formulating the Birkhoff-style characterizations referred to previously for equational, quasiequational and guasiequational classes. Guasiequational classes are characterized as those that are abstract and closed under subdirect intersections. Quasiequational classes are those that are directedly abstract and closed under subdirect intersections. Finally, equational classes are characterized as those that are closed under subdirect intersections and morphic images. The section concludes with some additional characterizations of these three classes involving the structure of the subcollection $\mathrm{ConSys}^{\mathsf{K}}(\mathcal{F})$ of the complete lattice $\mathrm{ConSys}(\mathcal{F})$. All of those additional results are based on the main characterizations described above.

1.3.2 Chapter 3

In Chapter 3 we start in earnest the study of the Leibniz hierarchy of π-institutions. Chapters 3-9 deal with the *semantic Leibniz hierarchy*. Here the classes are defined using properties of the Leibniz operator on theory families/systems of a π-institution. Chapters 11-**??**, on the other hand, deal with the *syntactic Leibniz hierarchy* in which classes are defined using collections of natural transformations satisfying specific definability properties. We shall see that "corresponding" classes in the two hierarchies may not coincide, but, nevertheless, the two hierarchies are closely connected - in fact may be seen as forming parts of a single hierarchy - and they are both modeled on the Leibniz hierarchy of sentential logics.

In Section 3.2, we study three properties. The first two are fundamental because they introduce concepts and terminology that play a critical role throughout the monograph. The third is used to establish classes of π-institutions at the very bottom of the hierarchy which abstract all other classes considered later in this and in subsequent chapters.

The first property if *systemicity*. A π-institution $\mathcal{I}$ is called *systemic* if every theory family of $\mathcal{I}$ is actually a theory system, i.e., if $\mathrm{ThFam}(\mathcal{I}) = \mathrm{ThSys}(\mathcal{I})$. Recalling from Chapter 2 that, given a theory family T of $\mathcal{I}$, $\overleftarrow{T}$ is the largest theory system included in T, $\mathcal{I}$ is systemic if and only if, for every theory family T, $\overleftarrow{T} = T$. Yet another characterization asserts that, for every Σ-sentence ϕ of $\mathcal{I}$, the least theory family $C(\phi)$ of $\mathcal{I}$ generated by ϕ contains all translates of ϕ under arbitrary signature morphisms. One of the reasons why systemicity plays such a critical role is that, for a systemic π-institution, it suffices to restrict attention to theory systems, i.e., one may take invariance under signature morphisms for granted.

The second property is *stability*. It may be thought of as the counterpart of systemicity when focus shifts from theory families to corresponding Leibniz congruence systems. A π-institution $\mathcal{I}$ is *stable* if, for all theory families T, $\Omega(\overleftarrow{T}) = \Omega(T)$. Of course, every systemic π-institution is stable, and this implication is proper. Both systemicity and stability *transfer*. This means that a π-institution $\mathcal{I} = \langle \mathbf{F}, C\rangle$ is systemic if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, every $\mathcal{I}$-filter family of $\mathcal{A}$ is a filter system. Similarly, $\mathcal{I}$ is stable if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$ and every $\mathcal{I}$-filter family T of $\mathcal{A}$, $\Omega^{\mathcal{A}}(\overleftarrow{T}) = \Omega^{\mathcal{A}}(T)$. These two transfer results are only the first of a host of, so-called, *transfer theorems* that are proved in the sequel for the majority of properties used to define classes in the Leibniz hierarchy. Having established the pattern and exhibited the main idea, we only mention such results briefly from now on, postponing the details for the main account in the relevant sections of the text.

The third property we study in Section 3.2 is *loyalty*. Unlike systemicity and stability, loyalty comes, as is typical for many subsequently introduced properties, in multiple flavors. To establish the pattern that will be followed in the presentation throughout, we introduce, first, the four versions, termed *family*, *left*, *right* and *system*. They may or may not be all different. So we study their properties, show which ones, if any, coincide, establish general implications between those that are not equivalent, and show, via examples, that these implications are proper, i.e., that no further collapsing of the subhierarchy based on these properties is possible.

A π-institution $\mathcal{I}$ is *family loyal* if, for all theory families T, T' of $\mathcal{I}$, $T \nless T'$ or $\Omega(T) \ngtr \Omega(T')$, or, equivalently, if it is not the case that $T < T'$ and $\Omega(T) > \Omega(T')$. If Ω, viewed as an operator mapping theory families to congruence systems, is either order preserving or order reflecting, then it is necessarily family loyal. So this property abstracts both monotonicity and reflectivity of Ω. Since both monotonicity and reflectivity play important roles in specifying classes in the Leibniz hierarchy, this observation provides partial justification for considering loyalty as a common abstraction. Here, as in all subsequently defined properties, once the family version is introduced, the other three versions follow by applying similar modifications. To obtain the *left version* one replaces, on the theory family side, all theory families by their arrow versions. So $\mathcal{I}$ is *left loyal* if, for all theory families T, T', $\overleftarrow{T} \nless \overleftarrow{T'}$ or $\Omega(T) \ngtr \Omega(T')$. To obtain the *right version*, a similar replacement is applied on the congruence system side. Thus, $\mathcal{I}$ is *right loyal* if, for all theory families T, T', $T \nless T'$ or $\Omega(\overleftarrow{T}) \ngtr \Omega(\overleftarrow{T'})$. Finally, the *system version* is obtained by imposing the same condition as in the family version, but restricting its application to theory systems, instead of insisting that it hold for all theory families. Accordingly, $\mathcal{I}$ is *system loyal* if, for all theory *systems* T, T' of $\mathcal{I}$, $T \nless T'$ or $\Omega(T) \ngtr \Omega(T')$.

Family loyalty properly implies stability. Moreover, family loyalty implies left loyalty, which, in turn, implies system and right loyalty, the latter two being equivalent properties. System loyalty, together with systemicity, imply family loyalty. That is, as is the case with virtually all properties introduced in the monograph, imposing systemicity has the effect of collapsing the entire four-class subhierarchy into a single class. This observation can be applied to obtain a backbone - or a bird's eye view - of the Leibniz hierarchy without worrying about the refinements and subdivisions due to the different flavors of each property. Section 3.2 concludes by showing that all three distinct versions of loyalty transfer, i.e., that a given π-institution has a certain loyalty property if the corresponding defining condition holds for all pairs of filter families (or systems) on arbitrary $\mathbf{F}$-algebraic systems.

In Section 3.3, we study *monotonicity properties*. A π-institution $\mathcal{I}$ is *family monotone* if, for all theory families $T, T', T \le T'$ implies $\Omega(T) \le \Omega(T')$, i.e., if the Leibniz operator on theory families is order preserving. In accordance with the general framework outlined above for loyalty, $\mathcal{I}$ is *left monotone* if, for all T, T', $\overleftarrow{T} \le \overleftarrow{T'}$ implies $\Omega(T) \le \Omega(T')$, *right monotone* if, for all T, T', $T \le T'$ implies $\Omega(\overleftarrow{T}) \le \Omega(\overleftarrow{T'})$ and *system monotone* if the same condition defining family monotonicity is restricted to theory systems, i.e., if the Leibniz operator on theory systems is order preserving. It is shown that family monotonicity implies stability. Most importantly, family and left monotonicity coincide as do system and right monotonicity. Following terminology inherited from sentential logics, we term π-institutions that satisfy family monotonicity *protoalgebraic* and those that satisfy the system version *prealgebraic*. Protoalgebraicity is equivalent to prealgebraicity plus stability. In particular, every protoalgebraic π-institution is prealgebraic, and this inclusion is proper. Both monotonicity properties transfer. Finally, pursuing connections with classes introduced in Section 3.2, we show that protoalgebraicity implies family loyalty, whereas prealgebraicity is sufficient for system loyalty.

In Sections 3.4 and 3.5, we study versions of a property called *complete monotonicity*. This is a property dual to complete order reflectivity, a property that characterizes truth equationality in the sentential framework. Given a sentential logic $\mathcal{S}$, complete order reflectivity stipulates that, for every collection $\mathcal{T} \cup \{T'\}$ of theories of $\mathcal{S}$, if $\bigcap_{T \in \mathcal{T}} \Omega(T) \subseteq \Omega(T')$, then $\bigcap \mathcal{T} \subseteq T'$. Since, in both the lattice of theories and that of congruences, meet and intersection coincide, but, on both theories and congruences, join is not the same as union, one may obtain two "dual" versions of complete order reflectivity. The first, following a set-theoretic approach, says that, for all $\mathcal{T} \cup \{T'\}$, $T' \subseteq \bigcup \mathcal{T}$ implies $\Omega(T') \subseteq \bigcup_{T \in \mathcal{T}} \Omega(T)$. The second, taking a lattice-theoretic point of view, asserts that, for all $\mathcal{T} \cup \{T'\}$, $T' \le \bigvee \mathcal{T}$ implies $\Omega(T') \le \bigvee_{T \in \mathcal{T}} \Omega(T)$, where the join in the hypothesis is taken in the complete lattice of theories of $\mathcal{S}$ and the one in the conclusion in the complete lattice of congruences on the formula algebra. In Section 3.4 we study an analog of the former property and in Section 3.5 an analog of the latter in the context of logics formalized as π-institutions. A few more details follow in the next two paragraphs.

In Section 3.4, we look at *complete* $\bigcup$*-monotonicity*, which is abbreviated as *c*$^\cup$*-monotonicity* or, simply, *c-monotonicity*. A π-institution is *family c*$^\cup$*-monotone* if, for every collection $\mathcal{T} \cup \{T'\}$ of theory families, $T' \le \bigcup \mathcal{T}$ implies $\Omega(T') \le \bigcup_{T \in \mathcal{T}} \Omega(T)$. *Left* and *right c*$^\cup$*-monotonicities* are obtained by replacing in the hypothesis and in the conclusion, respectively, every theory family occurring by its arrow version. Finally, *system c*$^\cup$*-monotonicity* is defined by the same condition as the family version, but applied exclusively to collections of theory systems. Family c$^\cup$-monotonicity implies stability, as does left c$^\cup$-monotonicity. Moreover, the family version is equivalent to the conjunction of the left and right versions and either of the latter implies system $c^\cup$-monotonicity. All four $c^\cup$-monotonicity properties transfer. And, whereas the left version is sufficient for protoalgebraicity, the system version implies only prealgebraicity.

In Section 3.5, we continue the study of complete monotonicity but switch from complete $\bigcup$-monotonicity to *complete* $\bigvee$*-monotonicity*, which is abbreviated as $c^\vee$*-monotonicity*. A π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is *family* $c^\vee$*-monotone* if, for every collection $\mathcal{T} \cup \{T'\}$ of theory families, $T' \leq \bigvee^{\mathcal{I}} \mathcal{T}$ implies $\Omega(T') \leq \bigvee^{\mathbf{F}}_{T \in \mathcal{T}} \Omega(T)$, where $\bigvee^{\mathcal{I}}$ denotes the join in the complete lattice of theory families of $\mathcal{I}$ and $\bigvee^{\mathbf{F}}$ the join in the complete lattice of congruence systems on $\mathbf{F}$. Again, following the general pattern, $\mathcal{I}$ is *left* $c^\vee$*-monotone* if, for all $\mathcal{T} \cup \{T'\}$, $\overleftarrow{T'} \leq \bigvee^{\mathcal{I}}_{T \in \mathcal{T}} \overleftarrow{T}$ implies $\Omega(T') \leq \bigvee^{\mathbf{F}}_{T \in \mathcal{T}} \Omega(T)$ and is *right* $c^\vee$*-monotone* if, for all $\mathcal{T} \cup \{T'\}$, $T' \leq \bigvee^{\mathcal{I}} \mathcal{T}$ implies $\Omega(\overleftarrow{T'}) \leq \bigvee^{\mathbf{F}}_{T \in \mathcal{T}} \Omega(\overleftarrow{T})$. Finally, $\mathcal{I}$ is *system* $c^\vee$*-monotone* if, for every collection $\mathcal{T} \cup \{T'\}$ of theory *systems* of $\mathcal{I}$, $T' \leq \bigvee^{\mathcal{I}} \mathcal{T}$ implies $\Omega(T') \leq \bigvee^{\mathbf{F}}_{T \in \mathcal{T}} \Omega(T)$. Again, either family or left $c^\vee$-monotonicity implies stability. The family version is equivalent to the conjunction of the left and right versions and either of those two implies system $c^\vee$-monotonicity. Left $c^\vee$-monotonicity implies protoalgebraicity and system $c^\vee$-monotonicity implies prealgebraicity.

Contrary to what the similarities of results pertaining to $c^\vee$-monotonicity with those of Section 3.4 on $c^\cup$-monotonicity may suggest, there are also significant differences between the two complete monotonicity properties. One instance concerns transfer theorems. Unlike $c^\cup$-monotonicity, $c^\vee$-monotonicity properties do not transfer in general. This is due to the fact that, unlike unions, joins do not commute with inverse surjective morphisms between algebraic systems. A second difference, which affords, perhaps, partial justification for introducing and discussing both types of properties in some detail, is that corresponding classes of π-institutions are incomparable. E.g., there exists a family $c^\vee$-monotone π-institution which is not family $c^\cup$-monotone and vice-versa.

In Section 3.6, we study *injectivity*. A π-institution $\mathcal{I}$ is *family injective* if, for all theory families T, T', $\Omega(T) = \Omega(T')$ implies $T = T'$, i.e., if the Leibniz operator is injective on theory families. It is *left injective* if, for all T, T', $\Omega(T) = \Omega(T')$ implies $\overleftarrow{T} = \overleftarrow{T'}$ and *right injective* if, for all T, T', $\Omega(\overleftarrow{T}) = \Omega(\overleftarrow{T'})$ implies $T = T'$. Finally, it is *system injective* if the Leibniz operator is injective on theory systems. Right injectivity is the strongest of the four injectivity properties and it implies systemicity. It is followed by family injectivity, then left injectivity, which implies system injectivity. System injectivity together with systemicity is equivalent to right injectivity, whereas, together with stability, which is weaker than systemicity, it implies left injectivity. All four injectivity properties transfer.

In Section 3.7, we turn to *reflectivity properties*. A π-institution $\mathcal{I}$ is *family reflective* if, for all theory families T, T' of $\mathcal{I}$, $\Omega(T) \leq \Omega(T')$ implies

$T \leq T'$, i.e., if the Leibniz operator on theory families is order reflecting. If, for all T, T', $\Omega(T) \leq \Omega(T')$ implies $\overleftarrow{T} \leq \overleftarrow{T'}$, then $\mathcal{I}$ is *left reflective*, whereas, if, for all T, T', $\Omega(\overleftarrow{T}) \leq \Omega(\overleftarrow{T'})$ implies $T \leq T'$, $\mathcal{I}$ is *right reflective*. *System reflectivity* stipulates the order reflectivity of the Leibniz operator on theory systems. It turns out that family or right reflectivity imply systemicity. This allows showing that the two are actually equivalent properties. They imply left reflectivity, which, in turn, implies system reflectivity. System reflectivity, coupled with stability, implies left reflectivity, whereas, together with systemicity, it becomes equivalent to family reflectivity. All four versions transfer. Section 3.7 ends by relating reflectivity with the injectivity properties, introduced in Section 3.6, and with the loyalty properties, introduced in Section 3.2. More precisely, it is shown that family/right, left and system reflectivity imply, respectively, right, left and system injectivity and that family/right, left and system reflectivity imply, respectively, family, left and system loyalty.

Section 3.8, the last section of Chapter 3, introduces *complete reflectivity properties*, abbreviated to *c-reflectivity*. These form a generalization of the reflectivity properties of Section 3.7. Complete reflectivity originates in the work of Raftery, where it is used to characterize truth equationality in the context of sentential logics. A π-institution is *family c-reflective* if, for every collection $\mathcal{T} \cup \{T'\}$ of theory families of $\mathcal{I}$, $\bigcap_{T \in \mathcal{T}} \Omega(T) \leq \Omega(T')$ implies $\bigcap \mathcal{T} \leq T'$. It is *left c-reflective* if, for all $\mathcal{T} \cup \{T'\}$, $\bigcap_{T \in \mathcal{T}} \Omega(T) \leq \Omega(T')$ implies $\bigcap_{T \in \mathcal{T}} \overleftarrow{T} \leq \overleftarrow{T'}$ and *right c-reflective* if, for all $\mathcal{T} \cup \{T'\}$, $\bigcap_{T \in \mathcal{T}} \Omega(\overleftarrow{T}) \leq \Omega(\overleftarrow{T'})$ implies $\bigcap_{T \in \mathcal{T}} \mathcal{T} \leq T'$. *System c-reflectivity* is defined using the same condition as family c-reflectivity restricted to collections of theory systems. As was the case with reflectivity, either family or right c-reflectivity implies systemicity and this enables showing that the family and right versions are equivalent. They imply left c-reflectivity, which, in turn, implies the system version. System c-reflectivity and systemicity are jointly equivalent to family c-reflectivity, whereas system c-reflectivity, augmented with stability, implies left c-reflectivity. All complete reflectivity properties transfer and, as is apparent from the relevant definitions, each version of c-reflectivity implies the corresponding reflectivity version.

1.3.3 Chapter 4

In Chapter 4, we visit weak prealgebraizability and weak algebraizability properties of π-institutions. These create a subhierarchy of π-institutions whose members roughly correspond to the weakly algebraizable logics in the sentential logic framework. Weak prealgebraizability classes arise when coupling family monotonicity with either of injectivity, reflectivity or complete reflectivity properties. Analogously, weak algebraizability results by combining system monotonicity with injectivity, reflectivity or complete reflectivity.

Before describing the versions of weak prealgebraizability and algebraizability in more detail, we mention, firstly, that the term "weak" refers to the use of monotonicity, as opposed to the stronger notion of equivalentiality, in the definitions, and remind, secondly, the reader of the hierarchy, established in Chapter 3, of the various flavors of injectivity, reflectivity and c-reflectivity properties, which assumed the form depicted in the diagram.

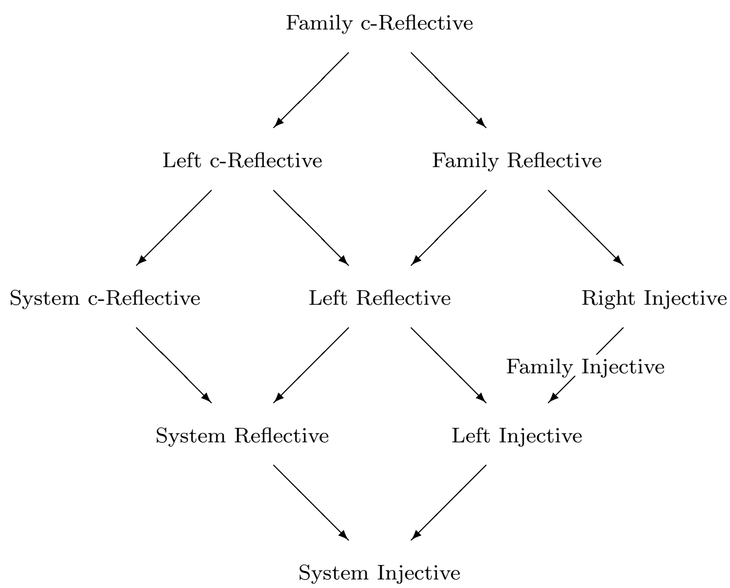

In Section 4.2, we define the classes of *weakly prealgebraizable π-institutions*. Each class results by imposing prealgebraicity (system monotonicity) and one of the ten flavors of injectivity, reflectivity and complete reflectivity shown in the preceding hierarchy. Since prealgebraicity is shared by all classes, the deciding factor in the subhierarchy is the type of injectivity, reflectivity or c-reflectivity imposed. Thus, a priori, one obtains ten potentially distinct classes whose hierarchy reflects that shown in the preceding diagram. We name the corresponding property "weak X prealgebraizability", or "WX prealgebraizability" for short, where the string X stands for one of SI, LI, FI, RI for system, left, family, right injectivity, respectively, SR, LR, FR for system, left, family reflectivity, respectively, or SC, LC, FC for system, left, family c-reflectivity, respectively.

In our first result, we show that prealgebraicity is sufficient to identify all system versions, which forces the collapsing of the classes of WSI, WSR and WSC prealgebraizable π-institutions. We call the corresponding property *WS prealgebraizability*. In what sets a pattern for subsequent work in this chapter, it is shown that WS prealgebraizability transfers and, further, a characterization is obtained via properties of the Leibniz operator $\Omega^{\mathcal{A}}$, viewed as a mapping between ordered sets, for arbitrary $\mathbf{F}$-algebraic systems $\mathcal{A}$. More precisely, it is shown that a π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is WS prealgebraizable iff, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$

is an order embedding. Next, it is shown that, in view of prealgebraicity, family reflectivity implies family c-reflectivity and this leads to the identification of WFR prealgebraizability and WFC prealgebraizability. Moreover, under protoalgebraicity, family injectivity implies family reflectivity. This enables showing that both WFR and WRI prealgebraizability are characterized as the conjunction of WFI prealgebraizability and systemicity and, hence, are identical properties. Both WFR and WFI prealgebraizability transfer. Moreover, the WFI version is characterized by the property that, for all $\mathcal{A}$, $\Omega^{\mathcal{A}}$ is a bijection on filter families, restricting to an order embedding on filter systems, whereas the WFR version is characterized by the condition that, for all $\mathcal{A}$, $\Omega^{\mathcal{A}}$ is an order isomorphism.

At this point, the hierarchy has been reduced to six classes, since, as it turned out, all three system classes are identical and the three family plus the WRI prealgebraizability collapse down to two classes. The only classes not put under the microscope yet are those defined using the left versions of injectivity, reflectivity and c-reflectivity. We return to them after a short break that gives a glimpse of further possible reductions under special circumstances. Namely, it is proven that, under systemicity, the entire hierarchy collapses to a single class and that, under stability, it collapses down to two classes, as the only properties that can be distinguished are the family (but including also WRI prealgebraizability) from the remaining versions.

Returning to the left properties, Section 4.2 concludes by showing that all three transfer and by providing characterizations along the lines outlined previously, using $\Omega^{\mathcal{A}}$. More precisely, it is shown that $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is WLC (WLR, WLI, respectively) prealgebraizable iff, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is a left completely order reflecting (left order reflecting, left injective, respectively) surjection, restricting to an order embedding on theory systems.

In Section 4.3, we study versions of weak algebraizability. These combine protoalgebraicity (family monotonicity) with the various versions of injectivity, reflectivity and complete reflectivity. Since protoalgebraicity dominates prealgebraicity, it is clear that one obtains at least as many identifications between the ten apparent weak algebraizability properties as those established between corresponding weak prealgebraizability properties in Section 4.2. However, the situation under closer scrutiny turns out to be much more radical. Since protoalgebraicity is strong enough to yield stability, the emerging landscape was anticipated by the previously mentioned collapse of the weak prealgebraizability hierarchy down to two classes in the presence of stability. Similarly, under protoalgebraicity and, hence, stability, all three weak family algebraizability properties together with WRI algebraizability collapse to a single property, termed *WF algebraizability*. Further, all remaining six left and system versions also collapse to a single property we call *WS algebraizability*. Both of these properties transfer. Also, for both one may obtain Leibniz operator type characterizations. More specifically, $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is WS algebraizable iff it is stable and, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism, whereas it is WF algebraizable iff, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism.

Observing that the characterization of WF algebraizability is identical with that obtained for WF prealgebraizability, we conclude that the top classes in the weak prealgebraizability and weak algebraizability subhierarchies actually coincide. Thus, by fusing these two subhierarchies, one obtains a total of seven potentially distinct classes, which form the combined hierarchy depicted in the diagram.

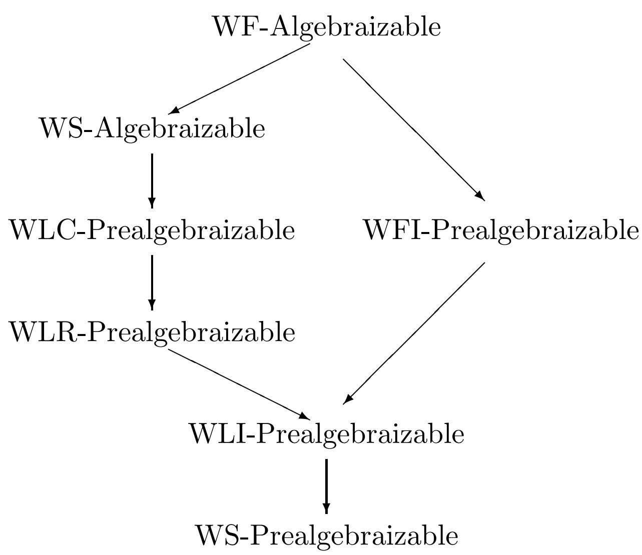

1.3.4 Chapter 5

In Chapter 5, we deal with classes of π-institutions that result from weakly prealgebraizable and weakly algebraizable π-institutions when the properties of prealgebraicity (system monotonicity) and protoalgebraicity (family monotonicity) are strengthened to preequivalentiality and equivalentiality, respectively. The strengthening, i.e., the passage from proto- (or pre-) algebraicity to (pre)equivalentiality, involves adding the condition of either family or system extensionality. Depending on which of these two properties is imposed, one obtains two parallel hierarchies, one on top of the other, both of which reflect the structure of the weak (pre)algebraizability hierarchy, described in Chapter 4.

In Section 5.2, we introduce and study *extensionality*. The definition requires the notion of subsystem of an algebraic system $\mathbf{F}$ generated by a given sentence family X, which is denoted by $\langle X \rangle$ and was introduced in Section 2.12. A π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is called *family extensional* if, for all sentence families X of $\mathbf{F}$ and all theory families T of $\mathcal{I}$, $\Omega(T) \cap \langle X \rangle^2 = \Omega^{\langle X \rangle}(T \cap \langle X \rangle)$.

It is called *system extensional* if the same condition holds, but T is quantified over all theory systems of $\mathcal{I}$, instead of ranging over arbitrary theory families. Since system extensionality specializes family extensionality, every family extensional π-institution is also system extensional. It is, moreover, the case that system extensionality, coupled with stability, implies family extensionality. Extensionality is very useful because, when satisfied, it causes certain properties that hold in a π-institution to be inherited by all its subinstitutions. For instance, under system extensionality, stability propagates from a π-institution $\mathcal{I}$ to all its subinstitutions $\mathcal{I}' \leq \mathcal{I}$. Additionally, system or family extensionality causes prealgebraicity or protoalgebraicity, respectively, to be inherited by all subinstitutions of a given π-institution. Both versions of extensionality transfer. The section closes by looking at *2-extensionality*, an apparently weaker condition than extensionality, which, however, turns out to be equivalent to it. A π-institution $\mathcal{I} = \langle \mathbf{F}, C\rangle$ is *family 2-extensional* if, for all $\Sigma \in |\mathbf{Sign}^\flat|$, all $\phi, \psi \in \mathrm{SEN}^\flat(\Sigma)$ and every theory family T of $\mathcal{I}$, $\langle \phi, \psi\rangle \in \Omega_\Sigma(T)$ if and only if $\langle \phi, \psi\rangle \in \Omega_\Sigma^{\langle \phi,\psi\rangle}(T \cap \langle \phi, \psi\rangle)$. *System 2-extensionality* is defined by the same condition in which T is quantified over theory systems. A π-institution is family/system extensional if and only if it is family/system 2-extensional, respectively.

In Section 5.3, we study *Leibniz commutativity*. The notion relies on the concepts of *extension* and *logical extension*. Given an algebraic system $\mathbf{F}$ and a sentence family X of $\mathbf{F}$, an *extension* is an algebraic system morphism of the form $\langle I, \alpha\rangle : \langle X\rangle \to \mathbf{F}$, where $\langle X\rangle$ is the algebraic subsystem of $\mathbf{F}$ generated by X and I is the identity functor on signatures. Given a π-institution $\mathcal{I} = \langle \mathbf{F}, C\rangle$, an extension $\langle I, \alpha\rangle : \langle X\rangle \to \mathbf{F}$ is called *logical*, denoted $\langle I, \alpha\rangle : \mathcal{I}^{\langle X\rangle} \to \mathcal{I}$, if, for every signature Σ and all $\Phi \subseteq \langle X\rangle_\Sigma$, $\alpha_\Sigma(C_\Sigma^{\langle X\rangle}(\Phi)) \subseteq C_\Sigma(\alpha_\Sigma(\Phi))$, where $C^{\langle X\rangle}$ is the restriction of C on $\langle X\rangle$, discussed in detail in Section 2.12. A characterization of this notion asserts that $\langle I, \alpha\rangle$ is logical if and only if α^{-1} preserves theory families, i.e., if $\alpha^{-1}(T) \in \mathrm{ThFam}(\mathcal{I}^{\langle X\rangle})$, for every $T \in \mathrm{ThFam}(\mathcal{I})$.

Logical extensions form the background for introducing the property of *Leibniz commutativity*, or, simply, *commutativity*. A π-institution $\mathcal{I} = \langle \mathbf{F}, C\rangle$ is called *family commuting* if the Leibniz operator on theory families commutes with logical extensions, i.e., if, for every sentence family X of $\mathbf{F}$, every logical extension $\langle I, \alpha\rangle : \mathcal{I}^{\langle X\rangle} \to \mathcal{I}$ and all $T' \in \mathrm{ThFam}(\mathcal{I}^{\langle X\rangle})$, $\alpha(\Omega^{\langle X\rangle}(T')) \leq \Omega(C(\alpha(T')))$. Applying the same condition, where T' ranges over all theory systems of $\mathcal{I}^{\langle X\rangle}$, defines *system commutativity*. A closely related concept is that of *inverse Leibniz commutativity*, or, simply, *inverse commutativity*. A π-institution $\mathcal{I}$ is *family inverse commuting* if, for every sentence family X, every logical extension $\langle I, \alpha\rangle : \mathcal{I}^{\langle X\rangle} \to \mathcal{I}$, and all $T \in \mathrm{ThFam}(\mathcal{I})$, $\alpha^{-1}(\Omega(T)) = \Omega^{\langle X\rangle}(\alpha^{-1}(T))$. The same condition, imposed on theory systems only, defines *system inverse commutativity*. The fact that injection morphisms $\langle I, j\rangle : \mathcal{I}^{\langle X\rangle} \to \mathcal{I}$ of subinstitutions into their parent institutions are logical extensions allows us to show that family/system inverse commutativity implies family/system extensionality, respectively. It is clear that the family version implies the system version and, as it turns out, the system version augmented by stability implies the family version. What is important for our purposes, and the reason why both direct and inverse commutativity properties are studied, is that under pre/proto-algebraicity, respectively, system/family commutativity is equivalent to system/family inverse commutativity. Moreover, in a result strengthening the relationship mentioned above, it is proven that family/system inverse commutativity and family/system extensionality, respectively, are actually equivalent properties. This section concludes by showing that both versions of inverse commutativity transfer.

In Section 5.4, we introduce *equivalentiality*. This is the section we have been preparing for by studying extensionality and commutativity in Sections 5.2 and 5.3, respectively. *Equivalentiality* is the result of coupling monotonicity with extensionality. Since each of those two properties comes in two flavors, there are, a priori, four possible versions of equivalentiality. *Family equivalentiality* combines protoalgebraicity with family extensionality. *System equivalentiality* keeps protoalgebraicity but uses system extensionality. *Family* and *system preequivalentiality* are defined analogously, but here one uses prealgebraicity instead of protoalgebraicity. Since protoalgebraicity is strong enough to imply stability, it turns out that family and system equivalentiality coincide. This property is referred to simply as *equivalentiality*. Thus, we get three properties in this hierarchy, namely, in decreasing order of potency, equivalentiality, family preequivalentiality and system preequivalentiality. Moreover, equivalentiality is equivalent to system preequivalentiality plus stability. All three properties transfer. There also exist characterizations of equivalentiality and preequivalentiality by conditions imposed on the Leibniz operator on filter families/systems, respectively, on arbitrary **F**-algebraic systems. Finally, as is clear by the corresponding definitions, equivalentiality dominates protoalgebraicity and preequivalentiality dominates prealgebraicity.

In Section 5.5, by replacing prealgebraicity by preequivalentiality, we obtain from the weak prealgebraizability hierarchy of Section 4.2 two parallel *prealgebraizability hierarchies*. The term "prealgebraizability" in both refers to the fact that preequivalentiality, as opposed to equivalentiality, is applied. In one of the two hierarchies, "family prealgebraizability" refers to the application of family preequivalentiality, whereas in "prealgebraizability", it is understood that (system) preequivalentiality is applied. The five classes in the first hierarchy are termed *XF prealgebraizable* and in the second *X prealgebraizable*, where X is one of the following strings, suggesting the imposition of an additional property on the Leibniz operator.

- LC for left completely reflective;
- LR for left reflective;

- FI for family injective;
- LI for left injective; and
- S for system (system completely reflective, system reflective or system injective, which are all equivalent in view of prealgebraicity).

It is shown that systemicity causes the total collapse of the hierarchy into a single class, whereas stability collapses the two family injectivity classes, FI and FIF prealgebraizability, and, also, all eight remaining classes and, therefore, leads to a 2-class hierarchy. Moreover, it is proven that all ten properties transfer. The remainder of this section is devoted to providing characterizations of each of the ten classes using order theoretic properties of the Leibniz operator viewed as a mapping from lattices of filters systems/families to lattices of congruence systems over arbitrary **F**-algebraic systems. We focus only on a couple of pairs to give a flavor of the type of results obtained, and refer the reader to the main text for a full account. A π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is FIF prealgebraizable if and only if, for all **F**-algebraic systems $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is a bijection commuting with inverse logical extensions, which restricts to an order embedding on filter systems. A similar characterization is obtained for FI prealgebraizability, but with a subtle important change: $\mathcal{I}$ is FI prealgebraizable if and only if, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is a bijection, which restricts to an order embedding commuting with inverse logical morphisms on filter systems. Analogously, for the left reflectivity classes, we get, on the one hand, that $\mathcal{I}$ is LRF prealgebraizable if and only if, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is a left order reflecting surjection commuting with inverse logical extensions, which restricts to an order embedding on filter systems, and, on the other, noting again the same subtle change, $\mathcal{I}$ is LR prealgebraizable if and only if, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is a left order reflecting surjection, which restricts to an order embedding commuting with inverse logical extensions on filter systems. Characterizations of the remaining six classes follow a similar pattern.

In Section 5.6, we switch from prealgebraizability to *algebraizability*. Dropping "pre" signifies using equivalentiality instead of the weaker preequivalentiality property. Equivalentiality encompasses protoalgebraicity and, under protoalgebraicity, only two classes of the ten potentially different ones are actually distinct. Accordingly, we get *family algebraizability*, or, simply, *F algebraizability*, when family injectivity is added, and *system algebraizability*, or, simply, *algebraizability*, when system injectivity is added. Family algebraizability is equivalent to algebraizability plus systemicity. Both properties transfer. Finally, $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is algebraizable if and only if it is stable and, for every **F**-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism commuting with inverse logical extensions, whereas $\mathcal{I}$ is family algebraizable if and only if, for all $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism commuting with inverse logical extensions.

1.3.5 Chapter 6

The motivating force behind the considerations in this chapter is the observation that, since for a π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$, with $\mathbf{F} = \langle \mathbf{Sign}^\flat, \mathrm{SEN}^\flat, N^\flat \rangle$, $\Omega(\overline{\varnothing}) = \nabla^{\mathbf{F}} = \Omega(\mathrm{SEN}^\flat)$, no π-institution without theorems can satisfy any of the injectivity, reflectivity of complete reflectivity properties introduced in Chapter 3. The question naturally arises whether, in that case, the existence of theory families with empty components is the only reason causing the lack of these properties or whether the π-institution in question would still not satisfy them even if theory families with empty components were in some way "discarded" or "bypassed". We choose two ways in which this circumvention may be accomplished, and study the various flavors of injectivity, reflectivity and complete reflectivity properties that result.

In Section 6.2, we introduce and study the relation of rough equivalence between theory families of a π-institution. Let $\mathbf{F} = \langle \mathbf{Sign}^\flat, \mathrm{SEN}^\flat, N^\flat \rangle$ be an algebraic system and $\mathcal{I} = \langle \mathbf{F}, C \rangle$ a π-institution based on $\mathbf{F}$. Given a theory family T of $\mathcal{I}$, we define the *rough companion* (*rough associate* or *rough representative*) $\widetilde{T}$ of T as the theory family resulting from T by replacing all empty Σ-components of T by the corresponding set $\mathrm{SEN}^\flat(\Sigma)$ of Σ-sentences. We say that two theory families T and T' are *roughly equivalent*, written $T \sim T'$, if $\widetilde{T} = \widetilde{T'}$. The rough equivalence class of T is denoted by $\widetilde{[T]}$ and $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ denotes the collection of all rough equivalence classes. When one considers the restriction of rough equivalence on theory systems, the corresponding rough equivalence class is denoted by $\widetilde{\lfloor T \rfloor}$ and the collection of all these classes by $\widetilde{\mathrm{ThSys}}(\mathcal{I})$. Reasoning with rough equivalence classes is one way of bypassing theory families with empty components. An alternative way is to ignore those theory families that have at least one empty component. This is accomplished by considering the collections $\mathrm{ThFam}^{\lightning}(\mathcal{I})$ and $\mathrm{ThSys}^{\lightning}(\mathcal{I})$ of all theory families and theory systems, respectively, none of whose components is empty.

The usefulness of rough equivalence in considering properties of the Leibniz operator stems from the fact that, for every theory family T, $\Omega(T) = \Omega(\widetilde{T})$. As a consequence, the Leibniz operator is constant on each rough equivalence class. It is fairly obvious that the rough companion $\widetilde{T}$ of a theory family T is the maximum element in the class $\widetilde{[T]}$. However, even if T happens to be a theory system, $\widetilde{T}$ may not be one. On the other hand, it can be shown that, even in that case, $\widetilde{\lfloor T \rfloor}$ has a maximum element, which, of course, does not coincide with $\widetilde{T}$. An unfortunate fact, when considering the operators $\overleftarrow{\ }$ and $\widetilde{\ }$ in the same context is that, even if two theory families T and T' are roughly equivalent, the same may not hold for $\overleftarrow{T}$ and $\overleftarrow{T'}$. On the positive side, if $\mathcal{A} = \langle \mathbf{A}, \langle F, \alpha \rangle \rangle$ is an $\mathbf{F}$-algebraic system and T is an $\mathcal{I}$-filter family of $\mathcal{A}$, we do have $\alpha^{-1}(\widetilde{T}) = \widetilde{\alpha^{-1}(T)}$. This implies that the action of α^{-1} preserves rough equivalence, i.e., if T and T' are $\mathcal{I}$-filter families of $\mathcal{A}$, with $T \sim T'$, then $\alpha^{-1}(T) \sim \alpha^{-1}(T')$, the latter being roughly equivalent theory families of $\mathcal{I}$.

In Section 6.3, we look at some notions combining systemicity with rough equivalence. They form a hierarchy weakening systemicity in the absence of theorems. In the presence of theorems, however, all concepts considered coincide. We say that a π-institution $\mathcal{I}$ is *roughly systemic* if, for every theory family T, $\overleftarrow{T}$ is roughly equivalent to T, i.e., $\overleftarrow{T} \sim T$. We say $\mathcal{I}$ is *narrowly systemic* if, for every theory family T in $\mathrm{ThFam}^{\nleq}(\mathcal{I})$ (i.e., with all components nonempty), $\overleftarrow{T} = T$. Finally, we say that $\mathcal{I}$ is *exclusively systemic* if, for all $T \in \mathrm{ThFam}^{\nleq}(\mathcal{I})$, such that $\overleftarrow{T} \in \mathrm{ThSys}^{\nleq}(\mathcal{I})$, $\overleftarrow{T} = T$. Systemicity is the strongest of these four conditions, followed by rough and narrow systemicity, which are incomparable in strength, and each of these two implies exclusive systemicity. Moreover, as mentioned previously, exclusive systemicity in the presence of theorems implies systemicity and, therefore, in that case, the entire hierarchy collapses to a single class.

In Section 6.4, we formalize and study various versions of rough injectivity, resulting by combining injectivity of the Leibniz operator with rough equivalence. The easiest to grasp is rough family injectivity. A π-institution $\mathcal{I}$ is *roughly family injective* if, for all theory families T, T', $\Omega(T) = \Omega(T')$ implies $T \sim T'$. *Rough left injectivity* results by replacing in the conclusion of the implication defining rough family injectivity T and T' by $\overleftarrow{T}$ and $\overleftarrow{T'}$, respectively. *Rough right injectivity* arises by a similar replacement in the hypothesis. Finally, *rough system injectivity* imposes the same condition as the family version, but restricts its application to theory systems. Rough right injectivity implies rough systemicity, but the converse fails in general. The rough injectivity hierarchy turns out to be more complex than the injectivity hierarchy studied in Section 3.6. There, it was shown that right injectivity implies family injectivity, which implies left injectivity, which, in turn, implies system injectivity, giving rise to a linear injectivity hierarchy. On the other hand, in the rough case, it is shown that rough right injectivity implies rough family injectivity, which implies the system version, and, in addition, rough left injectivity also implies the system version. Moreover, rough right injectivity is equivalent to rough system injectivity plus rough systemicity. Rough system injectivity, supplemented with stability, implies rough left injectivity. Each of the four rough injectivity properties, together with the availability of theorems, is equivalent to the corresponding injectivity property. The section concludes by establishing that all four rough injectivity properties transfer and by providing characterizations of rough family and rough system injectivity via the Leibniz operator Ω, viewed as a mapping from $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ and $\widetilde{\mathrm{ThSys}}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 6.5, we switch to a different version of injectivity properties, the overarching motivation still remaining that of bypassing theory families with empty components. *Narrow family injectivity* is defined by imposing the injectivity of the Leibniz operator on $\mathrm{ThFam}^{\nmid}(\mathcal{I})$, i.e., by stipulating that, for all $T, T' \in \mathrm{ThFam}^{\nmid}(\mathcal{I})$, $\Omega(T) = \Omega(T')$ implies $T = T'$. *Narrow left injectivity* replaces T, T' in the conclusion by $\overleftarrow{T}, \overleftarrow{T'}$, respectively, whereas *narrow right injectivity* applies the same replacement in the hypothesis. Finally, *narrow system injectivity* enforces the same condition as that of narrow family injectivity, but restricts its scope on theory systems in $\mathrm{ThSys}^{\nmid}(\mathcal{I})$. Narrow right injectivity implies exclusive systemicity, but does not imply any of the stronger versions of rough or narrow systemicity. With narrow injectivity, we recover the linearity of the injectivity hierarchy that was lost in passing to rough injectivity. That is, narrow right injectivity implies narrow family injectivity, which implies narrow left injectivity, which, in turn, implies the system version. Moreover, narrow system injectivity, supplemented by narrow systemicity, implies narrow right injectivity. It turns out that narrow family injectivity is equivalent to rough family injectivity. On the other hand, the two left injectivity properties, narrow left and rough left injectivity, ar incomparable, i.e., none implies the other. Some order is regained when looking at the right versions, where rough right injectivity implies narrow right injectivity. This order is maintained at the system level in which rough system injectivity also implies narrow system injectivity. As was the case with rough injectivity, each narrow injectivity property, supplemented with the existence of theorems, is equivalent to the corresponding injectivity property. Moreover, all four narrow injectivity properties transfer. Finally, the family and system versions have characterizations in terms of the injectivity of Ω, viewed as a mapping from $\mathrm{ThFam}^{\nmid}(\mathcal{I})$ and $\mathrm{ThSys}^{\nmid}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Sections 6.4 and 6.5, we looked at the rough and narrow injectivity hierarchies. Following this paradigm, in Sections 6.6 and 6.7, we introduce and study the rough and narrow reflectivity properties and, then, in Sections 6.8 and 6.9, the rough and narrow complete reflectivity properties.

In Section 6.6, we turn to rough reflectivity. Once more, the family version is the easiest to describe. A π-institution is called *roughly family reflective* if, for all theory families T, T', $\Omega(T) \leq \Omega(T')$ implies $\widetilde{T} \leq \widetilde{T'}$. *Rough left reflectivity* results by replacing T, T' in the conclusion by $\overleftarrow{T}$, $\overleftarrow{T'}$, respectively. *Rough right reflectivity* applies the same change in the hypothesis. Finally, *rough system reflectivity* imposes the same implication as the family version, but only on theory systems. Rough right reflectivity implies rough systemicity. It also implies rough family reflectivity, which implies rough system reflectivity. Rough left reflectivity also implies the system version. Rough right reflectivity is actually equivalent to the system version plus rough systemicity. On the other hand, rough system reflectivity and stability imply rough left reflectivity. It is straightforward to see, based on the relevant defining conditions, that each of the four rough reflectivity versions implies the corresponding rough injectivity version. Furthermore, each rough reflectivity version, supplemented by the existence of theorems, is equivalent to the corresponding reflectivity property. The section concludes with a proof that all four rough reflectivity properties transfer and with characterizations of rough family and rough system reflectivity in terms of the Leibniz operator, viewed as a mapping from $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ and $\widetilde{\mathrm{ThSys}}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 6.7, we look at narrow reflectivity properties. These constitute alternatives to rough reflectivity when dealing with reflectivity properties while attempting to bypass theory families with empty components. A π-institution is *narrowly family reflective* if, for all theory families T, T' in $\mathrm{ThFam}^{\lightning}(\mathcal{I})$, $\Omega(T) \leq \Omega(T')$ implies $T \leq T'$. As before, *narrow left reflectivity* results by replacing T, T' in the conclusion by $\overleftarrow{T}$, $\overleftarrow{T'}$, respectively, and *narrow right reflectivity* by performing the same replacement in the hypothesis instead. Finally, *narrow system reflectivity* stipulates that, for all $T, T' \in \mathrm{ThSys}^{\lightning}(\mathcal{I})$, $\Omega(T) \leq \Omega(T')$ implies $T \leq T'$. Narrow family reflectivity implies exclusive systemicity. As was the case with narrow injectivity properties, narrow reflectivity properties also align into a linear hierarchy. The strongest is narrow right reflectivity, followed by narrow family reflectivity, then by the left version and, at the tail, by narrow system reflectivity. The weakest one, narrow system reflectivity, supplemented by narrow systemicity, implies narrow right reflectivity. The relationships between corresponding rough and narrow versions of reflectivity follow those established in Section 6.5 between corresponding rough and narrow injectivity properties. First, rough family and narrow family reflectivity are equivalent. On the opposite end, the left versions turn out to be incomparable. Somewhere in between, for both the right and system versions, it turns out that the rough property implies the narrow one. Not surprisingly, each narrow reflectivity property implies the corresponding narrow injectivity property. Moreover, a given narrow reflectivity property is equivalent to the corresponding reflectivity property in the presence of theorems. All four narrow reflectivity properties transfer. Finally, characterizations are provided of narrow family and narrow system reflectivity in terms of the Leibniz operator seen as a mapping from $\mathrm{ThFam}^{\lightning}(\mathcal{I})$ and $\mathrm{ThSys}^{\lightning}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 6.8, we turn to complete reflectivity (c-reflectivity) properties starting with rough complete reflectivity. A π-institution $\mathcal{I}$ is *roughly family c-reflective* if, for every collection $\mathcal{T} \cup \{T'\}$ of theory families, $\bigcap_{T \in \mathcal{T}} \Omega(T) \leq \Omega(T')$ implies $\bigcap_{T \in \mathcal{T}} \widetilde{T} \leq \widetilde{T}'$. The *left version* results by replacing each theory family by its arrow counterpart in the conclusion, whereas the *right one* by applying the same change in the hypothesis instead. Finally, the *system version* stipulates that the same condition as that defining the family version applies, but $\mathcal{T} \cup \{T'\}$ is allowed to range over collections of theory systems instead of arbitrary theory families. Paralleling the rough reflectivity hierarchy, rough right c-reflectivity implies rough family c-reflectivity, which implies rough system c-reflectivity, while the left version also implies the system version. In fact, rough right c-reflectivity is equivalent to rough system c-reflectivity plus rough systemicity, whereas rough system c-reflectivity, together with stability, imply rough left c-reflectivity. It is clear that each rough c-reflectivity property generalizes the corresponding rough reflectivity property. It is also not difficult to show that each rough c-reflectivity property, in the presence of theorems, coincides with the corresponding c-reflectivity property. All four rough c-reflectivity properties transfer and, as before, characterizations may be formulated of the family and system versions in terms of the Leibniz operator, perceived as a mapping from $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ and $\widetilde{\mathrm{ThSys}}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

Section 6.9 deals with narrow complete reflectivity. A π-institution $\mathcal{I}$ is *narrowly family c-reflective* if, for every collection $\mathcal{T} \cup \{T'\} \subseteq \mathrm{ThFam}^{\nmid}(\mathcal{I})$, $\bigcap_{T\in\mathcal{T}} \Omega(T) \leq \Omega(T')$ implies $\bigcap \mathcal{T} \leq T'$. Once more, the *left version* arises by replacing all theory families in the conclusion by their arrow counterparts and, similarly, the *right version* by performing the same change in the hypothesis. *Narrow system c-reflectivity* imposes the same condition as the family version, but restricted to collections $\mathcal{T} \cup \{T'\} \subseteq \mathrm{ThSys}^{\nmid}(\mathcal{I})$. As with narrow reflectivity, the narrow c-reflectivity hierarchy is linear. The right version is the strongest, followed by the family, then the left and, finally, the system version. In addition, narrow system c-reflectivity, together with narrow systemicity, implies the right version. Comparisons between the rough c-reflectivity and the narrow c-reflectivity classes also follow the pattern revealed for corresponding reflectivity properties. In accordance, rough family and narrow family c-reflectivity are equivalent, rough left and narrow left c-reflectivity are incomparable, whereas the rough right and rough system versions imply, respectively, the narrow right and narrow system versions. As with their rough counterparts in Section 6.8, all four narrow c-reflectivity properties coincide with the corresponding c-reflectivity properties in the presence of theorems. Furthermore, all four narrow c-reflectivity properties transfer. The family and system versions have characterizations via the Leibniz operator seen as a mapping from $\mathrm{ThFam}^{\nmid}(\mathcal{I})$ and $\mathrm{ThSys}^{\nmid}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$, analogous to the ones obtained for both narrow injectivity and narrow reflectivity.

The last section of the chapter, Section 6.10, contains some characterizations of the property of a π-institution possessing theorems. This is closely connected to the overarching ideas governing the properties investigated in Sections 6.2-6.9, which aimed at rectifying the "pathologies" introduced by the absence of theorems. The availability of theorems is characterized by the injectivity of the Frege equivalence family operator, as well as by both the injectivity and the complete reflectivity of the Lindenbaum equivalence family operator, both applied to the collection of theory families of the π-

institution. These operators were introduced in Section 2.11. Possession of theorems transfers to the collections of all $\mathcal{I}$-filter families over arbitrary **F**-algebraic systems.

1.3.6 Chapter 7

In Chapter 7, we further pursue our endeavor of making properties in the lower bottom of the algebraic hierarchy suitable for the study of π-institutions that do not have theorems. Similarly to Chapter 6, we employ rough equivalence and narrowness to achieve this goal, but, unlike in Chapter 6, the focus here is on monotonicity and complete monotonicity properties, rather than on injectivity, reflectivity and complete reflectivity properties.

In Section 7.2, we define a *stability hierarchy*, which serves, in the sequel, to formalize properties of some of the classes in the monotonicity and complete monotonicity hierarchies. Recall that a π-institution $\mathcal{I}$ is *stable* if, for all theory families $T \in \mathrm{ThFam}(\mathcal{I})$, $\Omega(\overleftarrow{T}) = \Omega(T)$. Weakening this notion, we call $\mathcal{I}$ *narrowly stable* if the same equation holds, provided $T \in \mathrm{ThFam}^{\nmid}(\mathcal{I})$, i.e., the scope is restricted to theory families all of whose components are nonempty. A further weakening insists that the same equation hold for all $T \in \mathrm{ThFam}^{\nmid}(\mathcal{I})$, such that $\overleftarrow{T} \in \mathrm{ThSys}^{\nmid}(\mathcal{I})$, i.e., it further restricts the scope of the quantification to theory families all of whose components are nonempty and whose arrow counterparts also have all components nonempty. Clearly, stability implies narrow stability, which, in turn, implies the last property, which is termed *exclusive stability*. It is shown that both implications are strict.

In Section 7.3, we study the *rough monotonicity hierarchy*. Recall that, given a π-institution $\mathcal{I}$ and a theory family T of $\mathcal{I}$, $\widetilde{T}$ denotes the *rough companion* of the theory family T, which is the theory family resulting from T by replacing all empty Σ-components of T by $\mathrm{SEN}^\flat(\Sigma)$. Two theory families T and T' are *roughly equivalent* if they have the same rough companion. This is equivalent to saying that if T and T' differ at some signature Σ, they one has an empty Σ-component, whereas the other has $\mathrm{SEN}^\flat(\Sigma)$ as its Σ-component. A π-institution $\mathcal{I}$ is *roughly family monotone* if, for all theory families $T, T' \in \mathrm{ThFam}(\mathcal{I})$, $\widetilde{T} \leq \widetilde{T'}$ implies $\Omega(T) \leq \Omega(T')$. *Rough left monotonicity* results by replacing T, T' in the hypothesis by $\overleftarrow{T}$, $\overleftarrow{T'}$, respectively, and *rough right monotonicity* by applying the same replacement in the conclusion. *Rough system monotonicity* stipulates that the original implication hold, for all $T, T' \in \mathrm{ThSys}(\mathcal{I})$. It turns out that rough left monotonicity implies both rough family and rough right monotonicity and that each of the latter two implies the system version. Additionally, the strongest version, rough left monotonicity, is equivalent to the weakest, system, version, together with stability. Recall from Section 3.3 that family and left monotonicity are equivalent and this property was termed *protoalgebraicity*.

Recall also, from the same section, that system and right monotonicity are equivalent and this property was called *prealgebraicity*. Protoalgebraicity implies rough left monotonicity, whereas prealgebraicity implies rough right monotonicity. Tighter connections can be established under some fairly general hypotheses. For non almost inconsistent π-institutions, protoalgebraicity is equivalent to rough family or rough left monotonicity, coupled with the availability of theorems. Similarly, for π-institutions having a theory family $T \neq \mathrm{SEN}^\flat$, with $\overleftarrow{T} \neq \overline{\varnothing}$, prealgebraicity is equivalent to rough right or rough system monotonicity, supplemented with the availability of theorems. All four rough monotonicity properties transfer. E.g., $\mathcal{I}$ is roughly family monotone if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A} = \langle \mathbf{A}, \langle F, \alpha \rangle \rangle$ and all $\mathcal{I}$-filter families $T, T' \in \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A})$, $\widetilde{T} \leq \widetilde{T'}$ implies $\Omega^{\mathcal{A}}(T) \leq \Omega^{\mathcal{A}}(T')$. Both rough family and rough system monotonicity can be characterized using properties of the Leibniz operator viewed as a mapping from $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ and $\widetilde{\mathrm{ThSys}}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 7.4, we switch from rough monotonicity to *narrow monotonicity properties*. These constitute an alternative approach to bypassing theory families and theory systems with one or more empty components. We say that a π-institution $\mathcal{I}$ is *narrowly family monotone* if, for all theory families T, T', with all components nonempty, $T \leq T'$ implies $\Omega(T) \leq \Omega(T')$. The *left version* results by replacing T, T' by $\overleftarrow{T}$, $\overleftarrow{T'}$, respectively, in the hypothesis and the *right version* by performing the same replacement in the conclusion instead. *Narrow system monotonicity* stipulates that, for all $T, T' \in \mathrm{ThSys}^{\not\emptyset}(\mathcal{I})$, $T \leq T'$ implies $\Omega(T) \leq \Omega(T')$. Narrow left monotonicity implies narrow family monotonicity, which implies narrow system monotonicity, while the latter is also a consequence of narrow right monotonicity. Narrow left monotonicity is strong enough to yield exclusive stability, which, however, is the weakest of the three stability versions studied in Section 7.2. Under narrow systemicity, introduced in Section 6.3, the narrow monotonicity hierarchy collapses to a single class. Protoalgebraicity implies narrow left monotonicity and prealgebraicity implies the right version. In this case as well, tighter connections are possible under additional, fairly general, hypotheses, as was the case with rough monotonicity properties. Namely, under the hypothesis that $\mathcal{I}$ is not almost inconsistent, protoalgebraicity is equivalent to narrow left or narrow family monotonicity, coupled with the existence of theorems. And, provided that $\mathcal{I}$ possess a theory system $T \neq \overline{\varnothing}, \mathrm{SEN}^\flat$, prealgebraicity is equivalent to narrow right or narrow system monotonicity, together with the availability of theorems. Of central interest here is whether and how the rough monotonicity properties are related to the narrow monotonicity properties. In comparing the two hierarchies, we discover that the two family versions are equivalent, whereas each of the three remaining rough monotonicity properties implies the corresponding narrow monotonicity property. All four narrow monotonicity properties transfer. Finally, characterizations of the family and the system versions may be formulated in terms of the Leibniz operator seen as a mapping from $\mathrm{ThFam}^{\sharp}(\mathcal{I})$ and $\mathrm{ThSys}^{\sharp}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 7.5, we return to roughness, but study complete monotonicity (c-monotonicity) instead of monotonicity properties. *Rough family c-monotonicity* stipulates that, for all collections $\mathcal{T} \cup \{T'\} \subseteq \mathrm{ThFam}(\mathcal{I})$, $\widetilde{T'} \leq \bigcup_{T\in\mathcal{T}} \widetilde{T}$ implies $\Omega(T') \leq \bigcup_{T\in\mathcal{T}} \Omega(T)$. *Rough left c-monotonicity* and *rough right c-monotonicity* result by replacing in the hypothesis and in the conclusion, respectively, all theory families by their arrow versions. *Rough system c-monotonicity* imposes the same condition as does the family version, but restricts its applicability on collections $\mathcal{T} \cup \{T'\}$ consisting of theory systems. Here, it turns out that each of the left, family and right versions implies the system version. Moreover, rough left c-monotonicity is equivalent to the conjunction of rough system c-monotonicity and stability. It is also the case that, under stability, the rough family and rough right c-monotonicity properties coincide and that, under rough systemicity, the entire rough c-monotonicity hierarchy collapses to a single class. From the definitions, it is obvious that each of the four rough c-monotonicity properties implies the corresponding rough monotonicity version. It is also the case that each c-monotonicity property implies its rough c-monotonicity counterpart. Once more, for non almost inconsistent π-institutions, family (left c-monotonicity, respectively) is equivalent to the conjunction of rough family (rough left, respectively) c-monotonicity and the existence of theorems. Furthermore, if $\mathcal{I}$ possesses a theory family $T \neq \mathrm{SEN}^\flat$, such that $\overleftarrow{T} \neq \overline{\varnothing}$, then system (right, respectively) c-monotonicity is equivalent to rough system (right, respectively) c-monotonicity plus the existence of theorems. All four rough c-monotonicity properties transfer and one may, in this case also, recast the family and system versions in terms of properties of the Leibniz operator seen as a mapping from $\widetilde{\mathrm{ThFam}}(\mathcal{I})$ and $\widetilde{\mathrm{ThSys}}(\mathcal{I})$, respectively, to $\mathrm{ConSys}^*(\mathcal{I})$.

In Section 7.6, we switch from rough versions of c-monotonicity to narrow versions of the same property. A π-institution $\mathcal{I}$ is called *narrowly family c-monotone* if, for all collections $\mathcal{T} \cup \{T'\} \subseteq \mathrm{ThFam}^{\sharp}(\mathcal{I})$, $T' \leq \bigcup_{T\in\mathcal{T}} T$ implies $\Omega(T') \leq \bigcup_{T\in\mathcal{T}} \Omega(T)$. In the *left version*, all theory families are replaced in the hypothesis by their arrow counterparts and, in the *right version*, the same change is applied in the conclusion. The *system version* stipulates that the implication above hold for all collections $\mathcal{T} \cup \{T'\} \subseteq \mathrm{ThSys}^{\sharp}(\mathcal{I})$. Each of the left, family and right versions implies the system version. Moreover, each of the four c-monotonicity versions implies the corresponding narrow c-monotonicity version. As was the case in relating rough and narrow monotonicity classes in Section 7.4, rough family c-monotonicity is equivalent to narrow family c-monotonicity, whereas each of the other three rough c-monotonicity properties implies the corresponding narrow c-monotonicity version. From the definitions, it is clear that a narrow c-monotonicity property implies its narrow monotonicity counterpart, the latter being a special-ization of the former. All four narrow c-monotonicity properties transfer. In closing, both the family and the system versions have characterizations in terms of properties of the Leibniz operator perceived as a mapping from $\mathrm{ThFam}^{\nmid}(\mathcal{I})$ and $\mathrm{ThSys}^{\nmid}(\mathcal{I})$, respectively, into $\mathrm{ConSys}^*(\mathcal{I})$.

1.3.7 Chapter 8

In Chapter 8, we undertake the study of *regularity*. Roughly speaking, it is the property stipulating that, whenever two sentences belong to a theory family of a given π-institution, they must be identified modulo the Leibniz congruence system relative to that theory family. When, in addition to regularity, availability of theorems is also postulated, the property of *assertionality* is obtained. Assertionality strengthens complete reflectivity and, as a result, it can be used to strengthen (weak) (pre)algebraizability properties. These strengthenings and their associated hierarchies are under the microscope in Sections 8.4-8.7. The classes of π-institutions obtained here are among the most powerful classes in the semantic hierarchy of π-institutions, i.e., satisfy the strongest properties and are included in most of the other classes in the hierarchy.

In Section 8.2, we introduce *regularity*. As was the case with other properties in preceding chapters, regularity comes in four different versions. Once more, we begin from the easiest to describe, the family version. A π-institution $\mathcal{I}$ is *family regular* if, for all theory families T, all signatures Σ and all Σ-sentences ϕ an ψ, if $\phi, \psi \in T_\Sigma$, then $\langle \phi, \psi \rangle \in \Omega_\Sigma(T)$. *Left regularity* results by replacing T in the hypothesis by $\overleftarrow{T}$, *right regularity* by performing the same replacement in the conclusion instead, whereas *system regularity* stipulates that the implication hold for all theory systems T. Family regularity is the strongest of the four properties, followed by right regularity, which implies left regularity, which, in turn, implies the system version. Thus, regularity properties are stratified into a linear hierarchy. System regularity plus stability imply left regularity, and right regularity plus stability yield family regularity. It follows that, under stability, the four-class hierarchy is reduced to two classes. On the other hand, system regularity plus systemicity clearly yield family regularity, whence, systemicity leads to a collapse of the regularity hierarchy into a single class. The family, left and system versions have elegant characterizations in terms of the Suszko operator and one of its variants. E.g., a π-institution $\mathcal{I}$ is family regular if and only if, for every signature Σ and all Σ-sentences ϕ and ψ, $\langle \phi, \psi \rangle \in \widetilde{\Omega}^{\mathcal{I}}_\Sigma(C(\phi, \psi))$, where $C(\phi, \psi)$ is the least theory family of $\mathcal{I}$ containing ϕ and ψ. All four regularity properties transfer. For instance, with regards to the right version, $\mathcal{I}$ is right regular if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, all $\mathcal{I}$-filter families T of $\mathcal{A}$, all signatures Σ in $\mathcal{A}$ and all Σ-sentences ϕ and ψ, $\phi, \psi \in T_\Sigma$ implies $\langle \phi, \psi \rangle \in \Omega^{\mathcal{A}}_\Sigma(\overleftarrow{T})$. The other three transfer results are formalized similarly.

Finally, the family and system versions may be characterized by the property that the filter family (system, respectively) in any reduced matrix family (system, respectively) is at most a singleton, in the sense that it consists of components with at most one element each.

In Section 8.3, we study *assertionality*. This is the property resulting from regularity by adding the requirement that theorems exist. Accordingly, four versions of assertionality are a priori obtained, depending on which of the four versions of regularity is postulated. They are termed *family*, *right*, *left* and *system assertionality* and, based on the hierarchy of regularity properties of Section 8.2, these also form a linear hierarchy, with the family version at the top, followed by the right, then the left and, finally, the system version at the bottom of the hierarchy. Assertionality is characterized by asserting, roughly speaking, that each theory family is fully determined by its Leibniz congruence system as the equivalence class of any theorem. Even though, a priori, there are four assertionality versions, there is a reduction holding without proviso. More precisely, it can be shown that right assertionality implies systemicity and this entails that right and family assertionality are equivalent. This property implies left assertionality, which, in turn, implies the system version. Moreover, the latter supplied with systemicity, implies family assertionality. By the definitions, it is clear that each assertionality version implies the corresponding regularity version. What is, however, more interesting, albeit not much more challenging to demonstrate, is that each assertionality property implies the corresponding complete reflectivity (c-reflectivity) property (see Section 3.8). So the assertionality properties may be viewed as further strengthening the hierarchy of reflectivity and c-reflectivity properties, studied in Sections 3.7 and 3.8. All three different assertionality properties transfer. Again, indicative of the flavor, a π-institution $\mathcal{I} = \langle \mathbf{F}, C \rangle$ is, e.g., left assertional if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, the π-institution $\langle \mathcal{A}, C^{\mathcal{I},\mathcal{A}} \rangle$ is left assertional, meaning that, on the one hand, the least $\mathcal{I}$-filter family of $\mathcal{A}$ has all components nonempty and, on the other, that, for all $\mathcal{I}$-filter families T of $\mathcal{A}$, all signatures Σ and all Σ-sentences ϕ and ψ, such that $\phi, \psi \in T_\Sigma$, one has $\langle \phi, \psi \rangle \in \Omega^{\mathcal{A}}_\Sigma(T)$. The section concludes with characterizations of the family and system versions, analogous to the ones provided in the conclusion of Section 8.2 for regularity. Namely, it is shown that $\mathcal{I}$ is family (system) assertional if and only if the filter family (system, respectively) of every reduced matrix family (system, respectively) is a singleton (i.e., consists of singleton components).

In Sections 8.4-8.7, we take advantage of the role of assertionality in strengthening of c-reflectivity to obtain strengthened versions of weak (pre)-algebraizability and (pre)algebraizability properties. The first two are obtained by combining assertionality properties with pre- or protoalgebraicity, whereas the latter are obtained by using (pre)equivalentiality instead.

In Section 8.4, we look at *regular weak prealgebraizability* properties. These result from adding to prealgebraicity (i.e., system monotonicity) a version of assertionality. Since there are three distinct versions of assertionality, one obtains three distinct corresponding versions of regular weak prealgebraizability. A π-institution $\mathcal{I}$ is *regularly weakly family* (RWF) *prealgebraizable* if it is prealgebraic and family assertional. It is *regularly weakly left* (RWL) *prealgebraizable* if it is prealgebraic and left assertional and it is *regularly weakly system* (RWS) *prealgebraizable* if it is prealgebraic and system assertional. Since the distinguishing feature between these three properties is the type of assertionality imposed, the assertionality hierarchy immediately yields that RWF prealgebraizability implies RWL prealgebraizability, which, in turn, implies RWS prealgebraizability. Equally clear from the definitions is the fact that RWF/L/S prealgebraizability implies, respectively, family/left/system assertionality. Additionally, the fact that each assertionality property implies its c-reflectivity counterpart entails that RWF/L/S prealgebraizability implies, respectively, WF/L/SC prealgebraizability (see Section 4.2). All three versions of regular weak prealgebraizability transfer. The section concludes with characterizations of the three versions based on the Leibniz operator viewed as a mapping between ordered sets of filter families/systems and congruence systems. To provide a flavor, we look at RWF prealgebraizability. The characterization states that $\mathcal{I}$ is RWF prealgebraizable if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism, such that, for all $T \in \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A})$, $T/\Omega^{\mathcal{A}}(T)$ is a singleton.

In Section 8.5, we study *regular weak algebraizability*. The properties here are obtained from the regular weak prealgebraizability properties of Section 8.4 by upgrading prealgebraicity to protoalgebraicity. Accordingly, a π-institution $\mathcal{I}$ is *regularly weakly family* (RWF) *algebraizable* if it is protoalgebraic and family assertional, it is *regularly weakly left* (RWL) *algebraizable* if it is protoalgebraic and left assertional and it is *regularly weakly system* (RWS) *algebraizable* if it is protoalgebraic and system assertional. Notice that, since these properties constitute enhancements of the properties of Section 8.4, the right version has been absorbed within the family version. Here, however, protoalgebraicity, which, unlike prealgebraicity, implies stability, forces, in addition, the identification of the left and the system versions. Thus, there are only two distinct regular weak algebraizability properties, regular weak family (equivalently, right) algebraizability being the strongest and regular weak system (equivalently, left) algebraizability the weakest of the two. In comparing this two-step hierarchy with that of regular weak prealgebraizability properties, we discover that the two family versions coincide, whereas regular weak system algebraizability implies regular weak left prealgebraizability. As a consequence, the combined regular weak (pre)algebraizability hierarchy consists of four classes that are linearly ordered. Moreover, essentially due to the fact that assertionality properties imply c-reflectivity properties, each of the two regular weak algebraizability classes are included in the corresponding weak algebraizability classes. Both regular weak algebraizability properties transfer. Finally, both have characterizations in terms of the Leibniz operator seen as a mapping between ordered sets. E.g., $\mathcal{I}$ is RWS algebraizable if and only if, for every $\mathbf{F}$-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphim, such that, for all $T \in \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A})$, $T/\Omega^{\mathcal{A}}(T)$ is a singleton.

In Section 8.6, we introduce *regular prealgebraizability* properties. These are obtained by combining assertionality properties with preequivalentiality. Recalling that preequivalentiality is obtained by adding system extensionality to prealgebraicity, an alternative point of view is that regular prealgebraizability is obtained from regular weak prealgebraizability, studied in Section 8.4, by adding system extensionality. A π-institution $\mathcal{I}$ is *regularly family* (RF) *prealgebraizable* if it is preequivalential and family assertional, it is *regularly left* (RL) *prealgebraizable* if it is preequivalential and left assertional and it is *regularly system* (RS) *prealgebraizable* if it is preequivalential and system assertional. Based on the linear hierarchy of assertionality properties, we obtain a linear hierarchy of regular prealgebraizability properties, with RF prealgebraizability at the apex, followed by RL prealgebraizability, while RS prealgebraizability is at the bottom. Since preequivalentiality strengthens prealgebraicity, RF/L/S prealgebraizability implies, respectively, RWF/L/S prealgebraizability. Moreover, since each version of assertionality implies the corresponding c-reflectivity version, RF/L/S prealgebraizability implies, respectively, family/ left c-reflective/ system prealgebraizability (see Section 5.5). All three versions transfer. Finally, characterization theorems may be formulated for each of the three properties in terms of the Leibniz operator viewed as a mapping between ordered sets. To provide, once more, a preview, we mention the form this characterization takes in the case of regular left prealgebraizability. A π-institution $\mathcal{I}$ is regularly left prealgebraizable if and only if, for every $\mathbf{F}$-algebraic system, $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiSys}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order embedding, commuting with inverse logical extensions, such that, for all $T \in \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A})$, $\overleftarrow{T}/\Omega^{\mathcal{A}}(T)$ is a singleton.

In Section 8.7, the last section of Chapter 8, we turn to the study of *regular algebraizability* properties, which combine equivalentiality with assertionality. Equivalentiality forms a common strengthening of both protoalgebraicity and preequivalentiality. Even though one obtains, a priori, three versions of regular algebraizability, only two are distinct. We say that $\mathcal{I}$ is *regularly family* (RF) *algebraizable* if it is equivalential and family assertional, *regularly left* (RL) *algebraizable* if it is equivalential and left assertional, and *regularly system* (RS) *algebraizable* if it is equivalential and system assertional. Regular left and regular system algebraizability coincide and, as a result, the regular algebraizability hierarchy consists of the class of RF algebraizable π-institutions and its proper subclass of RS algebraizable π-institutions. In comparing regular algebraizability with regular prealgebraizability properties, we discover that the two family versions are equivalent and that regular system algebraizability implies regular left prealgebraizability. Further, in comparing regular algebraizability with regular weak algebraizability properties, we obtain, based on equivalentiality's dominant position over protoalgebraicity, that RF/S algebraizability implies, respectively, RWF/S algebraizability. In the ultimate comparison between subhierarchies, based on the fact that assertionality implies c-reflectivity, we obtain that RF/S algebraizability implies, respectively, F/S algebraizability. The section closes with the same type of theorems as previous sections. Namely, it is shown that both versions of regular algebraizability transfer from a π-institution to the filter families/systems over arbitrary **F**-algebraic systems and characterizations of both versions are obtained in terms of the Leibniz operator perceived as a mapping between ordered sets. The family version, e.g., asserts that $\mathcal{I}$ is regularly family algebraizable if and only if, for every **F**-algebraic system $\mathcal{A}$, $\Omega^{\mathcal{A}} : \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A}) \to \mathrm{ConSys}^{\mathcal{I}*}(\mathcal{A})$ is an order isomorphism commuting with inverse logical extensions, such that, for all $T \in \mathrm{FiFam}^{\mathcal{I}}(\mathcal{A})$, $T/\Omega^{\mathcal{A}}(T)$ is a singleton.

1.3.8 Chapter 9

In Chapter 9, we undertake the study of finitarity properties of weakly family algebraizable π-institutions. Here we draw inspiration by the analysis of corresponding properties of algebraizable sentential logics.

According to the theory of algebraization of sentential logics, a, not necessarily finitary, algebraizable sentential logic $\mathcal{S}$ is algebraized via an equivalence that relates its consequence relation with the equational consequence of a generalized quasivariety K. The relation of equivalence is established via a possibly infinite set of defining equations $E(x)$ in a single variable x, which serve to translate formulas into equations, and a possibly infinite set $\Delta(x,y)$ of equivalence formulas in two variables x and y, which serve to translate equations into formulas. Besides constituting interpretations between the two consequences, they should be mutually inverse in a specific sense. In examining the relationships between the various finitarity conditions that may hold, namely, $\mathcal{S}$ finitary, $\mathcal{S}_{\mathsf{K}}$ (the equational deductive system induced by K) finitary, $E(x)$ finite and $\Delta(x,y)$ finite, one may show that they are related by the implications depicted in the following diagram (see p. 137 in Section 3.4 of [86]).

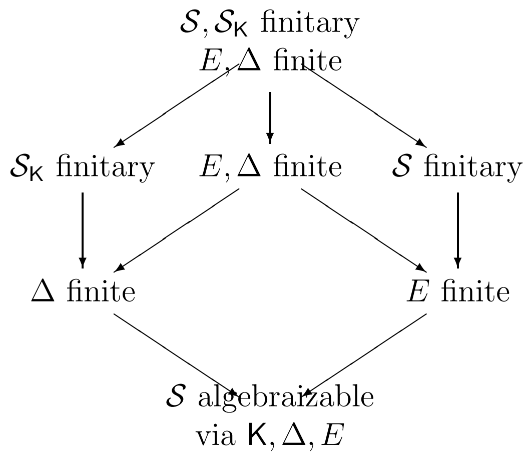

In the framework of sentential logics, roughly speaking, syntactic and semantic properties, i.e., those imposing the existence of transformations, such as $E(x)$ and $\Delta(x, y)$, satisfying certain properties, and those defined by order-theoretic properties of the Leibniz operator go hand-in-hand, in a tight correspondence. This is not the case in the framework of logics formalized as π-institutions. So in this chapter, the goal is to translate the sentential finitarity conditions to corresponding semantic properties and to establish an analogous hierarchy for weakly family algebraizable π-institutions. We also use examples from the sentential framework, recasting them as π-institutions, to obtain logical systems that serve to separate the classes of π-institutions specified by these finitarity properties.

In Section 9.2, the concept of *π-structure* is introduced, which abstracts that of a π-institution by removing the requirement of structurality. For π-structures, and, hence, also for π-institutions, the *finitary companion* is constructed, which is the π-structure over the same base algebraic system that has the largest finitary closure family included in the closure family of the given π-structure. *Locally finitely generated theory families* are defined and they are used to characterize those sentence families of a π-structure that are theory families of its finitary companion. These turn out to be exactly those sentence families that are unions of directed collections of locally finitely generated theory families of the given π-structure.

In Section 9.3, we investigate under which provisos, if any, the properties that define weak family algebraizability, i.e., protoalgebraicity and family reflectivity, are inherited by the finitay companion from the original π-structure and vice-versa. It is shown, first, that protoalgebraicity and family reflectivity are propagated from the finitary companion up to the parent π-structure unconditionally. On the other hand, the reverse inheritance requires additional conditions. To this end, the concept of *continuity* of the Leibniz and of the inverse Leibniz operator are introduced. The latter, of course, makes sense only if the π-institution under consideration is such that its Leibniz operator is an isomorphism, e.g., when it is weakly family algebraizable, which is precisely the case we focus on. If the Leibniz operator is continuous, it is easy to see that the π-institution is protoalgebraic. So continuity of the Leibniz operator actually strengthens protoalgebraicity. If, in addition to continuity, finiteness of the signature category is postulated, then the finitary companion is also protoalgebraic. Finally, it is shown that, if a π-institution, with a finite category of signatures, is weakly family algebraizable and both its Leibniz and inverse Leibniz operators are continuous, then its finitary companion is also weakly family algebraizable.

In Section 9.4, we undertake a detailed study of the interrelationships of the four finitarity properties pertaining to weakly family algebraizable π-institutions. These are the finitarity of the π-institution itself, the finitarity of its equational counterpart, the continuity of the Leibniz operator and the continuity of the inverse Leibniz operator, which is well-defined precisely because the π-institution is assumed to be weakly family algebraizable. These four properties are appropriate abstractions in the semantical institutional context of the properties of an algebraizable sentential logic being finitary, of its equivalent algebraic semantics being a quasivariety, of the set of equivalence formulas being finite and of the set of defining equations being finite, respectively. The close analogy is reflected in the fact that the results and hierarchy obtained here parallel the ones that hold for the corresponding properties in the sentential context. Our results come, as do their sentential counterparts, in dual pairs. In the first, it is shown that, for a weakly family algebraizable π-institution $\mathcal{I}$, the finitarity of $\mathcal{I}$ implies the continuity of its inverse Leibniz operator and, dually, the finitarity of the equational π-structure $\mathcal{Q}^{\mathsf{K}}$ induced by $\mathsf{K} := \mathrm{AlgSys}(\mathcal{I})$ implies the continuity of the Leibniz operator itself. Next, it is shown that, under weak family algebraizability, the finitarity of $\mathcal{I}$ and the continuity of the Leibniz operator imply that the equational counterpart is also finitary. Dually, the finitarity of the equational counterpart and the continuity of the inverse Leibniz operator imply that $\mathcal{I}$ itself is finitary. These implications lead to the following conditional equivalences, all applying to weakly family algebraizable π-institutions. For continuous Leibniz and inverse Leibniz operators, a π-institution is finitary if and only if its algebraic counterpart is finitary. For a finitary π-institution, its counterpart is finitary if and only if its Leibniz operator is continuous. Finally, if the algebraic counterpart of a π-institution is finitary, then the π-institution itself is finitary if and only if its inverse Leibniz operator is continuous. These outcomes lead to a finitarity hierarchy for weakly family algebraizable π-institutions paralleling the hierarchy depicted above for sentential logics.

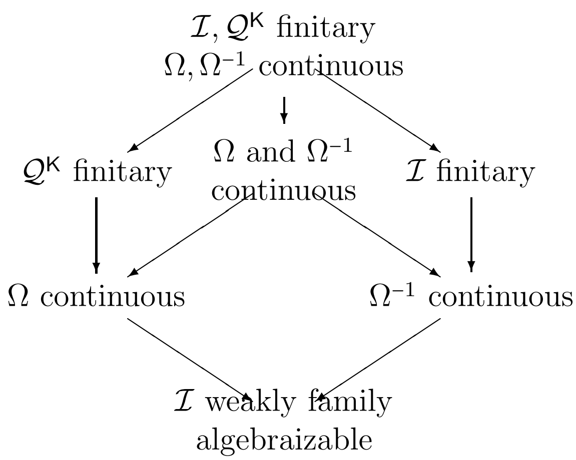

What remains to be done is separate the classes of π-institutions constituting the finitarity hierarchy. For this task, given the analogies established with the sentential framework, we seek inspiration from the realm of sentential logics.

In Section 9.5, we revisit three sentential logics that serve in separating the classes that form the finitarity hierarchy in the sentential framework. The classes related by the vertical arrows are separated by Łukasiewicz's infinite valued logic. This is a non-finitary, semantically defined sentential logic. It is algebraizable with a non-finitary equivalent algebraic semantics. On the other hand, both sets of defining equations and equivalence formulas are finite. The classes connected by the southeast arrows are separated using a finitary logic introduced by Dellunde and defined via a Hilbert calculus. It is regularly algebraizable via a singleton set of defining equations but a necessarily infinite set of equivalence formulas. Finally, the classes related by the southwest arrows of the diagram are separated using a non-finitary logic semantically defined by Raftery. This logic has a finitary equivalent algebraic semantics (actually a variety) and is algebraized via a finite set of equivalence formulas but a necessarily infinite set of defining equations. Even though we could certainly rely on well-written accounts from the literature to simply refer to these logics, we chose to recount all details, based on those original references. The Introduction to Chapter 9 and the main body contain more information, as well as appropriate references.

In Section 9.6, the three sentential logics of Section 9.5 are recast as π-institutions, according to the general procedure outlined in Section 1.1. The resulting π-institutions serve, in turn, in separating the corresponding classes appearing in the finitarity hierarchy of weakly family algebraizable π-institutions. Further evidencing the analogies described between the two finitarity hierarchies, the π-institution encapsulating Łukasiewicz's logic separates the classes of π-institutions connected via vertical arrows, the one incorporating Dellunde's logic separates classes along the southeast arrows, while the one arising from Raftery's logic separates classes related by the southwest arrows in the institutional finitarity hierarchy.

1.4 A Very Concise Summary of Contents

In Chapter 2, we introduce the basic definitions and fundamental results of algebra and logic and some indispensable notions and results pertaining to their interaction. These form the necessary background and the prerequisites for the general theory of algebraization of logics formalized as π-institutions that is presented in the monograph.

In Chapter 3, we introduce fundamental classes of the semantic Leibniz hierarchy. The term semantic alludes to the fact that they are defined purely by properties of the Leibniz operator on the complete lattices of the theory families or theory systems of π-institutions. Very central to our studies throughout, partly because they equip us with indispensable terminology regarding crucial properties, are the classes of systemic and stable π-institutions. At the bottom center of the hierarchy lie the loyalty properties. These simultaneously abstract monotonicity properties, on the one side, and reflectivity properties, on the other side. On the monotonicity side, we study monotonicity and two kinds of complete monotonicity, complete $\bigcup$-monotonicity, using the union operation, and complete $\bigvee$-monotonicity, using the join operation. In crossing over to the reflectivity side, we pass through, and study, injectivity properties. On the other side, we look, first, at reflectivity and, finally, at complete reflectivity properties. In Chapter 3, we not only define various flavors of each of these properties and compare their various strengths, but we also investigate the relations across those different kinds of properties. On the way, we also present many concrete examples, some of which are reused throughout the monograph to illustrate concepts, but, also - and mainly - to separate classes in the various hierarchies.

In Chapter 4, we study weak prealgebraizability and weak algebraizability properties. Weak prealgebraizability arises by combining prealgebraicity (system monotonicity) with one of the ten possible versions of injectivity, reflectivity or complete reflectivity. On the other hand, weal algebraizability results when combining protoalgebraiciy (family monotonicity) with one of those ten versions. Taking into account the combined hierarchy of injectivity, reflectivity and complete reflectivity properties, established in Chapter 3, we obtain a hierarchy of ten potentially different classes of weak prealgebraizability and a similar one consisting of ten potentially different classes of weak algebraizability. However, it is shown that the weak prealgebraizability hierarchy collapses down to six classes, whereas the one of weak algebraizability down to only two. Moreover, the top classes in the two hierarchies are identical. Therefore, when the two hierarchies are merged, a combined hierarchy consisting of seven distinct classes is obtained. The chapter includes, inter alia, characterizations of these seven classes using the Leibniz operator perceived as a mapping from the lattice of filter families to the poset of congruence systems over arbitrary algebraic systems.

In Chapter 5, we study the hierarchies of prealgebraizable and of algebraizable π-institutions. We look, first, at the property of extensionality and the seemingly weaker property of 2-extensionality and show that they are equivalent. Roughly speaking, extensionality relates Leibniz congruence systems of theories of an institution with those of corresponding theories of subinstitutions. Then, we look at the closely related properties of (Leibniz) commutativity and inverse (Leibniz) commutativity. These two properties are equivalent under monotonicity and, moreover, inverse commutativity is equivalent to extensionality. By combining monotonicity with extensionality properties, we build the hierarchy of equivalential π-institutions. Depending on which of the available versions of monotonicity or extensionality are imposed, three versions of equivalentiality arise, namely, equivalentiality, family preequivalentiality and (system) preequivalentiality in decreasing strength. By combining versions of preequivalentiality with injectivity, reflectivity or complete reflectivity properties, the ten classes of the prealgebraizability hierarchy are obtained. Similarly, by combining equivalentiality with injectivity properties (which are, in the presence of equivalentiality, equivalent to corresponding reflectivity or complete reflectivity properties), we get the two classes of algebraizable π-institutions.

In Chapter 6, we look at classes of the Leibniz hierarchy lying below the classes of injective, reflective and completely reflective π-institutions, which were introduced in Chapter 3. The motivating observation is that, if a π-institution satisfies injectivity or, a fortiori, reflectivity or complete reflectivity, then it must possess theorems. Thus, π-institutions without theorems are automatically excluded from consideration in contexts where these properties are postulated or studied. To bypass this hurdle, we define and study weakened versions of injectivity, reflectivity and complete reflectivity that can accommodate absence of theorems, but are equivalent to injectivity, reflectivity and complete reflectivity, respectively, in the presence of theorems. For each of those three properties, we study the rough versions and the narrow versions and carefully compare them to the original versions, as well as to each other, to obtain the hierarchies of injectivity, rough injectivity, narrow injectivity, reflectivity, rough reflectivity and narrow reflectivity, and c-reflectivity, rough c-reflectivity and narrow c-reflectivity classes of π-institutions. Roughly speaking, roughness identifies two theory families if their Σ-components are either equal or one is $\varnothing$ and the other is $\mathrm{SEN}^\flat(\Sigma)$. Those turn out to have identical Leibniz congruence systems. On the other hand, narrowness excludes from consideration altogether theory families with at least one empty component.

In Chapter 7, we continue the study of properties of π-institutions obtained by combining properties lying at the bottom of the Leibniz hierarchy with rough equivalence, on the one hand, and with narrowness, on the other. As opposed to Chapter 6, which considered properties lying below injectivity, reflectivity and complete reflectivity, this chapter undertakes the study of properties lying below monotonicity and complete monotonicity (c-monotonicity) properties. In a nutshell, roughly monotone and roughly c-monotone π-institutions form super classes, respectively, of the classes of monotone and c-monotone π-institutions. Additionally, narrowly monotone and narrow c-monotone π-institutions encompass respectively, roughly monotone and roughly c-monotone ones. By studying all four versions of each of these properties, we obtain a mixed hierarchy of rough and narrow monotonicity and rough and narrow c-monotonicity properties.

In Chapter 8, we study properties obtained by combining pre- or protoalgebraicity or (pre)equivalentiality, on the one hand, with assertionality, on the other. The latter, a property that strengthen complete reflectivity asserts, roughly speaking, that a π-institution has theorems and, in addition, each of its theory families is determined by its associated Leibniz congruence system as the equivalence class of a theorem. The chapter starts with the study of regularity, a property similar to assertionality, except that it does not require existence of theorems. It holds when any two sentences belonging to a theory family are identified modulo the Leibniz congruence system relative to that theory family. Assertionality properties are formalized next. The hierarchy they form and its interrelationships with the classes of the regularity hierarchy are explored in detail. Prealgebraicity, coupled with asserrtionality, gives rise to regular weak prealgebraizability, strengthening the classes of weak prealgebraizability properties. Protoalgebraicity, together with assertionality, leads to regular weak algebraizability properties. This hierarchy strengthens both regular weak prealgebraizability and weak algebraizability properties. Preequivalentiality and assertionality give rise to regular prealgebraizability, which strengthens both regular weak prealgebraizability and prealgebraizability. The chapter concludes with the study of regular algebraizability, which combines equivalentiality with assertionality. The classes of this hierarchy form subclasses of both those consisting of regularly prealgebraizable and those consisting of algebraizable π-institutions.

Chapter 9 starts with the introduction of the finitary companion of a π-institution. It is the largest finitary π-institution below the given one in the $\leq$ ordering of π-institutions based on the same algebraic system. The focus is on those properties defining weak family algebraizability, namely protoalgebraicity and family reflectivity. We investigate under which conditions, if any, those properties are passed from a π-institution to its finitary companion and vice-versa. In the second part, the focus shifts to the study of finitarity properties of weakly family algebraizable π-institutions. This class of π-institutions is chosen because, on its members, the Leibniz operator is an isomorphism and, hence, it makes sense to consider the inverse Leibniz operator. The four finitarity properties under investigation are the finitarity of the π-institution itself, the finitarity of its algebraic counterpart and the continuity of the Leibniz operator and of the inverse Leibniz operator. The implications holding between these properties give rise to the finitarity hierarchy of weakly family algebraizable π-institutions. The chapter also revisits some examples of sentential logics and formalizes them as π-institutions. The latter are then used to separate the various classes in the finitarity hierarchy. The three examples are Łukasiewicz's infinite valued logic, Dellunde's logic and a logic due to Raftery.

1.5 Further Reading

This is the first attempt to systematize the body of knowledge gathered over the years concerning the algebraization of logics formalized as π-institutions. However, for the readers interested in learning much more about the origins, history, concepts, results and developments in algebraic logic as applied to deductive systems, i.e., "abstract algebraic logic", there are a few excellent sources available that have served well over the years in educating the second and third generations of "abstract algebraic logicians".

Starting tangentially to the subject, but of interest, since they provide a comprehensive study of logical calculi and of institutions, respectively, the latter being the precursors of π-institutions used here, are the monographs by Wójcicki [34] and Diaconescu [79].

Two of the first sources that played a critical role in establishing and solidifying the discipline in its present form were the seminal "Memoirs" monograph of Blok and Pigozzi [35], in which algebraizable logics were introduced, and the pioneering monograph of Font and Jansana [52], in which generalized matrices were studied in a systematic way and the notion of Tarski congruence and accompanying reduced class of generalized matrices and underlying class of algebras were defined and studied in detail.

More at the textbook, rather than at the research, level, are the books of Czelakowski [64] and the more recent textbook by Font [86]. These are the only two books, to my knowledge, that are focused on systematically treating and presenting the most important results in the abstract setting. It goes, of course, without saying, that they both contain a plethora of concrete examples that have been studied in the literature, showcasing various aspects of the general theory and exemplifying the wide reach of its applicability.

Apart from research monographs and books, a few surveys have also appeared that provide overviews of, and/or details on, significant parts of the theory. Among them are [40], [68], [69, 80] and [90].

Finally, there have been a few, as far as I am aware, Ph.D. Dissertations which have dealt, either in their introductions or in their main corpus, with expositions and/or overviews of significant parts of the theory. Among them, some that have helped my own understanding and enhanced and/or diversified my point of view of various aspects of the theory are, in chronological order, those of Herrmann [43], Elgueta [47], Rebagliato [49], Dellunde [51], Gyuris [60], Martins [70], Russo [78], Albuquerque [85] and Moraschini [87].

The algebraization of logics formalized as π-institutions may be said to have started with the Ph.D. Dissertation by the author [97] (see, also, [98]), under the influence of preceding unpublished work by Zinovy Diskin [46] (see, also, [51]), which had been communicated to Professor Don Pigozzi, the author's Ph.D. Dissertation advisor, and used with Zinovy's kind permission and encouragement.