Saarland University
Faculty of Mathematics and Computer Science

Bachelor's Thesis

A Constructive Analysis of First-Order Completeness Theorems in Coq

Author
Dominik Wehr

Supervisor
Prof. Gert Smolka

Advisors
Dominik Kirst
Yannick Forster

Reviewers
Prof. Gert Smolka
Prof. Bernd Finkbeiner

Submitted: 15th July 2019

Eidesstattliche Erklärung

Ich erkläre hiermit an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe.

Statement in Lieu of an Oath

I hereby confirm that I have written this thesis on my own and that I have not used any other media or materials than the ones referred to in this thesis.

Einverständniserklärung

Ich bin damit einverstanden, dass meine (bestandene) Arbeit in beiden Versionen in die Bibliothek der Informatik aufgenommen und damit veröffentlicht wird.

Declaration of Consent

I agree to make both versions of my thesis (with a passing grade) accessible to the public by having them added to the library of the Computer Science Department.

Saarbrücken, 15th July, 2019

Abstract

First-order logic stands out among systems of similar expressivity as its deduction systems can be shown to be complete with regards to naive semantic accounts of validity. A deduction system is complete if it can prove every semantically valid formula. Historically, many proofs of first-order completeness have relied on non-constructive reasoning principles. In this thesis, we analyze multiple completeness theorems for variants of Gentzen's natural deduction and sequent calculus with regards to model and game semantics for first-order logic to determine which non-constructive principles are required to prove them.

In the first half of this thesis, we constructively analyze completeness theorems for the $\forall, \rightarrow, \bot$-fragment of first-order logic with regards to different notions of Tarski and Kripke models. We show that Veldman exploding models, which treat $\bot$ positively, and minimal models, which treat $\bot$ as an arbitrary logical constant, admit fully constructive completeness proofs. We also demonstrate the non-constructivity of completeness with regards to standard Tarski and Kripke models by relating them to the stability of provability. We derive tight characterizations of the requirements for multiple variants of standard completeness by identifying the principle of double-negation elimination and two different formulations of Markov's principle with the stability of provability restricted to different classes of theories.

The second half of the thesis is concerned with dialogue game semantics for full intutionistic first-order logic. We first give generic and fully constructive completeness proofs with regards to formal intuitionistic E- and D-dialogues. We then derive the completeness of the full intuitionistic sequent calculus with regards to formal first-order dialogues from this general result.

The analyses of this thesis are carried out in the calculus of inductive constructions and have been formalized in the interactive proof assistant Coq.

Acknowledgements

I want to thank both advisors of my thesis, Dominik Kirst and Yannick Forster. They let me explore the topic on my own accord, entertaining all of my whims and ideas while always acting as a voice of reason whenever I was at risk of losing sight of the thesis' direction. I want to thank Dominik in particular for making me realize that formal and philosophical inquiry do not have to be at odds with each other.

I also want to thank Professor Smolka for supervising this thesis. The courses and seminars offered by his chair made up for most of the electives of my undergraduate education and with good reason: The passion he displays for each aspect of the subjects he lectures on are both admirable and inspiring.

I thank Professor Finkbeiner who was my academic advisor during my undergraduate studies. His wealth of academic experience and willingness to critically question decisions I was about to make has helped me tremendously throughout these three years. I also thank him for agreeing to be the second reviewer of my thesis.

Further, I want to thank my family and friends. In particular, I want to thank Simon Spies for sharing my passion for logic and type theory. I also want to thank my family for their unwavering support of my academic aspirations, including graciously financing my studies in Saarbrücken. They were always there to cheer me on and back me up when I felt stuck or lost.

Lastly, I thank Dominik, Yannick and Simon for proof-reading this thesis and giving such helpful suggestions.

Contents

Chapter 1

Introduction

First-order logic has undoubtedly established itself as the canonical language of formal mathematical arguments. To give an example, Zermelo-Fraenkel set theory, the mathematical foundation most mathematicians and computer scientists work in, is traditionally formalized as a first-order theory. Even in fields outside the mathematical realm, such as linguistics and analytic philosophy, first-order logic enjoys widespread use. This predominant position is certainly not without basis as first-order logic is sufficiently expressive for most of mathematics while still having an extremely simple syntax.

The intrigue of logics originates in their semantics. A semantics is any formal or informal method of imbuing the formulas of a logical language with meaning. This is usually done by giving conditions for a formula A to be considered valid, which we denote by $\vDash A$. As an example, a possible account of $\vDash A \to B$ would be "in any universe in which A holds, B holds as well". When regarded only as a source of meaning, semantics should be simple and grant an immediate intuitive grasp on domain of the language they describe. However, this singular focus on the elegant characterization of a formula's meaning often leaves them unable to clarify another important aspect of logic: proof. Take as an example the semantic account of $\vDash A \to B$ from above. Knowing that $A \to B$ means that B holds in any universe in which A holds unambiguously characterizes $A \to B$, but how would one go about proving that $A \to B$ holds in the first place? This question cannot be answered on the basis of our semantic account without appeals to mathematical intuition.

Deduction systems are designed to give an unambiguous answer to this question. At first glance, deduction systems are a formalization of our intuitive understanding of valid mathematical arguments. They thus consist of multiple rules, each describing a valid argumentative step. A proof in a deduction system then consists of a finite chain of its rules, originating in axioms that are taken to always be evidently true and leading to the desired statement. Such deductive arguments usually establish that a finite list of formulas Γ "entails" a formula A, meaning that evidence for A can be deductively obtained from Γ. This fact is usually denoted by $\Gamma \vdash A$. An example of a deduction rule is given below.

$$\frac{\Gamma, A \vdash B}{\Gamma \vdash A \to B}$$

This rule expresses $A \to B$ can be deduced from Γ if B can be deduced from the list Γ extended with A. On its own, this rule isn't any more illuminating than the semantic account of $\vDash A \to B$. However, all rules of a deduction system taken together give a clear account of what constitutes a proof of $\Gamma, A \vdash B$. In fact, this notion of proof does not rely on any mathematical intuition, as the correctness of a deductive argument can be verified by simply checking if each of its deductive steps is in line with its rules. For this reason, deduction systems are often said to give a "purely syntactic" account of a logic. Of course, deduction systems are not completely detached from semantic accounts of validity, as their construction if often guided by a more intuitive semantics that serves as a heuristic for the admissibility of its rules.

Both semantics and deduction systems can be, and often are, specified in terms of an already existing mathematical framework. In that case, the underlying framework is called the "meta-logic" and the system that is being defined the "object-logic". If a deduction system $\Gamma \vdash A$ and a semantics $\Gamma \vDash A$ are defined in a meta-logic, their relationship can be analyzed formally. The two properties always considered in such analyses are soundness and completeness. A deduction system is said to be "sound" with regards to a semantics if everything that can be syntactically deduced in the system also holds semantically ($\Gamma \vdash A$ implies $\Gamma \vDash A$). This can be taken as the assertion that the deduction system maintains the validity characterized by the semantics: From valid sentences, only further valid sentences can be deduced. The opposite property is called "completeness". That is, a deduction system is complete with regards to a semantic account if everything that is semantically valid can be deduced ($\Gamma \vDash A$ implies $\Gamma \vdash A$). Note that soundness or completeness on their own each are fairly weak properties. For example, a deduction system that can not deduce anything is sound, while a deduction system that can deduce everything is complete. If, however, a deduction system is both sound and complete with regards to a semantics then the deduction system gives an accurate, purely syntactic characterization of the semantics' notion of validity. First-order logic is unique among systems of similar expressivity as its deduction systems can be shown to be both sound and complete with regards to rather naive (and thus very intuitive) semantic accounts.

In this thesis, we carry out multiple constructive analyses of completeness proofs for first-order logic. Constructivism is a philosophical position on what constitutes a convincing mathematical argument. Although the precise accounts of constructivism vary, they all agree on the demand that the witnesses of existential proofs should always be computable. This stands in opposition to the orthodox account of mathematical argument which allows existential statements to be derived from the impossibility of non-existence. As the constructivist demand on existential proof is thus not guaranteed by the classical systems usually employed in mathematics, constructivists work in systems of their own. These systems are often arrived at by removing deduction rules, such as the principle of the excluded middle, from classical systems. As a result, they can prove fewer theorems, which is often considered an indication of constructive systems being "weaker" than their classical counterparts. However, the opposite case can be made as well: While the constructive system may prove fewer formulas, the formulas it does prove have stronger epistemological backing, as all existential proofs that were made to arrive at them have explicit witnesses one can point to, thus making the system stronger. This line of argument can be taken even further: some statements that can be proven both classically and constructively are arguably only meaningful in a constructive setting. Examples of this can even be found in this thesis: In later chapters, we prove that certain completeness statements entail the classical principle of double negation elimination. In a classical system, such proofs are a triviality as the principle is already asserted by the system itself. However, in our constructive setting where these classical principles are not present, these proofs are an expression of the inherent classicality of these variants of first-order completeness and thus much more meaningful. As constructive systems always constitute subsystems of classical systems, one may be lead to believe that they are badly suited to expressing inherently classical concepts, such as semantic accounts of classical logics. However, this is not the case. For example, we give a very natural constructive account of a classical semantics in Chapter 3.

There are two distinct mathematical undertakings that are commonly deemed a "constructive analysis". The first one can be seen as a consequence of the close link between proof and computation in constructive settings: All constructive proofs have "computational content". That is, there always is an algorithm underlying the structure of constructive mathematical arguments. The resulting type of constructive analysis is an inquiry into these underlying programs, which can give further insight into constructively valid theorems. The second kind of mathematical undertaking is "constructive reverse mathematics" [34]. Reverse mathematics [60] is a broad field of mathematics that aims to find the minimal assumptions required to prove important mathematical theorems. If carried out from a constructivist perspective, the assumptions under scrutiny tend to be non-constructive reasoning principles, such as the principle of excluded middle or Markov's principle. While both approaches tend to come from a constructivist perspective, the analysis of computational content is concerned with constructive proofs, while constructive reverse mathematics usually works with non-constructive proofs. The notion of constructive analysis we adhere to in this thesis is closer to that of constructive reverse mathematics, although we use a different mathematical foundation than the second-order arithmetic usually employed by that community. More specifically, we analyze the minimal non-constructive assumptions required for various first-order completeness theorems.

The history of first-order completeness is rich enough to warrant its textbook. We thus restrict ourselves to a partial account, only covering proofs relating directly to the completeness theorems discussed in this thesis. The first proof of classical first-order completeness was given by Kurt Gödel in 1929 [23]. His proof was highly syntactic and made use of classical reasoning principles. 20 years later, Leon Henkin gave a revised, but still partially classical, account of first-order completeness [28], which has since become the standard presentation of this result due to its elegance. Saul Kripke proposed his model semantics for modal logic as a suitable semantics for intuitionistic first-order logic in 1965 [41]. In the same paper, he gave a classical proof of completeness with regards to Kripke semantics. A great advance in constructive completeness proofs was made by Wim Veldman in 1976, who gave a proof with regards to Kripke models which "treat negation positively" [65] that did not rely on the principle of excluded middle. However, his proof still required the use of the non-constructive fan theorem to handle $\vee$ and $\exists$. Jean-Louis Krivine applied this insight to classical first-order logic in 1996 [42], resulting in a proof of completeness for a similarly modified classical first-order semantics. As $\exists$ and $\vee$ do not have to be handled explicitly in classical settings, his proof was fully constructive.

We also consider dialogues as an intuitionistic semantics in this thesis. In his PhD thesis of 1961, Kuno Lorenz was the first to prove completeness with regards to various dialogue semantics for first-order logic [46], including one for intuitionistic logic. Walter Felscher later gave a modernized, constructive account of first-order completeness with regards to intuitionistic dialogues in 1985 [15], which has since become the canonical variant of this proof.

1.1 Contributions

The contents of this thesis can be split into two distinct halves. The first half analyzes completeness theorems for the $\forall, \rightarrow, \bot$-fragment of first-order logic with regards to various notions of Tarski and Kripke models. In that half

- we perform a unified analysis of the constructivity of Tarski and Kripke completeness for standard, exploding and minimal models. We thereby relate the various known results [29, 30, 38, 53, 4, 59] and highlight the symmetries of said results between the two semantics.

- we describe a constructive, Henkin-style theory extension procedure that is suitable for notions of validity with or without $\bot$, based on previous work by Herbelin and Ilik [29] and Schumm [59].
- we reframe the constructive analysis in terms of the stability of provability, thereby more clearly separating the non-constructivity of the completeness theorems from the specific formulations of the non-constructive principles (such as Markov's principle) we relate it to.

The second half is concerned with completeness proofs for intuitionistic dialogue semantics. In that half

- we present a novel way of formalizing dialogues as state transition systems which is better suited for type-theoretic settings than the traditional account of dialogues and thus allow for simpler proofs.
- we adopt the generalized completeness proofs for classical E-dialogues with finite rule sets by Sørensen and Urzyczyn [61] to intuitionistic E-dialogues with infinite rule sets
- we demonstrate how to derive the completeness of the full sequent calculus with regards to intuitionistic first-order E-dialogues from the general completeness result.
- we give a constructive proof of the equivalence of generalized intuitionistic E- and D-dialogues with enumerable rule sets based on a clear intuition in terms given by our formalization of dialogues as state transition systems.

The results of this thesis have been formalized in the interactive proof assistant Coq. The an explorable version of the source code of the formalization can be viewed online under `https://www.ps.uni-saarland.de/~wehr/bachelor/coq/toc.html`. To aid the exploration of the formalization, each definition and theorem is linked to its formalized counterpart.

1.2 Overview

We now remark on the structure of the remaining thesis. After closing this chapter with a brief overview of constructive type theory as a foundation of mathematics, we proceed by giving preliminary definitions in Chapter 2. These definitions cover central concepts of synthetic computability theory (Section 2.2), the syntax and a natural deduction system for first-order logic (Sections 2.3 and 2.4), and elaborations on the stability of provability and its relation to non-constructive principles (Section 2.5).

The main chapters of this thesis can be split into two halves. In the first half (Chapters 3 and 4), we constructively analyze completeness theorems for the $\forall, \rightarrow, \bot$-fragment of first-order logic with regards to different notions of Tarski and Kripke models, beginning with Tarski models in Chapter 3. For this, we give a generalized variant of Henkin's theory extension procedure in Section 3.2. We then constructively analyze completeness theorems with regards to different notions of Tarski models. We demonstrate the non-constructivity of completeness with regards to standard models by relating it to the stability of provability in Section 3.3.2. We then show that Veldman exploding models (Section 3.3.3), which treat $\bot$ positively, and minimal models (Section 3.3.4), which treat $\bot$ as an arbitrary logical constant, admit fully constructive completeness proofs. In Chapter 4, we derive symmetric results for Kripke models, meaning that exploding and minimal Kripke models allow for constructive completeness proofs (Section 4.3.1), while standard Kripke completeness can be related to the stability of provability (Section 4.3.2). We close the chapter by demonstrating that, as the Kripke completeness results were in terms of a normal sequent calculus, one can use them to derive a semantic proof normalization procedure in Section 4.4.

The second half of the thesis is concerned with dialogue semantics. We first present a novel way of formalizing dialogues and use it to give a completeness proof with regards to formal intuitionistic E-dialogues (Section 5.2). We then derive the completeness of the full intuitionistic sequent calculus with regards to formal first-order E-dialogues from this general result (Section 5.3). Further, we prove the completeness of generalized intuitionistic D-dialogues with enumerable rule sets (Section 5.4), which entails the equivalence between such E- and D-dialogues and the completeness of the full intuitionistic sequent calculus with regards to first-order D-dialogues.

We close the thesis in Chapter 6 by discussing our results (Section 6.1) as well as related and future work (Sections 6.2 and 6.3). We also remark on the Coq formalization that was developed alongside this thesis (Section 6.4).

1.3 On Constructive Type Theory

Type theories, such as Martin-Löf type theory [51] or the calculus of inductive constructions (CIC) [10], can be used as constructive foundations for mathematics. More precisely, they embody the strictest notion of constructivity, championed by mathematicians such as Heyting [32] and Bishop [6]. We also adapt this notion of constructivism for this thesis. As such, phrases like "constructive" and "non-constructive" should always be understood accordingly. The constructive analyses in this thesis are carried out in the calculus of inductive constructions. The remainder of this section gives an introduction to constructive type theory sufficient to enable a reader unfamiliar with type theory to understand the rest of this thesis.

At the heart of type theory is the concept of types. As an example, the natural numbers can be defined as the type given below.

$$\mathbb{N} : \mathbb{T} \;:=\; 0 \mid S\,n \qquad (n : \mathbb{N})$$

This definition establishes that the expression 0 is a member of the type of natural numbers $\mathbb{N}$, which is usually denoted by $0 : \mathbb{N}$. Additionally, the constructor S can be applied to an expression $n : \mathbb{N}$ to obtain another natural number $S\,n : \mathbb{N}$. This way of defining the natural numbers corresponds to that of Peano arithmetic, $0 : \mathbb{N}$ representing the number 0 and $S : \mathbb{N} \to \mathbb{N}$ denoting the successor function. The number 3, for example, is represented as $S\,(S\,(S\,0))$.

Another important feature of type theory are its built-in notions of function and computation. Functions can be defined in terms of λ-abstractions: the expression $(\lambda n.\ n + n) : \mathbb{N} \to \mathbb{N}$ defines the function that doubles a natural number. The computational behavior of terms is defined by reduction rules. For example, β-reduction describes the computation step of "inserting" arguments into functions. For instance, β-reduction allows $(\lambda n.\ n + n)\,m$ to be reduced to $m + m$.

Constructive type theory radically differs from other mathematical foundations, such as ZF set theory, by representing its propositions as types, just as it does with the objects that are reasoned about. The idea behind this approach is often called the "Curry-Howard-Isomorphism" and is summed up by the popular slogan "Proofs are Programs; Propositions are Types!". A simple example of a proposition represented as a type is that of the disjunction: Given two types A and B representing propositions, one can define a type representing their disjunction as below.

$$A \vee B : \mathbb{P} \;:=\; L\,a \mid R\,b \qquad (a : A, b : B)$$

A proposition-as-a-type is proven by giving an expression of that type. In the case of $A \vee B$, that can be done whenever a member (proof) of type A or B is already known. This exactly corresponds to the usual introduction rules for disjunctions found in natural deduction systems.

The attentive reader will have noticed a small difference between the definitions of $\mathbb{N}$ and $A \vee B$. The type of $\mathbb{N}$ is given as $\mathbb{T}$, whereas $A \vee B : \mathbb{P}$. These types of types are usually called type universes. Slightly simplifying, CIC has two universes: the impredicative universe of propositions $\mathbb{P}$ and the predicative universe of computational types $\mathbb{T}$. While there are important technical reasons for this distinction, it also improves the readability of type signatures. For example, the type of predicates on natural numbers is $\mathbb{N} \to \mathbb{P}$, that is, functions mapping natural numbers to propositions. An example for such a predicate is $\lambda n.\ n < 10 : \mathbb{N} \to \mathbb{P}$.

Types in CIC are not restricted to be parametric over other types. Types can also be defined dependent on terms (this is usually called "dependent types"). For example, the type corresponding to the existential quantifier can only be defined as a dependent type.

$$\exists p : \mathbb{P} \ := \ E\,x\,H \qquad (p : X \to \mathbb{P}, x : X, H : p\,x)$$

The type $\exists p$ is defined in terms of a predicate $p : X \to \mathbb{P}$. It can only be proven using the function E, which takes as arguments the witness $x : X$ as well as a proof $H : p\,x$ that x indeed satisfies the predicate p. More complicated propositions, such as equality or eveness of a number, can be represented as inductive types by using even more advanced features of type theory.

Function types take on a new role when interpreting types as propositions as well. If A and B are propositions, then $A \to B$ is the type corresponding to the proposition "A implies B". Thus, a proof of $A \to B$ is a function that takes a proof of A and transforms it into a proof of B. CIC and other powerful type theories like it have a second kind of function: dependent functions. As demonstrated in the previous paragraph, types can depend on terms. It can therefore be useful to be able to refer to arguments that have already been passed in when giving a function type. As an example, consider the function type $\forall x : X,\ x = x$ for a suitably defined predicate $(=) : X \to X \to \mathbb{P}$. This dependent function type describes a function that takes a value $x : X$ as an argument and returns a value (proof) of the type $x = x$. Thus, if we had a function of that type $r : \forall x : X,\ x = x$, we could obtain a proof of $x' = x'$ for a specific $x' : X$ with the expression $r\,x'$. As the notation we have used for dependent functions already suggests, they are the type-theoretic equivalent to the universal quantification. Note also that dependent functions cannot be defined inductively in CIC, such as $A \vee B$ or $\exists p$, but are primitives of the calculus instead.

Chapter 2

Preliminaries

Before we begin the analyses, we establish some preliminary notions. After briefly remarking on the type theory we use in Section 2.1, we elaborate on synthetic computability theory (Section 2.2) which we employ in Chapters 3 and 4. We then give the syntax of first-order logic (Section 2.3) and a natural deduction system for its $\forall, \rightarrow, \bot$-fragment (Section 2.4). We close this chapter by connecting the stability of deduction to various non-constructive principles (Section 2.5), which forms the basis for the constructive analysis in subsequent chapters.

2.1 Type Theory

All results of this thesis have been formalized using the interactive proof assistant Coq. Coq's logic is based on a constructive type theory called the "calculus of inductive constructions" (CIC) [10]. Matching the formalization, we assume a type theoretic framework in our writing.

The calculus of inductive constructions distinguishes between two kinds of universes. The impredicative universe $\mathbb{P}$ and an infinite hierarchy of predicative universes $\mathbb{T}_0 \subseteq \mathbb{T}_1 \subseteq \cdots$. For the sake of readability, we usually omit the index on $\mathbb{T}$. Traditionally, types representing propositions are defined as members of $\mathbb{P}$ whereas the types representing computational data are placed somewhere in the hierarchy of $\mathbb{T}$. Thus, predicates are functions of type $X_1 \rightarrow \cdots \rightarrow X_n \rightarrow \mathbb{P}$ with $X_i : \mathbb{T}$.

Throughout this thesis, we make use of multiple types that are well known in the type-theory community. Their definitions are given in Fig. 2.1. The boolean truth values true and false are represented by the type $\mathbb{B}$. The type $\mathbb{N}$ contains natural numbers in the usual Peano representation with $S : \mathbb{N} \rightarrow \mathbb{N}$ taking the role of the successor function. Tuples of fixed size are represented by product types of the form $X_1 \times \cdots \times X_n$. The sum type $X + Y$ may either contain a value of type X or a value of type Y. Lists of values of type X are represented by the type $\mathcal{L}(X)$. For the sake of readability, we write $[x_1, ..., x_n]$ for $x_1 :: ... :: x_n :: []$. The operation of concatenating two lists xs and ys is be denoted by $xs +\!\!+\, ys$. We also use a list-comprehension, denoted by $[\, f\, x_1 \dots x_n \mid x \in xs_1, \dots, x_n \in xs_n, p\, x_1 \dots x_n \,]$, yielding a list containing $f\, x_1 \dots x_n$ for all combinations of elements $x_i \in xs_i$ such that the boolean test $p\, x_1 \dots x_n$ holds. The type $\mathcal{O}(X)$ denotes optional values, either containing a value $x : X$, denoted by $\ulcorner x \urcorner$, or nothing at all, denoted by $\emptyset$. The type $\{x : X \mid p\, x\}$, often referred to as a Σ-type in the literature, can be seen as the subtype of type X which contains all values x satisfying the predicate $p : X \to \mathbb{P}$.

$$
\begin{aligned}
\mathbb{B} : \mathbb{T} &:= \mathsf{true} \mid \mathsf{false} & & \\
\mathbb{N} : \mathbb{T} &:= 0 \mid S\, n & & (n : \mathbb{N}) \\
X_1 \times \cdots \times X_n : \mathbb{T} &:= (x_1, \dots, x_n) & & (x_1 : X_1, \dots, x_n : X_n) \\
X + Y : \mathbb{T} &:= L\, x \mid R\, y & & (x : X, y : Y) \\
\mathcal{L}(X) : \mathbb{T} &:= x :: xs \mid [] & & (x : X, xs : \mathcal{L}(X)) \\
\mathcal{O}(X) : \mathbb{T} &:= \ulcorner x \urcorner \mid \emptyset & & (x : X) \\
\{x : X \mid p\, x\} : \mathbb{T} &:= E\, x\, H & & (x : X, p : X \to \mathbb{P}, H : p\, x)
\end{aligned}
$$

Figure 2.1: Type definition

While sets are not primitive to type theory, they can be represented as predicates. For example, sets of natural numbers are represented by predicates $X : \mathbb{N} \to \mathbb{P}$. A number $n : \mathbb{N}$ is considered a member of such a set X if $X\, n$ holds. To ease readability, we write $n \in X$ in for $X\, n$ and $2^{\mathbb{N}}$ for the type $\mathbb{N} \to \mathbb{P}$ when treating predicates as sets. Defining the usual set-theoretic operations on them is straightforward: The union of two sets $X \cup Y$ is given by $(\lambda x.\, x \in X \vee x \in Y)$ and their intersection $X \cap Y$ by $(\lambda x.\, x \in X \wedge x \in Y)$. Further, a function $f : A \to B$ can be applied to a set X as $f\, X := (\lambda b.\, \exists a.\, b = f\, a \wedge a \in X)$.

2.2 Synthetic Computability Theory

The results of computability theory have traditionally been established in terms of formalized models of computation, such as Turing machines or the λ-calculus. However, working in such models in a strict formal setting such as Coq can be tedious: They usually operate on a very low level of abstraction which means most interesting results can only be established after proving many intermediate results that one would simply gloss over on paper. Constructive foundations offer an attractive alternative to this approach: As any function that can be defined constructively is computable, many results of computability theory can be established in terms of these intrinsic computations. Thus, the computability of a certain function can be proven by simply giving an implementation of it in CIC. This approach is often called "synthetic computability theory" [56, 3]. We use this section to introduce core concepts of synthetic computability theory as adapted to CIC by Forster, Kirst and Smolka [20].

The first notion are concerned with are decidable predicates. That is, predicates which can be decided by a boolean function. We have mentioned previously that constructive settings such as CIC do not allow unrestricted classical reasoning. However, instances of classical principles such as that of the excluded middle can still be derived for any decidable predicate. Thus decidability of predicates becomes of great interest even in works not directly concerned with computability theory, such as ours.

Definition 2.1 (Decidable predicates) *A predicate $P : X \to \mathbb{P}$ is called decidable if there exists a boolean decider $d : X \to \mathbb{B}$ such that $d\,x = \texttt{true}$ iff $P\,x$ for every x.*

Fact 2.1 *Let $P : X \to \mathbb{P}$ be a decidable predicate. Then either $P\,x$ or $\neg P\,x$ can be derived for any x.*

Types with decidable equality are called discrete types. They are of interest as many important properties, such as list membership, are decidable for all discrete types.

Definition 2.2 (Discrete types) *A type X is called discrete if the predicate there is a boolean decider for the predicate $\lambda\,(x, y).\; x = y : X \times X \to \mathbb{P}$*

Along with decidablity, (recursive) enumerability also constitutes a standard concept of computability theory. A predicate / problem is said to be enumerable if there exists a possibly divergent procedure that outputs all values for which the predicate holds. However, the calculus of inductive constructions does not allow functions to diverge as this would cause inconsistencies in its notion of provability. Hence, the synthetic adaption of this notion is slightly different: Instead of a continuously running procedure, the enumeration is performed by a step-indexed enumerator $e : \mathbb{N} \to \mathcal{L}(X)$ which generates a growing list. This means, $e\,n$ is required to be a prefix of $e\,(S\,n)$. Intuitively, these enumerators map a number n to all the values that would have been output if the procedure "ran for n steps". This notion can be extended to types, which are considered enumerable if all of their values can be enumerated.

Definition 2.3 (Enumerability)

1. *A function $e : \mathbb{N} \to \mathcal{L}(X)$ is called an enumeration if $\forall n.\, \exists xs.\; e\,(S\,n) = e\,n +\!\!+\, xs$.*
2. *A predicate $P : X \to \mathbb{P}$ is called enumerable if there is an enumeration $e : \mathbb{N} \to \mathcal{L}(X)$ with $(\exists n.\; x \in\; e\,n)$ iff $P\,x$ for every $x : X$.*
3. *A type X is called enumerable if there exists an enumeration $e : \mathbb{N} \to \mathcal{L}(X)$ such that for every $x : X$ there is an n with $x \in e\,n$.*

2.3 First-Order Logic

2.3.1 Syntax

The syntax of first-order logic is split into two categories, each represented as a type. The terms $\mathfrak{T}$ represent the objects that are being reasoned about. They usually consist of function applications and constant symbols as well as the variables bound by the quantifiers. The formulas $\mathfrak{F}$ contain the first-order statements. These are built up from the logical connectives, the first-order quantifiers, the logical constants, and the predicates. The predicates are applied to terms, connecting the two categories.

While the selection of connectives and quantifiers in first-order languages is largely consistent throughout the literature, the choices of atomic formulas and term language can differ wildly. For example, the language of ZF set theory has a term language consisting entirely of variables while Peano arithmetic requires a constant 0, the successor function S and two binary functions $+$ and $\cdot$. However, the completeness proofs we explore in this thesis are unaffected by the concrete choices made in these matters. We therefore use a syntax that is generalized over signatures Σ. A signature $(\mathcal{F}, \mathcal{P}, |-|)$ consists of a type of function symbols $\mathcal{F}$, a type of predicate symbols $\mathcal{P}$ and a function $|-|$ assigning an arity to each function- or predicate symbol. Term constants and logical constants are represented by 0-ary functions and predicates, respectively.

Definition 2.4 (**First-order syntax**)

$$t : \mathfrak{T} ::= x \mid f\, t_1 \dots t_{|f|} \qquad f : \mathcal{F}, x : \mathbb{N} \quad \textbf{\textit{terms}}$$
$$\varphi, \psi : \mathfrak{F} ::= \dot{\top} \mid \dot{\bot} \mid P\, t_1 \dots t_{|P|} \mid \varphi \dot{\rightarrow} \psi \mid \varphi \dot{\land} \psi \mid \varphi \dot{\lor} \psi \mid \dot{\forall}\varphi \mid \dot{\exists}\varphi \qquad P : \mathcal{P} \quad \textbf{\textit{formulas}}$$

The syntactic elements of the formulas are topped by a dot, making it easier to distinguish them from their meta-logical counterparts. As the syntax does not contain negation, we write $\dot{\neg}\varphi$ to denote $\varphi \dot{\rightarrow} \dot{\bot}$.

The majority of the proofs in this thesis are only concerned with a fragment $\mathfrak{F}_F$ of first-order logic. While the terms remain unchanged, the formulas are restricted to $\dot{\forall}$, $\dot{\rightarrow}$, and $\dot{\bot}$ as well as the predicates.

Definition 2.5 ($\dot{\forall}$, $\dot{\rightarrow}$, $\dot{\bot}$**-fragment**)

$$\varphi, \psi : \mathfrak{F}_F ::= \dot{\bot} \mid P\, t_1 \dots t_{|P|} \mid \varphi \dot{\rightarrow} \psi \mid \dot{\forall}\varphi \qquad P : \mathcal{P} \quad \textbf{\textit{formulas}}$$

Many definitions and lemmas concerned with the syntax of formulas do not distinguish between the concrete connectives occurring in them. To enhance readability of these kinds of statements, we employ a notation to abstract over similar syntactic elements: we write $\varphi \,\dot{\Box}\, \psi$ to denote any of the binary connectives $\dot{\to}$, $\dot{\land}$, $\dot{\lor}$. $\dot{\nabla}\varphi$ means either of the quantifiers $\dot{\forall}$ or $\dot{\exists}$. Finally, $\dot{\mp}$ denotes any of the logical constants $\dot{\bot}$ and $\dot{\top}$.

Discreteness and enumerability extend from Σ to the first-order languages. As we make use of both properties throughout this thesis, we from now on always assume that the signature Σ is discrete and enumerable. Further, an enumeration of type $\mathbb{N} \to \mathfrak{F}$ is often be more convenient than the enumeration $e : \mathbb{N} \to \mathcal{L}(\mathfrak{F})$. Such an enumeration can be defined in terms of e as $\varphi_n := \mathsf{nth}\, n\, (e\, n)$ and we make use of it throughout this thesis.

Fact 2.2

1. *If $\mathcal{F}$ and $\mathcal{P}$ are discrete, so are $\mathfrak{F}$ and $\mathfrak{F}_F$.*
2. *If $\mathcal{F}$ and $\mathcal{P}$ are enumerable, so are $\mathfrak{F}$ and $\mathfrak{F}_F$.*

2.3.2 De Bruijn Formulas and Substitutions

The attentive reader might have already noticed the peculiar lack of a binding variable for both $\dot{\forall}$ and $\dot{\exists}$. This stems from the fact that we employ de Bruijn binders instead of the more common named binders. In the late 1960s, de Bruijn developed a machine-checkable formal system for mathematics called AUTOMATH [11]. He originally put forward de Bruijn binders as a machine-friendly representation of the terms of the typed λ-calculus of AUTOMATH [12]. Because of its machine-friendliness, this paradigm is well suited for the formal treatment of syntax with binders in general. For this reason, we have chosen to employ de Bruijn binders in our formalization as well.

In de Bruijn formulas, variables are represented by natural numbers instead of the usual strings of characters. In a closed formula, these numbers denote how many quantifiers have to be skipped to arrive at their binding quantifier (see Fig. 2.2 for a visual example). If a variable exceeds the amount of quantifiers in the formula, it is taken to refer to a free variable identified with the remainder of that number. In the formula $\dot{\forall} P\, 1 \,\dot{\to}\, \dot{\forall} P\, 2$, for example, 1 and 2 refer to the same free variable identified with the number 0.

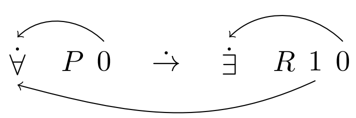

Figure 2.2: A de Bruijn formula

When using de Bruijn binders, many technical problems that usually arise when working with syntax containing binders are greatly simplified or even outright trivialized. For instance, α-equivalence, the fact that $\forall x.\ R\,x\,z$ and $\forall y.\ R\,y\,z$ are equivalent but $\forall x.\ R\,x\,z$ and $\forall x.\ R\,x\,y$ are not, requires reasoning about bound and free occurrences of variables when working with named binders. However, all members of an α-equivalence class are represented by the same de Bruijn formula. For example, both $\forall x.\ R\,x\,z$ and $\forall y.\ R\,y\,z$ are translated to $\dot\forall R\,0\,n$ for some fixed n corresponding to z. Thus reasoning about α-equivalence is reduced to reasoning about syntactic equality in a de Bruijn setting.

Another syntactic problem resolved by employing de Bruijn binders is that of a capture-free substitution. Consider the valid statement about natural numbers $\forall n.\ \exists m.\ n \neq m$. When naively instantiating this formula with the free variable m, one arrives at $\exists m.\ m \neq m$. The term that took the place of n now refers to the variable m bound by the existential quantifier instead of the free variable m, leading to an invalid statement. This phenomenon is called "variable capture" and can lead to inconsistencies in a deduction system as we just demonstrated.

The most convenient way of avoiding capture is working with a capture-free substitution, a notion of substitution that makes variable capture impossible. Defining a capture-free substitution for de Bruijn formulas turns out to be very simple. First, let us reexamine what leads to the capture in the previous example. The seemingly harmless term m that was supposed to replace n turned into a capturing formula when "moving below" $\exists m$. One way of resolving this would be to rename the binder of $\exists m$ and all of its bound variables to something different, such as $\exists x$. This method is usually called "α-renaming" and would lead m to not be captured. It is much more obvious how to handle substitutions moving below a universal quantifier in a de Bruijn setting. As stated previously, each variable "counts" the number of universal quantifiers one has to skip to arrive at their binder. Hence, when moving below a binder, every variable occurring in a substitution has to be increased by one to account for the additional binder introduced "above it". If this is done, all references made by variables in substitutions are always maintained, thereby avoiding variable capture. The formal definition of substitution based on this idea is given below.

Definition 2.6 (De Bruijn substitution)

1. *A substitution is a function $\sigma : \mathbb{N} \to \mathfrak{T}$ mapping the free variables to terms.*
2. *A term t is instantiated with a substitution σ, denoted $t[\sigma]$, as follows:*

$$x[\sigma] := \sigma\,x \qquad (f\,t_1 \ldots t_{|f|})[\sigma] := f\,(t_1[\sigma]) \ldots (t_{|f|}[\sigma])$$

3. *We denote the shifting substitution $\lambda x.\ S\,x$ by $\uparrow$. We write $\uparrow t$ instead of $t[\uparrow]$. This*

operation is lifted to the level of substitutions with $(\uparrow\sigma)\,x := \uparrow(\sigma\,x)$.

4. *A substitution* σ *can be extended by a term* t*, denoted by* (t,σ)*, as follows:*

$$(t,\sigma)\,x := \begin{cases} t & \text{if } x = 0 \\ \sigma\,y & \text{if } x = S\,y \end{cases}$$

5. *A formula* φ *is instantiated with a substitution* σ*, denoted* $\varphi[\sigma]$*, as follows:*

$$(P\,t_1\ldots t_{|P|})[\sigma] := P\,(t_1[\sigma])\ldots(t_{|P|}[\sigma]) \qquad \dot{\mp}[\sigma] := \dot{\mp}$$

$$(\varphi\,\dot{\Box}\,\psi)[\sigma] := \varphi[\sigma]\,\dot{\Box}\,\psi[\sigma] \qquad (\dot{\nabla}\varphi)[\sigma] := \dot{\nabla}(\varphi[0,\uparrow\sigma])$$

6. *We denote the identity substitution* $\lambda x.\ x$ *with* ι*. In a slight abuse of notation, we write* $s[t]$ *or* $\varphi[t]$ *for the single term substitution* (t,ι).

This definition of substitutions avoids capture, as references to free variables are always "corrected" by the lifting operation $\uparrow$ when moving below quantifiers. For an example, take the formula $\forall n.\ \exists m.\ n \neq m$ we used to illustrate variable capture. The formula can be represented in the syntax of our object logic by $\dot{\forall}\dot{\exists}\ \dot{\neg}E\,1\,0$. The de Bruijn analogue of instantiating $\exists m.\ n \neq m$ with n an application of the single term substitution $(\dot{\exists}\ \dot{\neg}E\,1\,0)[0]$. Evaluating this yields $\dot{\exists}\ \dot{\neg}E\,1\,0$ again, thereby avoiding variable capture. A detailed evaluation of this example can be seen below.

$$(\dot{\exists}\,\dot{\neg}E\,1\,0)[0] = (\dot{\exists}\,\dot{\neg}E\,1\,0)[0,\iota] = \dot{\exists}\,(\dot{\neg}E\,1\,0)[0,\uparrow(0,\iota)] = \dot{\exists}\,(\dot{\neg}E\,1\,0)[0,1,\uparrow\iota] = \dot{\exists}\,\dot{\neg}E\,1\,0$$

2.3.3 Contexts and Theories

When defining deduction systems or semantics, one often deals with collections of formulas which are assumed to already hold. Usually, these collections are guaranteed to be finite, in which case they are called contexts. We represent them as members of the type of lists of formulas $\mathcal{L}(\mathfrak{F})$ and denote them with Γ. However, we also need collections of potentially infinite size in Chapter 3. We call these theories and denote them by $\mathcal{T}$. They are represented by sets of formulas of the type $2^{\mathfrak{F}}$. We write $X \subseteq Y$ whenever $\varphi \in Y$ for all $\varphi \in X$, regardless of whether X and Y, respectively, are a context or a theory.

2.3.4 Fresh Variables

The notion of fresh variables and closed formulas has proven useful when reasoning about formulas and substitutions. A variable x is called fresh for term t or formula φ, denoted by $x\#t$ and $x\#\varphi$, if the variable x is never referred to in t and φ. A formula φ or term t is considered to be closed if no it doesn't refer to any free variables. The formal definitions of these notions is given below. Note that quantifiers again induce a variable shift to maintain the correct notion of reference.

Definition 2.7 (Fresh variables)

1. *The variable x being fresh for the term t is characterized inductively by $x\#t$:*

$$U_y^{\mathfrak{T}}\ \frac{x \neq y}{x\#y} \qquad\qquad U_f^{\mathfrak{T}}\ \frac{x\#t_1 \quad \dots \quad x\#t_{|f|}}{x\#f\, t_1 \dots t_{|f|}}$$

2. *The variable x being fresh for the formula φ is characterized inductively by $x\#\varphi$:*

$$U_{\dot{\mp}}^{\mathfrak{F}}\ \frac{}{x\#\dot{\mp}} \qquad\qquad U_{\dot{\Box}}^{\mathfrak{F}}\ \frac{x\#\varphi \quad x\#\psi}{x\#\varphi \mathbin{\dot{\Box}} \psi} \qquad\qquad U_{\dot{\nabla}}^{\mathfrak{F}}\ \frac{S\, x\#\varphi}{x\#\dot{\nabla}\varphi}$$

$$U_P^{\mathfrak{F}}\ \frac{x\#t_1 \quad \dots \quad x\#t_{|n|}}{x\#P\, t_1 \dots t_{|n|}}$$

3. *The variable x being fresh for the context Γ is defined as $x\#\Gamma := \forall \varphi \in \Gamma.\ x\#\varphi$.*

As any formula or context can only reference a finite amount of variables, one can always compute a boundary above which no variable is referenced by the formula or context.

Fact 2.3 *For any formula φ and context Γ, one can find a variable x such that any variable $y \geq x$ is fresh for φ or Γ.*

An important operation on formulas is closing them by capturing their free variables with universal quantifiers. We write $\dot{\forall}^n\, \varphi$ for the formula φ prefixed by n universal quantifiers. This operation captures all references to the free variables up to n and decreases all other references to free variables by n respectively. This is formalized by the fact below. When applied with an n found by the procedure of Fact 2.3, this formula transformation can be used to close terms. We denote this closing operation by $\dot{\forall}^*\, \varphi$.

Fact 2.4 *For any formula φ and variable x such that any variable $y \geq x$ is fresh for φ, any variable $z \geq x - n$ is fresh for $\dot{\forall}^n\, \varphi$.*

A common problem one encounters when working with substitutions is having to prove that two seemingly different substitutions σ_1 and σ_2 result in the same formula when applied to a certain formula φ. As the initial formula φ can only refer to a finite amount of free variables, it suffices to show that the two substitutions σ_1 and σ_2 agree on those variables to deduce that the resulting formulas are equal overall. This intuition is made formal by Fact 2.5 below, which is formulated in a very generic manner using a decidable predicate distinguishing between the fresh and the "relevant" variables. By specializing it to appropriate predicates, one arrives at Corollary 2.6, which gives the specific instances that we require throughout this thesis.

Fact 2.5 *Let p be a decidable predicate, φ a formula and σ_1, σ_2 substitutions, such that $\neg p\,x \to \sigma_1\,x = \sigma_2\,x$ and $p\,x \to x\#\varphi$. Then $\varphi[\sigma_1] = \varphi[\sigma_2]$.*

Corollary 2.6 *Let φ be a formula and σ_1, σ_2 be substitutions. Then*

1. *If $x\#\varphi$ and $\sigma_1\,y = \sigma_2\,y$ for all $y \neq x$ then $\varphi[\sigma_1] = \varphi[\sigma_2]$,*
2. *If all $x \leq y$ are fresh for φ and $\sigma_1\,z = \sigma_2\,z$ for all $z < x$ then $\varphi[\sigma_1] = \varphi[\sigma_2]$,*
3. *If φ is closed then $\varphi[\sigma_1] = \varphi[\sigma_2]$.*

Sometimes one has to prove that a variable is fresh for the result of a substitution instantiation. With similar reasoning as above, it suffices to show that the variable is fresh for all of the terms that the free variables referenced by the formula map to.

Fact 2.7 *Let φ be a formula, σ a substitution and y a variable such that $x\#\varphi$ or $y\#\sigma\,x$ for all x. Then $y\#\varphi[\sigma]$.*

2.4 Natural Deduction

We work with multiple different deduction systems throughout this thesis. As multiple interesting considerations go into defining these systems, we elaborate on some of the decisions we made in these matters with the example of the deduction system that is employed in Chapter 3.

$$\textsc{Ctx}\,\frac{\varphi \in \Gamma}{\Gamma \vdash \varphi} \qquad \text{II}\,\frac{\Gamma, \varphi \vdash \psi}{\Gamma \vdash \varphi \mathbin{\dot{\to}} \psi} \qquad \text{IE}\,\frac{\Gamma \vdash \varphi \mathbin{\dot{\to}} \psi \qquad \Gamma \vdash \varphi}{\Gamma \vdash \psi} \qquad \textsc{AllI}\,\frac{\uparrow\Gamma \vdash \varphi}{\Gamma \vdash \dot{\forall}\varphi}$$

$$\textsc{AllE}\,\frac{\Gamma \vdash \dot{\forall}\varphi}{\Gamma \vdash \varphi[t]} \qquad \textsc{Exp}\,\frac{\Gamma \vdash_E \dot{\bot}}{\Gamma \vdash_E \varphi} \qquad \textsc{Peirce}\,\frac{}{\Gamma \vdash_C ((\varphi \mathbin{\dot{\to}} \psi) \mathbin{\dot{\to}} \varphi) \mathbin{\dot{\to}} \varphi}$$

Figure 2.3: Natural deduction

The system we define is a variant of Gentzen's natural deduction [21, 22] on the syntactic fragment $\mathfrak{F}_F$. Its rules are given in Fig. 2.3. As we require different variants of natural deduction in the course of this thesis, we parameterize the system over two flags S and B. These parameters restrict the usage of the rules (Exp) and (Peirce). The first flag S characterizes the system's classicality: if it is C, the classical (Peirce) rule may be used, otherwise it is intuitionistic, which is denoted by an I. The second flag B describes how the system treats $\dot{\bot}$. If it is E, the system incorporates the (Exp) rule, otherwise $\dot{\bot}$ is treated as a logical constant, denoted by L. Many properties hold for any choice of S and B or only depend on the choice of one of them. To reduce visual clutter in such situations, we do not specify any of the components that can be freely chosen, such as in $\vdash_C$ and $\vdash_E$ in the (Peirce) and (Exp) rules below.

The rules of $\vdash$ are straight-forward: Any assumed formula can be deduced (Ctx). An implication holds if its conclusion can be deduced when assuming its premise (II) and can be used to deduce its conclusion whenever its premise holds (IE). A universally quantified property can be proven by demonstrating that it holds in a context in which the binding variable does not occur (AllI) and can be used to derive that property for any term t (AllE). If a contradiction $\dot\perp$ can be deduced, everything else can be deduced as well (Exp). The classical axiom (Peirce) states that any formula, which holds whenever it implies another formula, holds. While this is a somewhat uncommon method of ensuring classicality, it does not rely on $\dot\perp$ being contradictory and therefore allows for the classicality of the system and its treatment of $\dot\perp$ to be specified separately.

While $\vdash$ is defined in terms of finite contexts Γ, we also need to be able to reason about provability in terms of possibly infinite theories $\mathcal{T}$. Instead of defining a separate deduction system for theories, we define provability under theories in terms of provability under finite contexts. We extend the usage of the specifying parameters SB to this notion of provability.

Definition 2.8 (Provability under theories) *A formula φ is provable under a theory $\mathcal{T}$, written as $\mathcal{T} \vdash \varphi$, if there exists a finite context $\Gamma \subseteq \mathcal{T}$ with $\Gamma \vdash \varphi$.*

This definition captures the notion of provability under theories: Consider a natural deduction system for possibly infinite theories. It will still only admit finite derivations as valid proofs. However, a finite derivation can only make use of finitely many formulas from the theory $\mathcal{T}$ under which it was derived. This finite collection of formulas can be regarded as a context $\Gamma \subseteq \mathcal{T}$, thus yielding a proof according to the above definition.

The rules of deduction systems can usually be lifted to the level of provability under theories. We need two such lifted rules in Chapter 3.

Fact 2.8 *The following rules can be proven admissible for $\vdash$.*

$$\text{II}\ \frac{\mathcal{T} \cup \{\varphi\} \vdash \psi}{\mathcal{T} \vdash \varphi \dot\to \psi} \qquad\qquad \text{Cut}\ \frac{\mathcal{T} \cup \{\varphi\} \vdash \psi \quad \mathcal{T} \vdash \varphi}{\mathcal{T} \vdash \psi}$$

There are two kinds of weakening properties which all deduction systems discussed in this thesis posses. The first kind is weakening under assumption extension: A proof remains valid under assumptions that subsume the proof's original assumptions. The second kind is weakening under substitution: A proof remains valid if a substitution is applied to its context and claim.

Fact 2.9 *The following rules can be proven admissible for* $\vdash$.

$$\textsc{Weak}\,\frac{\Gamma' \vdash \varphi \quad \Gamma' \subseteq \Gamma}{\Gamma \vdash \varphi} \qquad \textsc{WeakT}\,\frac{\mathcal{T}' \vdash \varphi \quad \mathcal{T}' \subseteq \mathcal{T}}{\mathcal{T} \vdash \varphi} \qquad \textsc{WeakS}\,\frac{\Gamma \vdash \varphi}{\Gamma[\sigma] \vdash \varphi[\sigma]}$$

Most of the rules of $\vdash$ can be found identically throughout the literature. A noteworthy exception to this is (AllI), as there are multiple equally viable formulations of this rule for systems with de Bruijn binders. We call the variant we chose the "de Bruijn" version. It uses a context shift in the premise $\uparrow\Gamma \vdash \varphi$ to ensure that the variable originally bound by the universal quantifier does not refer to any free variables in the context. An alternative, especially popular in the programming languages community, is called "locally nameless" [2]. Here, the bound variable is replaced by a variable x in the premise $\Gamma \vdash \varphi[x]$. The variable x must be fresh for Γ and $\dot\forall\varphi$ to rule out variable capture. Generally, there is no clearly superior variant of the (AllI) rule. Our decision in this matter is motivated by a fairly technical reason: The proofs of the weakening properties in Fact 2.9 are simpler with the de Bruijn version. However, the locally nameless variant proves essential in a few other situations throughout this thesis. The following lemma allows us to freely switch between these two variants as needed.

Lemma 2.10 *Let* Γ *be a context,* φ *a formula and* x *a variable fresh for* Γ *and* $\dot\forall\varphi$*. Then*

$$\uparrow\Gamma \vdash \varphi \quad \textit{iff} \quad \Gamma \vdash \varphi[x]$$

Proof

- $\rightarrow$ Using Fact 2.9.3 with the substitution $[x]$ as $(\uparrow\Gamma)[x] = \Gamma$.
- $\leftarrow$ Using Fact 2.9.3 with the substitution $\sigma\, y := \begin{cases} 0 & \text{if } x = y \\ \uparrow y & \text{otherwise} \end{cases}$. □

A property of $\vdash$ that is crucial for the analysis is its enumerability. Both provability under arbitrary finite contexts and provability under enumerable theories can be enumerated. This proof also serves as an example how to prove enumerability of inductive predicates, such as deduction systems.

Lemma 2.11 *Let* Γ *be a context and* $\mathcal{T}$ *be an enumerable theory.*

1. *The predicate* $\lambda\varphi.\, \Gamma \vdash_{SB} \varphi$ *is enumerable*
2. *The predicate* $\lambda\varphi.\, \mathcal{T} \vdash_{SB} \varphi$ *is enumerable*

Proof We prove this by giving enumerations $e : \mathbb{N} \to \mathcal{L}(\mathfrak{F}_F)$.

1. We simultaneously define enumerations e_Γ of $\lambda\varphi.\ \Gamma \vdash_{SB} \varphi$ for each Γ simultaneously. For the sake of readability, we refer to both the enumeration of terms and formulas as e.

$$
\begin{aligned}
e_\Gamma\, 0 &:= \Gamma \\
e_\Gamma\, (S\, n) &:= e_\Gamma\, n && +\!\!+ \\
&\qquad [\, \varphi \dot{\to} \psi \mid \varphi \in e\, n, \psi \in e_{(\Gamma,\varphi)}\, n\,] && +\!\!+ \\
&\qquad [\, \psi \mid \psi \in e\, n, \varphi \in e_\Gamma\, n, (\varphi \dot{\to} \psi) \in e_\Gamma\, n\,] && +\!\!+ \\
&\qquad [\, \dot\forall \varphi \mid \varphi \in e_{(\uparrow\Gamma)}\, n\,] && +\!\!+ \\
&\qquad [\, \varphi[t] \mid \dot\forall \varphi \in e_\Gamma\, n, t \in e\, n\,] && +\!\!+ \\
&\qquad [\, ((\varphi \dot{\to} \psi) \dot{\to} \varphi) \dot{\to} \varphi \mid S = C, \varphi \in e\, n, \psi \in e\, n\,] && +\!\!+ \\
&\qquad [\, \varphi \mid B = E, \bot \in e_\Gamma\, n, \varphi \in e\, n\,]
\end{aligned}
$$

   These enumerations have one clause per rule of $\vdash_{SB}$ in the recursive case. For example, the second clause corresponds to the (IE)-rule. It generates a list of all those ψ that have been enumerated by the formula enumeration $e\, n$ and for which there is some φ such that φ and $\varphi \dot{\to} \psi$ have already been enumerated by $e_\Gamma\, n$, meaning that $\Gamma \vdash_{SB} \varphi$ and $\Gamma \vdash_{SB} \varphi \dot{\to} \psi$ have already been established.

2. We assume an enumeration $e_\mathcal{T} : \mathbb{N} \to \mathcal{L}(\mathfrak{F}_F)$ which enumerates all formulas in $\mathcal{T}$. Using $e_\mathcal{T}$, one can construct an enumeration $e' : \mathbb{N} \to \mathcal{L}(\mathcal{L}(\mathfrak{F}_F))$ which enumerates all contexts $\Gamma \subseteq \mathcal{T}$. Then the predicate $\lambda\varphi.\ \mathcal{T} \vdash_{SB} \varphi$ can be enumerated by the following enumeration:

$$
e\, 0 := [] \qquad e\, (S\, n) := e\, n +\!\!+ \mathsf{concat}\, [\, [\, \varphi \mid \varphi \in e_\Gamma\, n\,] \mid \Gamma \in e'\, n\,]
$$

□

2.5 Constructive Analysis

This thesis is concerned with constructively analyzing various first-order completeness results. More precisely, our goal is finding the non-constructive principles that are required to prove first-order completeness. This places our work into the field of constructive reverse mathematics, the inquiry into the proving principles necessary to establish noteworthy theorems. The standard techniques for such analyses have already been established: To show that a completeness theorem requires a certain principle, one proves that one can deduce the principle from the theorem. One can obtain an even tighter characterization of the relationship between the theorem and the principle by proving their equivalence.

Our analysis proceeds by relating completeness results to the stability of $\vdash_{CE}$ with regards to theories. The stability of $\vdash_{CE}$ can in turn be identified with different non-constructive proving principles, based on which classes of theories it is defined for. In this section, we give a class-based account of the stability of $\vdash_{CE}$ as well as making precise its relation to the non-constructive principles we consider in this thesis.

2.5.1 Stability

In constructive settings, such as CIC, the classical principle of double negation elimination is not generally valid. However, specific instances of it are still provable. A proposition for which the principle of double negation elimination can be proven is called stable.

Definition 2.9 (**Stability**) *A proposition P is stable if $\neg\neg P \to P$ holds.*

Stability is transported along equivalences. That is, if two propositions are equivalent, their stability is as well.

Fact 2.12 *Let P and Q be equivalent. Then P is stable if and only if Q is.*

As our analysis reveals, standard completeness with regards to different classes of theories requires different non-constructive principles to be proven. To match this insight, we define the notion of C-stability as the stability of $\vdash_{CE}$ with regards to the class of theories C. This allows for very fine-grained analyses of the completeness results.

Definition 2.10 (**C-stability**) *Let $C : 2^{\mathfrak{F}_F} \to \mathbb{P}$ be a class of theories. C-stability states that for any theory $\mathcal{T}$ for which $C\,\mathcal{T}$ holds, the proposition $\mathcal{T} \vdash_{CE} \varphi$ is stable for any choice of φ.*

For some of the analyses, we consider theory classes which are closed under certain operations.

Definition 2.11 (**Closed C-stability**) *A C-stability is closed under an operation $f : \mathfrak{F}_F \to \mathfrak{F}_F$ if for every $\mathcal{T}$ in C, $f\,\mathcal{T}$ is in C as well.*

2.5.2 Double-Negation Elimination

The strongest non-constructive principle which we consider for our analysis is that of double negation elimination (DN). When CIC is extended with this principle, it behaves almost like a traditional classical system.

Definition 2.12 (Double-Negation elimination) *The principle of double-negation elimination states that any proposition P is stable.*

The DN principle is equivalent to the stability of $\vdash_{CE}$ on all theories, which we refer to as U-stability.

Lemma 2.13 *DN and U-stability are equivalent.*

Proof DN trivially entails the stability of any $\mathcal{T} \vdash_{CE} \varphi$. We show that the inverse holds as well. Assume that $\mathcal{T} \vdash_{CE} \varphi$ is stable for any choice of $\mathcal{T}$ and φ. We have to show that any proposition P is stable. For this, consider $\mathcal{T}_P := (\lambda\varphi.\ \varphi = \dot{\bot} \land P)$, which is a theory that contains $\dot{\bot}$ if P holds and is empty otherwise. By Fact 2.12, it suffices to show that $\mathcal{T}_P \vdash_{CE} \dot{\bot}$ is equivalent to P to deduce the stability of P. Certainly, $\mathcal{T}_P = \{\dot{\bot}\}$ if P holds, which means $\mathcal{T}_P \vdash_{CE} \dot{\bot}$ can be derived. Now assume that $\mathcal{T}_P \vdash_{CE} \dot{\bot}$. Per definition, this means there is a $\Gamma \subseteq \mathcal{T}_P$ such that $\Gamma \vdash_{CE} \dot{\bot}$. We prove that P holds by case distinction on the emptiness of Γ. It cannot be the case that Γ is empty, as this would mean $\vdash_{CE} \dot{\bot}$, which is impossible by the consistency of $\vdash_{CE}$ which we prove in Lemma 3.42. Thus Γ contains at least one formula. Per definition of $\mathcal{T}_P$, this can only be $\dot{\bot}$. However, $\dot{\bot} \in \mathcal{T}_P$ already implies that P holds. □

2.5.3 Synthetic Markov's principle

CIC requires all functions which can be defined in it to be total. This is important to ensure it gives rise to a sensible notion of proof: Consider a variant of CIC which allowed for partial function definitions. As implications are functions in type theoretic foundations, one could prove $A \to B$ for any propositions A and B by simply giving a function that is undefined on all arguments. Clearly, this would make the type theory inconsistent.

The CIC ensures this totality by requiring all recursive definitions to be structurally recursive. For example, when defining a function per recursion on a natural number Sn, one may only recurse on values that are "contained within" Sn, such as n. Notably, this means that the notion of computability in the calculus of inductive constructions is slightly more restricted than that posed by the Church-Turing thesis.

A very common procedure that is impacted by this restriction on functions is unbounded search. For example, when considering some function $f : \mathbb{N} \to \mathbb{B}$ for which it is already known that $\exists n.\ fn = \mathsf{true}$, some non-trivial ideas are required to prove that a linear searching function similar to

$$g\,n := \text{if } f\,n \text{ then } n \text{ else } g\,(S\,n)$$

may be defined because the knowledge of $\exists n.\ f\,n = \mathsf{true}$ guarantees that it terminates eventually.

One variant of Markov's principle, the synthetic Markov's principle (SMP), is also concerned with such unbounded search procedures. Intuitively, it states that classical evidence $\neg\neg\exists n.\ f\,n = \mathsf{true}$ for the termination of such a search procedure is already sufficient to be allowed to perform the search in CIC. This variant is called synthetic because it directly influences the notion of computation inside CIC by allowing more functions to be definend and thus computed.

Definition 2.13 (Synthetic Markov's principle) *The synthetic Markov's principle states that for any function $f : \mathbb{N} \to \mathbb{B}$, the proposition $\exists n.\ f\,n = \mathsf{true}$ is stable.*

If CIC is not extended with additional axioms, SMP cannot be proven in it [50, 55]. This is not surprising as the principle is unconstructive: It allows for the derivation of existential proofs without giving an explicit witness. It should be noted, however, that some schools of intuitionistic thought still regard Markov's principle as intuitionistically valid.

The stability asserted by the synthetic Markov's principle can be extended to all enumerable predicates on discrete types.

Lemma 2.14 *Let X be discrete and $P : X \to \mathbb{P}$ enumerable. Under the synthetic Markov's principle, $P\,x$ is stable for any $x : X$.*

Proof Let $P : X \to \mathbb{P}$ be enumerated by $e : \mathbb{N} \to \mathcal{L}(X)$. As X is discrete, the function $f := \lambda n.\ x \in e\,n$ can be defined and has a true point if and only if $P\,x$ holds. Given $\neg\neg P\,x$, we have to deduce $P\,x$. For this, it suffices to show that f has a true point. By applying the synthetic Markov's principle, this can be deduced by showing that f does not consist solely of false points. But if f was the constant false-function, $\neg P\,x$ would follow, leading to a contradiction with the assumption $\neg\neg P\,x$. □

The synthetic Markov's principle is equivalent to E-stability, the stability of $\vdash_{CE}$ on enumerable theories. It is easy to prove that SMP entails E-stability.

Corollary 2.15 *Under the synthetic Markov's principle, the proposition $\mathcal{T} \vdash_{CE} \varphi$ is stable for any enumerable $\mathcal{T}$.*

Proof Follows by Lemma 2.14 and the enumerability of $\mathcal{T} \vdash_{CE} \varphi$ (Lemma 2.11). □

The proof that E-stability entails SMP is analogous to the proof that U-stability entails DN. However, to be able to apply E-stability, we also have to establish that the theory we construct is enumerable.

Lemma 2.16 *Under the stability of $\mathcal{T} \vdash_{CE} \varphi$ for enumerable theories $\mathcal{T}$, the proposition $\exists n.\ f\,n = \mathsf{true}$ is stable for every function $f : \mathbb{N} \to \mathbb{B}$.*

Proof Fix such a function $f : \mathbb{N} \to \mathbb{B}$. We define the theory $\mathcal{T}_f$ which contains $\dot\bot$ whenever f has a true point and is empty otherwise as $(\lambda\varphi.\ \exists n.\ f\,n = \mathsf{true} \land \varphi = \dot\bot)$.

Certainly, $\mathcal{T}_f$ is enumerable by the following enumerator $e_f : \mathbb{N} \to \mathcal{L}(\mathfrak{F}_F)$.

$$e_f\,0 := [] \qquad e_f\,(S\,n) := \text{if } f\,n \text{ then } e_f\,n +\!\!+ [\dot\bot] \text{ else } e_f\,n$$

Thus, by Fact 2.12, it suffices to show that $\mathcal{T}_f \vdash_{CE} \dot\bot$ is equivalent to $\exists n.\ f\,n = \mathsf{true}$. As $\dot\bot \in \mathcal{T}_f$ whenever $\exists n.\ f\,n = \mathsf{true}$, $\mathcal{T}_f \vdash_{CE} \dot\bot$ can be derived when knowing a true point of f. Now assume $\mathcal{T}_f \vdash_{CE} \dot\bot$. Per definition, this means there is a $\Gamma \subseteq \mathcal{T}_f$ with $\Gamma \vdash_{CE} \dot\bot$. We perform a case distinction on Γ. If $\Gamma = []$, this means that $\vdash_{CE} \dot\bot$, which is impossible because $\vdash_{CE}$ is consistent as we later prove in Lemma 3.42. Thus there is a $\varphi \in \Gamma$ and as $\Gamma \subseteq \mathcal{T}_f$, also $\varphi \in \mathcal{T}_f$, whereby $\exists n.\ f\,n = \mathsf{true}$. □

2.5.4 Object Markov's principle

The synthetic Markov's principle was defined in terms of the innate notion of computation in CIC. The computational power of the calculus of inductive constructions is not completely fixed as it may be modified by additional assumptions. For example, it is consistent to assume the existence of a function that can decide any semi-decidable problem, such as the halting problem. From the inside of CIC, one can thus not clearly reason about the limits of computability within it.

A more definite notion of computation can be obtained by formalizing a model of computation, such as the lambda calculus or Turing machines, inside of CIC. As is the case for deduction systems, such models would be called object models of computation to demarcate between it and the meta-notion of computation of CIC.

Markov's principle can also be formulated in terms of such an object model of computation. As an example, assume a variant of Church's lambda calculus formalized in CIC, that of Forster and Smolka [18]. Markov's principle for this lambda calculus could be formulated as the stability of term normalization. For Turing machines, Markov's principle could be phrased as the stability of halting.

We call Markov's principle in terms of such a object Church-Turing notion of computation the object Markov's principle. Similarly to the synthetic Markov's principle, the object Markov's principle entails the stability of every predicate that can be enumerated in the object model of computation. As such an object model of computation is much more definite than that intrinsic to CIC, its termination or normalization can be encoded into a finite first-order theory [16]. One can thus prove that the object Markov's principle is equivalent to the stability of $\vdash_{CE}$ on finite theories.

We did not formalize the object Markov's principle as the formalization of a sufficiently powerful model of computation together with a proof that its termination can be encoded as a finite first-order theory is outside the scope of this thesis. We thus simply employ F-stability, the stability of $\vdash_{CE}$ on finite theories, as a substitute for a fully formalized object Markov's principle.

Definition 2.14 (Object Markov's principle) *The object Markov's principle states that the proposition $\mathcal{T} \vdash_{CE} \varphi$ is stable for any theory for which there exists a finite context Γ with $\psi \in \Gamma$ iff $\psi \in \mathcal{T}$ for every ψ.*

The object Markov's principle is a weaker non-constructive axiom than the synthetic principle. This is because provability under finite contexts is enumerable according to Lemma 2.11 and its stability is thus entailed by synthetic Markov's principle. We strongly believe that the inverse direction does not hold as the synthetic Markov's principle implicitly depends on the computational power of CIC, which might be impossible to capture in the object model of computation. Note however, that most enumerable predicates that do not depend on the computational power of CIC can still be enumerated in any Church-Turing object model of computation.

Corollary 2.17 *The synthetic Markov's principle entails object Markov's principle.*

Proof Follows by Lemma 2.14 and the enumerability of $\Gamma \vdash \varphi$ (Lemma 2.11). $\square$

Chapter 3

Tarski Semantics

The first proof of first-order completeness was given by Gödel in 1929 [23]. Its highly technical nature has lead to it being superseded by an alternative proof given by Henkin in 1949 [28]. Henkin's proof presents a much more elegant argument via maximally consistent theories. This chapter introduces a unified framework for a constructive analysis of multiple variants of his proof. Much of this chapter is based on a similar analysis by Herbelin and Ilik [29] and Schumm's proof of the completeness for a minimal implicational predicate logic [59].

We begin this chapter by giving a brief overview of the strategy of Henkin's proof of first-order completeness in Section 3.1. We then proceed by giving a detailed account of a generalization of Henkin's theory extension construction in Section 3.2. Using this construction, we analyze the constructivity of a Henkin-style completeness proof for three different notions of Tarski models in Section 3.3. We close the chapter by demonstrating some extensions of these completeness results in Section 3.4 and proving the soundness of the natural deduction systems with regards to the semantics discussed in chapter in Section 3.5.

3.1 An Overview of Henkin's Proof

We only consider the $\dot\forall$, $\dot\rightarrow$, $\dot\bot$-fragment $\mathfrak{F}_F$ on a discrete and enumerable signature Σ. Note that this can already be considered a full completeness proof, as any connective of $\mathfrak{F}$ can be defined in terms of $\dot\forall$, $\dot\bot$ and $\dot\rightarrow$ in a classical setting.

The semantics are provided by Tarski models. An interpretation $\mathcal{I}$ on a domain D provides a predicate interpretation $P^{\mathcal{I}} : D^{|P|} \to \mathbb{P}$ for every $P : \mathcal{P}$ and a function interpretation $f^{\mathcal{I}} : D^{|f|} \to D$ for every $f : \mathcal{F}$. Each interpretation, together with an assignment $\rho : \mathbb{N} \to D$, gives rise to a term interpretation $\_^{\rho} : \mathfrak{T} \to D$. A formula being satisfied by this interpretation under an assignment ρ is defined by the translation into the meta-logic $\rho \vDash \varphi$ given below.

$$\begin{aligned} \vDash :\ (\mathbb{N} \to D) &\to \mathfrak{F}_F \to \mathbb{P} \\ \rho \vDash \dot{\bot} &:= \bot \\ \rho \vDash P\, t_1 \ldots t_{|P|} &:= P^{\mathcal{I}}\, t_1{}^{\rho} \ldots t^{\rho}_{|P|} \\ \rho \vDash \varphi \dot{\to} \psi &:= \rho \vDash \varphi \to \rho \vDash \psi \\ \rho \vDash \dot{\forall} \varphi &:= \forall d : D.\ d, \rho \vDash \varphi \end{aligned}$$

We extend this notion to contexts $\rho \vDash \Gamma$ and theories $\rho \vDash \mathcal{T}$, meaning that each formula contained in them is satisfied. A formula is called valid under a theory, written as $\mathcal{T} \vDash \varphi$, if $\rho \vDash \mathcal{T} \to \rho \vDash \varphi$ for every $\mathcal{I}$ and ρ. We say that a theory $\mathcal{T}$ has a model if there are $\mathcal{I}$ and ρ such that $\rho \vDash \mathcal{T}$ under $\mathcal{I}$.

Central to Henkin's proof is the model existence theorem. That is, the fact that every closed theory $\mathcal{T}$ which is consistent (formally $\mathcal{T} \nvdash \dot{\bot}$) has a model. Deducing completeness from it is straightforward, which we demonstrate below. Note, however, that the proof relies on the classical principle of proof by contradiction and is therefore not constructive.

Theorem 3.1 *For any closed $\mathcal{T}$ and φ, $\mathcal{T} \vDash \varphi \to \mathcal{T} \vdash \varphi$.*

Proof Let $\mathcal{T} \vDash \varphi$. Proof by contradiction: Assume $\mathcal{T} \nvdash \varphi$. This is equivalent to $\mathcal{T} \cup \{\dot{\neg}\varphi\} \nvdash \dot{\bot}$ as $\vdash$ is a classical deduction system. Then, by the model existence theorem, there exist $\mathcal{I}$ and ρ with $\rho \vDash \mathcal{T} \cup \{\dot{\neg}\varphi\}$. But $\rho \vDash \varphi$ holds as well, as $\mathcal{T} \vDash \varphi$, a contradiction. □

The remainder of this section is dedicated to proving the model existence theorem, which constitutes the majority of Henkin's original proof as well. The approach for this is to extend the initial theory $\mathcal{T}$ into a consistent theory Ω with the properties that $\varphi \dot{\to} \psi \in \Omega$ iff $\varphi \in \Omega \to \psi \in \Omega$ and $\forall t.\ \varphi[t] \in \Omega$ iff $\dot{\forall}\varphi \in \Omega$ for all formulas φ and ψ. The model $\mathcal{I}$ and ρ such that $\rho \vDash \mathcal{T}$ can then be obtained from Ω in a straightforward manner.

The construction of Ω proceeds in two steps, each helping to establish one of the properties stated above. The first step of the construction is the addition of the so called Henkin axioms to the theory $\mathcal{T}$, yielding a new theory $\mathcal{H} := \bigcup_{n \in \mathbb{N}} \mathcal{H}_n$. These formulas have the shape of the locally nameless (AllI) rule $\varphi[x] \dot{\to} \dot{\forall}\varphi$ and assert that if $\varphi[x]$ holds for a variable x which is fresh for the theory, the statement $\dot{\forall}\varphi$ holds as well. In the definition of $\mathcal{H}$ below, $\varphi_n : \mathbb{N} \to \mathfrak{F}_F$ denotes a function enumerating all formulas.

$$\mathcal{H}_0 := \mathcal{T} \qquad \mathcal{H}_{n+1} := \mathcal{H}_n \cup \{\varphi_n[x] \dot{\to} \dot{\forall}\varphi_n\} \quad \text{with } x \text{ fresh for } \mathcal{H}_n$$

These axioms assist in proving that $(\forall t.\ \varphi[t] \in \Omega) \to \dot{\forall}\varphi \in \Omega$ for all φ. The reason why the initial theory $\mathcal{T}$ is required to be closed lies with this construction, as it guarantees the unlimited supply of fresh variables that is needed for the axioms.

In the second step, we extend $\mathcal{H}$ into a maximal consistent theory. That is, a consistent theory Ω with the property that any formula that could be added to Ω without leading to an inconsistency is a member of Ω already (formally, that $\varphi \in \Omega$ for any φ with $\Omega \cup \{\varphi\} \nvdash \dot{\bot}$). The technique for creating this theory is very simple: The construction scans through all of the formulas using the enumeration $\varphi_{\_}$ and adds all those that maintain the theory's consistency, yielding a theory $\Omega := \bigcup_{n\in\mathbb{N}} \Omega_n$.

$$\Omega_0 = \mathcal{H} \qquad\qquad \Omega_{n+1} = \begin{cases} \Omega_n \cup \{\varphi_n\} & \text{if } \Omega_n \cup \{\varphi_n\} \nvdash \dot{\bot} \\ \Omega_n & \text{otherwise} \end{cases}$$

As $\mathcal{H}$ is consistent and only formulas maintaining its consistency were added, Ω is consistent overall. Ω is maximal as well: any formula φ consistent with Ω is also consistent with the partial Ω_n when $\varphi_n = \varphi$. Therefore it has already been added to Ω. The maximal consistency imbues Ω with another important property: it is closed under $\vdash$. That is, if $\Omega \vdash \varphi$ then $\varphi \in \Omega$. Using this fact, one can deduce that $\varphi \dot{\to} \psi \in \Omega$ iff $\varphi \in \Omega \to \psi \in \Omega$. By making use of the Henkin axioms added during the previous step, $\forall t.\ \varphi[t] \in \Omega$ iff $\dot{\forall}\varphi \in \Omega$ can be proven as well.

The model for the initial theory $\mathcal{T}$ is given by the identity assignment $\iota\, x := x$ and the interpretation $\mathcal{I}$ on the domain of terms $\mathfrak{T}$. Its term interpretation $-^{\rho}$ is the same as $-[\rho]$ and predicates hold under it whenever they are in Ω. This kind of model is called a syntactic model.

$$t^{\rho} = t[\rho] \qquad P^{\mathcal{I}}\, t_1 \ldots t_{|P|} := P\, t_1 \ldots t_{|P|} \in \Omega$$

One can now show that satisfaction in the model $\iota \vDash \varphi$ and membership in Ω precisely coincide. This implies that $\mathcal{T}$ is indeed satisfied by the model, as $\mathcal{T} \subseteq \Omega$. Examination of the proof of the above equivalence also reveals the reason why the initial theory $\mathcal{T}$ is required to be consistent: It involves showing that $\dot{\bot} \notin \Omega$, which is derived from the consistency of Ω, in turn stemming from the consistency of $\mathcal{T}$.

Note that this whole proof of model existence has been constructive. Hence, the only part of this proof calling for closer constructive analysis is that of completeness in Theorem 3.1. This is done in Section 3.3.

3.2 Generalized Theory Extension

In Section 3.3, we present and analyze three completeness proofs with regards to different notions of Tarski model. This section lays out a method for constructing maximal theories that is be central to each of these three proofs.

There is one significant difference to the extension outlined in Section 3.1: We do not require the deduction system to have the (Exp) rule. To allow for this, we assign the role of $\dot{\bot}$ to an arbitrary closed formula $\widetilde{\bot}$. While this change impacts many of the extension's details, its overall structure is still similar to that in Section 3.1.

In generalizing $\dot{\bot}$ to $\widetilde{\bot}$ we are also forced to adjust our notion of consistency. A generalization of consistency suitable for the proofs of this chapter is consistent extension.

Definition 3.1 (Consistent extension) *A theory $\mathcal{T}'$ is a consistent extension of a theory $\mathcal{T}$, written as $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T}'$, if $\mathcal{T} \subseteq \mathcal{T}'$ and $\mathcal{T}' \vdash_C \widetilde{\bot}$ entails $\mathcal{T} \vdash_C \widetilde{\bot}$.*

Therefore, any theory $\mathcal{T}'$ is a consistent extension of a theory $\mathcal{T}$ if it adds new formulas to $\mathcal{T}$ but does not introduce any inconsistencies with regards to $\widetilde{\bot}$ which were not already present in $\mathcal{T}$. This can be viewed as a generalization of the traditional theory extension procedure we gave in Section 3.1, which extends any consistent theory $\mathcal{T}$ into a consistent theory Ω, thus $\mathcal{T} \subseteq_{\dot{\bot}} \Omega$

The final result of the constructions we give in this section is a theory Ω with properties similar to those of Ω in Section 3.1. Namely, given a closed theory $\mathcal{T}$ and a closed formula $\widetilde{\bot}$, we construct a theory $\mathcal{T} \subseteq_{\widetilde{\bot}} \Omega$ such that

- It is closed under deduction: If $\Omega \vdash_C \varphi$ then $\varphi \in \Omega$
- Membership distributes over implication: $\varphi \dot{\rightarrow} \psi \in \Omega$ iff $\varphi \in \Omega \rightarrow \psi \in \Omega$
- Membership distributes over universal quantification: $\forall t.\ \varphi[t] \in \Omega$ iff $\dot{\forall}\varphi \in \Omega$

The steps of the theory extension procedure always take the shape of a union over a family of theories indexed by the natural numbers.

Definition 3.2 (Union) *Let $\mathcal{T}_{\_} : \mathbb{N} \rightarrow 2^{\mathfrak{F}}$ be given. Then $\bigcup_{n:\mathbb{N}} \mathcal{T}_n := \lambda\varphi.\ \exists n.\ \varphi \in \mathcal{T}_n$.*

The theory families of the unions used in the theory extension procedure are always cumulative. That is, $\mathcal{T}_n \subseteq \mathcal{T}_m$ for any $n \leq m$. This is because they represent the "limits" of iterative processes continuously adding formulas to a theory. Two properties exhibited by all such limits are crucial to the proofs of this section. First of all, anything provable under a limit is already provable under the result of a finite number of steps. Secondly, to prove that the limit constitutes a consistent extension of the starting theory, it suffices to show that each iterative step does not change the consistency.

Lemma 3.2 *Let $\mathcal{T} : \mathbb{N} \to 2^{\mathfrak{F}}$ be such a cumulative theory family.*

1. *If $\bigcup_{n:\mathbb{N}} \mathcal{T}_n \vdash \varphi$ then there is an n such that $\mathcal{T}_n \vdash \varphi$*
2. *It suffices to show that $\mathcal{T}_n \subseteq_{\widetilde{\bot}} \mathcal{T}_{S\,n}$ to deduce that $\mathcal{T}_0 \subseteq_{\widetilde{\bot}} \bigcup_{n:\mathbb{N}} \mathcal{T}_n$*

Proof

1. Assume $\bigcup_{n:\mathbb{N}} \mathcal{T}_n \vdash \varphi$. Thus there is a $\Gamma \subseteq \bigcup_{n:\mathbb{N}} \mathcal{T}_n$ with $\Gamma \vdash \varphi$. Per definition of the union operation, for each $\psi \in \Gamma$, there is a step i such that $\psi \in \mathcal{T}_i$. Let m be the maximum of those steps. Then $\Gamma \subseteq \mathcal{T}_m$ as $\mathcal{T}_\_$ is cumulative and thus $\mathcal{T}_m \vdash \varphi$.
2. Assume that $\mathcal{T}_n \subseteq_{\widetilde{\bot}} \mathcal{T}_{S\,n}$ for all n. We know that $\mathcal{T}_0 \subseteq \bigcup_{n:\mathbb{N}} \mathcal{T}_n$ per definition of the union operation. We thus assume $\bigcup_{n:\mathbb{N}} \mathcal{T}_n \vdash_{CL} \widetilde{\bot}$ to deduce $\mathcal{T}_0 \vdash_{CL} \widetilde{\bot}$. With 1., there is an n such that $\mathcal{T}_n \vdash_C \widetilde{\bot}$. From this, we can deduce $\mathcal{T}_0 \vdash_C \widetilde{\bot}$ inductively by using our initial assumption. □

3.2.1 Exploding Theories

The generalized theory extension requires one more construction step than its usual presentation. We already stated that $\widetilde{\bot}$ is supposed to take the place of $\dot{\bot}$. However, $\widetilde{\bot}$ currently lacks the explosion principle, which is required throughout in later parts of the proof. The first construction step thus turns the input theory into an exploding theory.

Definition 3.3 (Exploding theory) *We call a theory $\mathcal{T}$ exploding, if $(\widetilde{\bot} \dot{\to} \dot{\forall}^* \varphi) \in \mathcal{T}$ for all φ.*

Thus, a theory is exploding if it contains all closed instances of the explosion principle. This means one can treat $\widetilde{\bot}$ as a normal $\bot$ when proving something under an exploding theory, as one can take the needed instances of the explosion principle as part of the starting context Γ. It is worth noting that the closing of the conclusion of the explosion axioms via $\dot{\forall}^*$ is what allows the instances to be used even after an application of the (AllI) rule. This is because the axioms exhibit the property $\uparrow(\widetilde{\bot} \dot{\to} \dot{\forall}^* \varphi) = \widetilde{\bot} \dot{\to} \dot{\forall}^* \varphi$.

The construction used to extend the initial theory $\mathcal{T}$ into an exploding theory is based on a similar construction found in Schumm's proof of completeness for a minimal implicational predicate logic [59]. For each formula φ, the explosion axiom $\widetilde{\bot} \dot{\to} \dot{\forall}^* \varphi$ is added to the theory. An important difference between his construction and ours is that we additionally close all of the conclusions of the explosion axioms as Construction 3.2 requires its input theory to be closed.

Construction 3.1 (**Exploding $\mathcal{E}_\mathcal{T}$**) *Given a theory $\mathcal{T}$ and a formula $\widetilde{\bot}$, we construct*

$$\mathcal{E}_0 := \mathcal{T} \qquad \mathcal{E}_{S\,n} := \mathcal{E}_n \cup \{\widetilde{\bot} \dot\to \dot\forall^* (\varphi_n)\} \qquad \mathcal{E}_\mathcal{T} := \bigcup_{n:\mathbb{N}} \mathcal{E}_n$$

The constructed $\mathcal{E}_\mathcal{T}$ is exploding and remains closed, which is crucial as we want to use it as the input for Construction 3.2, which only accepts closed inputs.

Fact 3.3 *Let $\mathcal{T}$ be a theory and $\widetilde{\bot}$ a formula. Then $\mathcal{E}_\mathcal{T}$ is exploding and closed whenever $\mathcal{T}$ is closed.*

The construction $\mathcal{E}_\mathcal{T}$ constitutes a consistent extension of $\mathcal{T}$. This stems from the fact that the explosion axioms can only be used if $\widetilde{\bot}$ has already been proven, which means provability of $\widetilde{\bot}$ itself is not affected by their addition.

Lemma 3.4 *Let $\mathcal{T}$ be a theory and $\widetilde{\bot}$ a formula. Then $\mathcal{E}_\mathcal{T}$ is a consistent extension of $\mathcal{T}$.*

Proof We prove this using Lemma 3.2.2. Hence we only have to show $\mathcal{E}_n \subseteq_{\widetilde{\bot}} \mathcal{E}_{S\,n}$ for all n. Assume $\mathcal{E}_{S\,n} \vdash_C \widetilde{\bot}$. As $\mathcal{E}_{S\,n} := \mathcal{E}_n \cup \{\widetilde{\bot} \dot\to \dot\forall^* (\varphi_n)\}$, there exists a $\Gamma \subseteq \mathcal{E}_n$ with $\Gamma \vdash_C \dot\neg(\widetilde{\bot} \dot\to \dot\forall^* (\varphi_n))$ by (II). With (Peirce), we can conclude $\Gamma \vdash_C \widetilde{\bot}$ and hence $\mathcal{E}_n \vdash_C \widetilde{\bot}$. □

As $\widetilde{\bot}$ now serves its purpose as a replacement for $\dot\bot$, a few well known proving principles and connectives which originally relied on a $\dot\bot$ can be recovered. We from now on write $\widetilde{\neg}\varphi$ for $\varphi \dot\to \widetilde{\bot}$ and $\widetilde{\exists}\varphi$ for $\widetilde{\neg}\dot\forall\widetilde{\neg}\varphi$. The principles are listed below. Notably, the first principle does not not require a exploding $\widetilde{\bot}$, however it guarantees that the explosion axioms work as characterized in the second principle. As explicitly stating exactly which concrete explosion axioms are used in a proof quickly gets out of hand, we from now on only point out when explosion axioms are required, leaving the specific selection to the formalization.

Fact 3.5 *Let Γ be a context with $\widetilde{\bot} \dot\to \dot\forall^* \varphi \in \Gamma$. The following rules can be shown.*

$$\text{GClose}\ \frac{\Gamma \vdash \dot\forall^* \varphi}{\Gamma \vdash \varphi} \qquad \text{GExp}\ \frac{\Gamma \vdash \widetilde{\bot}}{\Gamma \vdash \varphi} \qquad \text{GDN}\ \frac{\Gamma, \widetilde{\neg}\varphi \vdash \widetilde{\bot}}{\Gamma \vdash \varphi}$$

$$\text{GExE}\ \frac{\Gamma \vdash_C \widetilde{\exists}\psi \qquad \uparrow\Gamma, \psi \vdash_C \uparrow\varphi}{\Gamma \vdash_C \varphi} \qquad \text{GExI}\ \frac{\Gamma \vdash \psi[t]}{\Gamma \vdash \widetilde{\exists}\psi}$$

$$\text{GXM}\ \frac{\Gamma, \varphi \vdash_C \psi \qquad \Gamma, \widetilde{\neg}\varphi \vdash_C \psi}{\Gamma \vdash_C \psi}$$

3.2.2 Henkin Theories

The second step of the generalized theory extension procedure corresponds to the first step of the procedure from Section 3.1. It extends $\mathcal{E}$ into a Henkin theory.

Definition 3.4 (Henkin theory) *We call a theory $\mathcal{T}$ Henkin if from $(\forall t.\ \mathcal{T}' \vdash \varphi[t])$ one can deduce $\mathcal{T}' \vdash \dot\forall\varphi$ for any $\mathcal{T} \subseteq \mathcal{T}'$ and φ.*

The necessity of Henkin theories can only be explained in a satisfactory manner once maximal theories have been introduced. We therefore postpone this discussion until Definition 3.5. We also note that the theory $\mathcal{T}'$ in Definition 3.4 is required to ensure that being a Henkin theory is monotonous. That is, for all theories $\mathcal{T} \subseteq \mathcal{T}'$, $\mathcal{T}'$ is Henkin whenever $\mathcal{T}$ is.

The rendition of the Henkin construction we present here stems from Herbelin and Ilik [29]. For each formula φ, the Henkin axiom $\varphi[n] \dot\to \dot\forall\varphi$ for some variable n is added to the theory. Intuitively, the Henkin axioms state that if $\varphi[x]$ can be proven for a variable x which is fresh for the theory, $\dot\forall\varphi$ can be derived as well. Note that in the definition below, we simply use n as the fresh variable x as we assume the initial theory $\mathcal{T}$ to be closed and $\varphi_\_$ to be constructed in such a way, that φ_n only refers to variables strictly smaller than n.

Construction 3.2 (Henkin $\mathcal{H}_\mathcal{T}$) *Given a theory $\mathcal{T}$, we construct*

$$\mathcal{H}_0 := \mathcal{T} \qquad \mathcal{H}_{S\,n} := \mathcal{H}_n \cup \{\varphi_n[n] \dot\to \dot\forall\varphi_n\} \qquad \mathcal{H}_\mathcal{T} := \bigcup_{n:\mathbb{N}} \mathcal{H}_n$$

Although their name would lead one to believe otherwise, this rendition of the Henkin axioms originates from Herbelin and Ilik, not Henkin. We choose to maintain their terminology by referring to them as Henkin axioms. While Henkin employs a similar construction, the axioms he uses would correspond more closely to $(\widetilde\exists\varphi_n) \dot\to \varphi_n[n]$ in our setting. That is, instead of designating witnesses for universal statements, he generates concrete witnesses for each valid existential statement. While there is very little difference between the two choices of axioms, Herbelin's variant leads to a slightly more elegant proof of that $\mathcal{H}_\mathcal{T}$ is indeed Henkin.

Lemma 3.6 *$\mathcal{H}_\mathcal{T}$ is Henkin for any $\mathcal{T}$.*

Proof Let $\mathcal{H}_\mathcal{T} \subseteq \mathcal{T}'$ and $\forall t.\mathcal{T}' \vdash \varphi[t]$. As $\varphi_\_$ is an enumeration, there is an n with $\varphi_n = \varphi$. Then $\varphi[n] \dot\to \dot\forall\varphi \in \mathcal{H}_\mathcal{T} \subseteq \mathcal{T}'$. Thus $\mathcal{T}' \vdash \dot\forall\varphi$ can be deduced. ☐

Additionally, $\mathcal{H}_\mathcal{T}$ is a consistent extension of the input theory $\mathcal{T}$. Recall the locally nameless variant of the (AllI) rule.

$$\text{AllI}'\;\frac{\Gamma \vdash \varphi[x] \quad x\#\Gamma, \dot\forall\varphi}{\Gamma \vdash \dot\forall\varphi}$$

The Henkin axioms $\varphi_n[n] \dot{\rightarrow} \dot{\forall}\varphi_n$ can be seen as the (AllI′) rule lifted to the level of theories. As the (AllI′) rule is admissible in a finite setting, it should not be surprising that adding Henkin axioms to a closed theory does not introduce any new inconsistencies either. While this constitutes an intuitive explanation of why $\mathcal{H}_\mathcal{T}$ is a consistent extension of $\mathcal{T}$, the proof is of a very different, more technical, nature. However, it still takes advantage of the correspondence between the de Bruijn and locally nameless approach as characterized in Lemma 2.10.

Lemma 3.7 *Let $\mathcal{T}$ be a closed, exploding theory. Then $\mathcal{H}_\mathcal{T}$ is a consistent extension of $\mathcal{T}$.*

Proof We prove this using Lemma 3.2.2. Assume $\mathcal{H}_{S\,n} \vdash_C \widetilde{\bot}$. By Fact 2.8.1, this implies that $\mathcal{H}_n \vdash_C \widetilde{\neg}(\varphi_n[n] \dot{\rightarrow} \dot{\forall}\varphi_n)$. Then there exists a context $\Gamma \subseteq \mathcal{H}_n$ such that $\Gamma \vdash_C \widetilde{\neg}(\varphi_n[n] \dot{\rightarrow} \dot{\forall}\varphi_n)$. When extending Γ into Γ' by adding suitable explosion axioms and noting that $\widetilde{\neg}(\varphi_n[n] \dot{\rightarrow} \dot{\forall}\varphi_n) = (\widetilde{\neg}(\varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n)))[n]$, we can deduce:

$$\widetilde{\exists}\text{E}\;\dfrac{\text{DP}\;\dfrac{}{\Gamma' \vdash_C \widetilde{\exists}\varphi_n \dot{\rightarrow} \uparrow(\dot{\forall}\varphi_n)} \qquad \dfrac{\text{Ctx}\;\dfrac{}{\uparrow\Gamma, \psi \vdash_C \psi} \qquad \dfrac{\dfrac{\Gamma \vdash_C (\widetilde{\neg}(\varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n)))[n]}{\uparrow\Gamma \vdash_C \widetilde{\neg}(\varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n))}\;\text{Lemma 2.10}}{\uparrow\Gamma', \varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n) \vdash_C \widetilde{\neg}(\varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n))}\;\text{Weak}}{\uparrow\Gamma', \varphi_n \dot{\rightarrow} \uparrow(\forall\varphi_n) \vdash_C \widetilde{\bot}}\;\text{IE}}{\Gamma' \vdash_C \widetilde{\bot}}$$

This proof makes use of a crucial property of the formula enumeration φ_n: One can define it in such a way that for $n \leq x$, x is fresh for φ_n. As the proof of this is fairly technical, we have chosen to not explicitly include it in this thesis. However, it has been formalized in Coq. This property allows for the use of Lemma 2.10 as n is still fresh for Γ and φ_n. Additionally, the formula $\Gamma' \vdash_C \widetilde{\exists}\varphi_n \dot{\rightarrow} \uparrow(\dot{\forall}\varphi_n)$, which is being eliminated at the bottom of the proof tree, is the well known classical tautology of the drinker's paradox. ☐

3.2.3 Maximal Theories

The final step of the extension process completes $\mathcal{H}$ into a maximal theory. Maximal theories are those which contain all formulas that do not alter their consistency with regards to $\widetilde{\bot}$.

Definition 3.5 (Maximal theories) *We call a theory $\mathcal{T}$ maximal, if $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$ implies that $\varphi \in \mathcal{T}$.*

Maximal theories exhibit one crucial property: they are closed under deduction. That is, maximal theories contain exactly those formulas φ that can be deduced under them. We make use of this fact to establish that membership in the theory resulting from the extension procedure does indeed distribute over implication and universal quantification, as we claimed at the beginning of this section.

Lemma 3.8 *Let $\mathcal{T}$ be maximal. Then $\varphi \in \mathcal{T}$ iff $\mathcal{T} \vdash_C \varphi$.*

Proof We only consider the $\leftarrow$ direction as the inverse is trivial. Let $\mathcal{T} \vdash_C \varphi$. We deduce $\varphi \in \mathcal{T}$ from the maximality of $\mathcal{T}$. Assume $\mathcal{T} \cup \{\varphi\} \vdash_C \widetilde{\bot}$. Then $\mathcal{T} \vdash_C \widetilde{\bot}$ follows via (CUT). □

Having defined maximal theories, we now explain the necessity of the $\mathcal{H}_\mathcal{T}$ phase. Recall that we require the membership of Ω, the result of the extension procedure, to distribute over universal quantification. That is, that $\dot\forall\varphi \in \Omega$ whenever $\varphi[t] \in \Omega$ for all t. This property cannot be concluded by virtue of the maximality of Ω alone, as there exist maximal theories whose membership does not distribute over universal quantification. Indeed, $\varphi[t] \in \Omega$ for every term t and $\dot\neg\dot\forall\varphi \in \Omega$ do not give rise to an inconsistency. Intuitively, this is because not all of the members of a domain described by a formula have to correspond to a term. As a slightly informal example, take φ to be the property "is described by a term".

Careful analysis of Construction 3.3 could rule out the possibility of such contrary theories being constructed. However, establishing this fact would require an exceeding amount of effort and a very deliberate choice of φ_-. Compared to that, the $\mathcal{H}_\mathcal{T}$ construction is a much simpler way of ensuring the desired outcome by essentially "stacking the theory in our favor" without affecting its consistency.

The following maximality construction is based on the consistent union operation $\overset{?}{\cup}$ which is defined below. This is a rendition of the constructively dubious condition "if $\Omega_n \cup \{\varphi_n\} \nvdash \dot\bot$" from Section 3.1 into constructive type theory. The crucial idea behind its constructivity is that it places the onus of asserting that $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$ on anyone trying to prove that $\varphi \in \mathcal{T} \overset{?}{\cup} \{\varphi\}$ instead of on the construction process itself.

Definition 3.6 (Consistent union) *Let $\mathcal{T}$ be a theory and φ be a formula. We define consistent union as:*

$$\mathcal{T} \overset{?}{\cup} \{\varphi\} := \lambda\psi.\ \psi \in \mathcal{T} \vee (\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\} \wedge \psi = \varphi)$$

Ω can now easily be defined in terms of the consistent union. Note that as opposed to the previous two phases, the result of this constructions is dependent on the order in which the formulas are enumerated by φ_-, with different choices of φ_- leading to different Ω.

Construction 3.3 (Maximal $\Omega_\mathcal{T}$) *Let $\mathcal{T}$ be a theory. We define*

$$\Omega_0 := \mathcal{T} \qquad \Omega_{S\,n} := \Omega_n \overset{?}{\cup} \{\varphi_n\} \qquad \Omega_\mathcal{T} := \bigcup_{n:\mathbb{N}} \Omega_n$$

We begin by proving that $\Omega_{\mathcal{T}}$ constitutes a consistent extension of $\mathcal{T}$ as this is required in the proof of its maximality. Differing from the previous consistency proofs, this deduction relies on a property of the consistent union instead of the shape of the formulas added by the construction.

Lemma 3.9 *Let $\mathcal{T}$ be a theory.*

1. *If $\mathcal{T} \mathbin{\circlearrowright\!\!\!\!\cup} \{\varphi\} \vdash_C \psi$ then $\mathcal{T} \vdash \psi$ or $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$.*
2. *$\Omega_{\mathcal{T}}$ constitutes a consistent extension of $\mathcal{T}$.*

Proof

1. Let $\Gamma \subseteq \mathcal{T} \mathbin{\circlearrowright\!\!\!\!\cup} \{\varphi\}$ with $\Gamma \vdash_C \psi$. We prove $\Gamma \subseteq \mathcal{T}$ or $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$ per induction on Γ. The case $\Gamma = []$ is resolved by $[] \subseteq \mathcal{T}$. Hence let $\Gamma = \Gamma', \theta$. If the inductive hypothesis yields $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$, we are done. Hence let $\Gamma' \subseteq \mathcal{T}$. As $\theta \in \mathcal{T} \mathbin{\circlearrowright\!\!\!\!\cup} \{\varphi\}$, either $\theta \in \mathcal{T}$ or $\mathcal{T} \subseteq_{\widetilde{\bot}} \mathcal{T} \cup \{\varphi\}$, the claim follows.
2. Proof using Lemma 3.2.2. Hence we have to deduce $\Omega_n \vdash_C \widetilde{\bot}$ from $\Omega_{S\,n} \vdash_C \widetilde{\bot}$. Then the claim follows in either case of 1. □

The fact that $\Omega_{\mathcal{T}}$ constitutes a consistent extension of $\mathcal{T}$ directly implies its maximality.

Lemma 3.10 *$\Omega_{\mathcal{T}}$ is a maximal theory.*

Proof Let φ be a formula such that $\Omega_{\mathcal{T}} \subseteq_{\widetilde{\bot}} \Omega_{\mathcal{T}} \cup \{\varphi\}$. There exists an n with $\varphi_n = \varphi$. It suffices to show that $\varphi \in \Omega_{S\,n} \subseteq \Omega_{\mathcal{T}}$. Per definition, this is the case whenever $\Omega_n \subseteq_{\widetilde{\bot}} \Omega_n \cup \{\varphi\}$. Hence, let $\Omega_n \cup \{\varphi\} \vdash_C \widetilde{\bot}$. As $\mathcal{T} \subseteq \Omega_n$, it suffices to show that $\mathcal{T} \vdash_C \widetilde{\bot}$ per (WeakT). As $\mathcal{T} \subseteq_{\widetilde{\bot}} \Omega_{\mathcal{T}} \subseteq_{\widetilde{\bot}} \Omega_{\mathcal{T}} \cup \{\varphi\}$, we can show $\Omega_{\mathcal{T}} \cup \{\varphi\} \vdash_C \widetilde{\bot}$ to prove this. This follows from $\Omega_n \cup \{\varphi\} \vdash_C \widetilde{\bot}$ using (WeakT). □

The maximality of $\Omega_{\mathcal{T}}$ and Lemma 3.8 allows giving a precise characterization of the conditions for formulas of the shapes $\dot\forall\varphi$ and $\varphi \mathbin{\dot\to} \psi$ to be members of $\Omega_{\mathcal{T}}$. These are crucial for proving the correctness of the syntactic models we define in Section 3.3.

Lemma 3.11 *If $\mathcal{T}$ is Henkin, then $\forall t.\ \varphi[t] \in \Omega_{\mathcal{T}}$ if and only if $\dot\forall\varphi \in \Omega_{\mathcal{T}}$.*

Proof

$\to$ This follows from Lemma 3.8 and the Henkin property of $\mathcal{T}$.

$\leftarrow$ By Lemma 3.8, it suffices to show $\Omega_{\mathcal{T}} \vdash_C \varphi[t]$. This follows via (AllE). □

Lemma 3.12 *For any φ and ψ, $\varphi \dot{\to} \psi \in \Omega_{\mathcal{T}}$ if and only if $\varphi \in \Omega_{\mathcal{T}} \to \psi \in \Omega_{\mathcal{T}}$.*

Proof

$\to$ This follows from Lemma 3.8 and (IE).

$\leftarrow$ We prove this using the maximality of $\Omega_{\mathcal{T}}$. By (IE) we can assume a $\Gamma \subseteq \Omega_{\mathcal{T}}$ with $\Gamma \vdash_C \tilde{\neg}(\varphi \dot{\to} \psi)$. This allows us to deduce $\Omega_{\mathcal{T}} \vdash_C \tilde{\bot}$ by proving $\Gamma \vdash_C \varphi$ as this implies $\Omega_{\mathcal{T}} \vdash_C \psi$ per assumption. This is proven as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\Gamma, \tilde{\neg}\varphi, \varphi \vdash_C \tilde{\bot}}{\Gamma, \tilde{\neg}\varphi, \varphi \vdash_C \psi}\,\text{Exp}}{\Gamma, \tilde{\neg}\varphi \vdash_C \varphi \dot{\to} \psi}\,\text{II} \qquad \dfrac{\Gamma \vdash_C \tilde{\neg}(\varphi \dot{\to} \psi)}{\Gamma, \tilde{\neg}\varphi \vdash_C \tilde{\neg}(\varphi \dot{\to} \psi)}\,\text{Weak}}{\dfrac{\Gamma, \tilde{\neg}\varphi \vdash_C \tilde{\bot}}{\Gamma \vdash_C \varphi}\,\text{DN}}\,\text{IE}$$

3.2.4 Summary

Having established these characterizations, we can deduce the full generalized theory extension theorem. It follows by combining all of the constructions and lemmas of this section.

Theorem 3.13 *Let $\mathcal{T}$ be a closed theory, $\tilde{\bot}$ a closed formula. Then one can construct a theory Ω such that*

1. $\mathcal{T} \subseteq_{\tilde{\bot}} \Omega$
2. *$\varphi \in \Omega$ for all φ with $\Omega \vdash_C \varphi$*
3. *$\forall t.\ \varphi[t] \in \Omega$ if and only if $\dot{\forall}\varphi \in \Omega$ for all φ*
4. *$\varphi \dot{\to} \psi \in \Omega$ if and only if $\varphi \in \Omega \to \psi \in \Omega$ for all φ, ψ*

3.3 Constructive Analysis of Completeness Theorems

This section encompasses the constructive analysis of the completeness theorems. After laying out the variant of Tarski semantics we use, we begin by showing that completeness with regards to standard models is equivalent to stability of provability, from which we derive equivalences between different completeness theorems and non-constructive principles from Section 2.5. We then present two alternative notions of model which allow for constructive proofs of completeness.

3.3.1 Tarski Models

We first make precise the notion of interpretation that is necessary to define any Tarski semantics. An interpretation on some domain D fixes a way of translating the terms into values of D and describes the truth conditions for first-order formulas in that setting.

Definition 3.7 (Generalized interpretation) *For a given signature Σ, a generalized interpretation $\mathcal{I}$ on a domain D consists of a function interpretation $f^{\mathcal{I}} : D^{|f|} \to D$ for every function symbol $f : \mathcal{F}$, a predicate interpretation $P^{\mathcal{I}} : D^{|P|} \to \mathbb{P}$ for every predicate symbol $P : \mathcal{P}$ and a $\dot{\bot}$ interpretation $\bot^{\mathcal{I}} : \mathbb{P}$.*

An interpretation $\mathcal{I}$, together with an assignment $\rho : \mathbb{N} \to D$, gives rise to a term interpretation $-^{\rho} : \mathfrak{T} \to D$.

$$x^{\rho} := \rho\, x \qquad (f\, t_1 \ldots t_{|f|})^{\rho} := f^{\mathcal{I}}\, t_1^{\rho} \ldots t_{|f|}^{\rho}$$

The conditions of a formula φ being satisfied by an interpretation $\mathcal{I}$ under an assignment ρ are given by the recursively defined translation $\rho \vDash \varphi$.

$$\begin{aligned} \rho \vDash \dot{\bot} &= \bot^{\mathcal{I}} \\ \rho \vDash P\, t_1 \ldots t_{|P|} &= P^{\mathcal{I}}\, t_1^{\rho} \ldots t_{|P|}^{\rho} \\ \rho \vDash \varphi \dot{\to} \psi &= \rho \vDash \varphi \to \rho \vDash \psi \\ \rho \vDash \dot{\forall} \varphi &= \forall d : D.\ d, \rho \vDash \varphi \end{aligned}$$

This definition of interpretations differs significantly from that sketched in Section 3.1 and those usually found throughout the literature. Instead of $\dot{\bot}$ always being interpreted as $\bot$, each interpretation can choose its own $\bot$-interpretation $\bot^{\mathcal{I}}$. This means validity defined in terms of all interpretations is different from the usual notion of validity, as interpretations in which $\bot^{\mathcal{I}}$ holds are considered as well. To offset this, we define a constrained of notion validity, which only considers interpretations satisfying a given constraint predicate. This provides a unified framework for reasoning about various notions of validity.

Definition 3.8 (Constrained validity)

1. *A theory $\mathcal{T}$ or context Γ is said to be satisfied by an interpretation $\mathcal{I}$ under an assignment ρ, written as $\rho \vDash \mathcal{T}$ or $\rho \vDash \Gamma$ respectively, if $\forall \varphi \in \mathcal{T}.\ \rho \vDash \varphi$ or $\forall \varphi \in \Gamma.\ \rho \vDash \varphi$.*
2. *A formula φ is said to be valid under a theory $\mathcal{T}$ or context Γ and under a constraint $X : \mathcal{I} \to \mathbb{P}$ if $\rho \vDash \mathcal{T} \to \rho \vDash \varphi$ or $\rho \vDash \Gamma \to \rho \vDash \varphi$ for any interpretation $\mathcal{I}$ satisfying X and assignment ρ. This is written as $\mathcal{T} \vDash_X \varphi$ or $\Gamma \vDash_X \varphi$ respectively.*

This granular definition of validity lends itself well to comparing different notions of validity: Oftentimes, the relationship between two notions of validity can be arrived at by analyzing their characterizing constraints. One inter-constraint relationship of special significance is constraint subsumption. We say a constraint X subsumes another constraint Y, if every interpretation satisfying Y also satisfies X. Whenever a constraint is subsumed by another constraint, their associated validities subsume each other as well.

Lemma 3.14 *Let X subsume Y. Then $\mathcal{T} \vDash_X \varphi \to \mathcal{T} \vDash_Y \varphi$ for any $\mathcal{T}$ and φ.*

Proof Assume $\mathcal{T} \vDash_X \varphi$. Given an interpretation $\mathcal{I}$ satisfying Y and an assignment ρ with $\rho \vDash \mathcal{T}$, we have to show $\rho \vDash \varphi$. This follows per assumption, as $\mathcal{I}$ also satisfies X because Y subsumes Y. □

3.3.2 Standard Models

We begin with a constructive analysis of the completeness proof outlined in Section 3.1. As noted earlier, unconstrained validity differs from the notion of validity employed in that proof. Therefore, we use validity constrained to standard models as defined below, which we denote by $\vDash_S$.

Definition 3.9 (Standard models)

1. *An interpretation $\mathcal{I}$ is said to be classical if $\rho \vDash (((\varphi \dot\to \psi) \dot\to \varphi) \dot\to \varphi)$ holds for all formulas φ and ψ and assignments ρ.*
2. *An interpretation $\mathcal{I}$ is said to have a standard $\dot\bot$ if $\dot\bot^{\mathcal{I}}$ entails $\bot$.*
3. *An interpretation is said to constitute a standard model if it is classical and has a standard $\dot\bot$.*

By requiring standard models to have a standard $\dot\bot$, the standard notion of validity $\vDash_S$ is equivalent to the more common definition of validity that always interprets $\dot\bot$ as $\bot$, such as the one from Section 3.1.

The restriction to classical interpretations in our definition of standard models is a very important detail. As it stands, validity translates from a classical object logic into a constructive meta-logic. That means that soundness, the fact that any formula that can be deduced is indeed valid, cannot be established if the notion of validity simply considers all interpretations, as the classical reasoning of the object logic will not translate into some of them. However, this problem can be resolved by restricting the interpretations considered for validity to those which allow for classical reasoning. Note that this restriction is not required for the proof of completeness. However, it is desirable that the deduction systems we consider are both sound and complete with regards to the semantics discussed in this chapter, which is why we chose to implement this restriction. We give a unified proof of soundness with regards to all notions of model examined in this chapter in Section 3.5.

The first proof of completeness we analyze is for the deduction system $\vdash_{CE}$. That is, natural deduction with both the classical principle of (PEIRCE) and the standard notion of $\bot$ expressed by (EXP). These two properties combined give rise to the classical principle of double negation elimination $\Gamma \vdash_{CE} \dot\neg\dot\neg\varphi \dot\to \varphi$, which imbues the system with the property of refutation completeness. This property is crucial in the proof of completeness.

Fact 3.15 *It suffices to show $\mathcal{T} \cup \{\dot\neg\varphi\} \vdash_{CE} \dot\bot$ to prove $\mathcal{T} \vdash_{CE} \varphi$.*

We begin our analysis by proving the model existence theorem. This property is classically equivalent to the traditional statement of completeness. Model existence can be proven constructively using the generalized theory extension from the previous section. It states that every closed, consistent theory $\mathcal{T}$ has a model. That is, a standard model $\mathcal{I}$ and a substitution ρ such that $\rho \vDash \mathcal{T}$ under $\mathcal{I}$. The model we construct is a syntactic model. That is, a model based on an interpretation whose domain are the terms $\mathfrak{T}$ of first-order logic.

Definition 3.10 (Syntactic interpretation) *Let $\mathcal{T}$ be a closed theory. Further, let Ω be the theory constructed using Theorem 3.13 with $\widetilde{\bot} = \dot\bot$. The syntactic interpretation of $\mathcal{T}$ denoted $\mathcal{I}_\mathcal{T}$, is given by:*

$$f^{\mathcal{I}_\mathcal{T}}\, t_1 \dots t_{|f|} := f\, t_1 \dots t_{|f|} \qquad P^{\mathcal{I}_\mathcal{T}} t_1 \dots t_{|P|} := P\, t_1 \dots t_{|P|} \in \Omega \qquad \bot^{\mathcal{I}_\mathcal{T}} := \dot\bot \in \Omega$$

The syntactic model has the property the formula satisfaction in $\mathcal{I}_\mathcal{T}$ for consistent $\mathcal{T}$ exactly coincides with membership of Ω. This fact relies on the fact that the assignments in $\mathcal{I}_\mathcal{T}$ are substitutions and the distributivity of membership in Ω.

Lemma 3.16 *Let $\mathcal{T}$ be closed. Then $\sigma \vDash \varphi$ iff $\varphi[\sigma] \in \Omega$ under $\mathcal{I}_\mathcal{T}$ for any substitution σ.*

Proof By induction on φ. We distinguish four cases:

$P\, t \dots t_{|P|}$: $P\, t_1^\sigma \dots t_{|P|}^\sigma \in \Omega$ is equivalent to $(P\, t_1 \dots t_{|P|})[\sigma] \in \Omega$ as $s^\sigma = s[\sigma]$ for all s.

$\varphi \dot\to \psi$: Using the inductive hypothesis, this is exactly Theorem 3.13.3.

$\dot\forall\varphi$: Note that $\forall t.\ t, \sigma \vDash_S \varphi$ is equivalent to $\forall t.\ \sigma \vDash_S \varphi[t]$. Then, using the inductive hypothesis, this is exactly Theorem 3.13.4.

$\dot\bot$: We have to prove $\dot\bot \in \Omega$ iff $\dot\bot \in \Omega$. This holds trivially. □

This result enables us to establish the model existence theorem for standard models.

Theorem 3.17 *Let $\mathcal{T}$ be closed and consistent. Then there is a standard model $\mathcal{I}$ and an assignment ρ such that $\rho \vDash_S \mathcal{T}$.*

Proof We choose $\mathcal{I}_\mathcal{T}$ as $\mathcal{I}$ and the identity assignment $\iota\, x := x$ as ρ. It remains to show that $\mathcal{I}_\mathcal{T}$ is indeed a standard model and $\iota \vDash_S \mathcal{T}$. The latter directly follows from Lemma 3.16 as $\mathcal{T} \subseteq \Omega$ by Theorem 3.13.1. To show that $\mathcal{I}_\mathcal{T}$ is a standard model, we have to prove two properties.

- $\mathcal{I}_\mathcal{T}$ is classical. That is, $\rho \vDash (((\varphi \dot{\rightarrow} \psi) \dot{\rightarrow} \varphi) \dot{\rightarrow} \varphi)$ for any ρ, φ and ψ. This again follows by Lemma 3.16 and Theorem 3.13.2 as (Peirce) is an axiom of $\vdash_{CE}$.
- $\mathcal{I}_\mathcal{T}$ has a standard $\dot{\bot}$. We have to prove that $\dot{\bot} \in \Omega$ leads to a contradiction. Certainly, it implies $\Omega \vdash_{CE} \dot{\bot}$. But as $\mathcal{T} \subseteq_{\dot{\bot}} \Omega$, this means $\mathcal{T} \vdash_{CE} \dot{\bot}$ as well, contradicting the consistency of $\mathcal{T}$. □

The proofs from Section 3.2 onward have been fully constructive. However, deducing the traditional statement of completeness from the model existence theorem requires a non-constructive reasoning principle. This insight is encoded in the statement below, which is nearly identical to that of traditional completeness, save for the doubly-negated conclusion, thereby admitting a constructive proof.

Theorem 3.18 *For any closed theory $\mathcal{T}$ and closed formula φ, $\mathcal{T} \vDash_S \varphi$ entails $\mathcal{T} \vdash_{CE} \varphi$.*

Proof Assuming $\mathcal{T} \vDash_S \varphi$ and $\mathcal{T} \nvdash_{CE} \varphi$, we have to prove $\bot$. By Fact 3.15, our second assumption is equivalent to $\mathcal{T} \cup \{\dot{\neg}\varphi\} \nvdash_{CE} \dot{\bot}$, the consistency of $\mathcal{T} \cup \{\dot{\neg}\varphi\}$. As φ and $\mathcal{T}$ are closed, so is $\mathcal{T} \cup \{\dot{\neg}\varphi\}$. Hence Theorem 3.17 grants us a standard model $\mathcal{I}$ and an assignment ρ such that $\rho \vDash \mathcal{T} \cup \{\dot{\neg}\varphi\}$. Then $\rho \vDash \varphi$ as $\rho \vDash \mathcal{T} \subseteq \mathcal{T} \cup \{\dot{\neg}\varphi\}$. This is a contradiction to $\rho \vDash \dot{\neg}\varphi$! □

Strong completeness, that is completeness on theories, can thus be derived for any theory class C which is stable under $\vdash_{CE}$.

Corollary 3.19 *Under C-stability, $\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vdash_{CE} \varphi$ can be derived for any closed φ and closed $\mathcal{T}$ for which $C\,\mathcal{T}$ holds.*

Having established that C-stability entails strong completeness, we can thus deduce different completeness theorems in terms of the non-constructive principles of Section 2.5 based on their characterizations as C-stabilities.

Corollary 3.20 *Let $\mathcal{T}$ and φ be closed.*

1. *Under double-negation elimination, $\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ holds.*
2. *Under synthetic Markov's principle, $\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ holds if $\mathcal{T}$ is enumerable.*
3. *Under object Markov's principle, $\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ holds if $\mathcal{T}$ is finite.*

We now prove that the inverse holds as well. That is, that strong completeness implies the stability of $\vdash_{CE}$. To prove this, we first establish another important fact: the stability of $\vDash_S$. Recall that this notion of validity is defined solely in terms of standard models. Formula satisfaction in a standard model is stable by virtue of its classicality. The following lemma establishes that this property is propagated to the level of validity.

Lemma 3.21 *For any $\mathcal{T}$ and φ, $\neg\neg\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vDash_S \varphi$.*

Proof Assume $\neg\neg\mathcal{T} \vDash_S \varphi$. We have to show $\rho \vDash \varphi$ for any standard model $\mathcal{I}$ and assignment ρ such that $\rho \vDash \mathcal{T}$. As $\mathcal{I}$ is classical, $\rho \vDash_S ((\varphi \dot{\to} \dot{\bot}) \dot{\to} \varphi) \dot{\to} \varphi$ holds. Hence it suffices to deduce $\rho \vDash_S \varphi$ from $\neg\rho \vDash_S \varphi$ as $\mathcal{I}$ has a standard $\dot{\bot}$. This follows per explosion, as $\neg\rho \vDash_S \varphi$ implies $\neg\mathcal{T} \vDash_S \varphi$. □

This proof is an application of McCarty's more general result in [53]. Intuitively, the stability of $\vdash_{CE}$ is entailed by completeness because $\vdash_{CE}$ and $\vDash_S$ are equivalent by soundness and completeness, meaning that the stability of $\vDash_S$ is extended to $\vdash_{CE}$. This proof makes use of the soundness of $\vdash_{CE}$ and $\vDash_S$, which we prove in Section 3.5. Note also that this proof assumes completeness for open theories and formulas. We show in Section 3.4 how to extend the current completeness result to also cover open theories.

Lemma 3.22 *$\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ entails the stability of $\mathcal{T} \vdash_{CE} \varphi$ for any $\mathcal{T}$.*

Proof Assume completeness and $\neg\neg\mathcal{T} \vdash_{CE} \varphi$. Per assumption, it suffices to show $\mathcal{T} \vDash_S \varphi$. Using Lemma 3.21, this can be achieved by proving $\neg\neg\mathcal{T} \vDash_S \varphi$. Assume $\neg\mathcal{T} \vDash_S \varphi$. Per assumption, it suffices to contradict $\mathcal{T} \vdash_{CE} \varphi$. Using Lemma 3.39, this implies $\mathcal{T} \vDash_S \varphi$, which is a contradiction. □

With this, we have established that strong completeness with regards to $\vdash_{CE}$ and stability of $\vdash_{CE}$ are equivalent. This gives rise to multiple characterizations of the requirements of strong completeness.

Corollary 3.23

1. *$\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ for any $\mathcal{T}$ entails double-negation elimination*
2. *$\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ for any enumerable $\mathcal{T}$ entails the synthetic Markov's principle*
3. *$\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ for any finite $\mathcal{T}$ entails the object Markov's principle*

3.3.3 Exploding Models

While we have demonstrated that the traditional phrasing of completeness for first-order logic is inherently unconstructive, we have not yet established why this is the case. Careful analysis allows us to pin down the cause of this unconstructiveness: the consistency requirement in the model existence theorem. As consistency of a theory is undecidable, it is not constructively stable. The reason why consistency is required is very singular as well: It is only used to show that the syntactic interpretation gives rise to a standard model in Theorem 3.17. Roughly speaking, this means that the root cause of the non-constructivity is "the way standard models treat $\dot{\bot}$".

If this analysis is accurate, one would hope that a fully constructive proof could be obtained if the notion of model is changed to treat $\dot{\bot}$ differently. In the remainder of this and the next section, we show that this is indeed the case by presenting two alternative notions of model that allow for a constructive completeness proof.

The first alternative notion of model we consider is a generalization of standard models. This variant of Tarski models was first presented by Krivine [42] based on a similar variation on Kripke models by Veldman [65]. Instead of requiring $\bot^{\mathcal{I}}$ to imply $\bot$, we restrict our notion of validity to those models that still admit the explosion principle. We call these models "exploding models".

Definition 3.11 (**Exploding models**)

1. *An interpretation $\mathcal{I}$ is said to have an exploding $\dot{\bot}$ if $\bot^{\mathcal{I}} \rightarrow P^{\mathcal{I}}\, t_1 \dots t_{|P|}$ for all $t_1, \dots, t_{|P|}$.*
2. *An interpretation is said to constitute an exploding model if it is classical and has an exploding $\dot{\bot}$.*

We denote the notion of validity restricted to exploding models with $\vDash_E$. Note that this allows for models where $\bot^{\mathcal{I}} \leftrightarrow \top$ as long as $P^{\mathcal{I}} t_1 \dots t_{|P|}$ holds for any $t_1, \dots, t_{|P|}$. As claimed above, all exploding models indeed admit the explosion principle.

Fact 3.24 *Let $\mathcal{I}$ be an interpretation with an exploding $\dot{\bot}$ and ρ an assignment. Then for any formula φ, $\rho \vDash \dot{\bot} \dot{\rightarrow} \varphi$.*

Exploding models are a slight variation of the standard notion of models. Indeed, these two notions of model are classically equivalent. More precisely, $\mathcal{T} \vDash_E \varphi$ constructively entails $\mathcal{T} \vDash_S \varphi$, while the contrary only holds non-constructively. This mismatch should not be considered disappointing. Rather, if these notions of model were constructively equivalent, completeness with regards to them would be as well, extinguishing all hope for a constructive completeness proof with regards to $\vDash_E$. Indeed, we later show that the translation from $\vDash_S$ to $\vDash_E$ is equivalent to the stability of $\vdash_{CE}$ as well.

As might be expected from the similarities between $\vDash_S$ and $\vDash_E$, the completeness proofs with regards to them do not differ much either. We first prove model existence and then deduce strong completeness from that. The interpretation we chose for model existence is exactly the same as for standard models.

Theorem 3.25 *Let $\mathcal{T}$ be closed. Then there is an exploding model $\mathcal{I}$ and an assignment ρ such that $\rho \vDash \mathcal{T}$ holds.*

Proof Most of this proof is completely analogous to Theorem 3.17. The only difference is that it suffices to show that $\mathcal{I}_\mathcal{T}$ has an exploding $\dot\bot$ to deduce that it constitutes a model. Hence let $\dot\bot \in \Omega$. We have to show that $P\,s\,t \in \Omega$ for any terms s and t. Per Theorem 3.13.2, this is the case whenever $\Omega \vdash_{CE} P\,s\,t$, which follows by (Exp) as $\dot\bot \in \Omega$. □

As Theorem 3.25 does not require the theory $\mathcal{T}$ to be consistent, the statement of completeness does not need to be doubly-negated. This allows us to prove strong exploding completeness fully constructively.

Theorem 3.26 *For any closed theory $\mathcal{T}$ and closed formula φ, $\mathcal{T} \vDash_E \varphi$ entails $\mathcal{T} \vdash_{CE} \varphi$.*

Proof By Fact 3.15, it suffices to prove $\mathcal{T} \cup \{\dot\neg\varphi\} \vdash_{CE} \dot\bot$. As φ and $\mathcal{T}$ are closed, so is $\mathcal{T} \cup \{\dot\neg\varphi\}$. Hence Theorem 3.25 grants us an interpretation $\mathcal{I}$ and an assignment ρ such that $\rho \vDash \mathcal{T} \cup \{\dot\neg\varphi\}$. Then $\rho \vDash \varphi$ as $\rho \vDash \mathcal{T} \subseteq \mathcal{T} \cup \{\dot\neg\varphi\}$. Together with $\rho \vDash \varphi$, this entails $\rho \vDash \dot\bot$. Then $\mathcal{T} \cup \{\dot\neg\varphi\} \vdash_{CE} \dot\bot$. □

Establishing strong exploding completeness allows for a very elegant characterization of the relationship between $\vDash_S$ and $\vDash_E$. As stated previously, $\vDash_E$ trivially entails $\vDash_S$. However, one can show that being able to translate $\vDash_S$ into $\vDash_E$ is equivalent to the stability of $\vdash_{CE}$. Intuitively, this is because such a translation would allow deriving completeness for $\vDash_S$ from that of $\vDash_E$.

Lemma 3.27 *Let $\mathcal{T}$ be a theory and φ a formula.*

1. *$\mathcal{T} \vDash_E \varphi$ entails $\mathcal{T} \vDash_S \varphi$*
2. *$\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vDash_E \varphi$ is equivalent to stability of $\vdash_{CE}$ for closed $\mathcal{T}$ and φ*

Proof

1. Per Lemma 3.14, it suffices to show that any standard model is an exploding model. This is the case as any standard $\dot\bot$ constitutes an exploding $\dot\bot$ by explosion in the meta-logic.

2. We prove this by showing that it is equivalent to strong standard completeness.

   $\rightarrow$ If $\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vDash_E \varphi$, then we can conclude $\mathcal{T} \vdash_{CE} \varphi$ from Theorem 3.26.

   $\leftarrow$ If $\vdash_{CE}$ is stable, then per Theorem 3.18 $\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vdash_{CE} \varphi$. Then per Lemma 3.39, we can conclude $\mathcal{T} \vDash_E \varphi$. □

3.3.4 Minimal Models

As exploding models allow for models which satisfy $\dot{\bot}$, they constitute a non-standard semantics. While this difference is classically insignificant, it does make a difference constructively. This difference might impede ones willingness to wholeheartedly call Theorem 3.26 a "real" completeness theorem.

The second alternative notion of model does not share this shortcoming. However, it is much more radical. It does not impose any restrictions on $\bot^{\mathcal{I}}$ at all. That is to say, this notion of models treats $\dot{\bot}$ as simply another logical constant, stripping it of the explosion principle. This notion of models is called "minimal models". We denote validity restricted to minimal models by $\vDash_M$.

Definition 3.12 (Minimal models) *An interpretation is said to constitute a minimal model if it is classical.*

Minimal models are a radical departure from the previous two notions of model. As $\dot{\bot}$ has lost its explosion principle, their notions of validity do not align anymore. More precisely, whenever $\mathcal{T} \vDash_M \varphi$ it follows that $\mathcal{T} \vDash_S \varphi$ and $\mathcal{T} \vDash_E \varphi$. However, the converse is not the case: consider $\dot{\bot} \dot{\rightarrow} Q$. As $\vDash_S$ and $\vDash_E$ share the explosion principle, the formula is valid under both notions of model. However, $\vDash_M$ does not validate this formula. For an example of a minimal model in which it does not hold, take an interpretation with $\bot^{\mathcal{I}} = \top$ and $Q^{\mathcal{I}} = \bot$.

Lemma 3.28 *Let $\mathcal{T}$ be a theory and φ a formula. If $\mathcal{T} \vDash_M \varphi$ then $\mathcal{T} \vDash_S \varphi$ and $\mathcal{T} \vDash_E \varphi$.*

Proof This follows by Lemma 3.14 as standard models and exploding models also constitute minimal models as they are classical. □

As opposed to exploding models, minimal models constitute the canonical model notion for a first-order logic restricted to the connectives $\dot{\rightarrow}$ and $\dot{\forall}$. While the name of the 0-ary predicate $\dot{\bot}$ might be slightly misleading, this alone is no reason for disconcern. Consequently, the theorem of completeness for minimal models ought to be considered "real".

We have already established that validity under $\vDash_M$ differs from that of $\vDash_S$. This means that the deduction system for which we can prove completeness changes as well: As $\dot{\bot}$ has been demoted to a logical constant, the proof uses $\vdash_{CL}$ as its deduction system. As this system is not refutation complete, the proof strategy significantly differs as well. Overall, it is a constructive variant of Schumm's classical proof of completeness of minimal implicational predicate logic [59].

The differences between the proof strategies for $\vDash_S$ and $\vDash_M$ start with the construction of the syntactic interpretation. As $\vDash_M$ does not guarantee an exploding $\dot{\bot}$, it is not immediately clear which formula to choose for $\widetilde{\bot}$ in the application of Theorem 3.13. Ultimately, this role will be taken by the semantically valid formula φ that should be deduced in $\vdash_{CL}$. For this reason, the choice of $\widetilde{\bot}$ is left open in the definition of the interpretation.

Definition 3.13 (Syntactic interpretation) *Let $\mathcal{T}$ be a closed theory and $\widetilde{\bot}$ a closed formula. Further, let Ω be the theory constructed using Theorem 3.13. The syntactic interpretation of $\mathcal{T}$, denoted $\mathcal{I}_{\mathcal{T},\widetilde{\bot}}$, is given by:*

$$f^{\mathcal{I}_{\mathcal{T},\widetilde{\bot}}}\, t_1 \dots t_{|f|} := f\, t_1 \dots t_{|f|} \qquad P^{\mathcal{I}_{\mathcal{T},\widetilde{\bot}}} t_1 \dots t_{|P|} := P\, t_1 \dots t_{|P|} \in \Omega \qquad \bot^{\mathcal{I}_{\mathcal{T},\widetilde{\bot}}} := \dot{\bot} \in \Omega$$

As before, we can prove a correspondence between $\vDash$ and membership in Ω, as well as the model existence theorem. Both proofs are analogous to those for the exploding models, as the only difference between the two syntactic interpretations, the choice of $\widetilde{\bot}$, does not matter for them.

Fact 3.29 *Let $\mathcal{T}$ be closed. Then $\sigma \vDash \varphi$ iff $\varphi[\sigma] \in \Omega$ under $\mathcal{I}_{\mathcal{T},\widetilde{\bot}}$.*

Fact 3.30 *Let $\mathcal{T}$ and φ be closed. Then there is a minimal model $\mathcal{I}$ and an assignment ρ such that $\rho \vDash \mathcal{T}$ holds.*

Using model existence, we can deduce strong minimal completeness. This proof crucially relies on the fact that $\widetilde{\bot}$ is chosen as φ.

Theorem 3.31 *For any closed theory $\mathcal{T}$ and closed formula φ, $\mathcal{T} \vDash_M \varphi$ entails $\mathcal{T} \vdash_{CL} \varphi$.*

Proof Assume $\mathcal{T} \vDash_M \varphi$. By Fact 3.30, there exists an interpretation $\mathcal{I}$ and an assignment ρ with $\rho \vDash \mathcal{T}$. Then $\rho \vDash \varphi$ as $\rho \vDash \mathcal{T}$ and hence $\varphi \in \Omega$. As $\mathcal{T} \subseteq_\varphi \Omega$, $\mathcal{T} \vdash_{CL} \varphi$ follows. □

3.4 Extending the Completeness Results

In the previous section, we have derived three different completeness results. However, these results differed from the more traditional statements of completeness as they were restricted to closed formulas and theories. This section demonstrates how to lift this restriction. We begin by deriving completeness for arbitrary finite contexts. After that, we give a brief account of how to extend the previous results to open theories and formulas.

3.4.1 Finite Completeness

We prove completeness with regards to finite contexts by reducing it to completeness of a closed formula under the empty context. The latter can be derived from the completeness results of Section 3.3. The reduction steps can be seen in Fig. 3.1. A context Γ and formula φ are first transformed into a single, equivalent formula $\Gamma \dot{\rightarrow} \varphi$, which is then closed using the closing operator $\dot{\forall}^*$ from Section 2.3.4. After proving its completeness, its unfolded version $\Gamma \vdash \varphi$ is deduced.

$$\begin{array}{ccccc} \Gamma \vDash_X \varphi & \longrightarrow & \vDash_X \Gamma \dot{\rightarrow} \varphi & \longrightarrow & \vDash_X \dot{\forall}^* (\Gamma \dot{\rightarrow} \varphi) \\ & & & & \downarrow \\ \Gamma \vdash \varphi & \longleftarrow & \vdash (\Gamma \dot{\rightarrow} \varphi) & \longleftarrow & \vdash \dot{\forall}^* (\Gamma \dot{\rightarrow} \varphi) \end{array}$$

Figure 3.1: Deriving finite completeness

The definition of the operation $\Gamma \dot{\rightarrow} \varphi$ as well as the properties needed to prove the correctness of this reduction are given below.

Definition 3.14 (Context reduction) *Let Γ be a context and φ a formula. We define*

$$[] \dot{\rightarrow} \varphi := \varphi \qquad (\Gamma, \psi) \dot{\rightarrow} \varphi := \psi \dot{\rightarrow} (\Gamma \dot{\rightarrow} \varphi)$$

Fact 3.32 *Let Γ be a context and φ a formula.*

1. $\Gamma \vDash_X \varphi \;\rightarrow\; \vDash_X \Gamma \dot{\rightarrow} \varphi$
2. $\vDash_X \varphi \rightarrow \vDash_X \dot{\forall}^* \varphi$
3. $\vdash \dot{\forall}^* \varphi \rightarrow \vdash \varphi$
4. $\vdash \Gamma \dot{\rightarrow} \varphi \;\rightarrow\; \Gamma \vdash \varphi$

Using these properties, we derive completeness for finite contexts. Notably, as this only requires completeness in terms of the empty context, which is finite, the object Markov's principle suffices for the proof of completeness.

Corollary 3.33 (Finite completeness) *Let Γ be a context and φ a formula.*

1. *Under the object Markov's principle, $\Gamma \vDash_S \varphi$ entails $\Gamma \vdash_{CE} \varphi$*
2. *$\Gamma \vDash_E \varphi$ entails $\Gamma \vdash_{CE} \varphi$*
3. *$\Gamma \vDash_M \varphi$ entails $\Gamma \vdash_{CL} \varphi$*

Indeed finite completeness and the translation from finite validity in $\vDash_S$ to finite validity in $\vDash_E$ can be proven equivalent to object Markov's principle. The proofs are completely analogous to those for general theories in Section 3.3.2 and Section 3.3.3.

Corollary 3.34

1. *If $\forall \Gamma \varphi.\, \Gamma \vDash_S \varphi \rightarrow \Gamma \vdash_{CE} \varphi$ then $\vdash_{CE}$ is stable under finite contexts.*
2. *If $\forall \Gamma \varphi.\, \Gamma \vDash_S \varphi \rightarrow \Gamma \vDash_E \varphi$ then $\vdash_{CE}$ is stable under finite contexts.*

3.4.2 Full Strong Completeness

Having derived finite completeness, we now move on to deriving strong completeness for open theories. A very elegant method for this is demonstrated by Herbelin and Ilik [29]. They also define their Henkin construction in terms of free variables which means their proof is restricted to theories $\mathcal{T}$ and formulas φ with infinitely many fresh variables. They solve this by doubling every free variable occurring in open $\mathcal{T}$ and φ, thereby guaranteeing that every uneven variable is fresh for the resulting theory and formula. As this transformation does not affect validity or provability, their result is thus extended to open theories and formulas.

While the same method can be used to extend our completeness results for standard and exploding models to open settings, it will not work for the minimal models. Recall that in their proof of completeness, we chose the formula φ as the $\dot{\bot}$ replacement $\widetilde{\bot}$. It is important that $\widetilde{\bot}$ is truly closed to ensure that the explosion axioms are not affected by variable shifts $\uparrow(\widetilde{\bot} \dot{\to} \dot{\forall}^* \psi) = \widetilde{\bot} \dot{\to} \dot{\forall}^* \psi$, meaning mere shifting of the free variables in φ does not suffice. We therefore follow a different approach, outlined in Fig. 3.2. Recall that the fragment $\mathfrak{F}_F$ was generalized over a signature Σ. We extend this signature by an infinite family of new constants c_n, arriving at the signature Σ_c. We then lift $\mathcal{T}$ and φ to Σ_c and close them by substituting every free variable n with c_n. After deducing the completeness of the modified $\mathcal{T}$ and φ, we revert them back their original form in Σ by replacing every c_n with the free variable n. As these transformation steps all maintain validity and provability, strong completeness can be deduced for open theories and formulas.

$$
\begin{array}{ccccc}
\mathcal{T} \vDash_X \varphi & \longrightarrow & \Uparrow\mathcal{T} \vDash_X \Uparrow\varphi & \longrightarrow & (\Uparrow\mathcal{T})[c_{\_}] \vDash_X (\Uparrow\varphi)[c_{\_}] \\
 & & & & \downarrow \\
\mathcal{T} \vdash \varphi & = & \Downarrow((\Uparrow\mathcal{T})[c_{\_}]) \vdash \Downarrow((\Uparrow\varphi)[c_{\_}]) & \longleftarrow & (\Uparrow\mathcal{T})[c_{\_}] \vdash (\Uparrow\varphi)[c_{\_}]
\end{array}
$$

Figure 3.2: Deriving full strong completeness

The definitions of the operations $\Uparrow$ and $\Downarrow$ are given below. Because $\Downarrow$ is supposed to map every c_n to the free variable n, it is defined carrying an index m that denotes how many universal quantifiers have to be skipped to arrive at the level of the free variables. Thus, $\Downarrow^m$ maps c_n to the variable $n + m$. However, we shorten $\Downarrow^0$ to $\Downarrow$ to aid readability.

Definition 3.15 (Lifting and dropping)

$$\begin{aligned} \Uparrow &: \mathfrak{T}^{\Sigma} \to \mathfrak{T}^{\Sigma_c} \\ \Uparrow x &:= x \\ \Uparrow (f\, t_1 \ldots t_{|f|}) &:= f_c\, (\Uparrow t_1 \ldots \Uparrow t_{|f|}) \end{aligned} \qquad \begin{aligned} \Downarrow &: \mathbb{N} \to \mathfrak{T}^{\Sigma_c} \to \mathfrak{T}^{\Sigma} \\ \Downarrow^m x &:= x \\ \Downarrow^m (f_c\, t_1 \ldots t_{|f|}) &:= f\, (\Downarrow^m t_1 \ldots \Downarrow^m t_{|f|}) \\ \Downarrow^m c_n &:= n + m \end{aligned}$$

$$\begin{aligned} \Uparrow &: \mathfrak{F}_F^{\Sigma} \to \mathfrak{F}_F^{\Sigma_c} \\ \Uparrow \dot{\bot} &:= \dot{\bot} \\ \Uparrow (P\, t_1 \ldots t_{|P|}) &:= P\ \Uparrow t_1 \ldots \Uparrow t_{|P|} \\ \Uparrow (\varphi \dot{\to} \psi) &:= \Uparrow \varphi \dot{\to} \Uparrow \psi \\ \Uparrow (\dot{\forall} \varphi) &:= \dot{\forall} (\Uparrow \varphi) \end{aligned} \qquad \begin{aligned} \Downarrow &: \mathbb{N} \to \mathfrak{F}_F^{\Sigma_c} \to \mathfrak{F}_F^{\Sigma} \\ \Downarrow^m \dot{\bot} &:= \dot{\bot} \\ \Downarrow^m (P\, t_1 \ldots t_{|P|}) &:= P\ \Downarrow^m t_1 \ldots \Downarrow^m t_{|P|} \\ \Downarrow^m (\varphi \dot{\to} \psi) &:= \Downarrow^m \varphi \dot{\to} \Downarrow^m \psi \\ \Downarrow^m (\dot{\forall} \varphi) &:= \dot{\forall} (\Downarrow^{S\,m} \varphi) \end{aligned}$$

The following facts are required to deduce full strong completeness. Note that their proofs are non-trivial spanning over 450 lines of Coq. However, due to their technical nature, we have decided to solely state them here.

Fact 3.35 *Let $\mathcal{T}$ be a theory, φ a formula and $X \in \{S, E, M\}$.*

1. $\mathcal{T} \vDash_X \varphi \to \Uparrow \mathcal{T} \vDash_X \Uparrow \varphi$
2. $\mathcal{T} \vdash \varphi \to (\Uparrow \mathcal{T})[\lambda n.\ c_n] \vDash_X (\Uparrow \varphi)[\lambda n.\ c_n]$
3. $\Downarrow ((\Uparrow \varphi)[\lambda n.\ c_n]) = \varphi$

Using the facts from above, we can deduce the full strong completeness results. For standard models, the C-stability has to be closed under the lifting operation $(\Uparrow -)[\lambda n.\ c_n]$ as the completeness proof from Section 3.3.2 is applied to the theory $(\Uparrow \mathcal{T})[\lambda n.\ c_n]$ to derive this result.

Fact 3.36 *Let $\mathcal{T}$ be a theory and φ a formula.*

1. *Under any $(\Uparrow -)[\lambda n.\ c_n]$ closed C-stability, $\mathcal{T} \vDash_S \varphi \to \mathcal{T} \vdash_{CE} \varphi$ for any $C\, \mathcal{T}$*
2. *$\mathcal{T} \vDash_E \varphi$ entails $\mathcal{T} \vdash_{CE} \varphi$*
3. *$\mathcal{T} \vDash_M \varphi$ entails $\mathcal{T} \vdash_{CL} \varphi$*

As all stability notions we singled out in Section 2.5.1 are closed under $(\Uparrow -)[\lambda n.\ c_n]$, we can derive the exact characterizations of strong completeness.

Corollary 3.37

1. *$\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vdash_{CE} \varphi$ for any $\mathcal{T}$ is equivalent to double negation elimination*
2. *$\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vdash_{CE} \varphi$ for any enumerable $\mathcal{T}$ is equivalent to the synthetic Markov's principle*
3. *$\mathcal{T} \vDash_S \varphi \rightarrow \mathcal{T} \vdash_{CE} \varphi$ for any finite $\mathcal{T}$ is equivalent to the object Markov's principle*

3.5 Soundness

In Section 3.3, we introduced three different notions of validity. We now give a unified proof of soundness covering all of them. In the context of completeness, a proof of soundness can be seen as a guarantee that the notions of model and validity were chosen correctly, thereby establishing its legitimacy.

We first establish an important fact about the interaction between assignments and substitution under formula satisfaction. Substitutions in the term under $\vDash_X$ can always be moved into the assignment instead.

Fact 3.38 *Let X be a constraint, $\mathcal{I}$ an interpretation, ρ an assignment and σ a substitution. Then $\rho \vDash_X \varphi[\sigma]$ iff $(\lambda x.\ (\sigma\, x)^\rho) \vDash_X \varphi$ for any φ.*

Proof We first prove $\forall \rho_1, \rho_2.\ \rho_1 \vDash_X \varphi$ iff $\rho_2 \vDash_X \varphi$ and use this to prove this lemma. Both proves are per formula induction of φ. ☐

Using this fact, one can now deduce soundness. As there are many similarities between the different variants of $\vdash$, all possible soundness statements can be proven at once. We require the models satisfying the constraint with regards to which soundness is being proven to fulfill certain properties representing the "features" of the deduction system $\vdash_{SB}$.

Lemma 3.39 *Let X be a constraint and SB a ND-variant such that all models satisfying X are classical if $S = C$ and all models satisfying X have exploding $\dot{\bot}$s if $B = E$. Further, let Γ be a context, $\mathcal{T}$ a theory and φ a formula.*

1. $\Gamma \vdash_{SB} \varphi \rightarrow \Gamma \vDash_X \varphi$
2. $\mathcal{T} \vdash_{SB} \varphi \rightarrow \mathcal{T} \vDash_X \varphi$

Proof

1. Proof per induction on $\Gamma \vdash_{SB} \varphi$. We only cover the more involved cases.

(Exp): As (Exp) requires $B = E$, we know any model satisfying X has an exploding $\dot{\bot}$. The claim then follows with Fact 3.24.

(Peirce): As (Peirce) requires $S = C$, the claim follows per assumption.

(AllI): Given an interpretation $\mathcal{I}$ and an assignment ρ, we have to show that $d, \rho \vDash_X \varphi$ for arbitrary d. Per inductive hypothesis $\uparrow \Gamma \vDash_X \varphi$. Hence it suffices to show that $d, \rho \vDash_X \uparrow \Gamma$. For a given $\uparrow \phi \in \uparrow \Gamma$, we know that $d, \rho \vDash_X \uparrow \phi$ iff $\rho \vDash_X \phi$ per Fact 3.38, the latter of which holds per assumption as $\phi \in \Gamma$.

(AllE): Given an interpretation $\mathcal{I}$ and an assignment ρ, we have to show that $\rho \vDash_X \varphi[t]$ for arbitrary t. By Fact 3.38, this is equivalent to $t^\rho, \rho \vDash_X \varphi$, which holds per inductive hypothesis.

2. If $\mathcal{T} \vdash_{SB} \varphi$ then there is $\Gamma \subseteq \mathcal{T}$ with $\Gamma \vdash_{SB} \varphi$. Using 1., this means $\Gamma \vDash_X \varphi$, which implies $\mathcal{T} \vDash_X \varphi$ as $\Gamma \subseteq \mathcal{T}$. □

The specific statements of soundness corresponding to the statements of completeness proven in the previous section can easily be deduced.

Corollary 3.40 *Let $\mathcal{T}$ be a theory and φ a formula.*

1. *$\mathcal{T} \vdash_{CE} \varphi$ entails $\mathcal{T} \vDash_S \varphi$*
2. *$\mathcal{T} \vdash_{CE} \varphi$ entails $\mathcal{T} \vDash_E \varphi$*
3. *$\mathcal{T} \vdash_{CL} \varphi$ entails $\mathcal{T} \vDash_M \varphi$*

An important consequence of soundness with regards to standard models is the consistency of $\vdash$. This proof is slightly more complicated than in other presentations of Tarski semantics, as we first need to establish the existence of a standard model.

Lemma 3.41 *There exists a standard model.*

Proof Consider the interpretation $\mathcal{I}$ defined on the domain of type **1** whose only member is $\star$.

$$f^{\mathcal{I}}\, t_1 \dots t_{|f|} := \star \quad P^{\mathcal{I}}\, t_1 \dots t_{|P|} := \top \quad \bot^{\mathcal{I}} := \bot$$

Per definition of $\bot^{\mathcal{I}}$, it has a standard $\dot\bot$. Further, one can prove per induction on φ that $\rho \vDash \varphi \vee \neg \rho \vDash \varphi$ for any formula φ and assignment ρ. One can thus prove that all instances of $(((\varphi \dot\rightarrow \psi) \dot\rightarrow \varphi) \dot\rightarrow \varphi)$ are satisfied under any assignment ρ. Thus $\mathcal{I}$ is classical and therefore a standard model. □

Lemma 3.42 *There is no derivation of $\vdash \dot\bot$.*

Proof Assume there was a derivation of $\vdash \dot\bot$. By Lemma 3.39, this means that $\vDash_S \dot\bot$. When instantiated with the standard model from Lemma 3.41, this yields a contradiction. □

3.6 Conclusion

In this chapter, we have analyzed the completeness of Tarski semantics for three different notions of models: standard, exploding and minimal. To this end, we have given a theory extension procedure Theorem 3.13, generalized from that of Henkin, which could be used in the model existence proofs for all three model notions. Based on this, we were able to establish that the completeness with regards to $\vDash_S$ cannot be deduced constructively, as it is equivalent to certain C-stabilities, among them the principle of double-negation elimination as well as the synthetic and object Markov's principle. The two non-standard notions of models allow for constructive completeness proof. These results can be summed up by Fig. 3.3. Normal arrows denote the usual implication, while dashed arrows are implications which are equivalent to the stability of $\vdash_{CE}$.

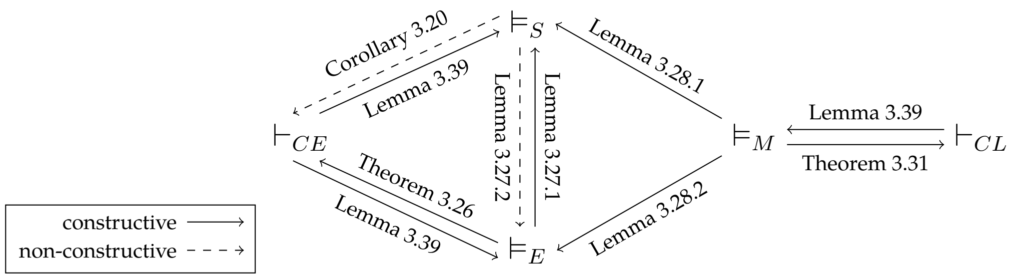

Figure 3.3: Results for classical completeness

Chapter 4

Kripke Semantics

In the previous chapter, we carried out a constructive analysis of completeness with regards to Tarski model semantics. This chapter will do the same for the canonical model semantics for intuitionistic first-order logic: Kripke models. Kripke models were originally developed to give a semantics to modal logic [40]. While a translation of intuitionistic logic into modal logic was already known, the first explicit rendering of Kripke models as a semantics for intuitionistic first-order logic, including a non-constructive completeness proof, was given by Kripke in 1965 [41]. However, Kreisel had already proven that the completeness of negative intuitionistic first-order logic (first-order logic restricted to results of Gödel's double-negation translation [25]) entails the non-constructive Markov's principle in 1962 [38]. As the notion of validity he uses in that proof is very general, the result extends to Kripke semantics as well. In 1976, Veldman proved that the non-constructive fan theorem suffices to prove completeness with regards to exploding Kripke models[65]. In this chapter, we carry out a constructive analysis of the completeness of multiple variants of Kripke semantics on the $\dot\forall$, $\dot\rightarrow$, $\dot\bot$-fragment $\mathfrak{F}_F$ of first-order predicate logic, based on a proof of the same by Herbelin and Lee [30].

We begin this chapter by defining Kripke models (Section 4.1) and the normal sequent calculus (Section 4.2), the deduction system employed in this chapter. We then conduct the constructive analysis of the completeness proofs in Section 4.3, first giving a constructive proof for minimal and exploding models and then demonstrating that completeness with regards to standard models requires the stability of $\vdash_{CE}$. We close the chapter by showing how to use the constructive completeness results for a syntactic normalization procedure in Section 4.4.

4.1 Kripke Models

Kripke models are a generalization of Tarski models that is suitable as the basis of semantics for modal and intuitionistic first-order logic. Instead of defining just one satisfaction relation $\rho \vDash \varphi$ for the whole model, Kripke models are defined in terms of a type $\mathcal{W}$ of possible worlds, each world $w : \mathcal{W}$ giving rise to a separate forcing relation $\rho \Vdash_w \varphi$. The worlds are not completely independent from each other. Instead, a world w can be reachable from a world v, denoted by $v \preccurlyeq w$. The forcing relation is monotonous with regards to $\preccurlyeq$: every formula which holds in v also holds in w.

Definition 4.1 (Kripke models) *A Kripke model $\mathcal{K}$ on a domain D consists of a type $\mathcal{W}$ of worlds preordered by a reachability relation $\preccurlyeq : \mathcal{W} \to \mathcal{W} \to \mathbb{P}$, a function interpretation $f^{\mathcal{K}} : D^{|f|} \to D$ for every $f : \mathcal{F}$, a predicate interpretation $P^{\mathcal{K}}_{-} : \mathcal{W} \to D^{|P|} \to \mathbb{P}$ for every $P : \mathcal{P}$ and a $\dot{\bot}$ interpretation $\bot^{\mathcal{K}}_{-} : \mathcal{W} \to \mathbb{P}$.*

The predicate and $\dot{\bot}$ interpretations are required to be monotonous with regards to $\preccurlyeq$. That is $P^{\mathcal{K}}_v\, d_1 \dots d_{|P|} \to P^{\mathcal{K}}_w\, d_1 \dots d_{|P|}$ and $\bot^{\mathcal{K}}_v \to \bot^{\mathcal{K}}_w$ for every $v \preccurlyeq w$.

Analogously to Tarski models (Definition 3.7), a Kripke model $\mathcal{K}$, together with an assignment $\rho : \mathbb{N} \to D$, gives rise to a term interpretation $-^{\rho} : \mathfrak{T} \to D$:

$$x^{\rho} := \rho\, x \qquad (f\, t_1 \dots t_{|f|})^{\rho} := f^{\mathcal{K}}\, t_1^{\rho} \dots t_{|f|}^{\rho}$$

The conditions of a formula φ being forced in the world v of a Kripke model $\mathcal{K}$ under an assignment ρ are given by the recursively defined translation into the meta-logic $\rho \Vdash_v \varphi$:

$$\begin{aligned} \rho \Vdash_v \dot{\bot} &:= \bot^{\mathcal{K}}_x \\ \rho \Vdash_v P\, t_1 \dots t_{|P|} &:= P^{\mathcal{K}}_v\, t_1^{\rho} \dots t_{|P|}^{\rho} \\ \rho \Vdash_v \varphi \dot{\to} \psi &:= \forall v \preccurlyeq w.\ \rho \Vdash_w \varphi \to \rho \Vdash_w \psi \\ \rho \Vdash_v \dot{\forall} \varphi &:= \forall d : D.\ d, \rho \Vdash_v \varphi \end{aligned}$$

Fact 4.1 *The forcing relation of any Kripke model $\mathcal{K}$ is monotonous with regards to $\preccurlyeq$. That is, for any ρ and φ, $\rho \Vdash_v \varphi$ implies $\rho \Vdash_w \varphi$ whenever $v \preccurlyeq w$.*

Kripke models are usually generalized even further by allowing each world $u : \mathcal{W}$ to choose its own domain D_w with the requirement that $D_v \subseteq D_w$ for all $v \preccurlyeq w$. Kripke models defined in this way require a slightly more complicated definition of the forcing relation for $\dot{\forall}\varphi$:

$$\rho \Vdash_v \dot{\forall} \varphi \ \text{ iff } \ \forall d : D_v, v \preccurlyeq w.\ d, \rho \Vdash_w \varphi$$

However, we have chosen to only consider Kripke models of constant domains D, allowing for the simpler definition of the forcing relation given in Definition 4.1. This suffices as the models constructed for the completeness proofs have constant domains. The completeness proofs can easily be adapted to Kripke models with varying domains.

Similar to the Tarski models of Chapter 3, our notion of Kripke models differs from that found throughout the literature. The forcing relation is usually defined as always interpreting $\dot{\bot}$ as $\bot$, whereas under our definition, each model can choose its own interpretation of $\dot{\bot}$ by means of $\bot^{\mathcal{K}}_{-}$.

As this chapter covers completeness proofs with regards to different notions of Kripke models, we define validity in terms of constraints $X : \mathcal{K} \to \mathbb{P}$ on models. As with constrained Tarski validity, subsumption of constraints implies subsumption of their associated notions of validity. These notions also extend to finite contexts Γ, as these can be treated as finite theories.

Definition 4.2 (**Constrained Kripke validity**)

1. *A theory $\mathcal{T}$ is said to be satisfied by a Kripke model interpretation $\mathcal{K}$ at world w under an assignment ρ, written as $\rho \Vdash_w \mathcal{T}$, if $\forall \varphi \in \mathcal{T}.\ \rho \Vdash_w \varphi$.*
2. *A formula φ is said to be valid under a theory $\mathcal{T}$ and under a constraint $X : \mathcal{K} \to \mathbb{P}$ if $\rho \Vdash_w \mathcal{T} \to \rho \Vdash_w \varphi$ holds for any Kripke model $\mathcal{K}$ satisfying X, any world w and assignment ρ. This is written as $\mathcal{T} \Vdash_X \varphi$.*

Fact 4.2 *Let C subsume C'. Then, for any φ and $\mathcal{T}$, $\mathcal{T} \Vdash_C \varphi$ entails $\mathcal{T} \Vdash_{C'} \varphi$.*

The three notions of Kripke models we consider are analogous to those of Chapter 3. However, we do not require the models to be classical as we did for the Tarski models, because we now discuss semantics for intuitionistic first-order logic.

Definition 4.3 (**Relevant notions of Kripke models**)

1. *A Kripke model $\mathcal{K}$ constitutes a standard model if $\bot^{\mathcal{K}}_u \to \bot$ for all $u : \mathcal{W}$. We denote standard validity by $\Vdash_S$.*
2. *A Kripke model $\mathcal{K}$ constitutes an exploding model if $\bot^{\mathcal{K}}_u \to P^{\mathcal{K}}_u\, d_1 \ldots d_{|P|}$ for any $P, d_1, \ldots, d_{|P|}$ and any u. We denote exploding validity by $\Vdash_E$.*
3. *Any Kripke model is a minimal model. We denote minimal validity by $\Vdash_M$.*

Fact 4.3 *For any exploding $\mathcal{K}$, $\rho \Vdash_w \varphi$ holds for any ρ and φ in any w with $\bot^{\mathcal{K}}_w$,*

The intuitionistic variants of the natural deduction system defined in Section 2.4 are sound with regards to the constrained Kripke semantics as defined above. If the natural deduction system allows for the use of the (Exp) rule, validity has to be constrained to at least only consider exploding models.

Theorem 4.4 *Let X be a constraint and and IB an ND-variant such that all models satisfying X are exploding if $B = E$. Then $\forall\, \mathcal{T}\, \varphi.\ \mathcal{T} \vdash_{IB} \varphi \to \Gamma \Vdash_X \varphi$.*

4.2 Normal Sequent Calculus

The completeness proofs of this section do not result in derivations in natural deduction but instead in sequents of the sequent calculus LJT, a normal variant of Gentzen's intuitionistic sequent calculus LJ [21, 22].

The rules of LJT can be seen in Fig. 4.1. There are two different kinds of derivations. Derivations with traditional conclusion $\Gamma \Rightarrow \varphi$ show that φ can be deduced from the context Γ. Secondly, the derivations resulting in conclusions with a focused formula $\Gamma ; \varphi \Rightarrow \psi$ show that ψ can be proven under the context Γ by using the focused formula φ. The distinction between conclusions with and without focused formula is only necessary to ensure the proofs derived in LJT are normal, a fact which we discuss further in Section 4.4. The rules are divided into left and right rules, depending on which side of $\Rightarrow$ they act on. Formulas in focus can be transformed by the left rules, either by instantiating universal quantifiers with (AllL) or by removing premises of implications with (IL). The right rules, (IR) and (AllR), correspond to the introduction rules (II) and (AllI) of the natural deduction system. The only way of bringing formulas into focus and discharging them again are the (Ctx) and (Ax) rules, respectively.

As with the natural deduction system, LJT is parameterized over a flag B, indicating whether the (Exp) rule may be used or not. We write $\Gamma \Rightarrow \varphi$ in the statements of lemmas and definitions whenever the choice of B is irrelevant.

$$\textsc{Ax}\frac{}{\Gamma ; \varphi \Rightarrow \varphi} \qquad \textsc{Ctx}\frac{\Gamma ; \varphi \Rightarrow \psi \quad \varphi \in \Gamma}{\Gamma \Rightarrow \psi} \qquad \text{IL}\frac{\Gamma \Rightarrow \varphi \quad \Gamma ; \psi \Rightarrow \theta}{\Gamma ; \varphi \dot{\rightarrow} \psi \Rightarrow \theta}$$

$$\text{IR}\frac{\Gamma, \varphi \Rightarrow \psi}{\Gamma \Rightarrow \varphi \dot{\rightarrow} \psi} \qquad \textsc{AllL}\frac{\Gamma ; \varphi[t] \Rightarrow \psi}{\Gamma ; \dot{\forall}\varphi \Rightarrow \psi} \qquad \textsc{AllR}\frac{\uparrow\Gamma \Rightarrow \varphi}{\Gamma \Rightarrow \dot{\forall}\varphi} \qquad \textsc{Exp}\frac{\Gamma \Rightarrow_E \dot{\bot}}{\Gamma \Rightarrow_E \varphi}$$

Figure 4.1: Normal Sequent Calculus LJT

LJT exhibits the same weakening properties as the natural deduction system from Section 2.4. That is, provability is maintained under context extensions and substitutions. Furthermore, one can always switch between the de Bruijn- and locally nameless variants of the (AllR) rule.

Fact 4.5] *The following rules can be shown admissible.*

$$\textsc{Weak}\frac{\Gamma' \Rightarrow \varphi \quad \Gamma' \subseteq \Gamma}{\Gamma \Rightarrow \varphi} \qquad \textsc{WeakS}\frac{\Gamma \Rightarrow \varphi}{\Gamma[\sigma] \Rightarrow \varphi[\sigma]}$$

Fact 4.6 *Let x be fresh for Γ and $\dot{\forall}\varphi$. Then $\uparrow\Gamma \Rightarrow \varphi$ if and only if $\Gamma \Rightarrow \varphi[x]$.*

The notions of deduction can be extended to theories $\mathcal{T}$ in a similar manner to $\vdash$.

Definition 4.4 (Provability under theories)

1. *We write $\mathcal{T} \Rightarrow \varphi$ to mean that there is a $\Gamma \subseteq \mathcal{T}$ with $\Gamma \Rightarrow \varphi$.*
2. *We write $\mathcal{T} ; \varphi \Rightarrow \psi$ to mean that there is a $\Gamma \subseteq \mathcal{T}$ with $\Gamma ; \varphi \Rightarrow \psi$.*

The deduction predicates of $\Rightarrow$ are enumerable. The enumeration procedures for this are completely analogous to that of the proofs in $\vdash_{SB}$ in Lemma 2.11.

Fact 4.7

1. *The predicate $(\lambda\varphi.\, \Gamma \Rightarrow \varphi)$ is enumerable for any Γ.*
2. *The predicate $(\lambda\varphi.\, \mathcal{T} \Rightarrow \varphi)$ is enumerable for enumerable $\mathcal{T}$.*

Proof We prove this by giving enumerations $e : \mathbb{N} \to \mathcal{L}(\mathfrak{F}_F)$.

1. We simultaneously define mutually recursive enumerations e_Γ and $e_{\Gamma;\varphi}$ for $\lambda\psi.\, \Gamma \Rightarrow \psi$ and $\lambda\psi.\, \Gamma ; \varphi \Rightarrow \psi$, respectively. For the sake of readability, we refer to both the enumeration of terms and formulas as e.

$$\begin{aligned}
e_\Gamma\, 0 &:= \Gamma \\
e_\Gamma\, (S\, n) &:= e_\Gamma\, n && + \\
&\quad [\, \psi \mid \varphi \in \Gamma, \psi \in e_{\Gamma;\varphi}\, n \,] && + \\
&\quad [\, \varphi \dot{\to} \psi \mid \varphi \in e\, n, \psi \in e_{(\Gamma,\varphi)}\, n \,] && + \\
&\quad [\, \dot{\forall}\varphi \mid \varphi \in e_{(\uparrow\Gamma)}\, n \,] && + \\
&\quad [\, \varphi \mid B = E, \bot \in e_\Gamma\, n, \varphi \in e\, n \,]
\end{aligned}$$

$$\begin{aligned}
e_{\Gamma;\varphi}\, 0 &:= [\, \varphi \,] \\
e_{\Gamma;\varphi\to\psi}\, (S\, n) &:= e_{\Gamma;\varphi}\, n + [\, \tau \mid \tau \in e_{\Gamma;\psi}\, n,\ \varphi \in e_\Gamma\, n \,] \\
e_{\Gamma;\dot{\forall}\varphi}\, (S\, n) &:= e_{\Gamma;\varphi}\, n + [\, \psi \mid t \in e\, n, \psi \in e_{\Gamma;\varphi[t]}\, n \,] \\
e_{\Gamma;\varphi}\, (S\, n) &:= e_{\Gamma;\varphi}\, n +
\end{aligned}$$

As with the enumerations for $\vdash$ in Lemma 2.11, each clause of the recursive cases of the enumerations corresponds to a deduction rule of $\Rightarrow$. Notably, the behavior of the enumerations for focused derivations $e_{\Gamma;\varphi}$ depends on the shape of φ.

2. We assume an enumeration $e_\mathcal{T} : \mathbb{N} \to \mathcal{L}(\mathfrak{F}_F)$ which enumerates all formulas in $\mathcal{T}$. Using $e_\mathcal{T}$, one can construct an enumeration $e' : \mathbb{N} \to \mathcal{L}(\mathcal{L}(\mathfrak{F}_F))$

which enumerates all contexts $\Gamma \subseteq \mathcal{T}$. Then the predicate $\lambda\varphi.\, \mathcal{T} \Rightarrow \varphi$ can be enumerated by the following enumeration:

$$e\,0 := [] \qquad e\,(S\,n) := e\,n +\!\!+ \,\mathsf{concat}\,[\,[\,\varphi \mid \varphi \in e_\Gamma\,n\,] \mid \Gamma \in e'\,n\,]$$

□

4.3 Constructive Analysis of Completeness Theorems

This section contains the constructive analysis of the completeness theorems with regards to Kripke models based on the proofs by Herbelin and Lee [30]. Similar to them, we only give completeness proofs in terms of finite contexts and explain how to modify those to obtain completeness for arbitrary theories. We begin by analyzing the slightly simpler proofs for exploding and minimal models. We then continue by giving a proof for the completeness of standard models as a slight variation on the previous proof. We close this section by demonstrating that completeness for standard models entails the stability of $\vdash_{CE}$.

4.3.1 Exploding and Minimal Models

The overall approach of these proofs is very different to that of completeness with regards to Tarski models. While both construct a suitable syntactic models (models whose domain are the terms $\mathfrak{T}$), the proof of classical completeness requires the construction of different models, depending on the theory and formula under scrutiny, whereas the proof for Kripke models only needs one model, called the universal Kripke model. To construct it, one simply chooses the contexts as the worlds $\mathcal{W}$ and then defines $P^{\mathcal{K}}$ and $\perp^{\mathcal{K}}$ in terms of provability under them. Importantly, this model is applicable to both variants of LJT, $\Rightarrow_L$ and $\Rightarrow_E$.

Definition 4.5 (Universal model) *The universal Kripke model $\mathcal{U}$ is a syntactic model. Its worlds $\mathcal{W}$ are the contexts $\mathcal{L}(\mathfrak{F}_F)$ with reachability $\preccurlyeq$ being defined as $\subseteq$. Furthermore*

$$P^{\mathcal{U}}_\Gamma\, t_1 \ldots t_{|P|} := \Gamma \Rightarrow P\, t_1 \ldots t_{|P|} \qquad \perp^{\mathcal{U}}_\Gamma := \Gamma \Rightarrow \dot{\perp}$$

The proof of completeness with regards to minimal or normal Kripke models now proceeds by proving a characteristic lemma about the universal Kripke model. This proof constitutes the majority of the completeness proof. It consists of two properties of the universal model, shown by simultaneous induction. Intuitively, these two statements allow one to convert between a formula being forced and the same formula being provable (1) and back (2).

Lemma 4.8 *Let Γ be a context, σ a substitution and φ a formula. In the universal Kripke model $\mathcal{U}$ the following hold.*

1. $\sigma \Vdash_\Gamma \varphi \to \Gamma \Rightarrow \varphi[\sigma]$
2. $(\forall\, \Gamma' \psi.\ \Gamma \subseteq \Gamma' \to \Gamma' ; \varphi[\sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi) \to \sigma \Vdash_\Gamma \varphi$

Proof We prove the following generalizations simultaneously by induction on φ.

1. $\forall\, \Gamma\, \sigma.\ \sigma \Vdash_\Gamma \varphi \to \Gamma \Rightarrow \varphi[\sigma]$
2. $\forall\, \Gamma\, \sigma.\ (\forall\, \Gamma'\, \psi.\ \Gamma \subseteq \Gamma' \to\ \Gamma' ; \varphi[\sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi) \to \sigma \Vdash_\Gamma \varphi$

Case $\dot\perp$ and $P\, t_1 \ldots t_{|P|}$:

1. In syntactic models, $t^\sigma = t[\sigma]$. Thus $\sigma \Vdash_\Gamma \varphi$ iff $\Gamma \Rightarrow \varphi[\sigma]$ in $\mathcal{U}$ for $\varphi = \dot\perp$ or $\varphi = P\, t_1 \ldots t_{|P|}$.
2. From $(\forall\, \Gamma'\, \psi.\ \Gamma \subseteq \Gamma' \to\ \Gamma' ; \varphi[\sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi)$ for $\varphi = \dot\perp$ or $\varphi = P\, t_1 \ldots t_{|P|}$, one can deduce $\Gamma \Rightarrow \varphi[\sigma]$ by choosing $\Gamma' = \Gamma$ and $\psi = \varphi[\sigma]$ since $\Gamma' ; \varphi[\sigma] \Rightarrow \varphi[\sigma]$ by the (Ax) rule. This is equivalent to $\sigma \Vdash_\Gamma \varphi$ as in 1.

Case $\varphi \dot\to \psi$:

1. Assuming $\forall \Gamma'.\ \Gamma \subseteq \Gamma' \to \sigma \Vdash_{\Gamma'} \varphi \to \sigma \Vdash_{\Gamma'} \psi$, one has to derive that $\Gamma \Rightarrow (\varphi \to \psi)[\sigma]$. Per (IR) and inductive hypothesis 1. for ψ it suffices to show $\sigma \Vdash_{\Gamma, \varphi[\sigma]} \psi$. Applying the inductive hypothesis 2. for φ and the assumption, it suffices to show $\forall \Gamma'\, \theta.\ \Gamma, \varphi[\sigma] \subseteq \Gamma' \to \Gamma' ; \varphi[\sigma] \Rightarrow \psi[\sigma] \to \Gamma' \Rightarrow \psi[\sigma]$, which holds per (Ctx).
2. Assuming $\forall\, \Gamma'\, \theta.\ \Gamma \subseteq \Gamma' \to \Gamma' ; (\varphi \dot\to \psi)[\sigma] \Rightarrow \theta \to \Gamma' \Rightarrow \theta$ one has to deduce $\forall\, \Gamma'.\ \Gamma \subseteq \Gamma' \to \sigma \Vdash_{\Gamma'} \varphi \to \sigma \Vdash_{\Gamma'} \psi$. Because of the inductive hypothesis 2. for ψ it suffices to show $\forall \Gamma''\, \theta.\ \Gamma' \subseteq \Gamma'' \to \Gamma'' ; \psi[\sigma] \Rightarrow \theta \to \Gamma'' \Rightarrow \theta$. By applying using the assumption, $\Gamma'' \Rightarrow \theta$ reduces to $\Gamma'' ; (\varphi \dot\to \psi)[\sigma] \Rightarrow \theta$. This follows by (IL), as the assumption $\sigma \Vdash_{\Gamma'} \varphi$ implies $\Gamma'' \Rightarrow \varphi[\sigma]$ per inductive hypothesis 1 and (Weak).

Case $\dot\forall \varphi$:

1. Assuming $\forall t.\ t, \sigma \Vdash_\Gamma \varphi$, one has to deduce $\Gamma \Rightarrow (\dot\forall \varphi)[\sigma]$. By (AR), it suffices to show $\uparrow\Gamma \Rightarrow \varphi[0, \uparrow\sigma]$. By Fact 4.6, this is equivalent to showing $\Gamma \Rightarrow \varphi[x, \sigma]$ for a variable x which is fresh for Γ and $\dot\forall \varphi$. This follows from the initial assumption and inductive hypothesis 1.
2. Assuming $\forall\, \Gamma'\, \psi.\ \Gamma \subseteq \Gamma' \to \Gamma' ; (\dot\forall \varphi)[\sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi$, one has to deduce $t, \sigma \Vdash_\Gamma \varphi$ for every term t. Per inductive hypothesis 2, it suffices to show $\forall\, \Gamma'\, \psi.\ \Gamma \subseteq \Gamma' \to \Gamma' ; \varphi[t, \sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi$. Using the assumption, $\Gamma' \Rightarrow \psi$ reduces to $\Gamma' ; (\dot\forall \varphi)[\sigma] \Rightarrow \psi$, which can be derived via (AL). ☐

The lemma and its proof exhibit the structure of normalization by evaluation, a method mainly employed in the study of programming languages and type systems [5, 13]: Given a language $\mathtt{L}$ and a computational denotational semantics $\mathcal{D}$, one defines the mutually recursive functions reflect $\uparrow : \mathtt{L}^n \to \mathcal{D}$ and reify $\downarrow : \mathcal{D} \to \mathtt{L}^n$ which translate between the denotational semantics and the normal terms of the language. In the case of this lemma, the language consists of the derivation of $\Rightarrow$, the denotational semantics is given by the model $\mathcal{U}$ and the two statements 1. and 2. correspond to the functions $\uparrow$ and $\downarrow$, respectively.

Using the lemma to derive the completeness of $\Rightarrow_L$ with regards to $\Vdash_M$ is fairly straightforward.

Theorem 4.9 *For any context* Γ *and formula* φ*,* $\Gamma \Vdash_M \varphi \to \Gamma \Rightarrow_L \varphi$*.*

Proof Assume $\Gamma \Vdash_M \varphi$. Then $\iota \Vdash_\Gamma \Gamma \to \iota \Vdash_\Gamma \varphi$ in $\mathcal{U}$ with $\iota\, x = x$. Then, by Lemma 4.8.1, it suffices to show $\iota \Vdash_\Gamma \Gamma$ to deduce $\Gamma \Rightarrow_L \varphi$. With Lemma 4.8.2, this can be achieved by proving $\forall\, \Gamma'\, \theta.\, \Gamma \subseteq \Gamma' \to \Gamma' ; \psi \Rightarrow_L \theta \to \Gamma' \Rightarrow_L \theta$ for all $\psi \in \Gamma$, which holds by (Ctx). □

Proving the completeness of $\Rightarrow_E$ with regards to $\Vdash_E$ requires one additional step: proving that $\mathcal{U}$ is an exploding model when defined in terms of $\Rightarrow_E$. This is required to instantiate $\Gamma \Vdash_E \varphi$ with $\mathcal{U}$ in the proof of completeness.

Lemma 4.10 *The universal Kripke model for* $\Rightarrow_E$ *is an exploding model.*

Proof Per definition of $\mathcal{U}$, we have to show that $\Gamma \Rightarrow_E \dot{\bot} \to \Gamma \Rightarrow_E P\, t_1 \ldots t_{|P|}$ for arbitrary Γ. This follows by the (Exp) rule. □

Theorem 4.11 *For any context* Γ *and formula* φ*,* $\Gamma \Vdash_E \varphi \to \Gamma \Rightarrow_E \varphi$*.*

Proof Completely analogous to Theorem 4.9. □

4.3.2 Standard Models

The approach of the previous section does not cover standard models as $\mathcal{U}$ is a non-standard model. Recall that standard models have the property that $\bot_u^{\mathcal{K}} \to \bot$ in every world u. This is not the case in $\mathcal{U}$. For example, $\bot_{[\dot{\bot}]}^{\mathcal{U}} = [\dot{\bot}] \Rightarrow \dot{\bot}$ can be derived and is therefore not contradictory. The completeness proof with regards to standard models thus uses a slight variation of the universal model $\mathcal{U}$: instead of taking the worlds to be all contexts, we restrict the worlds to the consistent contexts, thereby forcing $\bot_\Gamma^{\mathcal{U}}$ to be a contradiction for any world Γ.

Definition 4.6 (Consistent model) *The consistent Kripke model $\mathcal{C}$ is a syntactic model. Its worlds $\mathcal{W}$ are the consistent contexts $\{\Gamma : \mathcal{L}(\mathfrak{F}) \mid \Gamma \not\Rightarrow_E \dot\bot\}$ with reachability being defined as $\subseteq$. Furthermore*

$$P^{\mathcal{C}}_{\Gamma}\, t_1 \dots t_{|P|} := \Gamma \Rightarrow_E P\, t_1 \dots t_{|P|} \qquad \bot^{\mathcal{C}}_{\Gamma} := \Gamma \Rightarrow_E \dot\bot$$

Fact 4.12 *$\mathcal{C}$ constitutes a standard model.*

Similar to standard Tarski models, this proof of completeness requires the stability of $\Rightarrow_E$. Whereas one application of stability was sufficient to prove standard completeness with regards to Tarski models, the proof for standard Kripke models requires it to be used multiple times. However, these applications always take the shape of the same principle, given below.

Lemma 4.13 *Under the stability of $\Rightarrow_E$ for finite theories, when proving $\Gamma \Rightarrow_E \varphi$, one may assume the consistency of Γ.*

Proof Assume that $\Gamma \not\Rightarrow_E \dot\bot \to \Gamma \Rightarrow_E \varphi$. As $\Rightarrow_E$ is stable, it thus suffices to show $\Gamma \Rightarrow_E \varphi$ under the assumption that $\Gamma \not\Rightarrow_E \varphi$. From $\Gamma \not\Rightarrow_E \varphi$ it follows that $\Gamma \not\Rightarrow_E \dot\bot$, as $\Gamma \Rightarrow_E \varphi$ could be derived by (Exp) otherwise. Thus $\Gamma \Rightarrow_E \varphi$ per assumption. □

All of the non-constructive principles we consider in our constructive analysis can be used to deduce this lemma. For the principle of double negation elimination, this is trivial. The synthetic Markov's principle can derive it as we have shown in Fact 4.7 that $\Rightarrow_E$ is enumerable and thus stable under it. As the enumeration we gave in Fact 4.7 can be adapted to any Church-Turing model of computation, the object Markov's principle can deduce it as well.

Using this principle, one can deduce the same two properties about the consistent Kripke model $\mathcal{C}$ as before.

Lemma 4.14 *Let Γ be a consistent context, σ a substitution and φ a formula. Under the object Markov's principle, the consistent Kripke model $\mathcal{C}$ has the following properties.*

1. $\sigma \Vdash_{\Gamma} \varphi \to \Gamma \Rightarrow \varphi[\sigma]$
2. $(\forall\, \Gamma' \psi.\ \Gamma \subseteq \Gamma' \to \Gamma' ; \varphi[\sigma] \Rightarrow \psi \to \Gamma' \Rightarrow \psi) \to \sigma \Vdash_{\Gamma} \varphi$

Proof We prove the following generalizations simultaneously by induction on φ.

1. $\forall\, \Gamma\, \sigma.\ \sigma \Vdash_{\Gamma} \varphi \to \Gamma \Rightarrow_E \varphi[\sigma]$
2. $\forall\, \Gamma\, \sigma.\ (\forall\, \Gamma' \psi.\ \Gamma \subseteq \Gamma' \to \Gamma' ; \varphi[\sigma] \Rightarrow_E \psi \to \Gamma' \Rightarrow_E \psi) \to \sigma \Vdash_{\Gamma} \varphi$

Most cases are completely analogous to those in Lemma 4.8. We only prove the case that requires an application of Lemma 4.13.

Case $\varphi \dot{\rightarrow} \psi$:

1. Assuming $\forall \Gamma'.\ \Gamma \subseteq \Gamma' \rightarrow \sigma \Vdash_{\Gamma'} \varphi \rightarrow \sigma \Vdash_{\Gamma'} \psi$, one has to derive that $\Gamma \Rightarrow (\varphi \rightarrow \psi)[\sigma]$. Per (IR), it suffices to show that $\Gamma, \varphi[\sigma] \Rightarrow_E \psi[\sigma]$, for which one can assume $\Gamma, \varphi[\sigma] \not\Rightarrow_E \dot{\bot}$ per Lemma 4.13. This allows for the application of inductive hypothesis 1. for ψ, leaving one having to show $\sigma \Vdash_{\Gamma,\varphi[\sigma]} \psi$. Applying the inductive hypothesis 2. for φ and the assumption, it suffices to show $\forall \Gamma'\,\theta.\ \Gamma, \varphi[\sigma] \subseteq \Gamma' \rightarrow \Gamma' ; \varphi[\sigma] \Rightarrow \psi[\sigma] \rightarrow \Gamma' \Rightarrow \psi[\sigma]$, which holds per (Ctx). □

Deriving completeness from the lemma above is fairly straightforward. However, another application of Lemma 4.13 is required to do so.

Theorem 4.15 *Under the stability of* $\Rightarrow_E$, $\forall\, \Gamma\, \varphi.\ \Gamma \Vdash_S \varphi \rightarrow \Gamma \Rightarrow_E \varphi$.

Proof Assume $\Gamma \Vdash_S \varphi$. We may also assume $\Gamma \not\Rightarrow_E \dot{\bot}$ to prove $\Gamma \Rightarrow_E \varphi$ by Lemma 4.13. As Γ is consistent, $\iota \Vdash_\Gamma \Gamma \rightarrow \iota \Vdash_\Gamma \varphi$ in $\mathcal{C}$ with $\iota\, x = x$. Then, by Lemma 4.14.1, it suffices to show $\iota \Vdash_\Gamma \Gamma$ to deduce $\Gamma \Rightarrow_E \varphi$. With Lemma 4.14.2, this can be achieved by proving $\forall\, \Gamma'\, \theta.\ \Gamma \subseteq \Gamma' \rightarrow \Gamma' ; \psi \Rightarrow_E \theta \rightarrow \Gamma' \Rightarrow_E \theta$ for all $\psi \in \Gamma$, which holds by (Ctx). □

We have so far only presented this completeness proof in terms of finite contexts. However, it can be modified to prove completeness for C-theories under any C-stability which is closed under finite theory extensions and substitutions. For this, the universal Kripke model's wrolds would have to be chosen to be the consistent C-theories. The proof would then proceed analogously to the proof for finite contexts with one important modification: The proof can only be applied to theories with "enough" free variables. Consider the case of statement 1. for the universal quantifier. One has to prove that $\mathcal{T} \Rightarrow_E \dot{\forall}\varphi[\sigma]$ under the assumption that $\mathcal{T} \Rightarrow_E \varphi[\sigma']$ for any substitution σ'. Here, one has to choose σ' as x, σ for a variable x which is fresh for $\varphi[\sigma]$ and $\mathcal{T}$. However, many theories use every single variable, thus making this step impossible. The easiest way of avoiding this is to prove completeness only for theories with an infinite amount of fresh variables and then deriving the completeness for all C-theories by generating fresh variables in the theories in questions using the substitution $[\lambda x.\ 2 * x]$.

We have shown that by assuming the object Markov's principle, one can prove the completeness of LJT with regards to standard Kripke models and explained how to extend this to C-theories when assuming a suitable C-stability. We now go on to demonstrate that these assumptions are necessary. That is, that completeness entails the stability of $\vdash_{CE}$.

The proof of this entailment employs a double-negation translation procedure. Such a procedure translates classically provable formulas into intuitionistically provable, classically equivalent formulas. While there exist multiple different variants of such procedures throughout the literature, we have chosen to use the version that was originally given by Gödel [25] for this proof.

Definition 4.7 (Double-Negation Translation) *We define the double-negation translation of a formula φ^N as follows:*

$$(P\,t_1 \ldots t_{|P|})^N := \dot\neg\dot\neg P\,t_1 \ldots t_{|P|} \quad \dot\bot^N := \dot\bot \quad (\varphi \dot\to \psi)^N := \varphi^N \dot\to \psi^N \quad (\dot\forall\varphi)^N := \dot\forall\varphi^N$$

Fact 4.16 *For any formula φ and theory $\mathcal{T}$, $\mathcal{T} \vdash_{CE} \varphi$ iff $\mathcal{T}^N \vdash_{IE} \varphi^N$.*

The proof of the equivalence also makes use of the fact that proofs in $\Rightarrow_B$ can be translated into proofs in $\vdash_{IB}$, which we prove in the next section in Lemma 4.19.

Theorem 4.17 *Let C be a theory class closed under theory $-^N$. Then standard Kripke completeness $\mathcal{T} \Vdash_S \varphi \to \mathcal{T} \Rightarrow_E \varphi$ for all C-theories entails C-stability.*

Proof Assume Kripke completeness for all C-theories and $\neg\neg\mathcal{T} \vdash_{CE} \varphi$ for a $C\,\mathcal{T}$. We prove $\mathcal{T} \vdash_{CE} \varphi$ by applying the object double-negation in $\vdash_{CE}$ and double-negation translation, thus being left with proving $\mathcal{T}^N \vdash_{IE} \dot\neg\dot\neg\varphi^N$. As proofs in $\Rightarrow_E$ can be translated into proofs in $\vdash_{IE}$ (Corollary 4.20) and $\Rightarrow_E$ is complete for $\mathcal{T}^N$ per assumption, it suffices to show $\mathcal{T}^N \Vdash_S \dot\neg\dot\neg\varphi^N$. Thus, let $\mathcal{K}$ be a standard Kripke model on some domain D, ρ an environment in D and w a world of $\mathcal{K}$. We assume $\rho \Vdash_w \mathcal{T}^N$ and $\rho \Vdash_w \dot\neg\varphi^N$ to prove $\bot$ as $\mathcal{K}$ is a standard model. As we are deriving a contradiction, $\neg\neg\mathcal{T} \vdash_{CE} \varphi$ gives us $\mathcal{T} \vdash_{CE} \varphi$ and by double-negation translation and soundness of $\vdash_{IE}$ (Theorem 4.4) $\rho \Vdash_w \varphi^N$. This is a contradiction to $\rho \Vdash_w \dot\neg\varphi^N$ as $\mathcal{K}$ is a standard model. □

This can again be used to derive precise statements of different variants of standard Kripke completeness.

Corollary 4.18

1. $\mathcal{T} \Vdash_S \varphi \to \mathcal{T} \Rightarrow_E \varphi$ *for arbitrary $\mathcal{T}$ entails double-negation elimination*
2. $\mathcal{T} \Vdash_S \varphi \to \mathcal{T} \Rightarrow_E \varphi$ *for enumerable $\mathcal{T}$ entails the synthetic Markov's principle*
3. $\mathcal{T} \Vdash_S \varphi \to \mathcal{T} \Rightarrow_E \varphi$ *for finite $\mathcal{T}$ entails the object Markov's principle*

4.4 Semantic Normalization

Normal proofs are of great interest to proof theory, as they exhibit a very regular structure that makes analyzing them much simpler. This section demonstrates how to employ the completeness proofs of the previous section to define a semantic normalization procedure for proofs $\Gamma \vdash_{IB} \varphi$.

A proof is considered normal if it does not contain any eliminations of proofs ending in an introduction rule. Through the lens of the Curry-Howard isomorphism, this means that the term-representation of the proof is in a normal-form and can thus not be reduced further.

Definition 4.8 (Normal proofs in $\vdash$) *A proof $\Gamma \vdash \varphi$ is considered normal if it does not contain an application of an introduction rule, followed by an application of the corresponding elimination rule, that is, if it doesn't contain a subderivation of the shapes*

$$\text{IE}\,\dfrac{\text{II}\,\dfrac{\dfrac{\cdots}{\Gamma, \psi \vdash \varphi}}{\Gamma \vdash \psi \dot{\rightarrow} \varphi} \qquad \dfrac{\cdots}{\Gamma \vdash \psi}}{\Gamma \vdash \varphi} \qquad\qquad \text{AllE}\,\dfrac{\text{AllI}\,\dfrac{\dfrac{\cdots}{\uparrow\Gamma \vdash \varphi}}{\Gamma \vdash \dot{\forall}\varphi}}{\Gamma \vdash \varphi[t]}$$

An important property of LJT is that it only allows for normal proofs. Recall that it is divided between derivations with conclusions of the shapes $\Gamma \Rightarrow \varphi$ and $\Gamma\,;\psi \Rightarrow \varphi$. The left rules, which correspond to the elimination rules of the ND system, can only be applied to the focused formula of derivations $\Gamma\,;\psi \Rightarrow \varphi$. However, as the only way of bringing a formula into focus is by taking it from the context Γ, no application of a left rule will ever be preceded by a right rule.

Because all proofs in $\Rightarrow$ are inherently normal, they can be transformed into proofs in $\vdash$ which are normal as well.

Lemma 4.19 *For every proof $\Gamma \Rightarrow \varphi$ there exists a normal proof $\Gamma \vdash \varphi$.*

Proof We prove this by defining two mutually recursive proof transformation functions $f : \Gamma \Rightarrow \varphi \rightarrow \Gamma \vdash \varphi$ and $g : \Gamma\,; \psi \Rightarrow \varphi \rightarrow \Gamma \vdash \psi \rightarrow \Gamma \vdash \varphi$.

$$f\left[\text{Ctx}\,\dfrac{\Gamma\,;\varphi \Rightarrow \psi \quad \varphi \in \Gamma}{\Gamma \Rightarrow \psi}\right] := \; g\,[\,\Gamma\,;\varphi \Rightarrow \psi\,]\left[\text{Ctx}\,\dfrac{\varphi \in \Gamma}{\Gamma \vdash \varphi}\right]$$

$$f\left[\text{IR}\,\dfrac{\Gamma, \varphi \Rightarrow \psi}{\Gamma \Rightarrow \varphi \dot{\rightarrow} \psi}\right] := \text{II}\,\dfrac{f\,[\,\Gamma, \varphi \vdash \psi\,]}{\Gamma \vdash \varphi \dot{\rightarrow} \psi}$$

$$f\left[\text{AllR}\,\dfrac{\uparrow\Gamma \Rightarrow \varphi}{\Gamma \Rightarrow \dot{\forall}\varphi}\right] := \text{AllI}\,\dfrac{f\,[\,\uparrow\Gamma \vdash \varphi\,]}{\Gamma \vdash \dot{\forall}\varphi}$$

$$f\left[\textsc{Exp}\frac{\Gamma \Rightarrow_E \dot{\bot}}{\Gamma \Rightarrow_E \varphi}\right] := \textsc{Exp}\frac{f\,[\,\Gamma \vdash_{IE} \dot{\bot}\,]}{\Gamma \vdash_{IE} \varphi}$$

$$g\left[\textsc{Ax}\frac{}{\Gamma\,;\varphi \Rightarrow \varphi}\right][\,\Gamma \vdash \varphi\,] := \Gamma \vdash \varphi$$

$$g\left[\text{IL}\frac{\Gamma \Rightarrow \varphi \quad \Gamma\,;\psi \Rightarrow \theta}{\Gamma\,;\varphi \dot{\rightarrow} \psi \Rightarrow \theta}\right][\,\Gamma \vdash \varphi \dot{\rightarrow} \psi\,] := g[\,\Gamma\,;\psi \vdash \theta\,]\left[\text{IE}\frac{\Gamma \vdash \varphi \dot{\rightarrow} \psi \quad f\,[\,\Gamma \Rightarrow \varphi\,]}{\Gamma \vdash \psi}\right]$$

$$g\left[\textsc{AllL}\frac{\Gamma\,;\varphi[t] \Rightarrow \psi}{\Gamma\,;\dot{\forall}\varphi \Rightarrow \psi}\right][\,\Gamma \vdash \dot{\forall}\varphi\,] := g\,[\,\Gamma\,;\varphi[t] \Rightarrow \psi\,]\left[\textsc{AllE}\frac{\Gamma \vdash \dot{\forall}\varphi}{\Gamma \vdash \varphi[t]}\right]$$

The proofs produced by these functions are normal. Applications of the (IE) and (AllE) rules are only produced by the function g when applied to the (IL) and (AllL) rules, respectively. However, those applications are always preceded by the proof passed in as the second argument of g, which only consist of application of the (AllE, IE) and (Ctx) rules. Thus no application of an elimination rule will be preceded by a matching introduction rule, meaning all proofs produced by f are normal. □

Corollary 4.20 *For every proof $\mathcal{T} \vdash_{IB} \varphi$, there is a proof $\mathcal{T} \Rightarrow_B \varphi$.*

The first normalization procedure was given by Gentzen [21, 22] and took the shape of a direct manipulation of derivation trees. The translation procedure we have given above, together with the completeness proofs of Section 4.3 give rise to a normalization procedure as well. Contrasting Gentzen's syntactic approach, this procedure is considered a semantic normalization procedure, as the normalization takes place in the (semantic) completeness proofs of Section 4.3.

Corollary 4.21 *For any proof $\Gamma \vdash_{IB} \varphi$, there is a normal proof $\Gamma \vdash_{IB} \varphi$.*

Proof Let $\Gamma \vdash_{IB} \varphi$. Then per Theorem 4.4, $\Gamma \Vdash_C \varphi$ with $C = E$ if $B = E$ and $C = BL$ if $B = L$. With Theorem 4.9 or Theorem 4.11, $\Gamma \Rightarrow_B \varphi$ can be derived. Thus, with Lemma 4.19, one arrives at a normal proof of $\Gamma \vdash_{IB} \varphi$. □

The translation procedure from Lemma 4.19 can also be used to derive the soundness of $\Rightarrow_B$ from the soundness of $\vdash_{IB}$ with regards to Kripke models.

Corollary 4.22 *Let C be a constraint and and B a LJT-variant such that all models satisfying C are exploding if $B = E$. Then $\Gamma \Rightarrow_B \varphi \rightarrow \Gamma \Vdash_C \varphi$ for all Γ and φ.*

Proof Let $\Gamma \Rightarrow_B \varphi$. Per Lemma 4.19, this means $\Gamma \vdash_{IB} \varphi$. Thus, by Theorem 4.4, $\Gamma \Vdash_C \varphi$. □

4.5 Conclusion

In this chapter, we have analyzed the completeness of the normal sequent calculus LJT with regards to Kripke semantics with three different notions of models: standard, exploding and minimal. We have proven completeness with regards to exploding models $\Vdash_E$ and minimal models $\Vdash_M$ in a constructive manner. Further, we have demonstrated that completeness with regards to standard Kripke models $\Vdash_S$ cannot be proven constructively by showing that it entails the stability of $\vdash_{CE}$. Lastly, we have demonstrated how proofs of $\Rightarrow_B$ can be translated into normal proofs in $\vdash_{IB}$, thereby extracting a semantic normalization procedure from the constructive completeness proofs. The results are summed up by Fig. 4.2. Normal arrows denote the usual implication, while dashed arrows are implications which are equivalent to the stability of $\vdash_{CE}$.

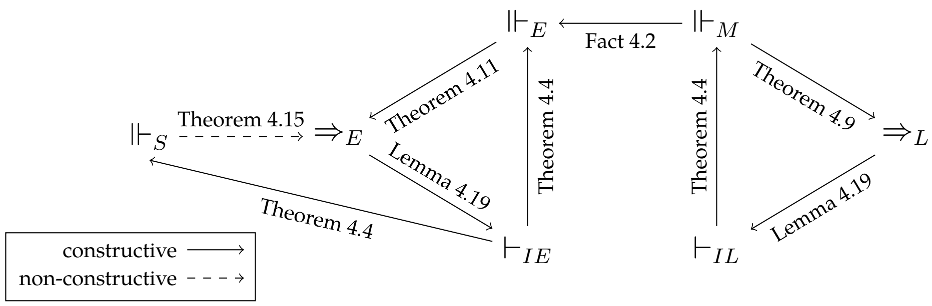

Figure 4.2: Results for Kripke completeness

Chapter 5

Dialogue Semantics

So far, we have analyzed two kinds of semantics: Tarski and Kripke model semantics. We showed that neither of their canonical formulations, the validity with regards to standard models, admit a constructive proof of completeness. In this section, we present a game semantics that does admit a constructive completeness proof for full intuitionistic first-order logic.

We begin this chapter with a rather informal overview of dialogues in Section 5.1. We then proceed by giving a formal account of abstract intuitionistic E-dialogues in the style of Sørensen and Urzyczyn [61], which we prove sound and complete in Section 5.2. Based on this abstract result, we deduce the completeness for full intuitionistic first-order logic from this abstract result (Section 5.3). We close the chapter by proving the soundness and completeness of enumerable abstract intuitionistic D-dialogues in Section 5.4.

5.1 An Overview of Dialogues

As dialogue semantics are not as well known as the semantics we have analyzed previously, we begin this chapter with an overview of their origins in Lorenzen's material dialogues as well as their contemporary presentation as formal dialogues.

5.1.1 Material Dialogues

Material dialogues were introduced by Paul Lorenzen in his talks "Logik und Agon" ("Logic and contest") of 1958 [48] and "Ein dialogisches Konstruktivitätskriterium" ("A dialogical criterion of constructivity") of 1959 [49]. They were originally developed to give a new, more understandable justification of intuitionistic validity.

He begins "Logik und Agon" by tracing back the origins of logic to the necessity of imposing rules on debates to withstand sophistic methods of argument. These rules transformed free debates into proper contests of wit (greek "agon" = struggle or contest). He contrasts these agonistic roots with the modern conception of logic as systems of rules describing how to derive true sentences from true sentences. He posits that these logical rules are viewed as a gift from god to man, thus eliminating the necessity of justifying them. However, he points out that Brouwer's critiques of classical logic have revealed that, curiously, god seems to have given different sets of these rules to different people. He deems the explanations of this phenomenon that had been provided at the time insufficient, blaming this state on the modern conception of logic as an unjustifiable gift from god. He thus puts forward his own justification of the intuitionistic conception of validity.

We follow Lorenzen's lead by giving a very informal dialogic account of logic. The contents of the following section are synthesized from both talks. A more formal presentation of dialogues is deferred to Section 5.1.2.

At the basis of his conception of intuitionistic logic, Lorenzen posits games whose outcomes can be observed and agreed upon. The example he gives is the derivability of sentences of a grammar as witnessed by derivations. However, he notes that any games with operational content ("operative Bedeutung"), even physical experiments, can take on this role. These observable games correspond to the atomic formulas of the logic (such as the predicates of first-order logic).

He then gives his justification of intuitonistic validity in terms of a contest between two players, the proponent P and the opponent O, which he calls the metagame. This contest always begins with P making some sort of claim. The most basic example of such a claim would be P claiming to be able to win one of the underlying games, for example by claiming that she can derive a specific sentence in the grammar. O can now challenge her on that claim, resulting in P winning the metagame if the outcome of that underlying game is in her favor.

Lorenzen now goes on to characterize the meaning of the logical connectives in terms of rules of the metagame. Consider the claim $\varphi \to \psi$. If P makes this claim, O can only challenge this claim by claiming φ herself. P now has two ways of responding to this. She may defend her original claim by claiming ψ or she can challenge O's claim of φ. For a more concrete example, consider the claim $\vdash s \to \vdash t$ where $\vdash s$ stands for s being derivable in a certain grammar. If P claims $\vdash s \to \vdash t$, O will challenge her by claiming $\vdash s$ herself. P can now respond by challenging O's claim, thereby forcing her to present a derivation of the sentence s in the grammar. Now P is forced to defend her claim of $\vdash t$. However, she can use the derivation of s that O presented her with to do this. The dialogic explanation of $\varphi \to \psi$ is thus the claim of "if φ is demonstrated to me, I can demonstrate ψ".

The dialogic accounts of the other connectives are similar. Claiming $\neg\varphi$ means claiming $\varphi \to \bot$ where $\bot$ is some underlying game that both players agree cannot be won. If one of the players of the metagame claims $\forall x.\ \varphi[x]$, her opponent may challenge $\varphi[t]$ for any object t of the domain of the underlying game. Dually, if a player claims $\exists x.\ \varphi[x]$, she can choose which object t she wants to defend the claim $\varphi[t]$ with once being challenged. The rules for $\varphi \wedge \psi$ and $\varphi \vee \psi$ are analogous to those of $\forall x.\ \varphi[x]$ and $\exists x.\ \varphi[x]$.

In "Ein dialogisches Konstruktivitätskriterium", Lorenzen also gives a dialogic account of inductive sets. Given a set X inductively defined by

$$\varphi[x] \to x \in X \qquad x \in X \wedge \psi[x, y] \to y \in X$$

the claim of $z \in X$ can either be defended by claiming $\varphi[z]$ or by choosing some x and claiming $x \in X \wedge \psi[x, z]$. If a player chooses to do the latter, she is also required to fix some natural number n, limiting how many times she is allowed to appeal to the inductive definition to prove her claim, thereby guaranteeing the justification of $z \in X$ to be well-founded.

Having given an account of all of the logical connectives, Lorenzen now moves on to justifying the intuitionistic position from the dialogical standpoint. For this, he introduces the concept of generally admissible claims ("allgemein-zulässig"). Generally admissible claims can be seen as claim-schemata which can be defended regardless of the choice of the specific claims as well as the underlying game. An example of a generally admissible claim is $(\varphi \to \psi) \to (\psi \to \theta) \to (\varphi \to \theta)$.

Lorenzen now goes on to explain why the classical principle of double negation elimination $\neg\neg\varphi \to \varphi$ is not generally admissible. For this, he considers P claiming the scheme for an underlying game φ for which only O knows how to win it. An example of such a game would be a sentence s of which only O knows how to derive it in the grammar. O thus challenges her by claiming $\neg\neg\varphi$. As P does not know how to win φ, she can only challenge O by claiming $\neg\varphi$. However, as O knows how to win φ, she can defend against P's challenge, leaving P in a losing position as she would now have to win the unwinnable $\bot$.

This explanation relies on an important rule of the metagame that Lorenzen never fully makes explicit in these two talks: A player may only ever defend against the last challenge that was leveled against her which she has not yet defended herself against. If one was to lift this restriction, P could win the previous metagame: Once O demonstrates how to win φ, P once again defends her original claim that $\neg\neg\varphi \to \varphi$ by winning φ. As this leaves O without anything to challenge, P has thus won the metagame of $\neg\neg\varphi \to \varphi$. Indeed, classical and intuitionistic dialogical accounts differ precisely by the absence or presence of this rule, respectively.

For a more indepth account of the history of dialogue semantics, we refer the reader to Krabbe's excellent essay on this topic [36].

5.1.2 Formal Dialogues

Formal dialogues were introduced by Lorenzen's student Kuno Lorenz in his PhD thesis of 1961 [46]. They differ from material dialogues as their claims are purely syntactical formulas, without atoms being represented by an underlying game. We now give a formal account of formal dialogues in the style that is found throughout the literature as exemplified by Felscher [15]. The definitions of this section have not been formalized in Coq, as we use a different presentation of dialogues in the completeness proofs of the next sections.

We begin by formally defining which attacks can be leveled against which formulas and how a player can defend against each of those attacks. These definitions are often referred to as the "local rules" in the literature. This should be viewed as more formal rendering of the account given in Section 5.1.1. Formally, attacks are members of the type $\mathcal{A}$ defined as below.

$$\mathcal{A} : \mathbb{T} := A_{\dot\bot} \mid A_{\dot\rightarrow}\, \varphi \mid A_L \mid A_R \mid A_{\dot\vee} \mid A_t\, \varphi \mid A_{\dot\exists} \qquad (\varphi : \mathfrak{F}, t : \mathfrak{T})$$

Which formula can be attacked by which attack is defined by the three-place attacks relation $\rhd : \mathfrak{F} \to \mathcal{A} \to \mathcal{O}(\mathfrak{F}) \to \mathbb{P}$. When $a \mid \psi \rhd \varphi$, we say a is an attack on φ and call ψ the admission. For each attack a, we define which formulas can be admitted to defend against a as a set of formulas $\mathcal{D}_a : 2^{\mathfrak{F}}$ called a's defense set. If $\varphi \in \mathcal{D}_a$, we say φ is a defense against a. The attack relation and defense sets are specified in Fig. 5.1. For the sake of readability, we write $a \rhd \varphi$ for $a \mid \emptyset \rhd \varphi$.

$$\begin{array}{rl}
A_{\dot\bot} \rhd \dot\bot & \mathcal{D}_{A_{\dot\bot}} = \{\} \\
A_{\dot\rightarrow}\, \psi \mid \ulcorner\varphi\urcorner \rhd \varphi \dot\rightarrow \psi & \mathcal{D}_{A_{\dot\rightarrow}\,\psi} = \{\psi\} \\
A_L \rhd \varphi \dot\wedge \psi & \mathcal{D}_{A_L} = \{\varphi\} \\
A_R \rhd \varphi \dot\wedge \psi & \mathcal{D}_{A_R} = \{\psi\} \\
A_{\dot\vee} \rhd \varphi \dot\vee \psi & \mathcal{D}_{A_{\dot\vee}} = \{\varphi, \psi\} \\
A_t\, \varphi \rhd \dot\forall\varphi & \mathcal{D}_{A_t\,\varphi} = \{\varphi[t]\} \\
A_{\dot\exists} \rhd \dot\exists\varphi & \mathcal{D}_{A_{\dot\exists}} = \{\varphi[t] \mid t : \mathfrak{T}\}
\end{array}$$

Figure 5.1: Attack relation and defense sets

Formal dialogues can be viewed as a two-player game between the proponent P and the opponent O. As in the material dialogues, P initially makes some claim φ. The goal of O is to attack the claim in such a way that P will be unable to defend it. P on the other hand aims to make O admit that φ does indeed hold.

Dialogues are formalized as sequences of player moves $\mathcal{L}((\mathcal{A} + \mathfrak{F}) \times \mathbb{N})$. Here, the value $(L\, a, n)$ represents a player attacking an admission made at position n and $(R\, \varphi, n)$ represents a player defending against an attack made against her at position n. As dialogues are turn-taking games, every even position will be a proponent move and every odd position will be an opponent move.

To define which sequences constitute dialogues, we first have to define what it means to admit a formula. There are two situations in which we say that a formula φ is admitted at position n of a dialogue: Firstly, if the value at position n is $(R\, \varphi, m)$ for some m, meaning a player defended against an attack at position m by admitting φ. Secondly, if it is $(L\, a, m)$ where ψ is admitted at position m and $a \mid \ulcorner\varphi\urcorner \rhd \psi$, meaning that φ was admitted by a player in the course of attacking ψ.

A sequence constitutes a dialogue if it fulfills the following three properties:

1. It begins with a tuple $(R\, \varphi, 0)$, representing the proponent's initial claim of the non-atomic formula φ.
2. Every attack move attacks an admission made by the other player. Formally: If the value at position n is $(L\, a, m)$ then a formula φ is admitted at position m with $a \mid \psi \rhd \varphi$ for some ψ. Further, m and n are of different parity.
3. Every defense move (except the proponent's initial claim) defends against an attack made by the other player. Formally: If the value at non-zero position n is $(R\, \varphi, m)$ then the value at position m is $(L\, a, k)$ with $\varphi \in \mathcal{D}_a$. Further, m and n are of different parity.

This account of dialogues alone does not yield a proper semantics of of first-order logic. We thus introduces different variants of dialogues, which further restrict what constitutes a valid move. The first of these variants we define are the classical D-dialogues. A dialogue is a classical D-dialogue if

(CD1) P may only admit an atomic formula after O has already admitted it. That is, if the atomic formula φ is admitted at the even position n, then there is an odd position $m < n$ at which φ is admitted as well.

(CD2) Any attack may be defended against at most once. That is, if the value at position n is $(R\, \varphi, m)$, then there is no position $n' < n$ with the value $(R\, \psi, m)$.

(CD3) Any admission made by P may only be attacked once. That is, if the value at an odd position n is $(L\, a, m)$, then there is no position $n' < n$ with the value $(L\, a, m)$.

The rule (CD1) is required as formal dialogues are not based on underlying games. Instead, the P's goal in a dialogue is now to make O admit her initial claim. Thus, P may only admit the atomic formulas that O already admitted to hold. Rule (CD2) means that players have to "commit" on their defenses. For example, a player defending the claim $\varphi \dot{\vee} \psi$ has to decide which side of the disjunction to choose and cannot go back on that choice one she announced it. The rule (CD3) only makes sense when considering the winning condition of dialogues: A dialogue is considered to be won by P if it ends in a move by P and O can't make any further moves. Thus, rule (CD3) is in place to guarantee that O cannot stall the dialogues indefinitely by repeatedly attacking the same admission made by P.

Intuitionistic D-dialogues can be defined in terms of classical D-dialogues. For this, one simply has to extend the rule set with the rule discussed at the end of the previous section.

(ID4) A player may only ever defend against the last attack she has not yet defended against. That is, if position n is $(R\,\varphi, m)$ there exists no position $m < m'$ of opposite parity of n with value $(L\,a, k)$ for which there exists no $n' < n$ with value $(R\,\psi, m')$.

The concrete formulation of the rule is very technical. A more intuitive account of this rule is that the open attacks against a player are stored in a stack. New attacks leveled against the player are pushed at the top of the stack. Under this rule, players may only defend against the attack at the top of the stack, thereby popping them.

In addition to the D-dialogues we have already defined, there is a second general variant of dialogues that also plays an important role in the study of formal dialogues: E-dialogues. Essentially, E-dialogues are D-dialogues with a harsh restriction placed upon the opponent: She can only react to the proponent's previous move. Maybe somewhat surprisingly, E-dialogues are not "easier" than D-dialogues: being able to win E- and D-dialogues is equivalent. The relationship between E- and D-dialogues is similar to that between arbitrary and cut-free derivations in derivation systems. As E-dialogues exhibit a simpler structure, meta-logical results, such as soundness and completeness, tend to be proven only for E-dialogues and then extended to D-dialogues on the basis of their equivalence. An E-dialogue can be defined in terms of D-dialogues, both classical and intuitionistic, by adding one more rule.

(E5) The opponent may only react to the proponent's previous move. That is, the value at any uneven position n has to be $(L\,a, n-1)$ or $(R\,\varphi, n-1)$.

We have now laid out all common variants of dialogues. We thus move on to discuss how to define the validity of a formula in terms of them. We have already defined what it means to win a dialogue. However, simply being able to win sometimes is not sufficient for a formula to be valid. Consider the formula $\dot{\top} \dot{\wedge} \dot{\bot}$. Certainly, the proponent can win this dialogue: After she claims the formula, O attacks with A_L, to which P responds with admitting $\dot{\top}$, leaving O without any further recourse, thus winning the dialogue. However, it is obvious that P was only able to win as O was not playing “smart”. Had O attacked with A_R instead, P would have lost. To avoid these situations, we require P to be able to win regardless of the actions taken by O. This possibility to always win is formally defined as the existence of a winning strategy.

A winning strategy is usually represented as a tree of moves $\mathcal{A} + \mathfrak{F}$. Every node at which the proponent has to make the next move, meaning every node on an odd level, has exactly one subtree. This means that the player has to choose and commit to which move she wants to perform in any situation that the opponent can force her into. Every node on an even level, meaning those in which the opponent makes the next move, have one subtree for every possible move the opponent could make at that point in the dialogue. For such a tree to constitute a winning strategy, every full path has to be a dialogue (of the specified variant) which is won by P. Thus, such a winning strategy can only be constructed if P can always react to any possible move of O in such a way that she can still win the dialogue.

Overall, we define intuitionistic validity of a formula φ as the existence of a winning strategy consisting of intuitionistic E-dialogues (or equivalently D-dialogues) which starts with P claiming φ. Classical validity is defined analogously.

5.2 Generalized Intuitionistic E-Completeness

Instead of giving a direct proof of the soundness and completeness of full intuitionistic first-order logic with regards to dialogues, we introduce an intermediate step inspired by work of Sørensen and Urzyczyn [61]. They prove soundness and completeness between classical E-dialogues generalized over the local rules and two special sequent calculi LKD and LKd. The central insight behind this proof strategy is that E-dialogues and sequent calculi share a lot of structure.

We begin this section by giving a formal account of dialogues that makes better use of the strengths of constructive type theory than the usual presentation in Section 5.2.1. We then continue by defining the dialogical sequent calculus LJD which is similar to the LKD of Sørenzen and Urzyczyn in Section 5.2.2. We end this section by proving that LJD is sound and complete with regards to generalized intuitionistic E-dialogues.

5.2.1 Generalized Intuitionistic E-Dialogues

The local rules of a dialogue are the rules that govern which formula may be attacked by which attack and which formula can be used to defend against which attack. We are going to define intuitionistic E-dialogues which are generalized over the specific choice of local rules.

We begin by defining rule sets. The rules for first-order logic we introduced in Section 5.1.2 can be seen as one such rule set. Notably, the set $\mathfrak{F}^*$ for first-order logic is the set of predicate formulas $\{P\, t_1 \ldots t_{|P|} \mid P : \mathcal{P}, t_1, \ldots, t_{|P|} : \mathfrak{T}\}$.

Definition 5.1 (Rule set) *A rule set $(\mathfrak{F}, \mathfrak{F}^*, \mathcal{A}, (\mathcal{D}_a)_{a \in \mathcal{A}}, \rhd)$ of local rules consists of a type of formulae $\mathfrak{F}$, a set of formulae $\mathfrak{F}^*$ considered atomic, a type of attacks $\mathcal{A}$, the defenses $\mathcal{D}_\_ : \mathcal{A} \to 2^{\mathfrak{F}}$, a family of sets of formulae indexed by the attacks, and the attacks relation $\lhd : \mathfrak{F} \to \mathcal{A} \to \mathcal{O}(\mathfrak{F}) \to \mathbb{P}$.*

Whenever $a \mid \psi \rhd \varphi$, we say "a is an attack on φ" and call ψ the admission. If $\varphi \in \mathcal{D}_a$, we call φ a defense against a.

We have outlined how formal dialogues are usually defined in the literature in Section 5.1.2. However, that presentation is not very well-suited for a formalization in Coq. Take, as an example, its reliance on accessing finite sequences by natural number indexes, a partial operation. As Coq only allows the definition of total functions, this could only be formalized through one of the many renditions of sequence indexing as a total function, for example by returning option values. However, this would introduce a considerable overhead, both in the statements of definitions and lemmas as well as the proofs.

To avoid these complications, we have opted for a presentation of dialogues that is better suited to a type theoretical setting. Instead of representing dialogues as sequences of moves, we formalize them as state transition systems. The type of states is $\mathcal{L}(\mathfrak{F}) \times \mathcal{A}$, the type of pairs of the list of the opponent's admissions and the last attack the opponent made (the opponent's challenge). Note that this simple state representation can only be used when defining E-dialogues as D-dialogues require more information to determine which moves are valid.

The proponent's moves are represented by the type $\mathcal{P}$. Its values are either $PA\, a$, denoting the proponent attacking one of the opponent's admissions with attack a, or $PD\, \varphi$, denoting the proponent defending against the opponent's last attack. We denote by $s \rightsquigarrow_p m$ that the proponent may perform move m in the state s.

Definition 5.2 (Proponent move)

1. *The type $\mathcal{P}$ describing the proponent's last move is defined as follows*

$$\mathcal{P} : \mathbb{T} := PA\, a \mid PD\, \varphi \qquad (a : \mathcal{A}, \varphi : \mathfrak{F})$$

2. *A formula φ is justified under opponent admissions A_o if $\varphi \in \mathfrak{F}^* \to \varphi \in A_o$. we take "justified $A_o\, \varphi$" for $\varphi : \mathcal{O}(\mathfrak{F})$ to mean "justified $A_o\, \psi$" if $\varphi = \ulcorner \psi \urcorner$ and $\bot$ otherwise.*

3. *The proponent's move relation $\rightsquigarrow_p$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{A} \to \mathcal{P} \to \mathbb{P}$ is defined by an inference system as follows*

$$\text{PA}\,\frac{\varphi \in A_o \quad a \mid \psi \rhd \varphi \quad \textit{justified } A_o\, \psi}{(A_o, c) \rightsquigarrow_p PA\, a} \qquad \text{PD}\,\frac{\varphi \in \mathcal{D}_c \quad \textit{justified } A_o\, \varphi}{(A_o, c) \rightsquigarrow_p PD\, \varphi}$$

Recall that in formal dialogues, the proponent may only admit atomic formulas if the opponent has admitted them already. Thus, we define the notion of justified formulas in such a way that non-atomic formulas are justified regardless of the opponent's admissions and atomic formulas are only justified if the opponent has already admitted them.

With this in mind, the definition of $\rightsquigarrow_p$ is fairly intuitive. The rule (PA) means the proponent may attack any opponent admission φ as long as the admission of that attack is justified. The rule (PD) allows the proponent to defend against the opponent's challenge by admitting a formula, so long as that formula is justified.

In E-dialogues, the opponent's moves are restricted to reactions to the proponent's previous move. We denote by the relation $s\,;\,m \rightsquigarrow_o s'$ that the opponent may react to the proponent move m by transforming the current state s into the next state s'.

Definition 5.3 (Opponent move) *The following inference system defines the opponent's move relation* $\rightsquigarrow_o$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{A} \to \mathcal{P} \to \mathcal{L}(\mathfrak{F}) \times \mathcal{A} \to \mathbb{P}$

$$\text{OA}\,\frac{c' \mid \psi \rhd \varphi}{(A_o, c)\,;\, PD\, \varphi \rightsquigarrow_o (\psi :: A_o, c')} \qquad \text{OD}\,\frac{\varphi \in \mathcal{D}_a}{(A_o, c)\,;\, PA\, a \rightsquigarrow_o (\varphi :: A_o, c)}$$

$$\text{OC}\,\frac{a \mid \ulcorner\psi\urcorner \rhd \varphi \quad c' \mid \theta \rhd \psi}{(A_o, c)\,;\, PA\, a \rightsquigarrow_o (\theta :: A_o, c')}$$

In a slight abuse of notation, $\emptyset :: A_o \triangleq A_o$ *and* $\ulcorner\varphi\urcorner :: A_o \triangleq \varphi :: A_o$.

If the proponent has just defended herself, the rule (OA) prescribes that the opponent may only react by attacking that defense. This replaces the challenge with that just issued by the proponent and also potentially extends the opponents admissions by the admission of that attack. If the proponent's last move was an attack on one of the opponent's admissions, the opponent can react in two different ways. She may defend herself against that attack by admitting a formula from the attack's defense set as described in rule (OD). According to (OC), if the proponent has made an admission along with her attack, the opponent may also counter by challenging that admission. The effects of a counter on the state are analogous to those of (OA), as both moves issue a new challenge.

Putting the definition of $\rightsquigarrow_p$ and $\rightsquigarrow_o$ together, it becomes clear that these definitions constitute a state transition system. Each round of the game constitutes one state transition. That is, starting in a state s, the proponent first chooses her move m according to $s \rightsquigarrow_p m$. The opponent then reacts to her move via $s\,;\,m \rightsquigarrow_o s'$, thus inducing transition from state s to state s'.

This conception of dialogues as state transition systems admits an elegant definition of winning state. Intuitively, a state s is winning if the proponent can choose a move m such that any state s' which the opponent can transition into via $s\,;\,m \rightsquigarrow_o s'$ is still winning. Vacuously, this definition means that any state in which the proponent can perform a move to which the opponent cannot respond is winning. This definition of winning states is a slight variation of the notion of strongly normalizing states in state transition systems: A state s is considered strongly normalizing if any state s' with $s \rightsquigarrow s'$ is strongly normalizing.

Definition 5.4 (Winning states) *A state s is winning if Win s can be derived in the following inference system*

$$\frac{s \rightsquigarrow_p m \quad \forall s'.\, s\,;m \rightsquigarrow_o s' \rightarrow \mathit{Win}\, s'}{\mathit{Win}\, s}$$

In the usual presentation of dialogues of Section 5.1.2, the proponent's initial challenge did not fit in neatly with the other moves of the game, as exemplified by the first and third dialogue conditions. This does not change if dialogues are formalized as state transition systems. Hence, we treat the first round of a dialogue differently in our definition of validity.

Definition 5.5 (E-Validity) *A formula φ is E-valid, denoted by $\Vdash_E \varphi$, if it is not atomic and for any attack $c \mid \psi \rhd \varphi$ on it, the resulting state $([\psi], c)$ is winning.*

That this definition aligns with the dialogues as defined in Section 5.1.2 requires some further elaboration. E-valid formulas cannot be atomic, just as the initial claim of the traditional account is required to be non-atomic. Now consider the winning criterion: In the traditional account, the first move by the opponent can only be an attack on the proponent's initial claim, as the proponent has not yet made any attack against the opponent that she could defend herself against. Thus any state corresponding to the sequence describing only the first round of the dialogue would be of the shape $([\psi], c)$ where c is the opponent's attack on the claim and $[\psi]$ might be the first admission the opponent is forced to make in the course of launching c.

5.2.2 Dialogical Sequent Calculus

The most remarkable aspect of the work by Sørensen and Urzyczyn [61] is the sequent calculus LKD they use in their soundness and completeness proofs. It is defined in such a way that a proof of a formula φ is exactly isomorphic to a winning strategy for the dialogue with the claim of φ. As their LKD is designed to easily facilitate completeness and soundness proofs for classical propositional logic, it is not directly suited for our proof. We thus adopt their sequent calculus into an intuitionistic version LJD which is suitable for first-order dialogues. Notably, it loses the usual finitary character of deduction systems, as winning strategies for first-order dialogues are often infinitely branching and thus a deduction system that aims to be isomorphic to them will have to be the same.

Deductions in LJD are defined by the inference system $\Rightarrow_D$: $\mathcal{L}(\mathfrak{F}) \rightarrow 2^{\mathfrak{F}} \rightarrow \mathbb{P}$ as given below in Fig. 5.2. Notably, the conclusions of these derivations are sets of formulas $2^{\mathfrak{F}}$ which could possibly be infinite. It thus differs strongly from the usual intuitionistic sequent calculus LJ whose characteristic property is its restriction to a conclusion of at most one formula. The system LJD has only two deduction rules (L) and (R). However, applications of these rules can have very different characteristics as LJD is defined in terms of a local rule set, just as the E-dialogues.

$$\text{L}\ \frac{\varphi \in \Gamma \quad \text{justified}\ \Gamma\, \psi \quad a \,|\, \psi \rhd \varphi \quad \forall\, \theta \in \mathcal{D}_a.\, \Gamma, \theta \Rightarrow_D \Delta \quad \forall a' \,|\, \tau \rhd \psi.\, \Gamma, \tau \Rightarrow_D \mathcal{D}_{a'}}{\Gamma \Rightarrow_D \Delta}$$

$$\text{R}\ \frac{\varphi \in \Delta \quad \text{justified}\ \Gamma\, \varphi \quad \forall a \,|\, \psi \rhd \varphi.\, \Gamma, \psi \Rightarrow_D \mathcal{D}_a}{\Gamma \Rightarrow_D \Delta}$$

Figure 5.2: System LJD

When instantiated with a logic's rule set, for example that for first-order logic, LJD is still very close to the traditional rules of the sequent calculus for that logic. For an example, consider the instantiations of the (R) rule for universal quantification and of the (L) rule for implications in Fig. 5.3. The (RAll) rule states that a claim set Δ containing a universal quantifier $\dot{\forall}\varphi$ can be proven by proving the set $\{\varphi[t]\}$ for every term t. While this rule differs from its usual finitary representation, it still captures well what it means to prove a universal quantifier. Analogously, the (LI) states that one may use an assumption of $\varphi \mathbin{\dot{\rightarrow}} \psi$, as long as the premise φ is justified and can be defended (which is intuitively the same as being able to prove $\{\varphi\}$), which leaves one proving the original claim under the additional assumption of the implication's conclusion.

$$\text{RAll}\ \frac{\dot{\forall}\varphi \in \Delta \quad \forall t.\, \Gamma \Rightarrow_D \{\varphi[t]\}}{\Gamma \Rightarrow_D \Delta}$$

$$\text{LI}\ \frac{\varphi \mathbin{\dot{\rightarrow}} \psi \in \Gamma \quad \text{justified}\ \Gamma\, \varphi \quad \Gamma, \theta \Rightarrow_D \Delta \quad \forall a' \,|\, \tau \rhd \varphi.\, \Gamma, \tau \Rightarrow_D \mathcal{D}_{a'}}{\Gamma \Rightarrow_D \Delta}$$

Figure 5.3: Instantiated LJD

As already stated, the derivations of LJD are isomorphic to winning strategies of E-dialogues. More precisely, each application of a derivation rule corresponds to one round of a winning strategy, or put another way, one application of the rule of the Win s derivation system. When regarded from the perspective of dialogues, the context Γ of LJD corresponds to the opponent's admissions and the set Δ corresponds to the defense set of the opponent's challenge. Thus, each of the derivation rules corresponds to one sort of proponent move. The (L) rule corresponds to the proponent attacking. The first three conditions of that rule are exactly the same as those of (PA) of the $\rightsquigarrow_p$ relation. The other two conditions essentially represent the proponent's ability to win after any possible opponent reaction, the fourth condition corresponding to the opponent defending via the (OD) rule and the fifth to an opponent's counter according to the (OC) rule. Analogously, the (R) rule corresponds to the proponent defending against the current challenge. Thus, she has to be able to win after the opponent issued a new challenge via (OA), as stated by the third condition of that rule.

5.2.3 Soundness and Completeness

Proving the soundness and completeness of LJD with regards to the generalized E-dialogues is fairly straight forward because of the isomorphism lined out in the previous section. Both proofs are essentially a more formal variant of the explanation given there.

Theorem 5.1 *For any formula φ, $\Rightarrow_D \{\varphi\}$ entails $\Vdash_E \varphi$.*

Proof As the context of $\Rightarrow_D \{\varphi\}$ is empty, its bottom-most rule application can only be of (R). Thus we know that φ is non-atomic, as it is justified under the empty context, and that $[\psi] \Rightarrow_D \mathcal{D}_a$ for any $c \,|\, \psi \rhd \varphi$. To prove $\Vdash_E \varphi$, it thus remains to show that Win $([\psi], c)$ for any $c \,|\, \psi \rhd \varphi$. We prove this by showing

$$\forall \Gamma, \Delta.\ \Gamma \Rightarrow_D \Delta \rightarrow \forall c.\ \Delta \subseteq \mathcal{D}_c \rightarrow \text{Win}\,(\Gamma, c)$$

per induction on the derivation $\Gamma \Rightarrow_D \Delta$.

(L) We have to show that for a challenge c with $\Delta \subseteq \mathcal{D}_c$, Win (Γ, c). We are given some $\varphi \in \Gamma$ and attack $a \,|\, \psi \rhd \varphi$ where ψ is justified. We know per the inductive hypotheses that $\forall \theta \in \mathcal{D}_a$. Win $(\theta :: \Gamma, c)$ and $\forall c' \,|\, \tau \rhd \psi$. Win $(\tau :: \Gamma, c')$. We thus choose $(\Gamma, c) \rightsquigarrow_p PA\, a$ as the proponent's move in constructing Win (Γ, c). Let $(\Gamma, c)\, ;\, PA\, a \rightsquigarrow_o s$ be an opponent response. We prove Win s by case distinction on that response.

(OD) Thus $s = (\theta :: \Gamma, c)$ for some $\theta \in \mathcal{D}_a$. Per the first inductive hypothesis, we can conclude Win s.

(OC) Thus $s = (\tau :: \Gamma, c')$ for some $c' \,|\, \tau \rhd \psi$. Per the second inductive hypothesis, we can conclude Win s.

(R) We have to show that for a challenge c with $\Delta \subseteq \mathcal{D}_c$, Win (Γ, c). We are given a justified $\varphi \in \Delta$. Per inductive hypothesis, we know $\forall c' \mid \psi \rhd \varphi$. Win $(\psi :: \Gamma, c')$. We thus choose $(\Gamma, c) \leadsto_p PD\,\varphi$ as the proponent's move. Let the opponent's response be $(\Gamma, c)\,;\, PD\,\varphi \leadsto_o s$. We know that this has to be an application of (OA) and thus $s = (\psi :: \Gamma, c')$ for some $c' \mid \psi \rhd \varphi$. We therefore have to show Win $(\psi :: \Gamma, c')$, which holds per inductive hypothesis. □

Theorem 5.2 *For any formula φ, $\Vdash_E \varphi$ entails $\Rightarrow_D \{\varphi\}$.*

Proof Per definition of $\Vdash_E \varphi$, we know that φ is not atomic and that for every attack $c \mid \psi \rhd \varphi$, Win $([\psi], c)$ holds. We prove $\Rightarrow_D \{\varphi\}$ with an application of the (R) rule. As φ is justified by virtue of being non-atomic, it remains to show that $\psi \Rightarrow_D \mathcal{D}_c$ for all $c \mid \psi \rhd \varphi$. We prove this by showing

$$\forall A_o, c.\ \text{Win}\,(A_o, c) \rightarrow A_o \Rightarrow_D \mathcal{D}_c$$

per induction on the winning strategy Win (A_o, c).

We are given a proponent move $(A_o, c) \leadsto_p m$ and know per inductive hypothesis that $A'_o \Rightarrow_D \mathcal{D}_{c'}$ for every $(A_o, c)\,;\, m \leadsto_o (A'_o, c')$. We perform a case distinction on $(A_o, c) \leadsto_p m$.

(PA) We know that $m = PA\,a$ for some $a \mid \psi \rhd \varphi$ where $\varphi \in A_o$ and ψ is justified. We thus apply the (L) rule. This leaves us with proving $A_o, \theta \Rightarrow_D \mathcal{D}_c$ for any defense $\theta \in \mathcal{D}_a$ and $A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$ for any counter $c' \mid \tau \rhd \psi$. Both follow per inductive hypothesis, as $(A_o, c)\,;\, PA\,a \leadsto_o (\theta :: A_o, c)$ per (OD) and $(A_o, c)\,;\, PA\,a \leadsto_o (\tau :: A_o, c')$ per (OC).

(PD) We know that $m = PD\,\varphi$ for a justified $\varphi \in \mathcal{D}_c$. We thus apply the (R) rule. This leaves us with proving $A_o, \psi \Rightarrow_D \mathcal{D}_{c'}$ for any attack $c' \mid \psi \rhd \varphi$. This follows per inductive hypothesis, as $(A_o, c)\,;\, PD\,\varphi \leadsto_o (\psi :: A_o, c')$ per (OA). □

5.3 Full Intuitionistics First-Order Completeness

The goal of this chapter is to demonstrate a constructive proof of completeness for the full syntax of intuitionistic first-order logic. In the previous section, we proved soundness and completeness of LJD with regards to generalized intuitionistic E-dialogues. When specializing said result to the local rules of first-order logic, one indeed obtains a proof of completeness of LJD for full intuitionistic first-order logic. However, this result is not fully satisfactory, as LJD does not exhibit the finitary character of "real deduction systems". In this section, we demonstrate how to translate between LJD and the intuitionistic sequent calculus LJ, thereby extending the soundness and completeness with regards to E-dialogues to a finitary deduction system.

5.3.1 Full Intuitionistic Sequent Calculus

The deduction system we use in these translations is a variant of the intuitionistic sequent calculs LJ. Importantly, the system LJ is cut-free, which allows translation to and from the cut-free LJD without first having to prove a cut-elimination theorem. The rules of the system LJ are given in Fig. 5.4.

$$\text{Ax}\,\frac{}{\Gamma, \varphi \Rightarrow_J \varphi} \qquad \text{Contr}\,\frac{\Gamma, \varphi, \varphi \Rightarrow_J \psi}{\Gamma, \varphi \Rightarrow_J \psi} \qquad \text{Weak}\,\frac{\Gamma \Rightarrow_J \psi}{\Gamma, \varphi \Rightarrow_J \psi}$$

$$\text{Perm}\,\frac{\Gamma, \psi, \varphi, \Gamma' \Rightarrow_J \theta}{\Gamma, \varphi, \psi, \Gamma' \Rightarrow_J \theta} \qquad \text{Exp}\,\frac{\Gamma \Rightarrow_J \dot\bot}{\Gamma \Rightarrow_J \varphi} \qquad \text{TR}\,\frac{}{\Gamma \Rightarrow_J \dot\top}$$

$$\text{IL}\,\frac{\Gamma \Rightarrow_J \varphi \quad \Gamma, \psi \Rightarrow_J \theta}{\Gamma, \varphi \dot\rightarrow \psi \Rightarrow_J \theta} \qquad \text{IR}\,\frac{\Gamma, \varphi \Rightarrow_J \psi}{\Gamma \Rightarrow_J \varphi \dot\rightarrow \psi} \qquad \text{AL}\,\frac{\Gamma, \varphi, \psi \Rightarrow_J \theta}{\Gamma, \varphi \dot\wedge \psi \Rightarrow_J \theta}$$

$$\text{AR}\,\frac{\Gamma \Rightarrow_J \varphi \quad \Gamma \Rightarrow_J \psi}{\Gamma \Rightarrow_J \varphi \dot\wedge \psi} \qquad \text{OL}\,\frac{\Gamma, \varphi \Rightarrow_J \theta \quad \Gamma, \psi \Rightarrow_J \theta}{\Gamma, \varphi \dot\vee \psi \Rightarrow_J \theta} \qquad \text{OR}_\text{L}\,\frac{\Gamma \Rightarrow_J \varphi}{\Gamma \Rightarrow_J \varphi \dot\vee \psi}$$

$$\text{OR}_\text{R}\,\frac{\Gamma \Rightarrow_J \psi}{\Gamma \Rightarrow_J \varphi \dot\vee \psi} \qquad \text{AllL}\,\frac{\Gamma, \varphi[t] \Rightarrow_J \psi}{\Gamma, \dot\forall \varphi \Rightarrow_J \psi} \qquad \text{AllR}\,\frac{\uparrow\Gamma \Rightarrow_J \varphi}{\Gamma \Rightarrow_J \dot\forall \varphi}$$

$$\text{ExL}\,\frac{\uparrow\Gamma, \varphi \Rightarrow_J \uparrow\psi}{\Gamma, \dot\exists \varphi \Rightarrow_J \psi} \qquad \text{ExR}\,\frac{\Gamma \Rightarrow_J \varphi[t]}{\Gamma \Rightarrow_J \dot\exists \varphi}$$

Figure 5.4: Full Sequent Calculus LJ

The management of the assumptions Γ by the system LJ differs from that exhibited by the other systems employed in this thesis. LJ contains the rules (Contr), (Weak) and (Perm) to explicitly drop, duplicate and move the formulas in Γ. These rules are required as the L-rules of LJ act by consuming the outermost formula of the context, which means it must be possible to move a copy of any assumption into this position to be able to properly deduce claims. Of course, this maneuver can be formalized, thus allowing to not further concern ourselves if explicit assumption management.

Fact 5.3 *One can prove the following rule admissible in the system LJ.*

$$\text{Ctx}\,\frac{\varphi \in \Gamma \quad \Gamma, \varphi \Rightarrow_J \psi}{\Gamma \Rightarrow_J \psi}$$

Just as the previous two deduction systems, LJ exhibits under context extension and substitution.

Fact 5.4 *One can prove the following rules admissible in the system LJ.*

$$\textsc{Weak}'\frac{\Gamma' \Rightarrow_J \varphi \quad \Gamma' \subseteq \Gamma}{\Gamma \Rightarrow_J \varphi} \qquad\qquad \textsc{WeakS}\frac{\Gamma \Rightarrow_J \varphi}{\Gamma[\sigma] \Rightarrow_J \varphi[\sigma]}$$

While the (AllR) and (ExL) rules of LJ are given in the de Bruijn style, their locally nameless variants are admissible as well.

Lemma 5.5 *Let x a variable fresh for Γ, ψ and $\dot{\nabla}\varphi$.*

1. $\uparrow\Gamma \Rightarrow_J \varphi$ *if and only if* $\Gamma \Rightarrow_J \varphi[x]$
2. $\uparrow\Gamma, \varphi \Rightarrow_J \uparrow\psi$ *if and only if* $\Gamma, \varphi[x] \Rightarrow_J \psi$

5.3.2 Translating between LJ and LJD

Before we begin with the translations, recall the rule-set for full first-order logic.

Definition 5.6 (First-order local rules) *The local rules of first-order logic are given by $(\mathfrak{F}, \mathfrak{F}^*, \mathcal{A}, \mathcal{D}_\_, \lhd)$. The set of atomic formulas $\mathfrak{F}^*$ is defined as*

$$\{P\, t_1 \ldots t_{|P|} \mid P : \mathcal{P}, t_1, \ldots, t_{|P|} : \mathfrak{T}\}$$

The type of attacks $\mathcal{A}$ as well as the attacks relation $\lhd$ and the defense sets $\mathcal{D}_\_$ are given below.

$$\mathcal{A} : \mathbb{T} \;:=\; A_{\dot{\bot}} \mid A_{\dot{\rightarrow}}\, \varphi \mid A_L \mid A_R \mid A_{\dot{\vee}} \mid A_t\, \varphi \mid A_{\dot{\exists}} \qquad (\varphi : \mathfrak{F}, t : \mathfrak{T})$$

$$\begin{array}{rl}
A_{\dot{\bot}} \rhd \dot{\bot} & \mathcal{D}_{A_{\dot{\bot}}} = \{\} \\
A_{\dot{\rightarrow}}\, \psi \mid \ulcorner\varphi\urcorner \rhd \varphi \dot{\rightarrow} \psi & \mathcal{D}_{A_{\dot{\rightarrow}}\, \psi} = \{\psi\} \\
A_L \rhd \varphi \dot{\wedge} \psi & \mathcal{D}_{A_L} = \{\varphi\} \\
A_R \rhd \varphi \dot{\wedge} \psi & \mathcal{D}_{A_R} = \{\psi\} \\
A_{\dot{\vee}} \rhd \varphi \dot{\vee} \psi & \mathcal{D}_{A_{\dot{\vee}}} = \{\varphi, \psi\} \\
A_t\, \varphi \rhd \dot{\forall}\varphi & \mathcal{D}_{A_t\, \varphi} = \{\varphi[t]\} \\
A_{\dot{\exists}} \rhd \dot{\exists}\varphi & \mathcal{D}_{A_{\dot{\exists}}} = \{\varphi[t] \mid t : \mathfrak{T}\}
\end{array}$$

We begin by proving that derivations in LJ can be translated into derivations of LJD. In the course of this translation, we make use of multiple properties exhibited by LJD which are analogous to those of finitary deduction systems.

Fact 5.6 *The following rule can be shown admissible for the system LJD.*

$$\text{WEAK}\frac{\Gamma' \subseteq \Gamma \quad \Gamma' \Rightarrow_D \Delta' \quad \Delta' \subseteq \Delta}{\Gamma \Rightarrow_D \Delta}$$

The first property is the usual weakening property of sequent calucli, allowing more formulas to be added on either side of the $\Rightarrow_D$ of a derivation without changing deducability.

Fact 5.7

1. *If* $\Gamma \Rightarrow_D \{\varphi\}$ *then* $\Gamma, \psi \Rightarrow_D \mathcal{D}_a$ *for all* $a \mid \psi \rhd \varphi$
2. *If* $\Gamma \Rightarrow_D \{\varphi\}$ *then* $\Gamma \Rightarrow_D \Delta$ *can be deduced for any* Δ *such that* $\Gamma' \Rightarrow_D \Delta$ *follows for any* $\Gamma \subseteq \Gamma'$ *with* φ *justified in* Γ' *and* $\forall a \mid \psi \rhd \varphi.\ \Gamma', \psi \Rightarrow_D \mathcal{D}_a$

To understand these two properties, recall the R-rule of LJD:

$$\text{R}\frac{\varphi \in \Delta \quad \text{justified}\ \Gamma\, \varphi \quad \forall a \mid \psi \rhd \varphi.\ \Gamma, \psi \Rightarrow_D \mathcal{D}_a}{\Gamma \Rightarrow_D \Delta}$$

Now consider an LJD derivation for a singleton set. To prove this claim, there are two possibilities for every path of the derivation: Either, it at some point applies the R-rule, thus "defending" the formula, or it proceeds by only ever applying the L-rule, thus being able to prove an claim Δ, not just the singleton set. With this in mind, consider the second property which states that when a singleton set can be deduced, a defense in $\mathcal{D}_a$ against each attack a against it can be deduced as well. This directly follows from the previous consideration as any singleton derivation will, on every of its paths, either apply the R-rule and thus be able to prove $\mathcal{D}_a$ or never apply the R-rule and thus prove any Δ, including $\mathcal{D}_a$. The third property is based on similar reasoning. Intuitively, it states that when any point at which the derivation applies the R-rule can also be continued with a derivation of Δ, Δ can be deduced overall.

Fact 5.8 *Assume a well-founded relation* $<$ *on formulas such that for any attack* $a \mid \psi \rhd \varphi$, $\psi < \varphi$ *and* $\theta < \varphi$ *for any* $\theta \in \mathcal{D}_a$. *Then* $\Gamma \Rightarrow_D \Delta$ *whenever there is a* φ *with* $\varphi \in \Gamma$ *and* $\varphi \in \Delta$.

This last property is the LJD equivalent of the (Ax) rule. It states that $\Gamma \Rightarrow_D \Delta$ can be deduced if Γ and Δ have a common element φ. When considered as a dialogue strategy, this property seems very obvious: The proponent can just attack the opponent's admission of φ and then mimic her actions until she has won. However, this only works for rule sets where attacks and defenses lead to "smaller" formulas for some well-founded notion of smaller. For a counter example, consider the rules-set on the formulas containing only the symbol $*$ with the only attack being $a \rhd *$ and its defense $\mathcal{D}_a = \{*\}$. Here, the proponent cannot win the game state $([*], a)$. If she were to attack $*$, the opponent could simply respond by admitting $*$, leading to the state $([*, *], a)$, which is essentially the same as $([*], a)$, thus leaving the proponent "stuck". This restriction can thus be viewed as a result of the asynchronous winning conditions of our formalization of dialogues: The opponent can win by forcing a dialogue to continue on forever, while the proponent can only win if the dialogue is finite.

Having established these properties, we can now prove that derivations of LJ can be transformed into derivations of LJD for the first-order rule-set.

Theorem 5.9 *Whenever* $\Gamma \Rightarrow_J \varphi$ *can be derived,* $\Gamma \Rightarrow_D \{\varphi\}$ *can be derived as well.*

Proof We prove the generalized statement

$$\forall \sigma.\, \Gamma \Rightarrow_J \varphi \to \Gamma[\sigma] \Rightarrow_D \{\varphi[\sigma]\}$$

per induction on the derivation $\Gamma \Rightarrow_J \varphi$. The cases for the L- and R-rules of the connectives are all very similar. We thus only give explicit proofs for $\dot\to$ and $\dot\exists$. For the sake of readability, we also leave out the substitution σ in all cases where it is not used.

(Ax) We have to prove that $\Gamma, \varphi \Rightarrow_D \{\varphi\}$. This follows from Fact 5.8, as the first-order rule-set is ordered by the subformula relation.

(Contr), (Weak), (Perm) All of these cases follow directly from (Weak).

(Exp) From $\Gamma \Rightarrow_D \{\dot\bot\}$ we have to derive $\Gamma \Rightarrow_D \{\varphi\}$. As $\{\dot\bot\}$ is a singleton, we know by Fact 5.7.1 and $A_{\dot\bot} \rhd \dot\bot$ that $\Gamma \Rightarrow_D \mathcal{D}_{A_{\dot\bot}}$. As $\mathcal{D}_{A_{\dot\bot}}$ is empty we can thus conclude $\Gamma \Rightarrow_D \{\varphi\}$ by (Weak).

(TR) We have to derive $\Gamma \Rightarrow_D \dot\top$. This follows by an application of the R-rule, as there are no attacks on $\dot\top$.

(IL) We assume $\Gamma \Rightarrow_D \{\varphi\}$ and $\Gamma, \psi \Rightarrow_D \{\theta\}$ to show $\Gamma, \varphi \dot\to \psi \Rightarrow_D \{\theta\}$. Per (Weak), $\Gamma \Rightarrow_D \{\varphi\}$ entails $\Gamma, \varphi \dot\to \psi \Rightarrow_D \{\varphi\}$. By Fact 5.7.2, it thus suffices to show that $\Gamma' \Rightarrow_D \{\theta\}$ for any $(\Gamma, \varphi \dot\to \psi) \subseteq \Gamma'$ with $\forall a \,|\, \tau \rhd \varphi.\, \Gamma', \tau \Rightarrow_D \mathcal{D}_a$ and φ justified in Γ', to deduce $\Gamma, \varphi \dot\to \psi \Rightarrow_D \{\theta\}$. We prove $\Gamma' \Rightarrow_D \{\theta\}$ by an application of the L-rule on the $\varphi \dot\to \psi \in \Gamma'$. As φ is justified in Γ' and $\forall a \,|\, \tau \rhd \varphi.\, \Gamma', \tau \Rightarrow_D \mathcal{D}_a$, this leaves us with proving $\Gamma', \psi \Rightarrow_D \{\tau\}$, which follows per assumption and (Weak).

(IR) We assume $\Gamma, \varphi \Rightarrow_D \{\psi\}$ to prove $\Gamma \Rightarrow_D \varphi \dot\to \psi$. This follows by an application of the R-rule.

(ExL) We assume $\forall\sigma.\ (\uparrow\Gamma,\varphi)[\sigma] \Rightarrow_D \{(\uparrow\psi)[\sigma]\}$ to prove $\Gamma, \dot{\exists}\varphi \Rightarrow_D \psi$. We thus apply the L-rule to $\dot{\exists}\varphi$, which leaves us proving $\Gamma, \varphi[t] \Rightarrow_D \psi$ for arbitrary terms t. This follows per assumption with the substitution $\sigma = [t]$.

(ExR) We assume $\Gamma \Rightarrow_D \{\varphi[t]\}$ to prove $\Gamma \Rightarrow_D \{\dot{\exists}\varphi\}$. We thus apply the R-rule. This leaves us proving $\Gamma \Rightarrow_D \{\varphi[s] \mid s : \mathfrak{T}\}$, which follows per assumption and (Weak). □

Corollary 5.10 *If $\Rightarrow_J \varphi$ then $\Vdash_E \varphi$ under first-order rules.*

We now turn our attention to transforming proofs in LJD into derivations in LJ. Between the two transformations of this section, this proof is the more interesting one as it converts infinitary derivations into finitary evidence.

The idea behind this proof is encoded in the statement we prove per induction:

$$\Gamma \Rightarrow_D \Delta \rightarrow \forall\varphi.\ (\forall\psi \in \Delta, \Gamma \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \psi \rightarrow \Gamma' \Rightarrow_J \varphi) \rightarrow \Gamma \Rightarrow_J \varphi$$

To understand this statement, recall the disjunction elimination rule (OE) of natural deduction systems.

$$\text{OE}\ \frac{\Gamma \vdash \varphi \dot{\vee} \psi \quad \Gamma, \varphi \vdash \theta \quad \Gamma, \psi \vdash \theta}{\Gamma \vdash \theta}$$

The statement $\Gamma \Rightarrow_D \Delta$ can be understood as "the disjunction over all formulas in Δ can be proven under Γ". The statement we prove per induction can thus be viewed as an adoption of the rule (OE) to the infinite disjunction over Δ: Any formula that can be derived from a proof for each element of Δ can be derived overall.

Before proving the full statement, we first prove a lemma.

Lemma 5.11 *Whenever $\forall\psi.\ (\forall\theta \in \mathcal{D}_a, (\Gamma,\tau) \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \rightarrow \Gamma' \Rightarrow_J \psi) \rightarrow \Gamma \Rightarrow_J \psi$ can be derived for all attacks $a \mid \tau \rhd \varphi$ for a Γ in which φ is justified, $\Gamma \Rightarrow_J \varphi$ can be derived.*

Proof Per case analysis on φ. Most cases are very similar, so we do not cover all of them.

$\varphi = \dot{\bot}$: Per the first assumption, we can conclude $\Gamma \Rightarrow_J \dot{\bot}$ by showing the vacuous statement $\forall\theta \in \mathcal{D}_{A_\bot}, \Gamma \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \rightarrow \Gamma' \Rightarrow_J \dot{\bot}$.

$\varphi = P\,t_1 \ldots t_{|P|}$: As $P\,t_1 \ldots t_{|P|} \in \Gamma$, as it is justified in Γ, the claim follows with the (Ctx) and (Ax) rules.

$\varphi = \psi \dot{\rightarrow} \theta$: We apply the (IR) rule and are left proving $\Gamma, \psi \Rightarrow_J \varphi$. This follows directly by applying the assumption with $A_{\dot{\rightarrow}}\ \psi$.

$\varphi = \dot\forall\psi$: After applying the (AR) and Lemma 5.5, we are left proving $\Gamma \Rightarrow_J \varphi[x]$ with x free in Γ and $\dot\forall\varphi$. This follows directly by applying the assumption with $A_x\,\varphi$. $\square$

Using this lemma, we can now derive the full translation theorem.

Theorem 5.12 *Whenever* $\Gamma \Rightarrow_D \{\varphi\}$ *can be derived,* $\Gamma \Rightarrow_J \varphi$ *can be derived as well.*

Proof We prove the generalized statement

$$\Gamma \Rightarrow_D \Delta \to \forall\varphi.\ (\forall\psi \in \Delta, \Gamma \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \psi \to \Gamma' \Rightarrow_J \varphi) \to \Gamma \Rightarrow_J \varphi$$

by induction on the derivation $\Gamma \Rightarrow_D \Delta$.

(R) There thus is a $\varphi \in \Delta$ which is justified under Γ. Per induction, we may assume $\forall\psi.\ (\forall\theta \in \mathcal{D}_a, (\Gamma,\tau) \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \psi) \to \Gamma \Rightarrow_J \psi$ holds for any attack $a\,|\,\tau \rhd \varphi$. We have to show that for any ψ with the property that $\forall\theta \in \Delta, \Gamma \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \psi$, one can derive $\Gamma \Rightarrow_J \psi$. We do this by applying said property for the $\varphi \in \Delta$, which leaves us proving that $\Gamma \Rightarrow_J \varphi$. This follows form Lemma 5.11.

(L) There thus is a $\varphi \in \Gamma$ and an attack $a\,|\,\psi \rhd \varphi$ with ψ being justified in Γ. Per induction, we may also assume that

1. For any $\tau \in \mathcal{D}_a$ and any ψ with $\forall\theta \in \Delta, (\Gamma,\tau) \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \psi$ one can deduce $\Gamma \Rightarrow_J \psi$
2. For any counter $c\,|\,\tau \rhd \psi$ one can deduce $\Gamma \Rightarrow_J \nu$ for any ν such that $\forall\theta \in \mathcal{D}_c, (\Gamma,\theta) \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \nu$

We have to prove that for any ν such that $\forall\theta \in \Delta, \Gamma \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \nu$ one can deduce ν. We proceed per case distinction on φ. As the cases for many connectives are very similar, we do not cover all of them.

$\varphi = \dot\bot$: We know that $\dot\bot \in \Gamma$. Then $\Gamma \Rightarrow_J \nu$ follows by the (Exp), (Ctx) and (Ax) rules.

$\varphi = \dot\top, \varphi = P\,t_1\,...\,t_{|P|}$: Both follow per meta-exfalso, as the assumed attack on them cannot exist.

$\varphi = \varphi \mathbin{\dot\to} \psi$: We know that the attack a has to be $A_{\dot\to}\ \psi\,|\,\ulcorner\varphi\urcorner \rhd \varphi \mathbin{\dot\to} \psi$. Applying the (Ctx) and (IL) rules leaves us proving $\Gamma \Rightarrow_J \varphi$ and $\Gamma, \psi \Rightarrow_J \nu$. The first claim follows as Lemma 5.11 applies by the second inductive hypothesis. The second claim follows by the first inductive hypothesis and the fact that $\forall\theta \in \mathcal{D}_c, (\Gamma,\theta) \subseteq \Gamma'.\ \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \nu$.

$\varphi = \dot{\exists}\varphi$: We know that a has to be the attack $A_{\dot{\exists}} \triangleright \dot{\exists}\varphi$. applying the (CTX) and (EL) rules as well as Lemma 5.5, we are left proving $\Gamma, \varphi[x] \Rightarrow_J \nu$ for an x which is fresh for Γ, $\dot{\exists}\varphi$ and ν. As $\varphi[x] \in \mathcal{D}_{A_{\dot{\exists}}}$, this claim follows by the first inductive hypothesis and the fact that $\forall\theta \in \mathcal{D}_c, (\Gamma, \theta) \subseteq \Gamma'. \Gamma' \Rightarrow_J \theta \to \Gamma' \Rightarrow_J \nu$. □

5.4 Generalized Intuitionistic D-Completeness

In Section 5.2, we have given an account of generalized, formal E-dialogues and showed them sound and complete, based on prior work by Sørensen and Urzyczyn [61]. In this section, we extend this approach to D-dialogues.

Recall that E- and D-dialogues differ in which moves the opponent is allowed to make. In E-dialogues, the opponent may only react to the proponent's previous move, while D-dialogues allow the opponent to react to moves made by the proponent in previous rounds. Intuitively, D-dialogues are "fairer" than E-dialogues as the moves the proponent and opponent may perform are more similar. It is thus slightly surprising that E- and D-dialogues are equivalent. That is, the proponent has a winning E-strategy if and only if she has winning D-strategy. This equivalence allows formal treatments of dialogues to focus only on the structurally simpler E-dialogues as most results can be extended to D-dialogues along this equivalence.

The distinction between E- and D-dialogues, although with different terminology, was first made by Lorenz [47]. In this paper, he already acknowledges the equivalence and gives a proof of it, although it's correctness is somewhat controversial [36]. Indeed, the proof of this equivalence is considered very complicated and thus appears relatively rarely in the literature, noteworthy exceptions being Felscher's [15] and Krabbe's [35] very formal renditions of it.

In the following section, we give a formal account of generalized intuitionistic D-dialogues (Section 5.4.1) and their equivalence to E-dialogues of the same nature (Section 5.4.2). For this, we prove the soundness and completeness of LJD with regards to D-dialogues on enumerable rule sets. We have chosen this approach, as opposed to direct translation of winning strategies, as it to leads to more intuitive proofs.

5.4.1 Generalized Intuitionistic D-Dialogues

Similarly to the formal E-dialogues of Section 5.2.1, we express formal D-dialogues in terms a state transition system. The symmetric nature of the rules of D-dialogues results in a very symmetric type for the state space: $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A})$. The first two components of these quadruples represent the proponent's open admissions (those, which the opponent may still attack) and the challenges issued by the opponent against her, which she has yet to defend herself against. The third and fourth components represent the same for the opponent.

As opposed to the formulation of proponent and opponent moves of Section 5.2.1, in which the proponent put forward a move and the opponent simply reacted to it, the move formulations for D-dialogues exhibit a stronger symmetry: The moves of both players are represented as binary relations on the state space.

We denote by $s \rightsquigarrow_p s'$ that the proponent may perform a move in state s, leading to a new state s'.

Definition 5.7 (Proponent moves) *The following inference system defines the proponent's move relation* $\rightsquigarrow_p$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \to \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \to \mathbb{P}$

$$\text{PD}\ \frac{\varphi \in \mathcal{D}_c \quad \textit{justified}\ A_o\, \varphi}{(A_p, c :: C_p, A_o, C_o) \rightsquigarrow_p (\varphi :: A_p, C_p, A_o, C_o)}$$

$$\text{PA}\ \frac{\varphi \in A_o \quad \textit{justified}\ A_o\, \psi \quad a \mid \psi \rhd \varphi}{(A_p, C_p, A_o, C_o) \rightsquigarrow_p (\psi :: A_p, C_p, A_o, a :: C_o)}$$

The (PD) rule states that the proponent may defend against the most recent open challenge by admitting a justified defense formula. This restriction of defenses to the most recent challenge ensures that the dialogues remain intuitionistic. The (PA) rule states that the proponent may attack any admission made by the opponent, possibly by admitting to a justified formula ψ herself if the attack a calls for it.

The legal opponent moves, given by the $\rightsquigarrow_o$ relation, are very similar to those of the proponent.

Definition 5.8 (Opponent moves) *The following inference system defines the opponent's move relation* $\rightsquigarrow_o$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \to \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A}) \to \mathbb{P}$

$$\text{OD}\ \frac{\varphi \in \mathcal{D}_a}{(A_p, C_p, A_o, a :: C_o) \rightsquigarrow_o (A_p, C_p, \varphi :: A_o, C_o)}$$

$$\text{OA}\ \frac{c \mid \psi \rhd \varphi}{(A_p +\!\!+\, \varphi :: A'_p, C_p, A_o, C_o) \rightsquigarrow_o (A_p +\!\!+\, A'_p, c :: C_p, \psi :: A_o, C_o)}$$

The rule (OD) governing the opponent's defense moves is nearly symmetric to that for the proponent. The only difference is that the opponent does not need any justifications for the formulas she admits, as any admission she makes already benefits the proponent. The (OA) rule, describing when and how she may attack the proponent's admissions. This rule differs from that of the proponent, as the attacked admission is removed from the list of the proponent's admissions, thereby assuring that every proponent admission is attacked at most once. This asymmetry in the rules balances out another rule asymmetry: The proponent can only win finite dialogues, which means the opponent has to be prevented from stalling a dialogue by repeatedly attacking the same proponent admission.

The definitions of winning states and the resulting D-validity are completely analogous to those for E-dialogues.

Definition 5.9 (Winning states) *A state s is winning if Win s can be derived in the following inference system*

$$\frac{s \rightsquigarrow_p s' \quad \forall s''.\, s' \rightsquigarrow_o s'' \rightarrow \mathit{Win}\, s''}{\mathit{Win}\, s}$$

Definition 5.10 (D-Validity) *A formula φ is D-valid, denoted by $\Vdash_D \varphi$, if it is not atomic and for any attack $c \,|\, \psi \rhd \varphi$ on it, the resulting state $([], [c], [\psi], [])$ is winning.*

5.4.2 Soundness and Completeness

We now prove the soundness and completeness of LJD with regards to generalized D-dialogues. As the LJD is also sound and complete with regards to generalized E-dialogues, this entails the E-D-equivalence.

We begin with completeness. We thus have to translate winning D-strategies into LJD derivations. Completeness is very easy to prove: As LJD derivations are isomorphic to winning E-strategies, this essentially constitutes a translation of a strategy suitable for a stronger opponent into a strategy suitable for a strictly weaker opponent. The proof is thus completely analogous to that of the E-completeness of LJD.

Theorem 5.13 *For any formula φ, $\Vdash_D \varphi$ entails $\Rightarrow_D \{\varphi\}$.*

Proof Per definition of $\Vdash_D \varphi$, we know that φ is not atomic and that for every attack $c \,|\, \psi \rhd \varphi$, Win $([], [c], [\psi], [])$ holds. We prove $\Rightarrow_D \{\varphi\}$ with an application of the (R) rule. As φ is justified by virtue of being non-atomic, it remains to show that $\psi \Rightarrow_D \mathcal{D}_c$ for all $c \,|\, \psi \rhd \varphi$. We prove this by showing

$$\forall A_o, c.\ \mathrm{Win}\, (A_p, c :: C_p, A_o, C_o) \rightarrow A_o \Rightarrow_D \mathcal{D}_c$$

per induction on the winning strategy Win $(A_p, c :: C_p, A_o, C_o)$.

We are given a proponent move $(A_p, c :: C_p, A_o, C_o) \rightsquigarrow_p s$ and know per inductive hypothesis that $A_o' \Rightarrow_D \mathcal{D}_{c'}$ for every $s \rightsquigarrow_o (A_p', c' :: C_p', A_o', C_o')$. We perform a case distinction on $(A_p, c :: C_p, A_o, C_o) \rightsquigarrow_p s$.

(PD) We know that $s = (\varphi :: A_p, C_p, A_o, C_o)$ for a justified $\varphi \in \mathcal{D}_c$. We thus apply the (R) rule. This leaves us with proving $A_o, \psi \Rightarrow_D \mathcal{D}_{c'}$ for any attack $c' \,|\, \psi \rhd \varphi$. This follows from the inductive hypothesis, as per the (OA) rule $(\varphi :: A_p, C_p, A_o, C_o) \rightsquigarrow_o (A_p, c' :: C_p, \psi :: A_o, C_o)$.

(PA) We know that $s = (\psi :: A_p, c :: C_p, A_o, a :: C_o)$ for some $a \,|\, \psi \rhd \varphi$ where $\varphi \in A_o$ and ψ is justified. We thus apply the (L) rule. This leaves us with proving $A_o, \theta \Rightarrow_D \mathcal{D}_c$ for any defense $\theta \in \mathcal{D}_a$ and $A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$ for any possible counter $c' \,|\, \tau \rhd \psi$. Both follow per inductive hypothesis, as per the (OD) $(\psi :: A_p, c :: C_p, A_o, a :: C_o) \rightsquigarrow_o (\psi :: A_p, c :: C_p, \theta :: A_o, a :: C_o)$ and $(\psi :: A_p, c :: C_p, A_o, a :: C_o) \rightsquigarrow_o (\psi :: A_p, c' :: c :: C_p, \tau :: A_o, a :: C_o)$ per (OC). □

The direction of soundness, converting LJD derivations into winning D-strategies, is much more complicated. Its strategy analogue, translating E-strategies into D-strategies, is why the proof of the E-D-equivalence is often considered very involved.

We start by giving an intuition why one should believe such a transformation is possible in the first place. For this, consider the following derivation in LJD.

$$\text{(L)}\ \frac{\varphi \dot{\to} \psi \in \Gamma \qquad \text{justified}\ \Gamma\, \varphi \qquad (1)\ \forall a \,|\, \theta \rhd \varphi.\ \Gamma, \theta \Rightarrow_D \mathcal{D}_a \qquad (2)\ \Gamma, \psi \Rightarrow_D \Delta}{\Gamma \Rightarrow_D \Delta}$$

As the derivation starts with an application of the (L) rule, any strategy that it might be translated into should start the round with the proponent attacking $\varphi \to \psi \in \Gamma$ accordingly. This leaves us with explaining how the proponent should react to the opponent moves following this attack. For the sake of simplicity, let us assume the state resulting from the proponent's attack is $([\varphi], [c], \Gamma, [A_{\to}\, \psi])$ for some open challenge c. This leaves room for two possible opponent moves: As the proponent has admitted φ, the opponent may choose to challenge her on that claim by attacking it with $a \,|\, \theta \rhd \varphi$. The proponent should now react by playing according to the translated strategy obtained from the proof (1). If the opponent instead chooses to defend against the attack by admitting ψ, the proponent should react according to the strategy obtained from (2). So far, this is no different from the translation of LJD derivations into E-strategies. However, consider again the case in which the defends against the proponent's attack on $\varphi \dot{\to} \psi$. The state resulting from that opponent move is $([\varphi], [c], \Gamma, [])$, the proponent admission φ still remaining in the state. That means that at any point in the future of this dialogue, the opponent may, instead of reacting to the preceding proponent move, choose to attack φ. However, this is not a problem to the proponent: She can just react according to (1). This is the crucial different between translating derivations into E- and D-strategies. As the E-opponent may only react to the previous proponent move, all "unused" derivations can be "forgotten" at the end of the round. This is not the case for D-dialogues: The unused derivations have to be remembered to be able to react to the opponent choosing to perform another reaction to the current proponent move later on in the dialogue.

For the proof of soundness, we introduce a new variant of dialogues we call S-dialogues ("S" for stack). These dialogues essentially represent the kinds of winning strategies for D-dialogues that can be obtained from the considerations of the previous paragraph and can thus be viewed as a "subset of D-dialogues". The state space of S-dialogues is given by the type $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A} \times \mathcal{A})$. As with the D-dialogues, the first two components represent the proponent's and opponent's open admissions respectively. The important change from the D-dialogues is the last component, a list of pairs of attacks, which is a combination of the list of C_p and C_o of open challenges of the D-dialogues into a stack of D of "deferred moves".

Another important change is that S-dialogues are not symmetric anymore. However, this time it is the proponent that has to react to the opponent's last challenge. The legal proponent moves in a state s reacting to the opponent challenge c is represented by the relation $s\,;\,c \rightsquigarrow_p s'$.

Definition 5.11 (Proponent moves) *The following inference system defines the proponent's move relation* $\rightsquigarrow_p$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A} \times \mathcal{A}) \to \mathcal{A} \to \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A} \times \mathcal{A}) \to \mathbb{P}$

$$\text{PD}\,\frac{\varphi \in \mathcal{D}_c \quad \textit{justified}\ A_o\,\varphi}{(A_p, A_o, D)\,;\,c \rightsquigarrow_p (\varphi :: A_p, A_o, D)}$$

$$\text{PA}\,\frac{\varphi \in A_o \quad \textit{justified}\ A_o\,\psi \quad a \mid \psi \rhd \varphi}{(A_p, A_o, D)\,;\,c \rightsquigarrow_p (\psi :: A_p, A_o, (a, c) :: D)}$$

Here, the (PD) rule allows the proponent to defend against the opponent's last challenge under the usual side conditions. The (PA) rule exemplifies the idea behind the deferred moves D: If the proponent chooses to attack one of the opponent's admissions with an attack a, the proponent's attack and the opponent's last challenge c are put on the stack together. This encodes an observation about the strategies derived from derivations with the translation we described previously: If the derivation $\Gamma \Rightarrow_D \mathcal{D}_c$ starts with an application of the (L) rule, meaning the proponent attacks an opponent admission with the attack a, the proof that explains to her how to continue defending against the current challenge c is $\forall \theta \in \mathcal{D}_a.\, \Gamma, \theta \Rightarrow_D \mathcal{D}_c$. Thus, she will only be able to "use" this information once the opponent has admitted some defense $\theta \in \mathcal{D}_a$. In S-dialogues, the proponent thus puts (a, c) on the stack of deferred moves, thereby signaling: "Only once the opponent defends against a will I be able to continue fending off c".

The opponent's moves thus signify her posing a challenge to c in a new state s' to the proponent, denoted by $s \rightsquigarrow_o s'\,;\,c$.

Definition 5.12 (Opponent moves) *The following inference system defines the opponent's move relation* $\rightsquigarrow_o$: $\mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A} \times \mathcal{A}) \to \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathfrak{F}) \times \mathcal{L}(\mathcal{A} \times \mathcal{A}) \to \mathcal{A} \to \mathbb{P}$

$$\text{OD}\,\frac{\varphi \in \mathcal{D}_a}{(A_p, A_o, (a,c) :: D) \rightsquigarrow_o (A_p, \varphi :: A_o, D)\,;\, c}$$

$$\text{OA}\,\frac{c \,|\, \psi \rhd \varphi}{(A_p +\!\!+\, \varphi :: A'_p, A_o, D) \rightsquigarrow_o (A_p +\!\!+\, A'_p, \psi :: A_o, D)\,;\, c}$$

There are two ways the opponent may pose the challenge. If she defends against the last attack a the proponent made against her via the (OD) rule, she reposes the challenge that was deferred alongside the attack a. If she instead chooses to attack one of the proponent's open admissions, this new attack is posed as the new challenge instead. To prevent the opponent from stalling the dialogue indefinitely by repeatedly attacking the same proponent admission, any proponent admission is removed once it has been attacked once.

The definitions of winning states of the transition system and S-validity are completely analogous to those E- and D-dialogues.

Definition 5.13 (Winning states) *A state* $s\,;\,c$ *is winning if* $\mathit{Win}\ s\,;\,c$ *can be derived in the following inference system*

$$\frac{s\,;\,c \rightsquigarrow_p s' \quad \forall s''c'.\ s' \rightsquigarrow_o s''\,;\,c' \to \mathit{Win}\ s''\,;\,c'}{\mathit{Win}\ s\,;\,c}$$

Definition 5.14 (S-Validity) *A formula* φ *is S-valid, denoted by* $\Vdash_S \varphi$*, if it is not atomic and for any attack* $c \,|\, \psi \rhd \varphi$ *on it, the resulting state* $([], [\psi], [])\,;\,c$ *is winning.*

As stated earlier, S-dialogues constitute a restricted variant of D-dialogues. Thus, D-validity subsumes S-validity.

Lemma 5.14 *Any S-valid formula* φ *is also D-valid.*

Proof Per definition, φ is atomic. Thus it remains to show that for any challenge $c \,|\, \psi \rhd \varphi$, Win $([], [\psi], [])\,;\,c$ entails Win $([], [c], [\psi], [])$. We show this by proving the generalized statement

$$\text{Win}\,(A_p, A_o, D)\,;\,c \to \text{Win}\,(A_p, c :: \pi_1 D, A_o, \pi_2 D)$$

per induction on Win $(A_p, A_o, D)\,;\,c$. As this proof is completely analogous to the other proofs of this chapter, we do not go into further detail. □

We now prove the soundness of LJD with regards to S-dialogues. The informal account of the translation from derivations into winning strategies is formalized by the statement we prove by induction:

$$\begin{aligned}\forall A_P, A_o, D, c.\ &(1)(\forall \varphi \in A_p, a \mid \psi \rhd \varphi.\ A_o, \psi \Rightarrow_D \mathcal{D}_a) \rightarrow\\ &(2)(\forall (a,c) \in D, \theta \in \mathcal{D}_a.\ A_o, \theta \Rightarrow_D \mathcal{D}_c) \rightarrow\\ &(3) A_o \Rightarrow_D \mathcal{D}_c \rightarrow \mathsf{Win}\,(A_p, A_o, D)\,;\,c\end{aligned}$$

Intuitively, this states that the state $(A_p, A_o, D)\,;\,c$ is winning if

(1) For every proponent admission $\varphi \in A_p$ and attack $a \mid \psi \rhd \varphi$ against it, the proponent has remembered a derivation $A_o, \psi \Rightarrow_D \mathcal{D}_a$ which explains how to defend the admission φ against said attack a

(2) For every deferred move $(a, c) \in D$ and every defense against the proponent's attack $\theta \in \mathcal{D}_a$, the proponent has remembered a derivation $A_o, \theta \Rightarrow_D \mathcal{D}_c$ which explains how to continue fending off the deferred challenge c once the opponent defends against a by admitting θ

(3) The proponent knows a derivation $A_o \Rightarrow_D \mathcal{D}_c$, telling her which move to make next in fending off the current challenge c

We will only be able to derive this translation for enumerable rule sets. That is, rule sets for which the attacks on any formula and the defenses against any attack can be enumerated.

Definition 5.15 (Enumerable rule-set) *A rule-set* $(\mathfrak{F}, \mathfrak{F}^*, \mathcal{A}, (\mathcal{D}_a)_{a\in\mathcal{A}}, \rhd)$ *is enumerable if there exist enumerations*

$$e_a : \forall\varphi.\ \mathbb{N} \rightarrow \mathcal{L}(\{(a,\psi) \mid a \mid \psi \rhd \varphi\}) \qquad \textit{and} \qquad e_D : \forall a.\ \mathbb{N} \rightarrow \mathcal{L}(\{\psi \mid \psi \in \mathcal{D}_a\})$$

While the proof proceeds by an inductive argument, there exists no obvious object to whose structure the argument follows. Intuitively, the proof is well-founded as each step of the induction removes one rule from the remembered derivation for c and only remembers new derivations obtained from that rule removal. For now, we do not concern ourselves further with this issue and just give the proof, assuming a suitable induction principle. Afterwards, we demonstrate how such an induction principle can be obtained.

Theorem 5.15 *Under any enumerable rule-set,* $\Rightarrow_D \varphi$ *implies* $\Vdash_S \varphi$.

Proof As the context of $\Rightarrow_D \varphi$ is empty, it can only start with an application of the (R) rule. Thus, we know that φ is non-atomic and for any attack $c \mid \psi \rhd \varphi$, $\psi \Rightarrow_D \mathcal{D}_c$.

We derive Win $([], [\psi], [])\,;\, c$ from this via the generalized statement

$$\begin{aligned}\forall A_P, A_o, D, c.\ (\forall \varphi \in A_p, a \mid \psi \rhd \varphi.\ A_o, \psi \Rightarrow_D \mathcal{D}_a) \to \\ (\forall (a, c) \in D, \theta \in \mathcal{D}_a.\ A_o, \theta \Rightarrow_D \mathcal{D}_c) \to \\ A_o \Rightarrow_D \mathcal{D}_c \to \text{Win}\,(A_p, A_o, D)\,;\, c\end{aligned}$$

which we prove per suitable induction.

We proceed by case distinction on $A_o \Rightarrow_D \mathcal{D}_c$.

(R) Then there is a $\varphi \in \mathcal{D}_c$ which is justified under A_o and for which we know that $\forall a \mid \psi \rhd \varphi.\ A_o, \psi \Rightarrow_D \mathcal{D}_a$. The proponent thus defends against the opponent's challenge, yielding $(A_o, A_p, D)\,;\, c \rightsquigarrow_p (A_o, \varphi :: A_p, D)$ by (PD). We proceed with a case distinction on the opponent's response $(A_o, \varphi :: A_p, D) \rightsquigarrow_o s\,;\, c'$.

(OD) Thus the opponent responds by defending against the last proponent attack a with $(A_o, \varphi :: A_p, (a, c') :: D') \rightsquigarrow_o (\theta :: A_o, \varphi :: A_p, D)\,;\, c'$ by admitting a $\theta \in \mathcal{D}_a$ for $D = (a, c') :: D'$, thereby reposing the deferred challenge c'.

Per assumption, there is a derivation $A_o, \theta \Rightarrow_D \mathcal{D}_{c'}$. As all assumed derivations, including that of $\forall a \mid \psi \rhd \varphi.\ A_o, \psi \Rightarrow_D \mathcal{D}_a$, can be lifted to the context A_o, θ by (Weak), the state $(\theta :: A_o, \varphi :: A_p, D)\,;\, c'$ can be won per inductive hypothesis.

(OA) Thus the opponent responds by attacking one of the admissions ψ by $(A_o, A_p' +\!\!+\, \psi :: A_p'', D) \rightsquigarrow_o (\tau :: A_o, A_p' +\!\!+ A_p'', D)\,;\, c'$ with an attack $c' \mid \tau \rhd \psi$ where $\varphi :: A_p = A_p' +\!\!+\, \psi :: A_p''$.

Per assumption, there is a derivation $A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$. As all assumed derivations, including that of $\forall a \mid \psi \rhd \varphi.\ A_o, \psi \Rightarrow_D \mathcal{D}_a$, can be lifted to the context A_o, θ by (Weak), the state $(\tau :: A_o, A_p' +\!\!+ A_p'', D)\,;\, c'$ can be won per inductive hypothesis.

(L) Then there is a $\varphi \in A_o$ and an attack $a \mid \psi \rhd \varphi$ such that ψ is justified under A_o and $\forall \theta \in \mathcal{D}_a.\ A_o, \theta \Rightarrow_D \mathcal{D}_c$ as well as $A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$ for any $c' \mid \tau \rhd \psi$. The proponent thus attacks one of the opponents admissions, yielding the state transition $(A_o, A_p, D)\,;\, c \rightsquigarrow_p (A_o, \psi :: A_p, (a, c) :: D)$. We proceed with a case distinction on the opponent's response $(A_o, \psi :: A_p, (a, c) :: D) \rightsquigarrow_o s\,;\, c'$.

(OD) Thus the opponent responds by defending against the last proponent attack a with $(A_o, \psi :: A_p, (a, c)) \rightsquigarrow_o (\theta :: A_o, \varphi :: A_p, D)\,;\, c$ by admitting a $\theta \in \mathcal{D}_a$, thereby reposing the deferred original c.

Per assumption, there is a derivation $A_o, \theta \Rightarrow_D \mathcal{D}_{c'}$. As all assumed derivations, including that of $\forall c' \mid \tau \rhd \psi.\ A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$, can be lifted to the context A_o, θ by (Weak), the resulting state $(\theta :: A_o, \varphi :: A_p, D) \,;\, c$ can be won per inductive hypothesis.

(OA) Thus the opponent responds by attacking one of the admissions ν by $(A_o, A'_p \mathbin{+\!\!+} \nu :: A''_p, (a, c) :: D) \rightsquigarrow_o (\tau :: A_o, A'_p \mathbin{+\!\!+} A''_p, (a, c) :: D) \,;\, c'$ with an attack $c' \mid \tau \rhd \nu$ where $\psi :: A_p = A'_p \mathbin{+\!\!+} \nu :: A''_p$.

Per assumption, there is a derivation $A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$. As all assumed derivations, including $\forall c' \mid \tau \rhd \psi.\ A_o, \tau \Rightarrow_D \mathcal{D}_{c'}$ and $\forall \theta \in \mathcal{D}_a.\ A_o, \theta \Rightarrow_D \mathcal{D}_c$, can be lifted to the context A_o, τ by applying (Weak), the resulting state $(\tau :: A_o, A'_p \mathbin{+\!\!+} A''_p, (a, c) :: D) \,;\, c'$ can be won per inductive hypothesis. □

Corollary 5.16 *Under any enumerable rule set, $\Rightarrow_D \varphi$ entails $\Vdash_D \varphi$.*

We now show how to derive a suitable induction principle for the proof above. Recall, that we claimed that the argument was well-founded as each inductive step removes one rule of one of the remembered derivations and only remembers new derivations obtained through that rule-removal. Intuitively, the argument thus proceeds per induction on the size of all of the remembered derivations. However, an LJD derivation may be of unbounded depth as our definition of rule sets in terms of possibly infinite defense sets allows for infinite branching. The size of LJD derivations thus needs to be measured in ordinal numbers, which allow going beyond finite numbers.

Because of time constraints, we were not able to formalize a type of ordinal numbers. We thus simply assume a type with suitable operations and properties. Note that such ordinals have already been formalized in Coq by Grimm [26]. However, technical reasons prevented us from integrating his work into our formalization.

Definition 5.16 (Ordinal numbers) *We assume a type $\mathbb{O}$ of ordinal numbers. This type possesses*

- *An element $O : \mathbb{O}$ and a successor function $S : \mathbb{O} \to \mathbb{O}$*
- *An addition operation $+ : \mathbb{O} \to \mathbb{O} \to \mathbb{O}$*
- *A well-founded preorder $<: \mathbb{O} \to \mathbb{O} \to \mathbb{P}$*
- *A countable supremum $\sup : (\mathbb{N} \to \mathbb{O}) \to \mathbb{O}$ such that $f\,n < \sup f$ for all n*

For rule sets in which the attacks and defense sets are enumerable, we can thus provide a suitable size function.

Lemma 5.17 *For any enumerable rule set one can define a size function on derivations suitable prove Theorem 5.15 by well-founded induction on the size of all remembered derivations.*

Proof With some slight abuses of notation, we define a size function σ as follows.

$$\sigma\,[] := 0 \qquad\qquad \sigma\,(H :: A) := S\,(\sigma\,H + \sigma\,A)$$

$$\sigma\,(H : \forall a \,|\, \psi \rhd \varphi.\, \Gamma, \psi \Rightarrow_D \mathcal{D}_a) := \sup\,(\lambda n.\ \sigma\,(H\,(e_a\,\varphi\,n)))$$

$$\sigma\,(H : \forall \theta \in \mathcal{D}_a.\, \Gamma, \theta \Rightarrow_D \mathcal{D}_c) := \sup\,(\lambda n.\ \sigma\,(H\,(e_D\,a\,n)))$$

$$\sigma\left[\text{L}\frac{H_1 : \forall\,\theta \in \mathcal{D}_a.\, \Gamma, \theta \Rightarrow_D \Delta \quad H_2 : \forall a' \,|\, \tau \rhd \psi.\, \Gamma, \tau \Rightarrow_D \mathcal{D}_{a'}}{\Gamma \Rightarrow_D \Delta}\right] := S\,(\sigma\,H_1 + \sigma\,H_2)$$

$$\sigma\left[\text{R}\frac{H : \forall a \,|\, \psi \rhd \varphi.\, \Gamma, \psi \Rightarrow_D \mathcal{D}_a}{\Gamma \Rightarrow_D \Delta}\right] := S\,(\sigma\,H)$$

By carefully studying the proof of Theorem 5.15, one can convince oneself that this induction can be seen as a well-founded size induction on $\sigma\,A$ where A contains all remembered proofs. □

As the rule set of first-order logic is enumerable, one can derive the soundness and completeness of the full sequent calculus LJ with regards to first-order intuitionistic D-dialogues.

Corollary 5.18

1. *For all* φ, $\Vdash_D \varphi$ *entails* $\Rightarrow_J \varphi$
2. *For all* φ, $\Rightarrow_J \varphi$ *entails* $\Vdash_D \varphi$

5.5 Conclusion

In this chapter, we discussed dialogues as a semantics for first-order logic. To this end, we gave a generalized completeness result for intuitionistic E-dialogues in the style of Sørensen and Urzyczyn [61]. We then derived the completeness for the full intuitionistic sequent calculus LJ with regards to intuitionistic first-order E-dialogues from said abstract result. We also proved a generalized completeness result for intuitionistic D-dialogues with enumerable rule sets, thereby proving the equivalence of E- and D-dialogues with enumerable rule sets and thus the completeness of the full intuitionistic sequent calculus LJ with regards to intuitionistic first-order D-dialogues. The results of this chapter, specialized to the rule set of first-order logic, are summed up by Fig. 5.5.

An interesting difference to the results of Chapter 3 and Chapter 4 is that formal dialogues, as we analyzed them in this chapter, admit constructive completeness

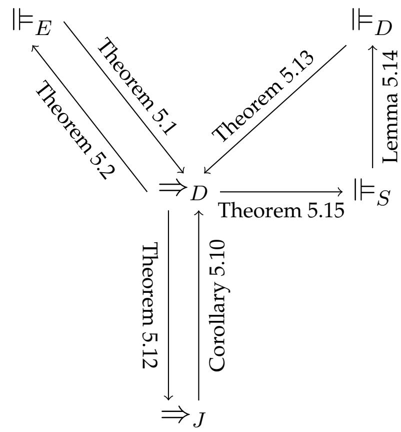

Figure 5.5: Results for Intuitionistic Dialogues

proofs, even on the full fragment of first-order logic, with regards to their canonical formulation. This is due to their syntactic nature: The notion of validity established by formal dialogues does not make any reference to domains and interpretations. Indeed, we have demonstrated in Section 5.2 that winning strategies for E-dialogues are isomorphic to derivations in the dialogical sequent calculus LJD. In that sense, formal dialogues can be thought of as "being less of a model semantics and more of a deduction system", which explains why it is easier to extract deductions from them and thus prove their completeness.

Chapter 6

Discussion

We close this thesis by summing up our results in a brief conclusion in Section 6.1. We then elaborate on related (Section 6.2) and future work (Section 6.3). We end the chapter with some remarks on the formalization which we developed alongside this thesis (Section 6.4).

6.1 Conclusion

In Chapter 3, we analyzed the constructivity of various completeness results for the $\dot\forall$, $\dot\rightarrow$, $\dot\bot$-fragment of classical first-order logic with regards to Tarski models. We showed that for standard models, in which $\dot\bot$ is never satisfied, completeness is equivalent to the stability of classical provability $\mathcal{T} \vdash_{CE} \varphi$ for certain classes of theories. From this, we deduced in particular that Tarski completeness

- on finite theories is equivalent to the object Markov's principle,
- on enumerable theories is equivalent to the synthetic Markov's principle
- on arbitrary theories is equivalent to double-negation elimination

As the theories mathematicians work with in practice are usually enumerable, the first two equivalences are the most interesting among the three. The third one should be regarded as hard statement on the infeasibility of completeness for arbitrary theories in any remotely constructive setting.

The non-constructivity of completeness stems from the requirement that no model may satisfy $\dot\bot$. We thus presented two ways of altering the strict notion of model which enable fully constructive completeness proofs. Exploding models, analogous to those introduced by Veldman [65], allow for models which satisfy $\dot\bot$ as long as they still satisfy all instances of the principle of explosion $\dot\bot \dot\rightarrow \varphi$. Minimal models simply do away with the special role $\dot\bot$ all-together, treating it like an arbitrary logical constant. It should be noted that the absence of $\dot\bot$ means deduction systems featuring the (Exp) rule are not sound with regards to minimal models.

In Chapter 4, we arrived at analogous results for Kripke models. Standard completeness for Kripke semantics also entails the stability of $\mathcal{T} \vdash_{CE} \varphi$. In particular, Kripke completeness

- on finite theories is equivalent to the object Markov's principle
- on enumerable theories entails the synthetic Markov's principle
- on arbitrary theories entails double-negation elimination

We did not formally prove the inverse directions of the last two statements, as we only presented a proof of completeness on finite theories. However, the proof could easily be modified to entail the other two equivalences as well. Similarly to Tarski semantics, Kripke semantics admit constructive completeness proofs if the notion of models is weakened to exploding or minimal models.

In Chapter 5, we analyzed completeness proofs for intuitionistic dialogue semantics, a game semantics modeling a debate about the validity of a formula. For this, we demonstrated a fully constructive proof of the completeness of full intuitionistic first-order logic with regards to both E- and D-dialogues. It might seem surprising that standard dialogues admit a constructive completeness proof for the full fragment while this is impossible for standard model semantics, such as Kripke models, even on smaller syntactical fragments. However, this can be traced back to the fact that winning strategies for dialogues are structurally very similar to deductions in the sequent calculus. One could interpret this as dialogues being "less of a model semantics and more of a deduction system".

6.2 Related Work

Prior Work Before discussing various works that have explored some of the same topics as this thesis, we want to begin by crediting the works on which this thesis' contents are directly based.

The adaption of synthetic computability theory to the calculus of inductive constructions we employ throughout this thesis is due to Forster, Kirst and Smolka [20]. In that paper, they formalize the undecidability of multiple variants of the Entscheidungsproblem in Coq. One of the points of future work they mention is the analysis of various completeness proofs, which has lead to this thesis. The formalization they developed alongside the paper also served as the starting point for our formalization. However, as we switched from the named binders they use to de Bruijn syntax, very little code is actually shared between the two formalizations.

Chapter 3 draws much of its content from the constructive account of Henkin's proof of first-order completeness by Herbelin and Ilik [29]. They show how to define classical Tarski semantics in a constructive setting by only considering classical models and give a clear, constructive account of the model existence theorem for systems with $\dot{\bot}$. They further demonstrate how to derive standard completeness for enumerable theories from Markov's principle and how to constructively obtain completeness for exploding models. Schumm's [59] classical completeness proof for the minimal predicate logic consisting only of predicates and implications are the source of the $\mathcal{E}$-construction and inspired the presentation of the model construction as multiple separate steps.

All completeness proofs as well as parts of the semantic normalization procedure of Chapter 4 stem from Herbelin and Lee's [30] discussion of semantic cut elimination for the $\dot{\forall}$, $\dot{\to}$, $\dot{\bot}$-fragment of first-order logic. They also already formalized their results in Coq, although we did not reuse any parts of their formalization.

The two presentations of dialogues we offer in Chapter 5 are of course based on the respective works by Lorenzen [48, 49] and Felscher [15]. The remarks on the historical background of dialogues we make in the rest of the chapter are informed by Krabbe's historical overview [36]. The approach of the generalized dialogue completeness proofs is based on work by Sørenzen and Urzyczyn [61] in which they present a generalized completeness proof of the classical finitary sequent calculus LKd logic with regards to E-dialogues.

Formalized completeness Because of its historical significance, the completeness of first-order logic has been formalized in many interactive theorem provers such as Isabelle/HOL [7, 57, 58], NUPRL [9, 64], Mizar [8], Lean [27] and Coq [30, 33]. We only go into detail for those which formalize constructive completeness proofs.

Constable and Bickford [9] give a constructive proof of completeness for the BHK-realizers of full intuitionistic first-order logic in NUPRL. For this, they represent the realizer of a formula φ as an intersection type of satisfaction over all models $\bigcap_M M \vDash \varphi$ which is inhabited by terms t which can be typed as $t : M \vDash \varphi$ for any model M. Their proof is fully constructive when realizers are restricted to be normal terms. To prove completeness without this restriction, Brouwer's fan theorem is required to first normalize the realizers.

In his PhD thesis [33], Ilik formalizes multiple constructive proofs of first-order completeness in Coq. His main focus in these proofs is analyzing their computational content to obtain normalization procedures. In Chapter 1, he gives a constructive completeness proof of exploding Boolean models for the $\dot{\forall}$, $\dot{\to}$, $\dot{\bot}$-fragment of classical first-order logic, very similar to that we give in Chapter 3. In Chapters 2 and 3, he proposes very non-standard, exploding Kripke models for full classical and intuitionistic first-order logic. As opposed to the traditional rendering of forcing as a recursive embedding into the meta-logic, Ilik defines it in terms of non-refutation and restricted non-refutation for classical and intuitionistic semantics, respectively. While these notions of model are far removed from their traditional accounts, they allow for constructive completeness proofs on the full syntax of first-order logic.

Constructive analysis of first-order model completeness Kreisel attributes the first proof that the standard completeness of intuitionistic first-order logic implies Markov's principle to Gödel in [37]. This proof is explained to be based on the reduction of primitive recursive relations to the provability of first-order formulas he developed as part of his first incompleteness proof [24]. However, Gödel never published this result.

The first published proof of this result is given four years later by Kreisel in [38], which he describes as an adoption of the proof relayed to him by Gödel simplified with several suggestions by Kleene. Notably, he uses a very general notion of semantic validity that also captures Kripke models, even though they are only introduced as an intuitionistic semantics in 1965. He proves two variants of the theorem. The first states that intuitonistic completeness on full first-order logic entails the stability of $\forall \alpha . \exists n.\ A(n, \alpha)$ for any primitive recursive relation $A(n, \alpha)$ between natural numbers and free choice sequences. For this, he essentially encodes $\exists \alpha . \forall n . \neg A(n, \alpha)$ as a first-order formula φ_A and shows that $\neg\neg\forall \alpha . \exists n.\ A(n, \alpha)$ implies $\vDash \dot{\neg}\varphi_A$. He obtains a derivation of $\vdash \dot{\neg}\varphi_A$ per completeness and extracts the information for $\forall \alpha . \exists n.\ A(n, \alpha)$ from it by means of Herbrand's theorem [31], more specifically a variant Kreisel proves in [37]. The second theorem he proves is, that when the primitive recursive relation $A(n)$ is restricted to be only on the natural numbers, the completeness of the negative formulas (those with all predicates under double-negation and without $\dot{\exists}$ or $\dot{\vee}$) still entails the stability of $\exists n . A(n)$. The crucial insight of this proof is that the formula φ_A can be reduced to a negative formula φ_A^N via Gödel's double-negation translation [25], while still maintaining the property that $\neg\neg\exists n.\ A(n)$ implies $\vDash \dot{\neg}\varphi_A^N$.

The proof of the Kreisel-Gödel theorem in [38] has since inspired a range of works deriving related non-constructivity results for different kinds of completeness. Especially noteworthy among these results are the Church-Turing thesis entailing that the set of intuitionistically valid first-order formulas in not enumerable [39, 45], the Church-Turing thesis and Markov's principle entailing that it is not arithmetically definable [52], the Church-Turing thesis entailing that pure intuitionistic first-order logic is incomplete [54] and a much more general result about the unprovability of regular logics in intuitionistic meta-theories [53].

By almost exclusively focusing our analysis on the completeness of the $\dot{\forall}$, $\dot{\rightarrow}$, $\dot{\bot}$-fragment of first-order logic, we were able to pinpoint how $\dot{\bot}$ prevents the constructivity of the result. However, we did not further concern ourselves with analysis of what $\dot{\exists}$ and $\dot{\vee}$ contribute to the non-constructivity of full completeness. Krivtsov [43, 44] does the exact opposite: He analyzes completeness with regards to exploding Tarski and Beth models, for full classical and intuitionistic first-order logic, respectively. His analysis reveals both completeness results to be equivalent to the weak fan theorem, thereby tightly characterizing the type of non-constructivity introduced when considering $\dot{\vee}$ and $\dot{\exists}$ for completeness. Espindola offers a similar, more abstract, result by showing that an abstract version of the model existence theorem for full first-order logic is equivalent to the boolean prime ideal theorem over intuitionistic ZF set theory [14].

Our approach to constructive analysis of completeness differs from that usually taken as we consider very specific notions of model, whereas these are often left more abstract in other works [38, 53, 14]. A very noteworthy work in this line of abstract analysis of completeness is that by Berardi [4]. He adopts a very general notion of models and interpretations to analyze which variants allow for constructive completeness proofs of classical first-order logic. He concludes that one needs to interpret $\dot{\neg}$ and $\dot{\rightarrow}$ in ways stronger than intuitonistic $\neg$ and $\rightarrow$ as well as $\dot{\vee}$ and $\dot{\exists}$ in ways weaker than the intuitonistic $\vee$ and $\exists$.

Completeness of first-order dialogues The first completeness proofs with regards to dialogues was given by Lorenz in PhD thesis [46]. These proofs were constructive and used the tableaux systems as deduction systems. Later, noteworthy constructive proofs of first-order completeness with regards to dialogues were given by Stegmüller [63], Felscher [15] and Krabbe [36]. Notably, Felscher and Krabbe prove completeness with regards to both E- and D-dialogues.

6.3 Future Work

Some of the results in this thesis were only discussed informally or not formalized fully. In particular, the equivalence between F-stability and the object Markov's principle as discussed in Section 2.5.4, the proof that a C-stability closed under substitution and extension entails Kripke completeness for theories of C we discuss after the proof of Theorem 4.15 and the ordinal number arithmetic involved in the proof of S-soundness of LJD in Section 5.4.2. It is of course desirable to formalize these results as well. Formalizing the extended Kripke completeness result should be straightforward, as it only constitutes a slight variation of the result in Chapter 4. Doing the same for the equivalence between F-stability and object Markov's principle, on the other hand, would certainly be more involved but could be eased by relying on previous work by Forster and Kunze [17]. Suitable ordinal numbers have already been formalized in Coq by Grimm [26]. Integrating his formalization into ours would be difficult, as he uses the ssreflect framework and we do not. However, it should be possible to adopt his approach to standard Coq.

This thesis focuses on the role of $\dot{\bot}$ in completeness proofs for the $\dot{\forall}$, $\dot{\rightarrow}$, $\dot{\bot}$-fragment of first-order logic. As mentioned in the previous section, Krivtsov [43, 44] has focused on the other difficulty of completeness: $\dot{\vee}$ and $\dot{\exists}$. He proves that various exploding completeness results for the full fragment are equivalent to the weak fan theorem. It would thus be of interest to try and connect these two results. Indeed, we conjecture that standard Tarski or Kripke completeness on a theory class C is equivalent precisely to C-stability and the weak fan theorem taken together. If this was the case, it would constitute a very insightful characterization of the requirements of completeness in a constructive setting.

In Section 5.1, we introduce material and formal dialogues. However, we then go on to only formally analyze the completeness of formal dialogues. The conclusion we arrived at based on that analysis was that formal dialogues lend themselves to a constructive completeness proof in their canonical formulation as they are structurally very close to sequent calculi. Material dialogues, however, seem to be much closer to model semantics. Recall that material dialogues were defined in terms of an underlying game, claims of the atomic formulas representing claims of "being able to win" certain constellations of said game. These underlying games can be formalized as a domain of game objects D together with a function interpretation $f^{\mathcal{I}} : D^{|f|} \to D$ for every $f : \mathcal{F}$. Which game constellations are claimed to be winnable by claiming a certain atomic formula would then be defined by a predicate interpretation $P^{\mathcal{I}} : D^{|P|} \to \mathbb{P}$ for every predicate $P : \mathcal{P}$. If defined with this notion of underlying game, material dialogues seem to be much more similar to regular model semantics, such as Tarski or Kripke models, than formal dialogues. This insight leads us to conjecture that their completeness proofs should thus exhibit the same constructivity properties as other model semantics. That is, standard material dialogue completeness should be equivalent to C-stability, while exploding and minimal completeness should be provable fully constructively on the $\dot\forall, \dot\to, \dot\bot$-fragment of first-order logic.

As we have mentioned already, the results of this thesis are very specific as we always reason about very concrete notions of model. Note, however, that the proofs of the fact that standard completeness entails stability are both very generic. Indeed, as laid out by McCarty [53], the proof we give for standard Tarski models can be generalized to arbitrary notions of semantic validity as long as they are sound and stable. Analogously, the proof we give for standard Kripke models can be extended to any sound notion of semantic validity with the property that $\Vdash \dot\neg\varphi$ iff $\neg \Vdash \varphi$. In this vein, it would be interesting to explore how far other results of this thesis can be generalized. An especially interesting question in this direction is if one can define a generalized notion of intuitionistic validity that is abstract enough to cover multiple intuitionistic semantics (such as Kripke models, Beth models, material dialogues) while still being specific enough to allow for a unified completeness proof in the style of that given in Chapter 4. This could yield a very insightful characterization of the non-constructivity of intuitionistic completeness.

6.4 Formalization

In this section we remark on the noteworthy aspects of the formalization that was developed alongside this thesis.

Synthetic computability theory has already been studied extensively in Coq. A lot of this work originates with a project by Forster et al. [19], which is mainly focused on deriving undecidability results by means of many-one reductions. As part of this ongoing project, a framework for synthetic computability theory has been developed which we also employ in our formalization. It contains common definitions, including those in Section 2.2, as well as lemmas and tactics that ease working with them.

As we explained in Section 2.3.1, the syntax we use is defined in terms of a signature Σ. As we want the signature to be inferred automatically, it is defined as a type class in Coq. The production rules of functions and predicates are represented by constructors taking a vector of terms to allow for the arities to be fixed by the signature. For example, the constructor for function application in the term language is $\mathsf{Func} : \forall \Sigma (f : \mathcal{F}_\Sigma).\ \mathsf{vector}\, \mathfrak{T}_\Sigma\, |f| \to \mathfrak{T}_\Sigma$. As Coq is unable to automatically generate a useful induction principle for nested inductive types, we had to provide and prove one manually. The variant for terms that has proven to be the most useful is given below.

$$\forall \Sigma (p : \mathfrak{T}_\Sigma \to \mathbb{P}).\ (\forall (f : \mathcal{F}_\Sigma)(v : \mathfrak{T}_\Sigma^{|f|}).\ (\forall t.\ t \in v \to p\, t) \to p\,(\mathsf{Func}\, f\, v)) \to \\ (\forall x : \mathbb{N}.\ p\, x) \to \forall t.\ p\, t$$

Stark et al. [62] have developed a tool called "AutoSubst 2" that eases working with de Bruijn syntax in Coq. Given a specification in higher-order abstract syntax, it generates the corresponding inductive types with de Bruijn binders. Additionally, it defines suitable substitution functions, proves that it constitutes a σ-calculus [1] and offers tactics that allow for automatic simplifications of expressions containing substitutions. We have used this tool to generate our syntax, including the generalization over signatures and the vector arguments. However, the facts about fresh variables and their interaction with substitutions can currently not be generated by AutoSubst2 and were therefore proven manually.

A related, important decision we made in the design of our first-order languages was to not formalize them as scoped syntax, even though AutoSubst2 could generate suitable code for this. Instead of allowing arbitrary free variables to occur in any term, scoped syntax limits the variables that may occur in a term to stay below a certain bound. In a type theory, this can be expressed with an index in the type of terms indicating the chosen bound. For example, $\mathfrak{T}_0$ would represent the closed terms, as only variables strictly less than 0 are allowed to occur in values of that type. A scoped first-order language could be very useful. For example, the axiom of replacement of ZF set theory is defined in terms of binary relations $\varphi(x, y)$, represented by formulas φ with two free variables. With scoped syntax, this could simply be represented by requiring that $\varphi : \mathfrak{F}_2$. There are two reasons for our decision against scoped syntax. First, recall the construction of the theory Ω for the Tarski completeness proofs. One of its steps, the Henkin construction $\mathcal{H}$, results in a theory which contains a free occurrence of every variable. Representing this with a scoped syntax is not impossible but would certainly introduce many complications. The second reason is that scoped syntax, in combination with the term signatures Σ, introduce the possibility of the term language being empty: simply choose a signature without function constants and formulas of scope 0. This is no problem for deduction systems, as the (AllI) rule raises the scope of all formulas in the subsequent deduction by one. However, as dialogue games are played on a fixed type of formulas, formulas like $\dot\neg(\dot\forall\dot\bot)$ would not be dialogically valid on these languages without terms, as the universal quantifier could not be attacked anymore, introducing a mismatch between deduction systems and semantics.

The proof of Tarski completeness requires many derivations in the object natural deduction system. To ease this type of proof, we developed a domain specific tactics language to work in the natural deduction system, inspired by that of Coq. An example, a proof of the drinker's paradox, in that language is given below.

```
Lemma DP {Sigma: Signature} phi:
  [] ⊢CE ∃̇(phi →̇ (∀̇ phi)[↑]).
Proof.
  oxm (∃̇ ¬̇ phi).
  - odestruct 0. oexists (var_term 0). ointros. oexfalso. oapply 1. ctx.
  - oexists (var_term 0). ointros. oindirect. oapply 2. oexists (var_term 0). ctx.
Qed.
```

The line count of the files of the formalization associated with each chapter is given below. It should be noted that the code associated with the Preliminaries was not fully written by us, but also contains syntax generated by AutoSubst 2 and parts of the library for synthetic computability theory developed by Forster, Kirst and Smolka in the course of [20], which accounts for about 1600 lines of code.

Chapter	**Specification**	**Proofs**
Preliminaries	1723	1606
Tarski Semantics	572	668
Kripke Semantics	321	202
Dialogue Semantics	467	532
Total	3083	3008

Bibliography

[1] Martin Abadi, Luca Cardelli, P. L. Curien, and J. J. Lévy. Explicit substitutions. *Journal of functional programming*, 1(4):375–416, 1991.

[2] Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. Mechanized metatheory for the masses: the p opl m ark challenge. In *International Conference on Theorem Proving in Higher Order Logics*, pages 50–65. Springer, 2005.

[3] Andrej Bauer. First steps in synthetic computability theory. *Electronic Notes in Theoretical Computer Science*, 155:5–31, 2006.

[4] Stefano Berardi. Intuitionistic completeness for first order classical logic. *The Journal of Symbolic Logic*, 64(1):304–312, 1999.

[5] Ulrich Berger and Helmut Schwichtenberg. An inverse of the evaluation functional for typed lambda-calculus. In [*1991*] *Proceedings Sixth Annual IEEE Symposium on Logic in Computer Science*, pages 203–211. IEEE, 1991.

[6] Errett Bishop. Foundations of constructive analysis. 1967.

[7] Jasmin Christian Blanchette, Andrei Popescu, and Dmitriy Traytel. Unified classical logic completeness. In *International Joint Conference on Automated Reasoning*, pages 46–60. Springer, 2014.

[8] Patrick Braselmann and Peter Koepke. Gödel's completeness theorem. *Formalized Mathematics*, 13(1):49–53, 2005.

[9] Robert Constable and Mark Bickford. Intuitionistic completeness of first-order logic. *Annals of Pure and Applied Logic*, 165(1):164–198, 2014.

[10] Thierry Coquand and Gérard Huet. *The calculus of constructions*. PhD thesis, INRIA, 1986.

[11] Nicolaas Govert de Bruijn. AUTOMATH, a language for mathematics, 1968.

[12] Nicolaas Govert de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. In *Indagationes Mathematicae (Proceedings)*, volume 75, pages 381–392. Elsevier, 1972.

[13] Peter Dybjer and Andrzej Filinski. Normalization and partial evaluation. In *International Summer School on Applied Semantics*, pages 137–192. Springer, 2000.

[14] Christian Espíndola et al. Semantic completeness of first-order theories in constructive reverse mathematics. *Notre Dame Journal of Formal Logic*, 57(2): 281–286, 2016.

[15] Walter Felscher. Dialogues, strategies, and intuitionistic provability. *Annals of pure and applied logic*, 28(3):217–254, 1985.

[16] Yannick Forster. First-order logic and Markov's principle. Draft, 2019. URL `https://www.ps.uni-saarland.de/~forster/drafts/MarkovFOL.pdf`.

[17] Yannick Forster and Fabian Kunze. Verified extraction from Coq to a lambda-calculus. In *Coq Workshop*, volume 2016, 2016.

[18] Yannick Forster and Gert Smolka. Weak call-by-value lambda calculus as a model of computation in Coq. In *International Conference on Interactive Theorem Proving*, pages 189–206. Springer, 2017.

[19] Yannick Forster, Edith Heiter, Dominik Kirst, Dominique Larchey-Wendling, and Gert Smolka. A library of formalised undecidable problems in Coq. URL `https://github.com/uds-psl/coq-library-undecidability`.

[20] Yannick Forster, Dominik Kirst, and Gert Smolka. On synthetic undecidability in Coq, with an application to the Entscheidungsproblem. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 38–51. ACM, 2019.

[21] Gerhard Gentzen. Untersuchungen über das logische Schließen. I. *Mathematische zeitschrift*, 39(1):176–210, 1935.

[22] Gerhard Gentzen. Untersuchungen über das logische Schließen. II. *Mathematische Zeitschrift*, 39(1):405–431, 1935.

[23] Kurt Gödel. *Über die Vollständigkeit des Logikkalküls*. 1929.

[24] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für mathematik und physik*, 38(1):173–198, 1931.

[25] Kurt Gödel. Zur intuitionistischen Arithmetik und Zahlentheorie. *Ergebnisse eines mathematischen Kolloquiums*, 4(1933):34–38, 1933.

[26] José Grimm. Implementation of three types of ordinals in Coq. 2013.

[27] Jesse Han and Floris van Doorn. A formalization of forcing and the consistency of the failure of the continuum hypothesis. In *International Conference on Interactive Theorem Proving*. Springer, 2019.

[28] Leon Henkin. The completeness of the first-order functional calculus. *The Journal of Symbolic Logic*, 14(3):159–166, 1949. ISSN 00224812.

[29] Hugo Herbelin and Danko Ilik. An analysis of the constructive content of Henkin's proof of Gödel's completeness theorem. Draft, 2016. URL `http://pauillac.inria.fr/~herbelin/articles/godel-completeness-draft16.pdf`.

[30] Hugo Herbelin and Gyesik Lee. Forcing-based cut-elimination for Gentzen-style intuitionistic sequent calculus. In *International Workshop on Logic, Language, Information, and Computation*, pages 209–217. Springer, 2009.

[31] Jacques Herbrand. *Recherches sur la théorie de la démonstration*. PhD thesis, Université de Paris, 1930.

[32] Arend Heyting. *Intuitionism: an introduction*, volume 41. Elsevier, 1966.

[33] Danko Ilik. *Constructive completeness proofs and delimited control*. PhD thesis, Ecole Polytechnique X, 2010.

[34] Hajime Ishihara. Reverse mathematics in bishop's constructive mathematics. *Philosophia Scientiæ. Travaux d'histoire et de philosophie des sciences*, (CS 6):43–59, 2006.

[35] Erik C. W. Krabbe. Formal systems of dialogue rules. *Synthese*, 63(3):295–328, 1985.

[36] Erik C. W. Krabbe. Dialogue logic. In *Handbook of the History of Logic*, volume 7, pages 665–704. Elsevier, 2006.

[37] Georg Kreisel. Elementary completeness properties of intuitionistic logic with a note on negations of prenex formulae. *The Journal of Symbolic Logic*, 23(3): 317–330, 1958.

[38] Georg Kreisel. On weak completeness of intuitionistic predicate logic. *The Journal of Symbolic Logic*, 27(2):139–158, 1962.

[39] Georg Kreisel and Anne S Troelstra. Formal systems for some branches of intuitionistic analysis. *Annals of mathematical logic*, 1(3):229–387, 1970.

[40] Saul A. Kripke. Semantical analysis of modal logic I. Normal modal propositional calculi. *Mathematical Logic Quarterly*, 9(5-6):67–96, 1963.

[41] Saul A. Kripke. Semantical analysis of intuitionistic logic I. In *Studies in Logic and the Foundations of Mathematics*, volume 40, pages 92–130. Elsevier, 1965.

[42] Jean-Louis Krivine. Une preuve formelle et intuitionniste du théorème de complétude de la logique classique. *Bulletin of Symbolic Logic*, 2(4):405–421, 1996.

[43] Victor N. Krivtsov. An intuitionistic completeness theorem for classical predicate logic. *Studia Logica*, 96(1):109–115, 2010.

[44] Victor N. Krivtsov. Semantical completeness of first-order predicate logic and the weak fan theorem. *Studia Logica*, 103(3):623–638, 2015.

[45] Daniel Leivant. Failure of completeness properties of intuitionistic predicate logic for constructive models. *Annales scientifiques de l'Université de Clermont. Mathématiques*, 60(13):93–107, 1976.

[46] Kuno Lorenz. *Arithmetik und Logik als Spiele*. PhD thesis, Christian-Albrechts-Universität zu Kiel, 1961.

[47] Kuno Lorenz. Dialogspiele als semantische Grundlage von Logikkalkülen. *Archive for Mathematical Logic*, 11(3):73–100, 1968.

[48] Paul Lorenzen. Logik und Agon. In *Atti del XII Congresso Internazionale di Filosofia*, volume 4, pages 187–194, 1960.

[49] Paul Lorenzen. Ein dialogisches Konstruktivitätskriterium. In *Proceedings of the Symposium on Foundations of Mathematics (Warsaw, 2 – 9 September 1959)*, pages 193–200, 1961.

[50] Bassel Mannaa and Thierry Coquand. The independence of Markov's principle in type theory. *Logical Methods in Computer Science*, 13, 2017.

[51] Per Martin-Löf and Giovanni Sambin. *Intuitionistic type theory*, volume 9. Bibliopolis Naples, 1984.

[52] Charles McCarty. Constructive validity is nonarithmetic. *The Journal of Symbolic Logic*, 53:1036–1041, 1988.

[53] Charles McCarty. Completeness and incompleteness for intuitionistic logic. *The Journal of Symbolic Logic*, 73(4):1315–1327, 2008.

[54] David Charles McCarty et al. Incompleteness in intuitionistic metamathematics. *Notre Dame journal of formal logic*, 32(3):323–358, 1991.

[55] Pierre-Marie Pédrot and Nicolas Tabareau. Failure is not an option. In *European Symposium on Programming*, pages 245–271. Springer, 2018.

[56] Fred Richman. Church's thesis without tears. *The Journal of symbolic logic*, 48 (3):797–803, 1983.

[57] Tom Ridge and James Margetson. A mechanically verified, sound and complete theorem prover for first order logic. In *International Conference on Theorem Proving in Higher Order Logics*, pages 294–309. Springer, 2005.

[58] Anders Schlichtkrull. Formalization of the resolution calculus for first-order logic. *Journal of Automated Reasoning*, 61(1-4):455–484, 2018.

[59] George F. Schumm. A Henkin-style completeness proof for the pure implicational calculus. *Notre Dame J. Formal Logic*, 16(3):402–404, July 1975.

[60] Stephen G. Simpson. Reverse mathematics. In *Proc. Symposia Pure Math*, volume 42, pages 461–471, 1985.

[61] Morten Heine Sørensen and Pawel Urzyczyn. Sequent calculus, dialogues, and cut elimination. *Reflections on Type Theory, λ-Calculus, and the Mind*, pages 253–261, 2007.

[62] Kathrin Stark, Steven Schäfer, and Jonas Kaiser. Autosubst 2: reasoning with multi-sorted de Bruijn terms and vector substitutions. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 166–180. ACM, 2019.

[63] Wolfgang Stegmüller et al. Remarks on the completeness of logical systems relative to the validity-concepts of p. lorenzen and k. lorenz. *Notre Dame Journal of Formal Logic*, 5(2):81–112, 1964.

[64] Judith Underwood. Aspects of the computational content of proofs. Technical report, Cornell University, 1994.

[65] Wim Veldman. An intuitiomstic completeness theorem for intuitionistic predicate logic 1. *The Journal of Symbolic Logic*, 41(1):159–166, 1976.