Decidable Fragments and Skolemization in Superintuitionistic and Modal Logics

Brian Burke Lewis
Stanford University
Undergraduate Honors Thesis

June 6, 2012

Chapter 1

ESkolemization in LC

1.1 Preliminaries

In the setting of classical logic, Skolemization allows one to remove strong quantifiers from formulas via the introduction of fresh function symbols in such a way that the new formla is equiderivable with the original one [16]. In intuitionistic logic, however, the process of Skolemization fails to be faithful: there are some underivable formulas in the calculus LJ whose Skolemizations are derivable. In fact, it is not difficult to see that the sequent

$$\forall x(Ax \wedge B) \Rightarrow (\forall x Ax \wedge B)$$

is underivable in LJ, whereas its Skolemization

$$\forall x(Ax \wedge B) \Rightarrow Ac \wedge B$$

is derivable in LJ. The work of Mints in [14] provides a complete characterization of the cases in intuitionistic logic for which Skolemization is faithful.

In [2], Baaz and Iemhoff define an alternative method called eSkolemization that is sound and complete for a class of formlas in LJ properly extending those for which ordinary Skolemization is intuitionistically sound and complete. The method makes use of an existence predicate, a device first considered by Scott in [15].

We recall some notation from [2]. Given a language $\mathcal{L}$, let $\mathcal{L}'$ extend $\mathcal{L}$ via the addition of the existence predicate, infinitely many Skolem functions for every arity, and infinitely many variables. Here, we assume that $\mathcal{L}$ has infinitely many constants, no variables, and no function symbols of arity ≥ 1.

As usual, given a set D, $\mathcal{L}_D$ denotes the language obtained from $\mathcal{L}$ by adding constants for all elements of D. $\mathcal{S}_{\mathcal{L}}$ denotes the collection of sequents in the language $\mathcal{L}$, and similarly for $\mathcal{L}'$. The eSkolem sequence of a formula A is the sequence of formulas $A = A_1, \ldots, A_n = A^s$ such that A_n does not contain any strong quantifiers and A_{i+1} is the result of replacing the first strong quantifier $QxB(x)$ in A_i by

$$Ef(y_1, \ldots, y_m) \to B(f(y_1, \ldots, y_m))\text{if } Q \text{ is } \forall, \text{ or by}$$

$$Ef(y_1, \ldots, y_m) \wedge B(f(y_1, \ldots, y_m))\text{if } Q \text{ is } \exists$$

where f is a Skolem function not occurring in A_i, and the weak quantifiers in the scope of which $QxB(x)$ occurs are exactly $Qy_1, \ldots, Qy_m$. This definition extends to sequents by treating $\Gamma \Rightarrow \Delta$ as $\bigwedge \Gamma \Rightarrow \bigvee \Delta$.

Define $\Sigma_{\mathcal{L}}$ as $\{\Gamma \Rightarrow Et \mid t \text{ a term in } \mathcal{L}, \Gamma \text{ a multi-set of formulas from } \mathcal{L}'\}$. The idea is that all terms from the original language exist at the root of any Kripke model, while terms involving the Skolem functions might exist only in worlds further along the model. LJE refers to the Gentzen-style calculus for the existence predicate introduced in [3]. The rules are exactly as in LJ, except that in the quantifier rules additional assumptions on the existence of terms must be made. LJE(Σ_L) is defined to be $\{S \in \mathcal{S}_{\mathcal{L}'} \mid \Sigma_{\mathcal{L}} \vdash_{\text{LJE}} S\}$. We write $\vdash_{\mathcal{L}} S$ when $S \in \text{LJE}(\Sigma_L)$, indicating that S is derivable under the assumption that all terms in $\mathcal{L}$ exist. $\Vdash^e$ denotes the usual notion of forcing in the language with the existence predicate; $\Vdash^e_{\mathcal{L}}$ denotes truth in all models in which Et is forced at all worlds for all $t \in \mathcal{L}$.

Baaz and Iemhoff show in [2] that eSkolemization is sound and complete for sequents whose strong quantifiers are all existential. Their proof yields as a corollary the decidability of $\vdash_{\text{LJ}}$ for sequents in which the only quantifiers are strong and existential. The aim of the present paper is to extend these results to the superintuitionistic logic LC. First we remind readers of the results and proofs from [2], [3].

1.2 Previous results

Proposition 1.2.1 (4.11 of [3]) *For every sequent S in $\mathcal{L}$ not containing E,* LJ $\vdash S$ *if and only if* LJE$(\Sigma_{\mathcal{L}}) \vdash S$.

Proof For the left-to-right direction, a proof in LJ of S can be modified by adding the appropriate $\Gamma \Rightarrow Et$ to the hypothesis of $R\exists$ and $L\forall$ rules or adding the appropriate Ey to the antecedent of $L\exists$ and $R\forall$ rules. Here, the proof in LJ is assumed to be cut-free, so it uses as terms only eigenvariables and terms from $\mathcal{L}$, so each of these added formulas belongs to $\Sigma_{\mathcal{L}}$. For the right-to-left direction, one replaces all instances of Et by $\top$ and checks that the resulting tree is a proof in LJ.

Proposition 1.2.2 (4.11 of [2]) *For quantifier-free closed sequents, the relations $\vdash_{\mathrm{LJE}}$ and $\vdash_{\mathcal{L}}$ are decidable.*

Theorem 1.2.3 (7.10 of [2]) *For all sets of closed sequents $\mathcal{S}$ and all closed sequents S in $\mathcal{L}'$, $\mathcal{S} \vdash_{\mathcal{L}} S$ if and only if $\mathcal{S} \Vdash^{e}_{\mathcal{L}} S$. In other words,* LJE($\Sigma_{\mathcal{L}}$) *is sound and complete with respect to the forcing relation $\Vdash^{e}_{\mathcal{L}}$.*

Baaz and Iemhoff use this completeness result to derive the following theorem.

Theorem 1.2.4 (8.10 of [2]) *Let $\exists x B(x, \bar{y})$ be a formula that occurs negatively in a closed sequent S in $\mathcal{L}$, is not in the scope of strong quantifiers, and where $\bar{y}$ are the variables of the weak quantifiers in the scope of which B occurs. If S' is the result of replacing in S the formula $\exists x B(x, \bar{y})$ by $Ef(\bar{y}) \wedge B(f(\bar{y}), \bar{y})$ where f is a fresh function symbol not in $\mathcal{L}$, then $\vdash_{\mathcal{L}} S$ if and only if $\vdash_{\mathcal{L}} S'$.*

The right-to-left direction is the difficult one. Given a model falsifying S, a model falsifying S' is constructed. In the course of this procedure, the worlds and accessibility relation from the first model carry over into the new one. Hence the proof of this theorem still works as long as the existence calculus is sound and complete with respect to $\Vdash^{e}_{\mathcal{L}}$ over the same class of frames $\mathcal{F}$ for which te original calculus is sound and complete with respect to the usual Kripke semantics.

Corollary 1.2.5 (8.9 of [2]) *For each closed sequent S in $\mathcal{L}'$ in which all strong quantifiers are existential, $\vdash_{\mathrm{LJ}} S$ iff $\vdash_{\mathcal{L}} S^{s}$. For the fragment of sequents without weak quantifiers and in which all strong quantifiers are existential, derivability in* LJ *is decidable.*

1.3 The tableaux calculi LCE and LCE(Σ_L)

In this section we define the system LCE, an analogue of the tableaux calculus LC extended by the existence predicate. The rules for LC match those of LCE with every instance of Et replaced by $\top$. LC is sound and complete over the class of Kripke models for which the accessibility relation is a linear order. The completeness proof is entirely analogous to the one presented here, and so we skip it. Recall that a Hilbert-style system characterizing this class of frames is obtained by adding $(p \to q) \vee (q \to p)$ to the standard intuitionistic Hilbert-style system.

As usual, terms—including variables—range over existing and non-existing elements, while the quantifiers range over existing objects only. In [8], Iemhoff provides a connection between linear models with constant domains and those with non-constant domains making use of the existence predicate. In this paper, the domains considered for LC are nested but not necessarily constant. The addition of the existence predicate lets us have constant domains in the semantics of LCE.

An indexed sequent is a sequent prefixed by a sequence of natural numbers, which is said to be the *index* of the sequent. Given a sequence σ and a natural number n not occurring in σ, $\sigma \star n$ is defined by placing n at the end of the sequence σ. We say that the sequence τ *extends* σ if τ is of the form $\sigma \star j_1 \star \cdots \star j_n$ for some natural number n (including the case $\sigma = \tau$, when $n = 0$). A tableau is defined to be a multi-set of indexed sequents. We say that a tableau U is *well-formed* if, whenever σ and τ are indices occurring in U, either τ extends σ or σ extends τ. In the rules of the system presented below, all tableaux are assumed to be well-formed.[1] Note that rule applications preserve well-formedness upward and downward.

The system LCE

$$\frac{}{U;\ \sigma\ \Gamma, P \Rightarrow P, \Delta}\ Ax \qquad\qquad \frac{}{U;\ \sigma\ \Gamma, \bot \Rightarrow \Delta}\ \bot$$

[1] In fact, indices can be eliminated entirely in the presentation that follows, resulting in a so-called hypersequent system akin to those found in [1] and [4]. Under this change, the linear ordering on a multiset of sequents can be read off of the order in which the sequents are written, rather than the explicit picture provided by indices.

$$\frac{U;\ \sigma\ \Gamma, A, B \Rightarrow \Delta}{U;\ \sigma\ \Gamma, A \wedge B \Rightarrow \Delta}\ L\wedge \qquad \frac{U;\ \sigma\ \Gamma \Rightarrow \Delta, A_0, A_1}{U;\ \sigma\ \Gamma \Rightarrow \Delta, A_0 \vee A_1}\ R\vee$$

$$\frac{U;\ \sigma\ \Gamma \Rightarrow A \qquad U;\ \sigma\ \Gamma \Rightarrow B}{U;\ \sigma\ \Gamma \Rightarrow A \wedge B}\ R\wedge \qquad \frac{U;\ \sigma\ \Gamma \Rightarrow A, \Delta \qquad U;\ \sigma\ \Gamma, B \Rightarrow \Delta}{U;\ \sigma\ \Gamma, A \to B \Rightarrow \Delta}\ L\to$$

$$\frac{U;\ \sigma\ \Gamma, A \Rightarrow \Delta \qquad U;\ \sigma\ \Gamma, B \Rightarrow \Delta}{U;\ \sigma\ \Gamma, A \vee B \Rightarrow \Delta}\ L\vee \qquad \frac{U;\ \sigma\ Ab, Eb, \Gamma \Rightarrow \Delta}{U;\ \sigma\ \exists x Ax, \Gamma \Rightarrow \Delta}\ L\exists$$

$$\frac{U;\ \sigma\ \forall x Ax, \Gamma \Rightarrow \Delta, Et \qquad U;\ \sigma\ At, \Gamma \Rightarrow \Delta}{U;\ \sigma\ \forall x Ax, \Gamma \Rightarrow \Delta}\ L\forall$$

$$\frac{U;\ \sigma\ \Gamma \Rightarrow \exists x Ax, Et, \Delta \qquad U;\ \sigma\ \Gamma \Rightarrow \exists x Ax, At, \Delta}{U;\ \sigma\ \Gamma \Rightarrow \exists x Ax, \Delta}\ R\exists$$

$$\frac{U;\ \sigma\ A, \Gamma \Rightarrow \Delta\ ;\ \tau\ A, \Pi \Rightarrow \Phi}{U;\ \sigma\ A, \Gamma \Rightarrow \Delta\ ;\ \tau\ \Pi \Rightarrow \Phi}\ Trans$$

$$\frac{U; T_0; \sigma \star n\ S; U_0' \qquad U; T_0; U_1; \sigma \star n_1 \star n\ S; U_1' \qquad \cdots \qquad U; T_0; U_k; \sigma \star n_1 \star \ldots \star n\ S}{U; \sigma\ \Gamma \Rightarrow \Delta, A \to B; \sigma \star n_1\ T_1; \sigma \star n_1 \star n_2\ T_2; \ldots; \sigma \star n_1 \star \ldots \star n_k\ T_k}\ R\to$$

$$\frac{U; T_0; \sigma \star n\ S; U_0' \qquad U; T_0; U_1; \sigma \star n_1 \star n\ S; U_1' \qquad \cdots \qquad U; T_0; U_k; \sigma \star n_1 \star \ldots \star n\ S}{U; \sigma\ \Gamma \Rightarrow \Delta, \forall Ax; \sigma \star n_1\ T_1; \sigma \star n_1 \star n_2\ T_2; \ldots; \sigma \star n_1 \star \ldots \star n_k\ T_k}\ R\forall$$

In the rule $L\exists$, b is taken to be a fresh variable. The transfer rule is applicable only when τ extends σ.

The rules $R \to$ and $R\forall$ merit special explanations. In each of these rules, n is taken to be a fresh natural number and σ extends all of the indices of sequents from U. T_0 is defined as the indexed sequent $\sigma\ \Gamma \Rightarrow \Delta, A \to B$ or $\sigma\ \Gamma \Rightarrow \Delta, \forall x Ax$ according to the rule. U_j is defined to be the tableau $\sigma \star n_1\ T_1; \cdots; \sigma \star n_1 \star \cdots \star n_j\ T_j$. U_j' is defined to be the tableau

$\sigma \star n_1 \star \cdots \star n_j \star n \star n_{j+1}\ T_{j+1}; \cdots ; \sigma \star n_1 \star \cdots \star n_j \star n \star \cdots \star n_k\ T_k$. In $R \rightarrow$, S is taken to be $\Gamma, A \Rightarrow B$. In $R\forall$, S is taken to be $\Gamma, Eb \Rightarrow Ab$, where b is taken to be a fresh variable.

We write LCE $\vdash U$ if the tableaux U is derivable in LCE. For a set of tableaux Σ, we say that S is derivable from Σ if S is derivable in LCE extended by axioms from Σ. Define LCE(Σ) to be the set of tableaux derivable in LCE extended by axioms Σ. As in the case of LJE, no particular term is assumed to exist. As a consequence, $\forall x Px \Rightarrow Pt$ is underivable, while $\forall x Px, Et \Rightarrow Pt$ is derivable. Ultimately, the original terms from $\mathcal{L}$ will be assumed to exist. This motivates the following definition:

$$\Sigma_{\mathcal{L}} \equiv_{def} \{U; \sigma\ \Gamma \Rightarrow Et \mid t \text{ a term in } \mathcal{L}, \Gamma \text{ a multiset in } \mathcal{L}', \\ \sigma \text{ a sequence of natural numbers}, U \text{ a tableau}\}$$

In the above definition, we take only tableaux U such that the resulting tableau is well-formed. We write $\vdash_{\mathcal{L}} S$ when S is derivable from $\Sigma_{\mathcal{L}}$ in LCE.

Proposition 1.3.1 *For every sequent S in $\mathcal{L}$, if $\vdash_{\mathcal{L}} S$ then* LC $\vdash S$.

Proof Replace all occurrences of Et in the proof of S by $\top$. This replacement transforms instances of LCE-rules to instances of LC-rules. Furthermore, the end sequent is not altered since S was assumed to be in $\mathcal{L}$ and so in particular does not contain the existence predicate.

Proposition 1.3.2 *For every sequent S in $\mathcal{L}$, if* LC $\vdash S$ *then* $\vdash_{\mathcal{L}} S$.

Proof The proof is as in [3]. The LC-proof of S may include Skolem terms—that is, terms containing functions from $\mathcal{L}' \backslash \mathcal{L}$. However, no quantifiers can bind any variables inside such terms because the goal sequent S is from $\mathcal{L}$ and the system LC satisfies a version of the subformula property. In this case, such terms can be replaced with constants from $\mathcal{L}$ without disrupting the proof. After performing these replacements, the terms that occur are either eigenvariables or terms from $\mathcal{L}$. As before, we simply add the appropriate $\Gamma \Rightarrow Et$ to the hypothesis of $R\exists$ and $L\forall$ rules or adding the appropriate Eb to the antecedent of $L\exists$ and $R\forall$ rules.

Proposition 1.3.3 *The relations* LCE $\vdash$ *and* $\vdash_{\mathcal{L}}$ *are decidable for quantifier-free sequents.*

Proof The system LCE satisfies a version of the subformula property. Namely, in an LCE derivation of $\Gamma \Rightarrow \Delta$, all formulas that occur in the derivation are either subformulas of Γ or of Δ or are instances of the existence predicate. This follows immediately by induction on the derivation. For quantifier-free sequents, none of the quantifier rules need be applied. But in this case the proof search terminates, as there are only finitely many relevant subformulas.

1.4 Semantics

We define a notion of forcing—called *existence forcing* and denoted $\Vdash^e$—for which LCE($\Sigma_{\mathcal{L}}$) is sound and complete. Following [2], a classical existence structure for $\mathcal{L}'_D$ is a pair (D, I_k) such that D is a set and I_k is a map from $\mathcal{L}'_D$ such that

- $I_k(E)$ is a non-empty unary predicate on D;
- For each n-ary predicate $P \in \mathcal{L}'_D$, $I_k(P)$ is an n-ary predicate on D;
- For each n-ary function $f \in \mathcal{L}'_D$, $I_k(f)$ is an n-ary function from D to D;
- $I_k(a) = a$ for every constant $a \in D$.

I_k is extended in the usual way to an interpretation of terms and then to all sentences. For a sentence $A \in \mathcal{L}'_D$, we write $(D, I_k) \models A$ if A holds in the structure (D, I_k), defined as is standard in the classical case.

A frame is a pair $(W, \leq)$ where W is a non-empty set and $\leq$ is a linear order on W with a root 0. An LC-existence model on a frame $F = (W, \leq)$ is a triple $K = (F, D, I)$ where D is a non-empty set called the domain and I is a collection $\{I_k \mid k \in W\}$ such that each (D, I_k) is a classical existence structure for $\mathcal{L}'_D$ that satisfies the persistency requirements: for all $k, l \in W$, for all predicates $P(\bar{x}) \in \mathcal{L}'_D$ (including E) and for all closed terms $\bar{t} \in \mathcal{L}'_D$, whenever $k \leq l$, it is the case that $(D, I_k) \models P(\bar{t}) \Rightarrow (D, I_l) \models P(\bar{t})$ and $I_k(\bar{t}) = I_l(\bar{t})$.

Note that $I_k(\bar{t}) = I_0(\bar{t})$ for all closed terms $\bar{t}$. Also note that the Kripke models defined in this way have constant domains, which is permissable because in the forcing notion to be defined quantification ranges over existing objects only, a subset of the domain that may vary.

Given an LC-existence model $K = (F, D, I)$, we define $\Vdash^e$ inductively on sentences of $\mathcal{L}'_D$. For predicates $P(\bar{x}) \in \mathcal{L}'$ and closed $\mathcal{L}'_D$ terms $\bar{t}$, we set

$$K, k \Vdash^e P(\bar{t}) \Leftrightarrow (D, I_k) \models P(\bar{t}).$$

We extend $\Vdash^e$ to all sentences $A \in \mathcal{L}'_D$ as follows:

$$K, k \nVdash^e \bot$$

$$K, k \Vdash^e A \wedge B \text{ iff } K, k \Vdash^e A \text{ and } K, k \Vdash^e B$$

$$K, k \Vdash^e A \vee B \text{ iff } K, k \Vdash^e A \text{ or } K, k \Vdash^e B$$

$$K, k \Vdash^e A \to B \text{ iff } \forall l \geq k : K, l \Vdash^e A \Rightarrow K, l \Vdash^e B$$

$$K, k \Vdash^e \exists x A(x) \text{ iff } \exists d \in D : K, k \Vdash^e Ed \wedge A(d)$$

$$K, k \Vdash^e \forall x A(x) \text{ iff } \forall d \in D : K, k \Vdash^e Ed \to A(d)$$

Note that, under this definition, the upwards persistency requirement extends to all sentences. We say that a sentence A is forced in a model K, denoted $K \Vdash^e A$, if for all nodes $k \in W$, $K, k \Vdash^e A$. For a formula $A(\bar{x})$ with free variables, we say that $K \Vdash^e A(x)$ if $K \Vdash^e A(\bar{a})$ for all $\bar{a} \in D$. We call K an $\mathcal{L}$-model if, for all terms t in the language $\mathcal{L}$, $K \Vdash^e Et$. Finally, we say that A is $\mathcal{L}$-forced, written $\Vdash^e_{\mathcal{L}} A$, when $K \Vdash^e A$ for all $\mathcal{L}$-models K.

Given a model $K = (W, \leq, D, I)$ and a (linearly ordered) sequence $\sigma = (\sigma_1, \ldots, \sigma_m)$ of natural numbers, we say that J is a K-interpretation of σ if J is a map from $\{\sigma_1, \ldots, \sigma_m\}$ into W satisfying $J(\sigma_i) \leq J(\sigma_j)$ if $i \leq j$.

We say that an indexed sequent $\sigma\ \Gamma \Rightarrow \Delta$ is forced in a model K if $K, J(\sigma_m) \Vdash^e \bigwedge \Gamma \to \bigvee \Delta$ for all K-interpretations J of σ. A tableau is forced at K if one of its constituent sequents is. We write $\Vdash^e_{\mathcal{L}} \sigma\ \Gamma \Rightarrow \Delta$ if $\sigma\ \Gamma \Rightarrow \Delta$ is forced in all $\mathcal{L}$-models, with the analogous definition for entire tableaux. We say that the sequent or tableau is $\mathcal{L}$-forced if this is the case.

1.5 Completeness

Theorem 1.5.1 (Soundness) *For any tableau U, if $\vdash_{\mathcal{L}} U$, then $\Vdash^e_{\mathcal{L}} U$.*

Proof The proof is by induction on the derivation of A. We show that the axioms of LCE($\Sigma_{\mathcal{L}}$) are $\mathcal{L}$-forced and that, if the hypotheses of some rule are $\mathcal{L}$-forced, then the conclusion of the rule is as well. We consider the rule $R \to$ and axioms from $\Sigma_{\mathcal{L}}$. Throughout the proof, let K be an arbitrary $\mathcal{L}$-model, $k \in W$ a world, and $\bar{a}$ a tuple from D.

Consider a sequent $\sigma\ \Gamma \Rightarrow Et$ from $\Sigma_{\mathcal{L}}$. In particular, t is a term in $\mathcal{L}$. Let $\bar{x}$ denote the free variables occurring in Γ. Since K is an $\mathcal{L}$-model, $K, k \Vdash^e Et$ for all $k \in W$. Hence $K, k \Vdash^e (\Gamma \Rightarrow Et)[\bar{a}/\bar{x}]$ for all $\bar{a} \in D$. As this holds for arbitrary $k \in W$, it follows that $K, J(\sigma_m) \Vdash^e (\Gamma \Rightarrow Et)[\bar{a}/\bar{x}]$ for all K-interpretations J. Hence $\Vdash^e_{\mathcal{L}} \sigma\ \Gamma \Rightarrow Et$, as claimed.

For $R \to$, suppose that the hypotheses of the rule are $\mathcal{L}$-forced. Since the rule's hypotheses are $\mathcal{L}$-forced, in each of the hypotheses some constituent sequent is $\mathcal{L}$-forced. If this is a sequent from U or T_0, then the conclusion is clearly $\mathcal{L}$-forced, since U and T_0 occur in the conclusion of the rule.

Similarly, if a sequent $\sigma \star \cdots \star n_i\ T_i$ from U_j is $\mathcal{L}$-forced, then it it holds under all K-interpretations J of $\sigma \star n_1 \star \cdots \star n_i$. This very sequent occurs in the conclusion of the rule, hence the corresponding sequent is $\mathcal{L}$-forced, so the entire conclusion is as well.

If a sequent of the form $\sigma \star n_1 \cdots \star n_j \star n\ \Gamma, A \Rightarrow B$ is $\mathcal{L}$-forced, we'll show that $\sigma\ \Gamma \Rightarrow A \to B$ is too. Let σ_m denote the last number in the sequence σ. Suppose the latter sequent is *not* $\mathcal{L}$-forced. Then there is some model K with a K-interpretation J of σ such that $K, J(\sigma_m) \nVdash^e \wedge\Gamma \to (A \to B)$. In particular, this means that there is some $l \geq J(\sigma_m)$ with $K, l \Vdash^e \wedge\Gamma$ and $K, l \nVdash^e A \to B$. Unraveling the latter implication, this means that there is some $p \geq l \geq J(\sigma_m)$ with $K, p \Vdash^e A$ and $K, p \nVdash^e B$. By the persistency requirement, $K, p \Vdash^e (\wedge\Gamma) \wedge A$. Define a K-interpretation L of $\sigma \star n_1 \cdots \star n_j \star n$ by making L agree with J on all of σ, setting $L(n_i) = L(\sigma_m)$ for $i \leq j$, and letting $L(n) = l$. By hypothesis, $K, L(n) \Vdash^e \Gamma, A \Rightarrow B$. Since $p \geq l$, it follows that if $K, p \Vdash^e (\wedge\Gamma) \wedge A$, then $K, p \Vdash^e B$. But this contradicts the choice of p.

Finally, suppose a sequent from U'_j is $\mathcal{L}$-forced. Without loss of generality, let this sequent be $\sigma \star \cdots \star n_j \star n \star \cdots \star n_i\ T_i$ for some $i > j$. We'll show that the sequent $\sigma \star n_1 \star \cdots \star n_i\ T_i$ from the rule's conclusion is $\mathcal{L}$-forced. Let J be a K-interpretation of $\sigma \star n_1 \star \cdots \star n_i$. Let L be the K interpretation of $\sigma \star \cdots \star n_j \star n \star \cdots \star n_i$ given by $L(m) = J(m)$ for $m \neq n$ and by $L(n) = L(n_j)$. By hypothesis, $K, L(n_i) \Vdash^e T_i$. But $L(n_i) = J(n_i)$, so $K, J(n_i) \Vdash^e T_i$. Hence $\sigma \star n_1 \star \cdots \star n_i\ T_i$ is $\mathcal{L}$-forced, as claimed.

Theorem 1.5.2 (Completeness) *For any sequent A, if* $\Vdash^e_{\mathcal{L}} A$ *then* $\vdash_{\mathcal{L}} 0\ A$.

Proof We argue by contraposition. Suppose that $\nvdash_{\mathcal{L}} 0\ A$. In this case, the proof search starting with $0\ A$ fails to terminate. In particular, there is some branch that remains open. In general, infinitely many rules will be applied along this branch; otherwise, we obtain a finite countermodel. Choose one such open branch $\mathcal{B}$ with the stipulation that if both branches resulting from the application of $L\forall$ or $R\exists$ rules remain open, the left branch—that is, the branch containing Et in the succedent—is chosen.

Using the chosen branch, we construct an $\mathcal{L}$-model $K = (W, \leq, D, I)$ such that $K, 0 \nVdash^e_{\mathcal{L}} A$. Hence the K-assignment sending all of σ to 0 witnesses the fact that $\nVdash^e_{\mathcal{L}} A$. Let W be the union of all indices σ occurring in the branch $\mathcal{B}$, and define $\sigma \leq \tau$ iff τ is an extension of σ. Define the domain D to be $\{t \mid t \text{ a term in } \mathcal{L}\} \cup \{b \mid b \text{ a term occurring in } \mathcal{B}\}$. Given a sequence σ occurring in $\mathcal{B}$, define the set

$$\Pi_\sigma := \{A \in \mathcal{L}'_D \mid A \text{ occurs in the antecedent of some sequent } \sigma\ T \text{ in } \mathcal{B}\}.$$

Similarly, let

$$\Lambda_\sigma := \{B \in \mathcal{L}'_D \mid B \text{ occurs in the succedent of some sequent } \sigma\ T \text{ in } \mathcal{B}\}.$$

Note that if $\sigma \leq \tau$, then $\Pi_\sigma \subseteq \Pi_\tau$, since in this case the transfer rule is applied so that the antecedents of sequents whose index is τ can be taken to contain all formulas from the antecedents of sequents whose index is σ. It remains to define I_σ. Let $I_\sigma(E)$ be the set $\{t \mid t \text{ a term in } \mathcal{L}\} \cup \{b \mid Eb \in \Pi_\sigma\}$. By our stipulation that $\mathcal{L}$ contain at least one constant, $I_\sigma(E)$ is non-empty. Given an n-ary predicate symbol $P \in \mathcal{L}'_D$ and a sequence $\bar{a} \in D^n$, let $\bar{a} \in I_\sigma(P)$ iff $P\bar{a}$ belongs to Π_σ. Given an n-ary function symbol $f \in \mathcal{L}'_D$ and a sequence $\bar{a} \in D^n$, let $I_\sigma(f)(\bar{a}) = f(\bar{a})$. Finally, for every constant $a \in D$, let $I_\sigma(a) = a$.

We show that $K = (W, \leq, D, I)$ is, in fact, an LC-existence model. As noted above, $\leq$ is a linear order on W. Let $\bar{t}$ be a term in $\mathcal{L}'_D$. Since $\Pi_\sigma \subseteq \Pi_\tau$ whenever $\sigma \leq \tau$, $D, I_\sigma \models P(\bar{t})$ implies that $D, I_\tau \models P(\bar{t})$. Moreover, $I_\sigma(\bar{t}) = I_\tau(\bar{t})$ by induction on the formation of the term $\bar{t}$. Also note that, because $K, \sigma \Vdash^e Et$ for all terms $t \in \mathcal{L}$, K is an $\mathcal{L}$-model.

We prove by induction on A that for A occuring in $\mathcal{B}$ with index σ,

$$A \in \Pi_\sigma \Leftrightarrow K, \sigma \Vdash^e A$$

$$A \in \Lambda_\sigma \Rightarrow K, \sigma \nVdash^e A$$

Consider the case where A is $B \to C$. First, suppose that $B \to C \in \Pi_\sigma$. We must show that $K, \sigma \Vdash^e B \to C$. So let $\tau \geq \sigma$ and suppose that $K, \tau \Vdash^e B$. By the induction hypothesis, $B \in \Pi_\tau$. At some point, the rule $L \to$ was applied to $\tau\ B \to C$, resulting in the branch containing either $C \in \Pi_\tau$ or $B \in \Lambda_\tau$. The latter is impossible, as then the branch would close since $B \in \Pi_\tau$ and $B \in \Lambda_\tau$. So $C \in \Pi_\tau$. Hence $K, \tau \Vdash^e C$ by the induction hypothesis, so $K, \sigma \Vdash^e B \to C$, as claimed.

Conversely, suppose that $K, \sigma \Vdash^e B \to C$. Then for all $\tau \geq \sigma$, if $K, \tau \Vdash^e B$ then $K, \tau \Vdash^e C$. We have to show that $B \to C \in \Pi_\sigma$. By hypothesis, $B \to C$ occurs in $\mathcal{B}$ with index σ. If $B \to C \in \Lambda_\sigma$, then after an application of $R \to$ we would have some $\tau \geq \sigma$ such that that $B \in \Pi_\tau$ and $C \in \Lambda_\tau$. By the induction hypothesis, $K, \tau \Vdash^e B$ and $K, \tau \nVdash^e C$, a contradiction. Hence $B \to C \in \Pi_\sigma$, as claimed.

Second, assume that $B \to C \in \Lambda_\sigma$. At some point, the $R \to$ rule was applied, resulting in part of the branch and an index τ extending σ such that $B \in \Pi_\tau$ and $C \in \Lambda_\tau$. By our induction hypothesis, $K, \tau \Vdash^e B$ and $K, \tau \nVdash^e C$. Hence $K, \sigma \nVdash^e B \to C$.

The other cases are similar. In particular, for the sequent $A \equiv \Gamma \Rightarrow \Delta$, $K, 0 \nVdash^e A$, since every formula from Γ is in Π_0 and every formula from Δ is in Λ_0. Hence $K, 0 \Vdash^e \bigwedge \Gamma$ but $K, 0 \nVdash^e \bigvee \Delta$, completing the proof.

1.6 ESkolemization

Theorem 1.6.1 *Let $\exists x B(x, \bar{y})$ be a formula that occurs negatively in a closed sequent S in $\mathcal{L}$, is not in the scope of strong quantifiers, and where $\bar{y}$ are the variables of the weak quantifiers in the scope of which B occurs. If S' is the result of replacing in S the formula $\exists x B(x, \bar{y})$ by $Ef(\bar{y}) \wedge B(f(\bar{y}), \bar{y})$ where f is a fresh function symbol not in $\mathcal{L}$, then $\vdash_{\mathcal{L}} \sigma\ S$ if and only if $\vdash_{\mathcal{L}} \sigma\ S'$.*

Proof The right-to-left direction proceeds exactly as in [2]. For the left-to-right direction, note that LCE $\vdash \sigma\ A \Rightarrow A'$ for any sequence σ and any sentence A. Hence $\vdash_{\mathcal{L}} S \Rightarrow \vdash_{\mathcal{L}} S'$, as claimed.

Corollary 1.6.2 *For each closed sequent S in $\mathcal{L}'$ in which all strong quantifiers are existential, $\vdash_{\mathcal{L}} S$ iff $\vdash_{\mathcal{L}} S^s$.*

Proof Let $S = S_1, \cdots, S_k = S^s$ be the eSkolem sequence for S. By the previous theorem, $\vdash_{\mathcal{L}} S_i$ iff $\vdash_{\mathcal{L}} S_{i+1}$.

Corollary 1.6.3 *For each closed sequent S in $\mathcal{L}$ in which all strong quantifiers are existential,* $\mathrm{LC} \vdash S$ *iff* $\vdash_{\mathcal{L}} \sigma\ S^s$.

Proof This follows immediately from the previous corollary, Proposition 3.1, and Proposition 5.3.

Corollary 1.6.4 *For each closed sequent S in $\mathcal{L}$ in which all quantifiers are strong and existential,* $\mathrm{LC} \vdash S$ *is decidable.*

Proof By the previous corollary, $\mathrm{LC} \vdash S$ iff $\vdash_{\mathcal{L}} \sigma\ S^s$ for some sequence σ. But the sequent $\sigma\ S^s$ is quantifier-free, so by Proposition 3.2 it is decidable whether or not $\vdash_{\mathcal{L}} \sigma\ S^s$ holds.

Chapter 2

Strong Quantifiers in Modal Logics

2.1 Preliminaries

In classical predicate logic, the derivability problem for formulas of the form $\forall x_1 \cdots \forall x_n \varphi(x_1, \ldots, x_n)$ is easily decidable: One simply instantiates the quantifiers by distinct fresh variables until one is left with the propositional formula $\varphi(a_1, \ldots, a_n)$, then checks whether this latter formula is a propositional tautology. This procedure extends to formulas whose prefixes in prenex normal form consist entirely of universal quantifiers.

First, we define positive and negative subformulas of a formula. A occurs positively in itself. A positive (respectively, negative) occurrence of A in B or C gives rise to a positive (negative) occurrence of A in $B \wedge C$, $B \vee C$, $\exists x B$, and $\forall x B$. A positive (negative) occurrence of A in B gives rise to a negative (positive) occurrence of A in $B \to C$ and $\neg B$. Finally, a positive (negative) occurrence of A in C gives rise to a positive occurrence of A in $B \to C$. An occurrence of the formula $\forall x A$ in B is called an occurrence of the universal quantifier in B, and its sign is defined to be the same as that of the formula $\forall x A$ in B. Similarly, an occurrence of the formula $\exists x A$ in B is called an occurrence of the existential quantifier in B, and its sign is defined to be the opposite of the sign of $\exists x A$ in B.

Equivalently, an occurrence of a universal quantifier is positive in a formula φ if it is in the scope of an even number of negations, treating the antecedents of conditionals in whose scope the particular universal quantifier occurs as contributing to this count. For an existential quantifier to be positive, it must occur in the scope of an odd number of negations, again with the caveat for conditionals. This coincides with the notion of strong quantifiers in [2].

Any formula whose only occurrences of quantifiers are positive is said to be a *pq-formula*. It is clear that *pq*-formulas are precisely those whose prefixes in prenex normal form contain only universal quantifiers.

As shown in [13] and later in [5], the aforementioned decidability result still holds for the class of *pq*-formulas in intuitionistic predicate logic, even though prenex normal form is unavailable and the interdefinability of quantifiers fails.

In this paper, decision procedures are presented for the class of *pq*-formulas in the minimal modal predicate logic **K** and in the logic **T**. In particular, it is shown that every falsifiable *pq*-formula can be falsified in a model with finitely many worlds, each of which has a finite domain. Upper bounds on the number of worlds and the sizes of their domains required in the falsifying model are provided as functions of the original formula.

2.2 Semantic tableaux

In the non-classical setting, one method of proof search is by semantic tableaux. In this section, we present the facts about tableaux systems that are necessary to the proof of the main theorem. For a more detailed account of these systems and an example implementation of one, see [6].

Each tableau begins with a formula to be falsified and from which a tree is constructed according to certain rules. If the formula is falsifiable, the method is intended to produce a countermodel for it. On the other hand, if the formula is not falsifiable, the method should produce a proof of it. It is shown in [18] that the tableaux procedure as described below is sound and complete.

More precisely, formulas prefixed with either T or F are said to be *signed* formulas, with the prefix indicating their intended truth values. A signed formula is said to be true when the ordinary formula it contains has the prefix as its truth value. For each connective, quantifier, and modal operator, there are two rules for developing the tableau, one for each sign. It is convenient to specify the rules by giving their sign and their main connective: $F\Box$, $T\wedge$, $F\forall$, etc. First, consider the case of propositional connectives.

If $T\neg X$ occurs at a node in a tableau, it may be replaced by FX. If $F\neg X$ occurs at a node in a tableau, it may be replaced by TX. This corresponds to the intuition that, under any valuation, the original signed formula is true if and only if its replacement is true. If $TX \wedge Y$ occurs at a node in a tableau, it may be replaced by TX and TY. Similarly, if $FX \vee Y$ occurs, it may be replaced by FX and FY. In either case, under any valuation, the original formula is true if and only if both of its replacements are true.

Under any valuation, $TX \vee Y$ is true if and only if one of TX or TY is true, and $FX \wedge Y$ is true if and only if one of FX or FY is true. To capture this reasoning, the following rules are introduced. If $TX \vee Y$ occurs at a node in a tableau, create a copy of the tableau in progress. Then, in the original tableau, replace the original formula by TX. In the copy, replace it by TY. An entirely analogous procedure handles the case of $FX \wedge Y$. Each such copy is said to be a *component* of the tableau.

A component is called *closed* it contains a node where both TX and FX occur for some formula X. A tableau is *closed* if each of its components is closed. A closed tableau starting with FX is a proof of X, since in this case the assumption that X could be made false leads to a contradiction regardless of the valuation. If no further rules apply to any formulas from any node of a component and the component is not closed, the component is said to be open. Dually, a tableau with at least one open branch is said to be open. An open tableau starting with FX furnishes a countermodel, with the valuation specified to the signs of the proposition letters that occur in the open component.

Together, these rules yield a decision procedure for propositional logic, since each rule decreases the total formula degree. In the modal case, there are two new types of signed formula. The first type, called *necessaries* or ν formulas, contains all signed formulas of the forms $T\Box X$ and $F\Diamond X$, while the second type, called *possibles* or π formulas, contains all signed formulas of the forms $F\Box X$ and $T\Diamond X$.

Let the *principal part* of such a formula be the result of removing the modal operator; e.g., the principal part of $T\Box X$ is simply TX. Denote the principal part of a ν formula by ν_0, and the principal part of a π formula by π_0. A necessary is true at a world ω if and only if its principal part is true at all worlds accessible from ω. A possible is true at a world ω if and only if its principal part is true at some world accessible from ω. This heuristic is built into the tableaux rules as follows.

When a π formula is to be made true at a world ω, we extend the current component by creating a new node σ, recording the fact that σ should be accessible from ω. The new node σ contains the principal component π_0 of the original π formula. The rule for ν formulas is passive, in the sense that it is applicable only after a new node has been created from the current one. When a ν formula is to be made true at a world ω, every successor node σ inherits the principal part ν_0. This corresponds to the semantic definition to be expected of $T\Box X$ and $F\Diamond X$.

Combining both modal rules, when a world σ is created by an application of the π rule, σ inherits the principal part π_0 of the original π formula, as well as every signed formula ν_0 for which there exists τ such that $\tau R \sigma$ and $\nu \in \tau$ both hold. This set is called the *trace* of the node σ, and is denoted $G(\sigma)$. For the initial node σ_0 in a tableau beginning with FX, define $G(\sigma_0)$ to be $\{FX\}$.

Closed components and tableaux are defined as before. Once again, closed tableaux are proofs, while open components of tableaux result in counter-models. The given system yields a decision procedure for the propositional modal logic **K**, and can be used to show that **K** has the effective finite model property. As before, each rule involves a reduction of formula degree, and the branching is bounded by the occurrences of modal operators in the original formula.

With the switch to first-order tableaux systems, decidability is lost. In the first-order case, there are two new types of signed formulas: the universals and the existentials. The universals consist of all formulas of the forms $T\forall x\phi(x)$ and $F\exists x\phi(x)$, while the existentials consist of all formulas of the forms $T\exists x\phi(x)$ and $F\forall x\phi(x)$. We extend the original language L by adding an infinite supply of constant symbols, called parameters. Proofs are of formulas in L, but may contain formulas from the extended language L^*.

The new rules are as follows. When the existential formula $T\exists x\phi(x)$ occurs at a node in a tableau, it may be replaced by $T\phi(c)$, where c is a constant symbol that does not occur anywhere in the present component of the tableau. Similarly, when the existential formula $F\forall x\phi(x)$ occurs at a node in a tableau, it may be replaced by $F\phi(c)$, with the same restrictions on c.

When the universal formula $T\forall x\phi(x)$ occurs at a node in a tableau, $\phi(t)$ may be added at the same node *for any closed term t*. The same holds, mutatis mutandis, in the case of the universal formula $F\exists x\phi(x)$. In [18] it is shown that the procedure without the rules for modalities results is a sound and complete proof search procedure for classical first-order logic.

Universals form a significant obstacle to decidability. These rules allow for the substitution of any closed term t, but there is often no clear strategy for which choice of term to make. Furthermore, a single universal formula may have to be analyzed multiple times during the course of the proof search.

2.3 Results

In the modal case, prenex normal form fails but quantifiers are interdefinable. For the present purposes, it is convenient to work with the smallest set of formulas containing all literals and closed under $\wedge$, $\vee$, $\forall$, $\exists$, $\Box$, and $\Diamond$. Such a formula is said to be in negative normal form.

Any *pq*-formula is equivalent to a formula in negative normal form not containing any occurrences of the existential quantifier, as can been seen by making the standard replacements of Boolean connectives and then pushing all negations inward using De Morgan's laws.

The tableau for φ begins by trying to make φ *false*, and every rule $F\wedge$, $F\vee$, $F\forall$, $F\Box$, and $F\Diamond$ has as its conclusion further formulas to be made false. The only rules that require switching the signs of formula are those for negation, which need only be applied at the propositional level by the choice of normal form. Hence no statement beginning with a universal quantifier must be made true at any node, preventing potentially ruinous applications of the rule for universals.

Theorem 2.3.1 *The class of all pq-formulas over* $\boldsymbol{K}$ *has the effective finite model property and the effective finite domain property.*

Proof Suppose that ψ is a *pq*-formula. Without loss of generality, take ψ to be in negative normal form. As discussed in the previous section, when a new node in a tableau is to be created, it must satisfy certain formulas in virture of its location in the tableau. With formulas in negative normal form, new nodes must be created whenever the $F\Box$ rule is applicable, and these nodes inherit other formulas from applications of the $F\Diamond$ rule.

Let $\lambda(\varphi)$ denote the number of occurrences of the modal operator $\Box$ in the formula φ. Next, associate to each node σ in a tableau the number $\Lambda(\sigma) := \sum_{\varphi\in G(\sigma)} \lambda(\varphi)$, representing the total number of occurrences of the modal operator $\Box$ in the trace of the node.

Lemma 2.3.2 *Whenever a node* σ *is created from a node* τ, $\Lambda(\sigma) < \Lambda(\tau)$.

The proof of this claim is straightforward. Since we are working in **K**, for each node σ other than the initial one there is exactly one node τ such that $\tau R\sigma$. Each signed formula $\varphi \in G(\sigma)$ is the principal part of (at least one) formula $\hat{\varphi}$ occurring at τ. In particular, it follows that $\lambda(\varphi) \leq \lambda(\hat{\varphi})$, since the initial modal operator is removed when passing from $\hat{\varphi}$ to φ. Moreover, note that, if α and β are distinct elements of $G(\sigma)$, $\hat{\alpha}$ and $\hat{\beta}$ must be distinct as well. Hence $\Lambda(\sigma) < \sum_{\varphi \in G(\sigma)} \lambda(\hat{\varphi})$, where the inequality is strict because σ is created from τ from an application of the $F\Box$ rule, which removes an initial $\Box$ from a particular formula. For every non-modal rule, the total number of occurrences of each modal operator in the starting formula is at least as large as the total number of occurrences of the modal operator in its conclusion. Since every $\hat{\varphi}$ at τ is the result of applying a sequence of non-modal rule applications to members of $G(\tau)$, it follows that the quantity $\sum_{\varphi \in G(\sigma)} \lambda(\hat{\varphi})$ is no greater than the total number of occurrences of the $\Box$ modality in the members of $G(\tau)$. This latter quantity is precisely $\Lambda(\tau)$, completing the proof of the lemma.

It follows that the quantity Λ decreases as the tree progresses. After proceeding to a depth of $\lambda(\psi)$, all nodes are guaranteed to have no occurrences of $\Box$ in their traces, at which point no further nodes need be created. The $F\Diamond$ rule is guaranteed to be inapplicable at this point. Accordingly, every subformula beginning with $\Diamond$ can be replaced by $\bot$, making the tableau procedure identical to the procedure for the classical case discussed in the introduction. If all of the tableau's branches have closed by this point, this constitutes a proof of ψ. If, on the other hand, at least one branch of the tableau has not closed by this point, this branch furnishes a countermodel for ψ.

Note that all branching in the tableau arises from applications of the $F\Box$ rule. Each node σ has no more than n successors, where n is the total number of occurrences of $\Box$ in the formula ψ. If ψ is a *pq*-formula falsifiable over **K**, it is falsifiable in a model of depth no more than $\lambda(\psi)$ and branching width no more than n. Together, these conditions imply that ψ is falsifiable in a model of size no more than $\sum_{i=0}^{\lambda(\psi)} n^i$. This establishes the effective finite model property for *pq*-formulas in **K**.

Finally, each constant of the domain comes from an application of the $F\forall$ rule. If m denotes the total number of occurrences of the universal quantifier in ψ, then each node can have no more than m applications of this rule, and hence at most m more constants than its immediate predeccessor. By the same reasoning as before, provided that ψ is falsifiable, it is falsifiable in a model with domains not exceeding the size $m\lambda(\psi)$.

Corollary 2.3.3 *The derivability problem for the class of pq-formulas over* ***K*** *is decidable.*

2.4 Concluding remarks

The situation for the modal logic **T** is quite similar. Although the definition of the trace of a node must be modified to account for the reflexivity of the accessibility relation, the quantity Λ still decreases as the tree progresses.

The decidability of *pq*-formulas in the modal logics **S4**, **S4.3**, and **S5** remain open problems. Via the Gödel-Tarski translation, the case of **S4** would extend the work of Mints from [13], in which the corresponding fragment of LJ is shown to be decidable. Similarly, the case of **S4.3** would entail the decidability of *pq*-formulas in LC, yielding as a corollary the main result from Chapter 1. These results require new techniques, as the transitivity of the accessibility relation obstructs the current paper's methods.

Chapter 3

A Survey of Decidability Results in Non-Classical Logics

3.1 Intuitionistic Theories

In [17], Smorynski provides a detailed summary of decidability results in the intuitionistic setting in addition to presenting several results of his own.

3.1.1 The monadic fragment

Kripke [9] showed that the fragment of intuitionistic predicate logic with two monadic predicates is undecidable by reducing it to the case of classical predicate logical with a single binary relation. Later, Maslov, Mints, and Orevkov showed syntactically that a single monadic predicate letter suffices for undecidability [12]. Following Smorynski, we present a semantic proof due to Gabbay, originally presented in [**?**].

The proof reduces the decidability of the fragment in question to that of the classical theory of a symmetric, reflexive binary relation, shown to be undecidable by Rogers. The reduction class consists of negations of formulas of the form $Q_1 z_1 \cdots Q_n z_n \bigwedge \left(\bigwedge_i \neg R x_i y_i \to \bigvee_j \neg R u_j v_j \right)$, where each Q_i is a quantifier. That these sentences constitute a reduction class for the classical theory is evident from putting them in conjunctive normal form.

Given $\neg A$, where A is of the form specified above, we associate it with a formula A' given by

$$Q_1 z_1 \cdots Q_n z_n \bigwedge \left(\bigwedge_i \neg(Px_i \wedge Py_i) \to \left(\bigvee_j \neg(Pu_j \wedge Pv_j) \vee \exists x Px \right) \right) \to (\exists x Px \vee \exists x \neg Px)\,.$$

Ulimately, the classical binary relation Rxy will be closely related to the intuitionistic formula $\neg(Px \wedge Py)$. The structure of A' is chosen to ensure that R is both reflexive and symmetric.

The goal is to prove that $\neg A$ is a theorem of the classical theory of a symmetric, reflexive binary relation if and only if A' is a theorem of the intuitionistic monadic predicate calculus on a single predicate letter, henceforth denoted M_1. From this, the undecidability of M_1 will be immediate.

For one direction, suppose that $\neg A$ is a theorem of the classical theory of a symmetric, reflexive binary relation. By way of contradiction, suppose further that A' is *not* a theorem of M_1. Let $(K, \leq, D, \Vdash)$ be a Kripke model and let $\alpha \in K$ be such that $\alpha \nVdash A'$. Since the main connective of A' is a conditional, the world α must satisfy the antecedent but not the consequent, $\alpha \nVdash \exists x Px \vee \exists x \neg Px$. Hence $\alpha \nVdash \exists x Px$ and $\alpha \nVdash \exists x \neg Px$.

From the intuitionistic model, we construct a classical model of a symmetric, reflexive binary relation, then show that this model satisfies A, contradicting our initial assumption.

The construction of the classical model M is as follows. Let the domain of M be the same as the domain of α, and let $M \models Rxy$ iff $\alpha \nVdash \neg(Px \wedge Py)$. The symmetry of R is evident in its syntax. Note that, since $\alpha \nVdash \exists x \neg Px$, $\alpha \nVdash \neg(Pc \wedge Pc)$ for any $c \in D\alpha$, demonstrating the reflexivity of R. We show that $M \models \bigwedge_i \neg Rx_i y_i \to \bigvee_j \neg Ru_j v_j$, the sentence A.[1]

Suppose $M \models \bigwedge_i \neg Rx_i y_i$. Then $\alpha \Vdash \bigwedge_i \neg(Px_i \wedge Py_i)$, so $\alpha \Vdash \bigvee_j \neg(Pu_j \wedge Pv_j) \vee \exists x Px$, since α satisfies the antecedent of the sentence A'. By construction, $\alpha \nVdash \exists x Px$, so $\alpha \Vdash \bigvee_j \neg(Pu_j \wedge Pv_j)$. In terms of the classical model M, this means that $M \models \bigvee_j \neg Ru_j v_j$. Hence $M \models A$, yielding the desired contradiction.

For the other direction, suppose that A' is a theorem of M_1 but that $\neg A$ is *not* a theorem of the classical theory under consideration. Let M be a countermodel for $\neg A$—that is, a model of A. Define a Kripke model by

$$K = \{0\} \cup \{\{x, y\} \mid M \models Rxy\}$$

[1] Technically, A may contain quantifiers, but these enter the meta-language and persist throughout the following discussion. A similar situation occurs in the proof of the converse.

$$\alpha \leq \beta \text{ iff } \alpha = 0 \text{ or } \alpha = \beta$$

$$\alpha \Vdash Px \text{ iff, for some } y,\ \alpha = \{x, y\}$$

$$\text{with } D\alpha \text{ equal to the domain of } M, \text{ for any } \alpha$$

We'll show that $0 \nVdash A'$, contradicting our original hypothesis. Suppose instead that $0 \Vdash A'$. Note that $0 \nVdash \exists x Px \vee \exists x \neg Px$, so 0 must **not** force the antecedent of the conditional A'. So 0 must not force one of the conjuncts: that is, for some $\beta \geq 0$, $\beta \Vdash \bigwedge_i \neg(Px_i \wedge Py_i)$, but $\beta \nVdash \bigvee_j \neg(Pu_j \wedge Pv_j) \vee \exists x Px$.

Here, the usefulness of the $\exists x Px$ term becomes apparent: If $\beta > 0$, then $\beta \Vdash \exists x Px$, so $\beta = 0$ must hold. Accordingly, $0 \Vdash \bigwedge_i \neg(Px_i \wedge Py_i)$ but $0 \nVdash \bigvee_j \neg(Pu_j \wedge Pv_j)$.

By construction, $M \models \neg Rxy$ iff $0 \Vdash \neg(Px \wedge Py)$. From the above, it follows that $M \models \bigwedge_i \neg Rx_i y_i$, but $M \not\models \bigvee_j \neg Ru_j v_j$. Hence $M \not\models A$, a contradiction. So $0 \nVdash A'$, but then it cannot be the case that A' is a theorem of M_1, as was assumed.

3.1.2 Further results

In [11], Lifshits proved the undecidability of the intuitionistic theories of equality and normal equality, the latter of which is obtained from the former by adding as an axiom $\neg\neg x = y \rightarrow x = y$. Smorynski presents alternative proofs of these results, making use of the undecidability of M_1. Each of these theories is able to interpret M_1, replacing the single monadic predicate Px of M_1 by $\forall y(x = y \vee \neg x = y)$. It is then shown that $M_1 \vdash A$ if and only if $E \vdash A'$, where A' is obtained from A by making this replacement. The method of proof is similar to that of the Maslov, Mints, Orevkov theorem, although the Kripke models involved in the case of equality are necessarily more complicated. Smorynski concludes with proofs that the intuitionistic theories of a successor function and of a dense linear order are both undecidable as well.

3.2 Variable Bounds

In order to avoid the undecidability of full classical predicate logic, restrictions are needed. One useful way of generating fragments is by permitting formulas in which fewer that a specified number of variables appear. These fragments are naturally ordered by inclusion, and their study reveals information about the border between decidability and undecidability. What follows is a whirlwind summary of some of these results, followed by a more in-depth discussion of a recent paper. Many of these results are well-known; see [7] for a comprehensive list of citations.

The one-variable fragment of classical predicate logic is decidable. In fact, this fragment is equivalent to the modal propositional logic **S5**. The two-variable fragment is also decidable. Notably, many propositional modal logics can be translated into this fragment, providing one explanation for their decidability. At the three-variable level, however, decidability in the classical case breaks down.

By the Gödel double-negation translation, the undecidability of the three variable fragment of intuitionistic predicate logic follows. Also immediate is the undecidability of the three-variable fragments of quantified modal logics, by their conservativity over classical predicate logic. On the other hand, the one-variable fragment of the modal logic **QS4**—and hence that of intuitionistic predicate logic, by means of the Gödel-Tarski translation—is decidable. Experience shows that the the two-variable fragments of various non-classical logics often end up being undecidable.

Gabbay and Shehtman proved the undecidability of the two-variable fragment of intuitionistic logic with constant domains, obtained by adding the so-called principle of constant domains $\forall x(P(x) \vee q) \to \forall x P(x) \vee q$ [7].[2] Kontchakov et al. [10] built upon this result by showing that the two-variable fragment of intuitionistic predicate logic *without* constant domains is likewise undecidable. This result forms the focus of this section.

The proof reduces the decidability of this fragment, henceforth denoted **QH(2)**, to that of an infinite tiling problem. One of the standard proofs for the aforementioned undecidability of the three-variable fragment of classical predicate logic involves showing that such a tiling problem can be encoded within this fragment. In this setup, the third variable is only used in a limited way. Moving to the intuitionistic setting gives enough freedom to cut this variable out entirely, as the authors show.

Precisely, the undecidable tiling problem is as follows: Given a finite set

[2]On p. 801, the authors offer a compelling way of thinking about their results: "Informally, the basic idea underlying the proofs is in viewing a 3-variable classical model as a 3-dimensional cube, while a 2-variable Kripke model is a family of 2-dimensional cubes parameterized by possible worlds, i.e., it is also a 3-dimensional structure. That is why sometimes we can embed a 3-cube into a 2-variable Kripke model."

T of tiles, each of whose four sides is assigned a certain color, decide whether or not T is able to tile $\mathbb{N} \times \mathbb{N}$; that is, decide whether there exists a function $\tau : \mathbb{N} \times \mathbb{N} \to T$ such that, for all $i, j \in \mathbb{N}$, the up-color of $\tau(i, j)$ is the down-color of $\tau(i, j+1)$ and the right-color of $\tau(i, j)$ is the left-color of $\tau(i+1, j)$. From such a T, the authors construct a sentence φ_T in the language of **QH(2)** such that φ_T is refutable in some Kripke model if and only if T tiles $\mathbb{N} \times \mathbb{N}$. The sentence φ_T is an implication $\psi_T \to \exists x(D(x) \to \bot)$, where ψ_T is the conjunction of six sentences that encode the tiling, roughly speaking.

$$\forall x \bigvee_{t \in T} \left(P_t(x) \wedge \bigwedge_{s \neq t} (P_s(x) \to \bot) \right) \tag{3.1}$$

$$\bigwedge_{right(t) \neq left(s)} \forall x \forall y \, (H(x, y) \wedge P_t(x) \wedge P_s(y) \to \bot) \tag{3.2}$$

$$\bigwedge_{up(t) \neq down(s)} \forall x \forall y \, (V(x, y) \wedge P_t(x) \wedge P_s(y) \to \bot) \tag{3.3}$$

$$\forall x \exists y H(x, y) \wedge \forall x \exists y V(x, y) \tag{3.4}$$

$$\forall x \forall y (V(x, y) \vee (V(x, y) \to \bot)) \tag{3.5}$$

$$\forall x \forall y \, (V(x, y) \wedge \exists x(D(x) \wedge H(x, y)) \to \forall y(H(x, y) \to \forall x(D(x) \to V(y, x)))) \tag{3.6}$$

Here, $V(x, y)$ (and $H(x, y)$) are interpreted as "y is the vertical (respectively, horizontal) successor of x," while $P_t(x)$ is interpreted as "tile t is at location x in the grid." Under these interpretations, (3.1) says that each location has exactly one tile associated with it, while (3.2) and (3.3) ensure that adjacent colors match. The meaning of (3.6) is more mysterious; its role becomes clear during the course of the proof. Again, ψ_T is taken to be the conjunction of (3.1)-(3.6), while φ_T is $\psi_T \to \exists x(D(x) \to \bot)$.

Suppose that φ_T is refutable. Then, there is a Kripke model with a world w where $w \models \psi_T$ but $w \not\models \exists x(D(x) \to \bot)$. Using the pieces (3.4), (3.5), and (3.6) of ψ_T and the fact that $w \not\models \exists x(D(x) \to \bot)$, one verifies that, for any a, b, c in the domain of the model at w, if $w \models H(a, b) \wedge V(a, c)$, then there is a d in the domain of the model at w such that $w \models H(c, d) \wedge V(b, d)$. This fact, combined with (3.4), yields that there are $a_{i,j}$ in the domain of w such that $w \models H(a_{i,j}, a_{i+1,j})$ and $w \models V(a_{i,j}, a_{i,j+1})$ for all $i, j \in \mathbb{N}$. The tiling τ is defined by $\tau(i, j) = t$ if and only if $w \models P_t(a_{i,j})$, and the remaining pieces of ψ_T are used to verify that this is, in fact, a successful tiling.

Conversely, if τ is a tiling of $\mathbb{N} \times \mathbb{N}$, the authors define a Kripke model refuting φ_T. The set of worlds is given as $\{w_0\} \cup (\mathbb{N} \times \mathbb{N})$, where every world is accessible from w_0 but all other worlds can only access themselves. It suffices to take constant domains of $\mathbb{N} \times \mathbb{N}$. For each world w, $w \models H(x, y)$ if and only if $(x, y) = ((i, j), (i + 1, j))$ for some $(i, j) \in \mathbb{N} \times N$, and similarly for $V(x, y)$. P_t is defined in the obvious manner. Finally, $D^{w_0} = \emptyset$, while $D^w = \{w\}$ for all $w \neq w_0$. The constructed model does, in fact, refute φ_T, as is easily verified.

Let **QS4(2)** denote the two-variable fragment of quantified **S4**. Coupled with the previous result, the Gödel-Tarski translation shows that validity in **QS4(2)** is also undecidable. In fact, the authors prove that, if **L** is a propositional modal logic having a Kripke frame in which there is a point that has infinitely many successors, then the two-variable fragment **QL(2)** is undecidable. The argument is analogous to the previous one, and shows that the two-variable fragments of quantified **K**, **K4**, **GL**, **S5**, **S4.2**, and **S4.3** are undecidable.

Recall the result from section 3.1 on the undecidability of monadic **QS4**. Kontchakov et al. conclude their paper by showing that the even the two-variable monadic fragment of many quantified modal logics—among them all the modal logics mentioned above—is undecidable. They conjecture that the monadic fragment of **QH(2)** is undecidable as well, but note that the techniques used in the modal case no longer apply.

3.3 Acknowledgements

I am extremely grateful to my advisor, Grigori Mints, for his endless patience and wisdom. No part of this thesis would have been possible without his support. I would also like to thank Professors Kannan Soundararajan and Eleny Ionel for their guidance and encouragement throughout my undergraduate experience.

To my family I owe everything.

Lastly, I would like to confess my indebtedness to the beautiful residents of XΘX, each of whom I bring with me.

Bibliography

[1] Arnon Avron. “Using Hypersequents in Proof Systems for Non-classical Logics.” *Annals of Mathematics and Artificial Intelligence.* 4:225-248, 1991.

[2] Matthias Baaz and Rosalie Iemhoff. “The Skolemization of existential quantifiers in intuitionistic logic.” *Annals of Pure and Applied Logic.* 142.1-3. 2006. pp. 269-295.

[3] Matthias Baaz and Rosalie Iemhoff. “Gentzen calculi for the existence predicate.” *Studia Logica.* 82.1. 2006. pp. 7-23.

[4] Matthias Baaz, Agata Ciabattoni, and Christian G. Fermüller. “Hypersequent calculi for Gödel logics - a survey.” *J. Log. Comput.*, 13(6):835-861, 2003.

[5] Gilles Dowek and Ying Jiang. “Eigenvariables, bracketing and the decidability of positive minimal predicate logic.” *Theoretical Computer Science.* 360.1-3. 2006. pp. 193-208.

[6] Melvin Fitting. “First-Order Modal Tableaux.” *Journal of Automated Reasoning* 4. 1988. pp. 191-213.

[7] D. M. Gabbay, V. B. Shehtman. “Undecidability of Modal and Intermediate First-Order Logics with Two Individual Variables.” *The Journal of Symbolic Logic.* Vol. 58, No. 3. 1993. pp. 800-823.

[8] Rosalie Iemhoff. “A Note on Linear Kripke Models.” *Journal of Logic and Computation.* Vol. 15 (4). 2005. 489-506.

[9] Saul Kripke. “Semantical analysis of intuitionistic logic, I.” *Formal systems and recursive functions.* J. N. Crossley and M. A. E. Dummett, eds. North-Holland, Amsterdam, 1965.

[10] Roman Kontchakov, Agi Kurucz, and Michael Zakharyaschev. "Undecidability of First-Order Intuitionistic and Modal Logics with Two Variables." *The Bulletin of Symbolic Logic.* Vol. 11, No. 3. 2005. pp.428-438.

[11] V.A. Lifshits. "Problem of decidability for some constructive theories of equalities." *Studies in constructive mathematics and mathematical logic, Part 1.* A.O. Slisenko, ed. 1969.

[12] S.Y. Maslov, G.E. Mints, and V.P. Orevkov. "Unsolvability of the constructive predicate calculus of certain classes of formulas containing only monadic predicate variables." *Soviet Mathematics Doklady.* Vol. 163. 1965. pp. 918-920.

[13] Grigori Mints. "Solvability of the problem of deducibility in LJ for a class of formulas not containing negative occurrences of quantifiers." *Five Papers on Logic and Foundations.* 1992. pp. 135-145.

[14] Grigori Mints. "The Skolem method in intuitionistic calculi." *Proceedings of Steklov Institute of Mathematics* 121 (1972) 73109.

[15] Dana Scott. "Identity and existence in intuitionistic logic." *Applications of Sheaves, Proc. Res. Symp. Durham 1977*, M. P. Fourman, et al. (eds), Vol. 753 of *Lecture Notes in Mathematics,* pp. 660696, Heidelberg, 1979.

[16] Thoralf Skolem. Logisch-kombinatorische Untersuchungen über die Erfüllbarkeit oder Beweisbarkeit mathematischer Sätze nebst einem Theorem über dichte Mengen, Skrifter utgitt av Videnskapsselskapet i Kristiania, I, Matematisk Naturvidenskapelig Klasse 4 (1920) 1993-2002.

[17] Craig Smorynski. "Elementary Intuitionistic Theories." *The Journal of Symbolic Logic.* Vol. 38, No. 1. 1973. pp. 102-134.

[18] Raymond Smullyan. *First-Order Logic.* Springer-Verlag. Berlin. 1968.

[19] Frank Wolter, Michael Zakharyaschev. "Decidable Fragments of First-Order Modal Logics." *The Journal of Symbolic Logic.* Vol. 66, No. 3. 2001. pp. 1415-1438.