Some Intuitions Behind Realizability Semantics for Constructive Logic: Tableaux and Läuchli countermodels.

James Lipton
Wesleyan University

Michael J. O'Donnell
The University of Chicago

Abstract

We use formal semantic analysis based on new constructions to study abstract realizability, introduced by Läuchli in 1970, and expose its algebraic content. We claim realizability so conceived generates semantics-based intuitive confidence that the Heyting Calculus is an appropriate system of deduction for constructive reasoning.

Well-known semantic formalisms have been defined by Kripke and Beth, but these have no formal concepts corresponding to constructions, and shed little intuitive light on the meanings of formulae. In particular, the completeness proofs for these semantics do not generate confidence in the sufficiency of the Heyting Calculus, since we have no reason to believe that every intuitively constructive truth is valid in the formal semantics.

Läuchli has proved completeness for a realizability semantics with formal concepts analogous to constructions. We argue in some detail that, in spite of a certain inherent inexactness of the analogy, every intuitively constructive truth is valid in Läuchli semantics, and therefore the Heyting Calculus is powerful enough to prove all constructive truths. Our argument is based on the postulate that a uniformly constructible object must be communicable in spite of imprecision in our language, and that the permutations in Läuchli's semantics represent conceivable imprecision in a language, while allowing a certain amount of freedom in choosing the particular structure of the language.

We give a detailed generalization of Läuchli's proof of completeness for the propositional part of the Heyting Calculus, in order to make explicit constructive and algebraic content.

In our treatment, we establish several new results about Läuchli models. We show how to extend the *sconing* and *gluing* constructions familiar from Kripke and Frame semantics and Topos theory, to Läuchli models, and use them to give an algebraic approach to countermodel construction. In particular, the Läuchli arguments are given without the restriction to the integers, $\mathbb{Z}$, as a group of permutations, which makes much of the coding scheme used in Läuchli's original paper transparent.

We also make use of a new propositions-as-types syntax for the Heyting calculus, with limited non-determinism, in which validity of formulae can be decided without loop-detection.

1 A Philosophical Introduction

This paper presents a semantic analysis of constructive reasoning that formalizes a set of explicit intuitions about constructions. We use that analysis to demonstrate that the usual formal systems for constructive propositional logic (the Heyting Propositional Calculus [Hey30] and equivalent formal systems) have precisely the right power for constructive reasoning. Kripke semantics already provides a technically very successful semantic analysis of constructive logic, with formal soundness and completeness proofs. We argue that these proofs do not justify a claim that a formal logical system is the right one for constructive reasoning, because Kripke's semantics are not formally based on the intuitions of constructive reasoning. Our approach, based on the realizability semantics of Kleene and Läuchli, formalizes the concept of a construction directly, so our soundness and completeness proofs yield genuine evidence that the formal rules of constructive logic precisely capture constructive reasoning.

Logical systems are ultimately based on philosophical positions. In particular, classical logic is based on the position that every meaningful assertion is definitely either *true* or *false*. Of course, we may not know the truth or falsehood of a given assertion α, but we may use in reasoning the fact that one or the other must hold. In contrast, constructive philosophy holds that an assertion is true only if a mental *construction* proves its correctness, and is false only if there is a mental *construction* of an absurdity from it. This more stringent concept of truth leads to a more conservative logic, with fewer theorems. For example, the classical law of the excluded middle $(\alpha \vee \neg\alpha)$ is not constructively valid. An *a priori* argument for this is that there may be formulae α such that neither α nor $\neg\alpha$ is proved by a construction. Constructive philosophy allows the possibility that there are formulae that are neither true nor false.

While negation displays the disagreement between constructive and classical logic in the most transparent way, there are also purely positive formulae on which the logics disagree. For example, $\alpha \vee (\alpha \Rightarrow \beta)$ is classically, but not constructively, valid. Excluded middle is just the special case where β denotes a patent falsehood. The constructive fallacy in arguments for this assertion is discussed in Section 2. If the only available connective is implication, we may still distinguish the logics by Peirce's law [Pei85], $(((\alpha \Rightarrow \beta) \Rightarrow \alpha) \Rightarrow \alpha)$, which is classically, but not constructively, valid. In all of these cases, constructive logic rejects the validity of the assertion in question, but does not affirm its negation.

An intriguing perspective on constructive reasoning is advanced by Gödel's double-negation translation of classical logic into intuitionistic logic. Since, in classical logic, the connectives $\vee$ and $\Rightarrow$ are eliminable in favor of conjunction and negation, we may view propositions containing these connectives as abbreviations for their equivalent $\wedge, \neg$-forms and take the position that disjunction and implication simply don't occur in classical logic at all. Then constructive reasoning can be viewed as an *extension* of classical reasoning by two genuinely new connectives. For details about the formal properties of this translation and related ones the reader may consult [Tv88].

A consequence of this translation and its properties is that if a classical proof of falsity exists in pure predicate logic then so does a constructive proof. Thus one cannot hope to make a case in favor of constructive reasoning by establishing the inconsistency of classical logic.

For each philosophical position about truth and falsehood, it is useful to design languages for expressing assertions, to discover appropriate systems of reasoning, and even to define these languages and reasoning systems formally. An appropriate system of reasoning must be *faithful* to its underlying philosophy—that is, every logical consequence derived by the system must be philosophically correct. Given faithfulness, it is desirable that a system of reasoning be as powerful as possible. Ideally, it should be *full*—that is, every logical consequence, expressible in the language, that follows from the philosophy should be derivable in the system. Well-designed systems of reasoning may often be seen to be faithful by direct intuitive inspection. Demonstrations of fullness are usually much more subtle and complex, requiring substantial mathematical tools. So, to show that a system is full, we need a formal description of the *semantics* of the language, capturing the crucial qualities of the philosophical position underlying the system of reasoning. A formal description also improves our understanding of faithfulness, even though faithfulness may be demonstrated more directly and intuitively. Feferman has defined related concepts of *faithfulness* and *adequacy* [Fef79, Bee80]. His concept of faithfulness is essentially the same as ours, but adequacy requires that the language be powerful enough to express the important properties of a system, while our fullness requires only that the true formulae in a given language be provable.

A mathematical semantic system typically identifies a class of structures, called *worlds* or *models*, representing the possible world states that logical assertions are intended to describe. In addition, a formal description of semantics must define the conditions under which a given assertion is *true* of a given model. The act of inferring a conclusion β as a logical consequence of assumption α is *valid* if β is true in every model in which α is true. A formal system for reasoning is *sound* with respect to a semantic system if every logical consequence derivable in the system is valid; it is *complete* if every valid logical consequence can be derived. Soundness is a formal analogue of faithfulness, and completeness is a formal analogue of fullness. The formal versions have the advantage that they are susceptible to mathematical proof. In order to use mathematical proofs of soundness and completeness to provide evidence for faithfulness and fullness, we must analyze the relationship between the formal semantics and the philosophy behind it. In particular, soundness implies faithfulness if validity in the formal semantics implies philosophical correctness. Similarly, completeness implies fullness if philosophical correctness implies formal validity.

The formal mathematical treatment of semantics for classical propositional logic is an intuitively direct formalization of its philosophy. The philosophy holds that every formula is either true or false, and that the truth or falsehood of a composite formula depends only the truth or falsehood of its components. So, a formal model simply assigns truth or falsehood to each atomic assertion. It is intuitively clear that every conceivable reality is represented by a model, and that every model represents a conceivable reality, so the well-known soundness and completeness proofs for classical logic provide completely satis-

α	F
β	T

Figure 1: Classical Truth-table Countermodel for $(\alpha \Rightarrow \beta) \Rightarrow (\beta \Rightarrow \alpha)$

factory demonstrations of faithfulness and fullness. In particular, the proof of completeness provides us with a truth-table countermodel for each unprovable formula α, and that countermodel yields a *completely concrete* interpretation of the atomic formulae under which α is clearly and intuitively false. For example, $(\alpha \Rightarrow \beta) \Rightarrow (\beta \Rightarrow \alpha)$ is not provable in classical logic. The countermodel is given by the truth table in Figure 1. If we interpret α as the proposition

> *in the diagram of Figure 1, the symbol "α" appears in the same row with the symbol "T",*

and β as the proposition

> *in the diagram of Figure 1, the symbol "β" appears in the same row with the symbol "T",*

it is intuitively clear that $\alpha \Rightarrow \beta$ is true, but $\beta \Rightarrow \alpha$ is false.

The situation with respect to constructive logic is less satisfactory. There are well-known proof systems that can easily be seen by intuitive inspection to be faithful. These proof systems are known to be sound and complete with respect to the semantics of Kripke [Kri65] and Beth [Bet59] The models of Kripke semantics consist of partially ordered collections of *possible worlds,* each of which has the form of a classical model in that it accepts or rejects every atomic assertion. Intuitively, however, acceptance of an assertion in a possible world is intended as an indication that the assertion is not just classically true, but *knowable* in that world, while rejection indicates only that the assertion is *not knowable,* not that its negation is knowable. So, each possible world represents a conceivable state of knowledge. World $\mathfrak{v}$ preceding world $\mathfrak{w}$ in the partial ordering of a model is intended to indicate that it is possible that the state of knowledge represented by $\mathfrak{v}$ will develop into that represented by $\mathfrak{w}$ over time.

Kripke semantics seems to formalize a certain temporal-epistemic philosophy of truth directly. If we take the position that an assertion is known precisely when we have a mental construction proving it, then a Kripke model may be read as describing a postulated temporal development of the collection of constructions that we possess. Because there is no immediately apparent way to extract a *construction* from an arbitrary Kripke model, however, we have no *a priori* reason to believe that *every* Kripke model represents a conceivable constructive reality. The mere fact that standard intuitionistic propositional proof systems are complete with respect to Kripke models does not clearly imply that they are full with respect to the philosophy of constructivism. The extensional correspondence between validity for Kripke models and formal provability might in some sense only be a coincidence, which may be very useful as a technical device, but which cannot be taken as an intuitive justification for the formal rules of proof.

In order to develop intuitively satisfactory mathematical semantics for constructive logic, we should find direct formal representatives for constructions, and define formally what it means for a construction to prove an assertion. Heyting spelled out an informal blueprint indicating what such a formal definition might entail in the so-called Brouwer-Heyting-Kolmogorov (BHK) interpretation [Tv88]. It is sketched out in definition 4.1 below. While such an interpretation does not pin down the meaning of the connectives, and indeed does not even rule out a classical interpretation, it has led to several interpretations based directly on the notion of a construction, in particular the so-called Curry-Howard interpretation, and the one that constitutes the main theme of this paper.

Kleene [Kle45, Kle59, KV65] also proposed a BHK-style semantics, called *realizability* semantics, in which constructions are represented by set-theoretic objects, called *realizers,* of appropriate types—particularly, a realizer proving $\alpha \Rightarrow \beta$ is a function from realizers for α to realizers for β. An object *realizing* a formula is the formal analogue to a construction proving an assertion. There are fundamental reasons to doubt the possibility of defining in any precise and effective way the class of constructions. To establish faithfulness through a mathematical proof of soundness, it is important to identify a class of realizers that clearly represent constructions. To establish fullness through a proof of completeness, we must find a class of realizers that clearly includes a representative for each possible construction. There is no need to find a single class of realizers corresponding precisely to constructions.

Kleene's realizability semantics allows the partial computable functions to act as realizers representing constructions. It is intuitively clear that every construction is computable, but Rose showed that the formula

$$((\neg\neg\gamma \Rightarrow \gamma) \Rightarrow (\neg\neg\gamma \vee \neg\gamma)) \Rightarrow (\neg\neg\gamma \vee \neg\gamma)$$

where

$$\gamma = (\neg\alpha \vee \neg\beta)$$

is true in Kleene's semantics, but not provable in Heyting Calculus [Ros53], so Heyting Calculus is not complete for Kleene's realizability. Läuchli achieved the first completeness result for formal constructive reasoning (in the first-order predicate calculus) with respect to a realizability semantics [Läu70], by allowing as realizers only functions that are invariant under certain permutations. We believe that Läuchli's analysis has the essential technical form required for a genuine demonstration of fullness, but two important components are lacking. First, Läuchli does not explain the intuitive connection between permutation invariance and constructions. Second, he proves completeness only of the theorems of formal constructive logic—the formulae that can be proved with no assumptions—while we want completeness of the logical consequence relation. Friedman proved another realizability completeness result (unpublished), using a simpler technical basis than Läuchli's, but not addressing any of the points above.

In this paper, we work out an intuitive basis for Läuchli's realizability semantics, showing how it can be derived from intuitions about the unambiguous communicability of a function in spite of certain ambiguities of language. Based on this intuitive explanation of Läuchli's semantics, we argue that completeness with respect to such formal semantics does in fact imply fullness with respect to a coherent philosophical position. We do not argue that the philosophical position identified here is correct, nor even that it is held by any particular constructivist or intuitionist, but only that it is a plausible position for which the Heyting calculus is an appropriate formal system.

We also underscore the crucial model-theoretic constructions involved in the proof of completeness, and extend Läuchli's proof to *logical consequences* in the propositional calculus, in addition to theorems (the argument for *fullness* for logical consequences depends on the intuitive form of the deduction rule). Finally, we examine the constructive content of the completeness proof. For the propositional calculus, the proof is completely constructive, but it depends on the decidability of validity in that language, and the sufficiency of *finite* countermodels for invalid formulae. In the predicate calculus, the definition of a countermodel for an invalid formula is completely constructive, but the proof that it is a countermodel requires the axiom of choice.

We assume familiarity with constructive formalism, but not necesarily with the specific systems used here. This paper is not intended to provide an introduction to the systems, nor on various intuitive philosophical bases for logics, but rather to illuminate connections between the two for a reader who already understands each in isolation. We repeat well-known and sometimes elementary details of formal definitions and proofs whenever we find them helpful to focus attention on their connection to the intuition we are seeking to highlight, or on special details of our completeness proof

2 Positive Constructive Propositional Logic

2.1 A Formal Notation for Propositional Formulae

In our study we restrict attention to the language of *positive propositional logic* because the most important semantic issues have to do with implication. Negation introduces complications that are technically easy to solve, but that obscure the essential insights. From now on, all references to propositional formulae, proofs, etc. are understood to be positive unless otherwise stated. The set of propositional formulae is denoted PF.

For convenience, we introduce several abbreviations into our printed notation. We systematically drop parentheses, giving precedence to $\wedge$, $\vee$, and $\Rightarrow$ in that order. The symbols $\wedge$, $\vee$, and $\Rightarrow$ associate to the right, so $\alpha \Rightarrow \beta \Rightarrow \gamma$ is an abbreviation for $(\alpha \Rightarrow (\beta \Rightarrow \gamma))$. In the case of $\wedge$ and $\vee$, the direction of association does not matter semantically in either classical or constructive logic, but the association of $\Rightarrow$ is significant in both logics.

2.2 A Formal Notation for Constructive Propositional Proofs

Proofs are the formal analogues of rigorously reasoned arguments. Formal descriptions of proofs are often treated as mere syntactic devices for enumerating true formulae. In constructive logic, we can get a deeper insight into proofs by regarding a proof formula as a syntactic object denoting a semantic construction. So proof formulae, as well as propositional formulae, have semantic content. The following notation for proofs is technically equivalent to the *natural deduction* notation of Fitch [Fit74], but uses *lambda-terms* [GLT89, How80] to exhibit the intended interpretation of a proof as denoting a construction, more directly and naturally than Fitch's notation. We use $\Gamma \vdash a\colon\alpha$ to mean that a is a proof formula proving the propositional formula α, possibly using assumptions in the collection Γ of labelled formulae, but no others. For notational simplicity, set braces are omitted in collections of assumptions, and commas are used for unions (so $x\colon\alpha, \Gamma \vdash b\colon\beta$ abbreviates $\{x\colon\alpha\} \bigcup \Gamma \vdash b\colon\beta$).

Definition 2.1 *Assume that there is an unlimited supply of* labels *for assumptions. The* (deterministic) proof formulae *are defined inductively as follows:*

(B) *If x is a label, then $x\colon\alpha \vdash x\colon\alpha$ for all formulae α;*

($\wedge$I) *If $\Gamma \vdash a\colon\alpha$ and $\Delta \vdash b\colon\beta$,*
then $\Gamma, \Delta \vdash \langle a, b\rangle\colon(\alpha \wedge \beta)$;

($\wedge$E) *If $\Gamma \vdash c\colon(\alpha \wedge \beta)$, then $\Gamma \vdash (\sigma_0 c)\colon\alpha$, and $\Gamma \vdash (\sigma_1 c)\colon\beta$;*

($\vee$I) *If $\Gamma \vdash a\colon\alpha$, then $\Gamma \vdash \langle 0, a\rangle\colon(\alpha \vee \beta)$, and $\Gamma \vdash \langle 1, a\rangle\colon(\beta \vee \alpha)$;*

($\vee$E) *If $\Gamma \vdash c\colon(\alpha\vee\beta)$, $\Delta \vdash d\colon(\alpha \Rightarrow \gamma)$, and $\Theta \vdash e\colon(\beta \Rightarrow \gamma)$, then $\Gamma, \Delta, \Theta \vdash (\chi cde)\colon\gamma$;*

($\Rightarrow$I) *If $x\colon\alpha, \Gamma \vdash b\colon\beta$ and Γ does not contain an assumption with label x, then*
$\Gamma \vdash (\lambda x\colon\alpha\ .\ b)\colon(\alpha \Rightarrow \beta)$;

($\Rightarrow$E) *If $\Gamma \vdash a\colon\alpha$, and $\Delta \vdash b\colon(\alpha \Rightarrow \beta)$, then $\Gamma, \Delta \vdash (ba)\colon\beta$;*

(K) *If $\Gamma \vdash a\colon\alpha$, then $\Gamma, \Delta \vdash a\colon\alpha$.*

The notions of subformulae *of proof formulae, and* free *and* bound *occurrences of labelled assumptions within proof formulae, are defined in the usual manner [Ste72].*

Like propositional formulae, proof formulae are not character strings, but abstract formal objects constructed from labels and various proof operators denoted by $\langle\cdot,\cdot\rangle$, $(\sigma_0\ \cdot)$, $(\sigma_1\ \cdot)$, $\langle 0,\cdot\rangle$, $\langle 1,\cdot\rangle$, $(\chi\ \cdot\cdot\cdot)$, $(\lambda\cdot : \alpha\ .\ \cdot)$, and $(\cdot\,\cdot)$. Parentheses are omitted in notations for proof formulae when no ambiguity results. Function applications associate to the left, so that xyz abbreviates $((xy)z)$. Notice that most of the rules for constructing proof formulae in Definition 2.1 come in the form ($\odot$I), or ($\odot$E), where $\odot$ is one of the operators $\wedge$, $\vee$, or $\Rightarrow$. The ($\odot$I) rule is the *introduction* rule for $\odot$, since it shows how to *introduce* $\odot$ into the head position of a formula in a proof. The ($\odot$E) rule is the *elimination* rule for $\odot$, because it shows how to *eliminate* $\odot$ as head symbol in a formula, in order

Assume Γ, Δ

$\ldots b \ldots$ (a proof of $\alpha \Rightarrow \beta$ from Γ)

$\alpha \Rightarrow \beta$

$\ldots a \ldots$ (a proof of α from Δ)

α

β, by *modus ponens*

Figure 2: Natural deduction proof corresponding to (ba)

to reason about its principal subformulae. The only exceptions to this structure are the *basis* rule (B) and the *weakening* rule (K).

Intuitively, $\langle a, b\rangle$ denotes the *ordered pair* of a and b, and σ_0 and σ_1 denote the associated selection functions. The proof formula $\langle 0, a\rangle$ marks a as a construction for the left-hand side of a disjunction, $\langle 1, a\rangle$ for the right. The proof operator χ denotes a conditional function, such that $(\chi\langle 0, c\rangle de) = (dc)$ and $(\chi\langle 1, c\rangle de) = (ec)$. In the terminology of the lambda-calculus, the label x on an assumption α is a *variable* of *type* α. The proof formula $(\lambda x\colon \alpha\ .\ b)$ denotes the function that, given a construction for α, assigns that construction as the value of x and returns the resulting value of b. Notice that x can be variously thought of as a label on an assumption, or as a variable, with no inconsistency. The proof operation (ba) denotes the application of the function b to argument a.

In the discussions that follow, we will abbreviate expressions of the form $\Gamma \vdash a\colon \alpha$ in several ways.

Definition 2.2 *$\alpha_0, \ldots, \alpha_m \vdash b_0\colon \beta_0, \ldots, b_n\colon \beta_n$ if and only if there exist labels $x_0, \ldots, x_m$ such that $x_0\colon \alpha_0, \ldots, x_m\colon \alpha_m \vdash b_0\colon \beta_0, \ldots, b_n\colon \beta_n$.*

$\Gamma \vdash \alpha$ if and only if there exists a proof formula a such that $\Gamma \vdash a\colon \alpha$. In this case, we say there is a constructive deduction *from Γ to α.*

$\vdash a\colon \alpha$ if and only if $\emptyset \vdash a\colon \alpha$.

$\vdash \alpha$ if and only if $\emptyset \vdash \alpha$. In this case, we say that α is a constructive theorem.

A constructive proof formula as defined in Definition 2.1 above may be read straightforwardly as a natural deduction proof. In particular, function application corresponds precisely to *modus ponens*: if $\Gamma \vdash b\colon (\alpha \Rightarrow \beta)$ and $\Delta \vdash a\colon \alpha$, then the proof (ba) of β from assumptions in Γ, Δ may be rewritten as in Figure 2. Lambda abstraction corresponds to the *deduction rule*: if $x\colon \alpha, \Gamma \vdash b\colon \beta$, then the proof $(\lambda x\colon \alpha\ .\ b)$ of $\alpha \Rightarrow \beta$ from assumptions in Γ may be rewritten as in Figure 3. The main peculiarity of the lambda notation is that it distinguishes different assumptions $x\colon \alpha$ and $y\colon \alpha$ of the same formula α. Such a distinction is superfluous for the purpose of deriving theorems, but it is certainly not harmful, and it leads to a unique association of a function with each proof of an implication.

Assume Γ

Assume α

$\ldots b \ldots$ (a proof of β from α, Γ)

β

$\alpha \Rightarrow \beta$, by the *deduction rule*

Figure 3: Natural deduction proof corresponding to $(\lambda x{:}\,\alpha \;.\; b)$

Pairing corresponds to introducing $\alpha \wedge \beta$ after proving α and β individually, and selection corresponds to detaching conjuncts. Marking a proof with 0 or 1 corresponds to introducing a disjunction after proving one of the disjuncts, and the conditional operator χ corresponds to proof by cases.

The proof formulae of Definition 2.1 were designed to represent as simply as is possible the structure of logical arguments and the structure of constructions of functions. In order to analyze the concept of proof formally, we consider a variant, called *nondeterministic proof formulae*, in which conditional forms (χabc) (i.e., proofs by cases) are decomposed into a two-level form, where the inner level associates a conditional-branch/case with the condition in which it is used, and the outer level combines the branches. Nondeterministic proof formulae denote partial proofs that are contingent upon the selection of certain branches of disjunctions. To accommodate such contingency, the relation $\vdash$ is extended to allow collections of *proof-formula*:*formula* pairs on the right-hand side, resulting in the form $x_0{:}\,\alpha_0, \ldots, x_m{:}\,\alpha_m \vdash b_0{:}\,\beta_0, \ldots, b_n{:}\,\beta_n$. The intended meaning of this form is that, whenever each of the x_is is replaced by a proof of the corresponding α_i, then at least one of the b_js denotes a constructive proof of the corresponding β_j. Note that the choice of b_i may depend on the particular proofs substituted for $x_0, \ldots, x_m$

Definition 2.3 *As in Definition 2.1, assume an unlimited supply of labels for assumptions. The* nondeterministic proof formulae *are defined inductively as follows:*

(nB) *If x is a label, then $x{:}\,\alpha \vdash x{:}\,\alpha$ for all formulae α;*

(n∧I) *If $\Gamma \vdash a{:}\,\alpha, \Phi$ and $\Delta \vdash b{:}\,\beta, \Psi$, then $\Gamma, \Delta \vdash \langle a, b\rangle{:}\,(\alpha \wedge \beta), \Phi, \Psi$;*

(n∧E) *If $\Gamma \vdash c{:}\,(\alpha \wedge \beta), \Phi$, then $\Gamma \vdash (\sigma_0 c){:}\,\alpha, \Phi$, and $\Gamma \vdash (\sigma_1 c){:}\,\beta, \Phi$;*

(n∨I) *If $\Gamma \vdash a{:}\,\alpha, \Phi$, then $\Gamma \vdash \langle 0, a\rangle{:}\,(\alpha \vee \beta), \Phi$, and $\Gamma \vdash \langle 1, a\rangle{:}\,(\beta \vee \alpha), \Phi$;*

(n∨E) *If $\Gamma \vdash c{:}\,(\alpha \vee \beta), \Phi$, then $\Gamma \vdash (\rho_0 c){:}\,\alpha, (\rho_1 c){:}\,\beta, \Phi$;*

(nN) *If $\Gamma \vdash a_0{:}\,\alpha, a_1{:}\,\alpha, \Phi$, then $\Gamma \vdash [a_0, a_1]_\alpha{:}\,\alpha, \Phi$;*

(n⇒I) *If $x{:}\,\alpha, \Gamma \vdash b{:}\,\beta$, and Γ does not contain an assumption with label x, then $\Gamma \vdash (\lambda x{:}\,\alpha \;.\; b){:}\,(\alpha \Rightarrow \beta)$;*

1. $x{:}\,\alpha \vdash x{:}\,\alpha$ (nB).

2. $x{:}\,\alpha \vdash x{:}\,\alpha, x{:}\,\beta$ 1, (nK).

3. $\vdash x{:}\,\alpha, (\lambda x{:}\,\alpha\ .\ x){:}\,(\alpha \Rightarrow \beta)$ 2, (n⇒IC).

4. $\vdash \langle 0, x\rangle{:}\,(\alpha \vee (\alpha \Rightarrow \beta)),$
   $(\lambda x{:}\,\alpha\ .\ x){:}\,(\alpha \Rightarrow \beta)$ 3, (n∨I).

5. $\vdash \langle 0, x\rangle{:}\,(\alpha \vee (\alpha \Rightarrow \beta)),$
   $\langle 1, (\lambda x{:}\,\alpha\ .\ x)\rangle{:}\,(\alpha \vee (\alpha \Rightarrow \beta))$ 4, (n∨I).

6. $\vdash [\langle 0, x\rangle, \langle 1, (\lambda x{:}\,\alpha\ .\ x)\rangle]{:}\,(\alpha \vee (\alpha \Rightarrow \beta))$ 5, (nN).

Figure 4: Constructively fallacious argument for $\alpha \vee (\alpha \Rightarrow \beta)$.

(n⇒E) *If* $\Gamma \vdash c{:}\,(\alpha \Rightarrow \beta), \Phi$, *and* $\Delta \vdash a{:}\,\alpha, \Psi$, *then* $\Gamma, \Delta \vdash (ca){:}\,\beta, \Phi, \Psi$*;*

(nK) *If* $\Gamma \vdash \Phi$, *then* $\Gamma, \Delta \vdash \Phi, \Psi$.

In principle, the subscript α in $[\cdot,\cdot]_\alpha$ is required to avoid ambiguity in nondeterministic proof formulae, but we omit it when it is clear from the context.

Notice that most of the rules in Definition 2.3 allow the same constructions as in Definition 2.1, with arbitrary additional *proof-formula*:*formula* pairs on the right-hand side. The rule (n⇒I) for introducing an implication sign is an exception, and it allows only one pair on the right-hand side. It is tempting to generalize this rule to allow additional pairs, just like the others:

(n⇒IC) *If* $x{:}\,\alpha, \Gamma \vdash b{:}\,\beta, \Phi$, *and* Γ *does not contain an assumption with label* x, *then* $\Gamma \vdash (\lambda x{:}\,\alpha\ .\ b){:}\,(\alpha \Rightarrow \beta), \Phi$.

Such a generalization is sound for classical logic, but not for constructive logic, as it allows the fallacious derivation of the law $\alpha \vee (\alpha \Rightarrow \beta)$ shown in Figure 4. From this constructively fallacious formula, all classically true formulae may be derived constructively. Notice how the label/variable x is used as an undischarged assumption in $\langle 0, x\rangle$, and inconsistently as if it had two different types (i.e., as if it labelled the two different formulae α and β) within $\langle 1, (\lambda x{:}\,\alpha\ .\ x)\rangle$. The basic idea of nondeterministic proof as embodied in (nB) allows the type-inconsistent usage in the second contingent proof in line (2), because the first contingent proof is type-consistent. Then, the generalized implication introduction rule (n⇒IC) allows the type-consistent use of the assumption x in the first contingent proof in line (2) to remain undischarged, because the type-inconsistent use in the second contingent proof *was* discharged. Thus, each contingent proof does one thing right, and one thing wrong, and the rules (n∨I) and (nN) allow them to be combined and treated as correct. It is intuitively clear that $[\langle 0, x\rangle, \langle 1, (\lambda x{:}\,\alpha\ .\ x)\rangle]$ does not denote a sensible assumption-free construction[1]

[1] We might try to blame the rule (nK) instead of (n⇒IC), but a slightly longer version of the proof uses (n∨E) instead of (nK)

1. $u\colon(\alpha \vee \beta) \vdash u\colon(\alpha \vee \beta)$ (nB).

2. $u\colon(\alpha \vee \beta) \vdash (\rho_0 u)\colon\alpha, (\rho_1 u)\colon\beta$ 1, (n$\vee$E).

3. $\vdash (\lambda u\colon(\alpha \vee \beta)\ .\ (\rho_0 u))\colon((\alpha \vee \beta) \Rightarrow \alpha),$
$(\lambda u\colon(\alpha \vee \beta)\ .\ (\rho_1 u))\colon((\alpha \vee \beta) \Rightarrow \beta)$ 2, (n$\Rightarrow$ICW).

For brevity, let a denote the first proof formula in line (3), and let b denote the second. Let c and d be proof formulae for the easily derived constructive tautologies of lines (4) and (5) below:

4. $\vdash c\colon((\alpha \vee \beta) \Rightarrow \alpha) \Rightarrow (\beta \Rightarrow \alpha)$

5. $\vdash d\colon((\alpha \vee \beta) \Rightarrow \beta) \Rightarrow (\alpha \Rightarrow \beta)$

Using these abbreviations, we continue:

6. $\vdash (ca)\colon(\beta \Rightarrow \alpha),$
$(db)\colon(\alpha \Rightarrow \beta)$ 3, 4 (n$\Rightarrow$E); 3, 5 (n$\Rightarrow$E)

7. $\vdash \langle 1, (ca)\rangle\colon(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha),$
$\langle 0, (db)\rangle\colon(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$ 6, (n$\vee$I) twice

8. $\vdash [\langle 1, (ca)\rangle, \langle 0, (db)\rangle]\colon(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$ 7, (nN)

Figure 5: Constructively fallacious argument for $(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$.

for $\alpha \vee (\alpha \Rightarrow \beta)$. Notice that we have not demonstrated that $\alpha \vee (\alpha \Rightarrow \beta)$ is constructively invalid, but merely that the obvious argument in its favor is constructively invalid. In fact, $\alpha \vee (\alpha \Rightarrow \beta)$ is not derivable by the rules of Definition 2.3, but we will not be able to demonstrate this until we have developed constructive semantics.

Another tempting generalization of (n$\Rightarrow$I) is

(n$\Rightarrow$ICW) *If $x\colon\alpha, \Gamma \vdash b_0\colon\beta_0, \ldots, b_n\colon\beta_n$, and Γ does not contain an assumption with label x, then*

$$\Gamma \vdash (\lambda x\colon\alpha\ .\ b_0)\colon(\alpha \Rightarrow \beta_0), \ldots, (\lambda x\colon\alpha\ .\ b_n)\colon(\alpha \Rightarrow \beta_n).$$

The constructive fallacy in (n$\Rightarrow$ICW) is a bit more subtle. It allows the abbreviated proof of $(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$ shown in in Figure 5. This formula constructively implies $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha)$, but not all classically valid formulae are proved, and the logic of (n$\Rightarrow$ICW) is intermediate between classical and constructive logic. The fallacious step is line (3). Notice that in line (2) a *single* object u, assumed to prove $(\alpha \vee \beta)$, was guaranteed to provide either a proof $(\rho_0 u)$ of α or a proof $(\rho_1 u)$ of β. In line (3), u was abstracted *separately* from each of $(\rho_0 u)$ and $(\rho_1 u)$, so in effect there are two *different* u's in line (3), each assumed to prove $\alpha \vee \beta$. There is no way to guarantee that either the first proves α or the second proves β, since it is quite possible that the first proves β and the second proves α. So, line (3) does not intuitively represent a pair of constructions at least one of which must be correct, and therefore does not fulfill the intent of the nondeterministic proof formulae. The remaining steps merely work the fallacy of line (3) into the form $(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$.

Intuitively, ρ_0 and ρ_1 are intended to denote *contingent* (i.e., *partial*, in the terminology of recursion theory) operators such that $(\rho_0\langle 0, a\rangle) = a$, $(\rho_1\langle 1, b\rangle) = b$, and $(\rho_0 c)$, $(\rho_1 d)$ do not denote proofs when c and d are not of the forms $\langle 0, a\rangle$ and $\langle 1, b\rangle$, respectively. The proof operator $[\cdot, \cdot]_\alpha$ denotes a contingent *nondeterministic* operator such that $[a, b]_\alpha = a$ if a denotes a proof of α but b does not, $[a, b]_\alpha = b$ in the opposite case, $[a, b]_\alpha$ is nondeterministically either of a and b when both denote proofs of α, and finally $[a, b]_\alpha$ does not denote a proof of α if neither of a, b does so. The usual conditional-form deterministic proof formula (χabc) proving γ is represented by the nondeterministic proof formula $[(b(\rho_0 a)), (c(\rho_1 a))]_\gamma$. Nondeterministic proof formulae in which $[\cdot, \cdot]_\alpha$, $(\rho_0 \cdot)$, and $(\rho_1 \cdot)$ appear in other forms than the preceding one do not translate directly into deterministic proof formulae, but it is easy to show that this does not change the essential power of the system.

Theorem 2.4 *There exists a collection of nondeterministic proof formulae, $a_0^{nd}, \ldots, a_n^{nd}$ such that $\Gamma \vdash a_0^{nd}\colon \alpha_0, \ldots, a_n^{nd}\colon \alpha_n$ if and only if there exists a deterministic proof formula a^d such that $\Gamma \vdash a^d\colon (\alpha_0 \vee \cdots \vee \alpha_n)$.*

To prove Theorem 2.4, we need a few lemmata.

Lemma 2.5 *Let $\alpha_0, \ldots, \alpha_m$ and $\beta_0, \ldots, \beta_n$ be two sequences of formulae, such that $\{\alpha_0, \ldots, \alpha_m\} \subseteq \{\beta_0, \ldots, \beta_n\}$ Let a^d be a deterministic proof formula such that $\Gamma \vdash a^d\colon (\alpha_0 \vee \cdots \vee \alpha_m)$. Then, there is a deterministic proof formula b^d such that $\Gamma \vdash b^d\colon (\beta_0 \vee \cdots \vee \beta_n)$.*

Proof: A straightforward generalization of the following example. We derive $(\beta \vee \gamma \vee \alpha)$ from $(\alpha \vee \beta)$. We begin by deriving $\alpha \Rightarrow (\beta \vee \gamma \vee \alpha)$:

1. $v\colon \alpha \vdash v\colon \alpha$ (B)
2. $v\colon \alpha \vdash \langle 1, v\rangle\colon (\gamma \vee \alpha)$ 1, ($\vee$I)
3. $v\colon \alpha \vdash \langle 1, \langle 1, v\rangle\rangle\colon (\beta \vee \gamma \vee \alpha)$ 3, ($\vee$I)
4. $\vdash (\lambda v\colon \alpha\ .\ \langle 1, \langle 1, v\rangle\rangle)\alpha \Rightarrow (\beta \vee \gamma \vee \alpha)$ 3, ($\Rightarrow$I)

Similarly, we can derive $\beta \Rightarrow (\beta \vee \gamma \vee \alpha)$.

5. $\vdash (\lambda w\colon \beta\ .\ \langle 0, \langle 0, w\rangle\rangle)\colon \beta \Rightarrow (\beta \vee \gamma \vee \alpha)$

Finally, we collect the results of the separate derivations using ($\vee$E)

6. $u\colon (\alpha \vee \beta) \vdash u\colon (\alpha \vee \beta)$ (B)
7. $u\colon \alpha \vee \beta \vdash (\chi u(\lambda v\colon \alpha\ .\ \langle 1, \langle 1, v\rangle\rangle)(\lambda w\colon \beta\ .\ \langle 0, \langle 0, w\rangle\rangle))\colon$
   $(\beta \vee \gamma \vee \alpha)$ 4, 5, 6 ($\vee$E)

$\square$ Lemma 2.5

Lemma 2.6 *If a^d is a deterministic proof formula such that $\Gamma \vdash a^d : \alpha$, then there is a nondeterministic proof formula a^{nd} such that $\Gamma \vdash a^{nd} : \alpha$.*

Proof: The Lemma follows by an easy induction on the construction of a^d. Every clause in the definition of deterministic proof formulae is replicated in the definition of nondeterministic proof formulae, except ($\vee$E), which constructs conditional proofs.

Assume that $\Gamma \vdash b^d : (\alpha \vee \beta)$, $\Delta \vdash c^d : (\alpha \Rightarrow \gamma)$, and $\Theta \vdash d^d : (\beta \Rightarrow \gamma)$. It suffices to show that there is a nondeterministic proof formula e^{nd} such that $\Gamma, \Delta, \Theta \vdash e^{nd} : \gamma$.

By our inductive assumption, there are nondeterministic proof formulae b^{nd}, c^{nd}, and d^{nd} such that

1. $\Gamma \vdash b^{nd} : (\alpha \vee \beta)$
2. $\Delta \vdash c^{nd} : (\alpha \Rightarrow \gamma)$
3. $\Theta \vdash d^{nd} : (\beta \Rightarrow \gamma)$

Our construction of e^{nd} continues as follows:

4. $\Gamma \vdash (\rho_0 b^{nd}) : \alpha,$
   $(\rho_1 b^{nd}) : \beta$ — 3, (n$\vee$E)
5. $\Gamma, \Delta, \Theta \vdash (c^{nd}(\rho_0 b^n d)) : \gamma,$ — 2, 4, (n$\Rightarrow$E)
   $(d^{nd}(\rho_1 b^n d)) : \gamma$ — 3, 4, (n$\Rightarrow$E)
6. $\Gamma, \Delta, \Theta \vdash [(c^{nd}(\rho_0 b^n d)), (d^{nd}(\rho_1 b^n d))] : \gamma$ — 5, (nN)

$\square$ Lemma 2.6

Proof of Theorem 2.4:

($\Longleftarrow$) Assume that $\Gamma \vdash a^d : (\alpha_0 \vee \cdots \vee \alpha_n)$. By Lemma 2.6, there is a nondeterministic proof formula a^{nd} such that $\Gamma \vdash a^{nd} : (\alpha_0 \vee \cdots \vee \alpha_n)$.

We now obtain $a_0^{nd}, \ldots, a_n^{nd}$ such that $\Gamma \vdash a_0^{nd} : \alpha_0, \ldots, a_n^{nd} : \alpha_n$ by repeated application of (n$\vee$E).

($\Longrightarrow$) The proof proceeds by induction on the length of the derivation $\Gamma \vdash a_0^{nd} : \alpha_0, \ldots, a_n^{nd} : \alpha_n$.

Basis Straightforward.

Induction The induction step proceeds by cases according to the last step in the derivation. We do the (n$\wedge$I) case for example, the others are closely analogous.

(n∧I) Assume that the last step in deriving $\Gamma \vdash a_0^{nd}\colon \alpha_0, \ldots, a_n^{nd}\colon \alpha_n$ uses (n∧I). We may assume without loss of generality that (n∧I) is used to create the $a_0^{nd}\colon \alpha_0$ term (the right-hand side of the nondeterministic form represents a set, so order is insignificant, and Lemma 2.5 allows us to reorder the disjuncts in the deterministic form at will). So, the derivation has the form:

$\vdots$

i. $\Gamma \vdash b^{nd}\colon \beta, a_1^{nd}\colon \alpha_1, \ldots, a_n^{nd}\colon \alpha_n$

$\vdots$

j. $\Gamma \vdash c^{nd}\colon \gamma, a_1^{nd}\colon \alpha_1, \ldots, a_n^{nd}\colon \alpha_n$

$\vdots$

k. $\Gamma \vdash \langle b^{nd}, c^{nd}\rangle\colon (\beta \wedge \gamma), a_1^{nd}\colon \alpha_1, \ldots, a_n^{nd}\colon \alpha_n$ — i, j, (n∧I)

where $a_0^{nd} = \langle b^{nd}, c^{nd}\rangle$ and $\alpha_0 = \beta \wedge \gamma$. Notice that by (nK), we may assume that the $a_i^{nd}\colon \alpha_i$'s for $i \neq 0$ and the elements of Γ are all present in lines (i) and (j). By our inductive hypothesis, there are deterministic proof formulae b^d and c^d such that

1. $\Gamma \vdash b^d\colon (\beta \vee \alpha_0 \vee \cdots \vee \alpha_n)$
2. $\Gamma \vdash c^d\colon (\gamma \vee \alpha_0 \vee \cdots \vee \alpha_n)$

As in Lemma 2.5, there must be a d^d for the easily provable tautology:

3. $\vdash d^d\colon (\beta \vee \alpha_0 \vee \cdots \vee \alpha_n) \Rightarrow (\gamma \vee \alpha_0 \vee \cdots \vee \alpha_n) \Rightarrow ((\beta \wedge \gamma) \vee \alpha_0 \vee \cdots \vee \alpha_n)$

Finally

4. $\Gamma \vdash (d^d b^d)\colon (\gamma \vee \alpha_0 \vee \cdots \vee \alpha_n) \Rightarrow ((\beta \wedge \gamma) \vee \alpha_0 \vee \cdots \vee \alpha_n)$ — 1, 3, (⇒E)
5. $\Gamma \vdash (d^d b^d c^d)\colon ((\beta \wedge \gamma) \vee \alpha_0 \vee \cdots \vee \alpha_n)$ — 2, 4, (⇒E)

as required.

□ Theorem 2.4

Corollary 2.7 *There exists a nondeterministic proof formula, a^{nd}, such that $\Gamma \vdash a^{nd}\colon \alpha$ if and only if there exists a deterministic proof formula, a^d, such that $\Gamma \vdash a^d\colon \alpha$.*

Proof: This is the $n = 0$ case of Theorem 2.4.

□ Corollary 2.7

Thanks to Theorem 2.4 and Corollary 2.7, various uses of $\vdash$ may be interpreted with respect to deterministic or nondeterministic proof formulae, as convenient, without affecting the meaning of the relation. In addition, nondeterministic proof formulae give a natural meaning to the form $\Gamma \vdash \alpha_0, \ldots, \alpha_n$, i.e., there exist $a_0, \ldots, a_n$ such that $\Gamma \vdash a_0\colon \alpha_0, \ldots, a_n\colon \alpha_n$. The form $\Gamma \vdash \Psi$, with collections of propositional formulae on each side of the $\vdash$, but no proof formulae, is called a *sequent.*

2.3 Sequent Calculus

The inductive definition of proof formulae in Definition 2.1 provides a very natural basis for realizability semantics, but it is not the most convenient for proving completeness results. For that purpose, a *sequent,* or *tableau,* formulation is more helpful. We regard sequent and tableau rules as alternate ways of deriving sequent relations of the form $\Gamma \vdash \Psi$, as discussed after Corollary 2.7. Except for $\Rightarrow$RW, which is an innovation from Kurtz, Mitchell and O'Donnell [KMO92], the following rules (Table 1) are equivalent to the Beth tableau rules [Bet59, Fit69], but presented in sequent notation [Pra65]. The rules are presented schematically about a horizontal line, with several sequent schemas above, and one below. The meaning is that one can derive an instance of the sequent below the line from corresponding instances of the sequents above the line.

Theorem 2.8 *If a sequent relation $\alpha_0, \ldots, \alpha_m \vdash \beta_0, \ldots, \beta_n$ is derivable by the rules in Table 1, then for all labels $x_0, \ldots, x_m$ there exist nondeterministic proof formulae $b_0, \ldots, b_n$ such that $x_0\!:\alpha_0, \ldots, x_m\!:\alpha_m \vdash b_0\!:\beta_0, \ldots, b_n\!:\beta_n$.*

$$(\mathrm{B})\ \Gamma, \alpha \vdash \alpha, \Psi$$

$$(\wedge\mathrm{L})\ \frac{\Gamma, \alpha, \beta, \alpha \wedge \beta \vdash \Psi}{\Gamma, \alpha \wedge \beta \vdash \Psi} \qquad (\wedge\mathrm{R})\ \frac{\begin{array}{c}\Gamma \vdash \alpha \wedge \beta, \alpha, \Psi \\ \Gamma \vdash \alpha \wedge \beta, \beta, \Psi\end{array}}{\Gamma \vdash \alpha \wedge \beta, \Psi}$$

$$(\vee\mathrm{L})\ \frac{\begin{array}{c}\Gamma, \alpha, \alpha \vee \beta \vdash \Psi \\ \Gamma, \beta, \alpha \vee \beta \vdash \Psi\end{array}}{\Gamma, \alpha \vee \beta \vdash \Psi} \qquad (\vee\mathrm{R})\ \frac{\Gamma \vdash \alpha \vee \beta, \alpha, \beta, \Psi}{\Gamma \vdash \alpha \vee \beta, \Psi}$$

$$(\Rightarrow\mathrm{L})\ \frac{\begin{array}{c}\Gamma, \alpha \Rightarrow \beta \vdash \alpha, \Psi \\ \Gamma, \beta, \alpha \Rightarrow \beta \vdash \Psi\end{array}}{\Gamma, \alpha \Rightarrow \beta \vdash \Psi} \qquad (\Rightarrow\mathrm{RS})\ \frac{\Gamma, \alpha \vdash \beta}{\Gamma \vdash \alpha \Rightarrow \beta, \Psi}$$

$$(\Rightarrow\mathrm{RW})\ \frac{\Gamma \vdash \alpha \Rightarrow \beta, \beta, \Psi}{\Gamma \vdash \alpha \Rightarrow \beta, \Psi}$$

Table 1: Constructive Sequent Rules

Proof: The proof is by induction on the length of derivation. The basis case is straightforward.

($\odot$R) Assume that the last step in the derivation is an instance of ($\wedge$R):

$$\frac{\begin{array}{c}\gamma_0, \ldots, \gamma_m \vdash (\alpha \wedge \beta), \alpha, \delta_0, \ldots \delta_n \\ \gamma_0, \ldots, \gamma_m \vdash (\alpha \wedge \beta), \beta, \delta_0, \ldots \delta_n\end{array}}{\gamma_0, \ldots, \gamma_m \vdash (\alpha \wedge \beta), \delta_0, \ldots \delta_n}$$

Fix labels $x_0, \ldots, x_m$. By our inductive assumption, there exist nondeterministic proof formulae a, b, c, d, $e_0, \ldots, e_n$, and $f_0, \ldots, f_n$ such that

1. $x_0\colon\gamma_0, \ldots, x_m\colon\gamma_m \vdash a\colon\alpha, c\colon(\alpha \wedge \beta), e_0\colon\delta_0, \ldots, e_n\colon\delta_n$

2. $x_0\colon\gamma_0, \ldots, x_m\colon\gamma_m \vdash b\colon\beta, d\colon(\alpha \wedge \beta), f_0\colon\delta_0, \ldots, f_n\colon\delta_n$

We continue:

3. $x_0\colon\gamma_0, \ldots, x_m\colon\gamma_m \vdash \langle a, b\rangle\colon(\alpha \wedge \beta), c\colon(\alpha \wedge \beta), d\colon(\alpha \wedge \beta),$
$e_0\colon\delta_0, \ldots, e_n\colon\delta_n, f_0\colon\delta_0, \ldots, f_n\colon\delta_n$ 1, 2, (n∧I)

4. $x_0\colon\gamma_0, \ldots, x_m\colon\gamma_m \vdash [\langle a, b, \rangle, [c, d]]\colon(\alpha \wedge \beta),$
$e_0\colon\delta_0, \ldots, e_n\colon\delta_n, f_0\colon\delta_0, \ldots, f_n\colon\delta_n$ 3, (nN) twice

5. $x_0\colon\gamma_0, \ldots, x_m\colon\gamma_m \vdash [\langle a, b\rangle, [c, d]]\colon(\alpha \wedge \beta),$
$[e_0, f_0]\colon\delta_0, \ldots, [e_n, f_n]\colon\delta_n$ 4, (nN) n times

The other (⊙R) rules translate similarly to the corresponding (n⊙R) rules.

(⊙L) Assume that the last step in the derivation was by an instance of (∧L):

$$\frac{\gamma_0, \ldots, \gamma_m, \alpha, \beta, (\alpha \wedge \beta) \vdash \delta_0, \ldots, \delta_n}{\gamma_0, \ldots, \gamma_m, (\alpha \wedge \beta) \vdash \delta_0, \ldots, \delta_n}$$

Fix labels $x_0, \ldots, x_m$, w, y, and z. By our inductive hypothesis, there must exist nondeterministic proof formulae $d_0, \ldots, d_n$ such that

k. $x_0\colon\gamma_0, \ldots x_m\colon\gamma_m, y\colon\alpha, z\colon\beta, w\colon(\alpha \wedge \beta) \vdash d_0\colon\delta_0, \ldots, d_n\colon\delta_n$

Now,

1. $v\colon(\alpha \wedge \beta) \vdash v\colon(\alpha \wedge \beta)$ (nB)

2. $v\colon(\alpha \wedge \beta) \vdash (\sigma_0 v)\colon\alpha$ 1, (n∧E)

3. $v\colon(\alpha \wedge \beta) \vdash (\sigma_1 v)\colon\beta$ 1, (n∧E)

Lines (2) and (3) enable us to replace occurrences of the label y by $(\sigma_0 v)$, and z by $(\sigma_1 v)$. Let d'_i be the result of such a substitution applied to d_i. By performing such a substitution on the proof formulae in each sequent in the derivation of (k), we convert it to a derivation of

k′. $x_0\colon\gamma_0, \ldots x_m\colon\gamma_m, w\colon(\alpha \wedge \beta) \vdash d'_0\colon\delta_0, \ldots, d'_n\colon\delta_n$

by simple syntactic substitution.

The other (⊙L) rules translate similarly to the corresponding (n⊙E) rules.

□ Theorem 2.8

The proof system of Theorem 2.8 can be modified in several ways that do not affect its power. For example, adding the following weakening rule does not increase the power of the system:

$$(\mathrm{K})\ \frac{\Gamma \vdash \Psi}{\Gamma, \Delta \vdash \Psi, \Theta}$$

Similarly, we can remove ($\Rightarrow$RW) without decreasing the power of the system.

Although the sequent rules above do not mention proofs directly, the construction in the proof of Theorem 2.8 shows that the rules provide a different structural form for the construction of proof formulae. While the rules of Definition 2.1 and Definition 2.3, and the ($\odot$R) rules of Theorem 2.8 build up proof formulae by combining one or two subproofs with a single new operator at the head, the ($\odot$L) rules of Theorem 2.8 build up proof formulae by substituting simple subproofs for the variables/labels at the leaves of other proofs.

Notice that not all proof formulae are generated by the sequent rules. In particular, the sequent rules do not generate proof formulae, such as $((\lambda x{:}\,\alpha\ .\ b)a)$, in which elimination rules are applied to remove operators previously created by introduction rules. Such proof formulae are superfluous, in the sense that they may always be replaced by more direct proof formulae that never eliminate what they have introduced. So, the converse of Theorem 2.8 holds, and every logical consequence derivable by proof formulae is also derivable by the sequent rules. This converse is tricky to prove directly, but it follows from the completeness proof of Section 10.

The sequent rules are presented so that, whenever the relations above the line are true, so is the one below, but they are more often used backwards in a goal-directed search for a derivation. Notice that all of the rules except (B) and ($\Rightarrow$RS) have the property that every formula in the relation below the line appears above the line as well. In all but (B) and ($\Rightarrow$RW), a particular formula below the line is chosen as the focus of attention, and its principal subformulae appear above the line. In ($\Rightarrow$RW), only the right-hand subformula appears above the line. So, when applied backwards in the search for a derivation, the rules may be regarded as decomposing formulae about their head symbols. It is tempting to make ($\Rightarrow$RS) and ($\Rightarrow$RW) look like the others, i.e., generalize them to ($\Rightarrow$RC) as follows:

$$(\Rightarrow\mathrm{RC})\ \frac{\Gamma, \alpha \vdash \alpha \Rightarrow \beta, \beta, \Psi}{\Gamma \vdash \alpha \Rightarrow \beta, \Psi}$$

This generalization embodies the same fallacy as the generalized rule (n$\Rightarrow$IC) for nondeterministic lambda-abstraction proposed after Definition 2.3. The sequent version of that fallacious derivation is shown in Figure 6. So, all theorems of classical propositional calculus are derivable using the sequent rules of Theorem 2.8 plus ($\Rightarrow$RC). It is the retention of Ψ above the line that causes the problem—the retention of $\alpha \Rightarrow \beta$ is benign but pointless. The intermediately powerful rule

$$(\Rightarrow\mathrm{RCW})\ \frac{\Gamma, \alpha \vdash \alpha \Rightarrow \beta_0, \beta_0, \ldots, \alpha \Rightarrow \beta_n, \beta_n}{\Gamma \vdash \alpha \Rightarrow \beta_0, \ldots, \alpha \Rightarrow \beta_n}$$

1. $\alpha \vdash \beta, (\alpha \Rightarrow \beta), \alpha \vee (\alpha \Rightarrow \beta), \alpha$ (B)

2. $\vdash (\alpha \Rightarrow \beta), \alpha \vee (\alpha \Rightarrow \beta), \alpha$ 1, ($\Rightarrow$RC)

3. $\vdash \alpha \vee (\alpha \Rightarrow \beta)$ 2, ($\vee$R).

Figure 6: Constructively fallacious sequent derivation for $\vdash \alpha \vee (\alpha \Rightarrow \beta)$.

is similarly equivalent to the fallacious rule (n$\Rightarrow$ICW) proposed for nondeterministic proof formulae.

So, it is critical to the soundness of constructive reasoning that, when we assume the left-hand side of an implication, we use it only to derive the corresponding right-hand side, rather than some other desired formula. This restriction makes each backwards application of ($\Rightarrow$RS) a crucial choice point in the search for a proof. Since no other rule discards any formula, we may apply those rules indiscriminately, delaying choices until we apply ($\Rightarrow$RS). Notice that this delaying tactic requires the use of multiple formulae on the right-hand side in order to deal with right-hand side disjunctions. Otherwise, the simpler form with a single formula on each right-hand side would suffice, as in Definition 2.1.

The closest we can come to ($\Rightarrow$RC), while maintaining constructive soundness, is to include both ($\Rightarrow$RS) and ($\Rightarrow$RW). These particular rules, with their redundancy, have been chosen to simplify the completeness proof. They allow the maximum possible work to be accomplished by indiscriminate backward application of rules other than ($\Rightarrow$RS), delaying the choice points as much as possible.

3 Kripke Models

Kripke [Kri65] interprets constructive logic as a modal system, in which $\wedge$ and $\vee$ behave classically, but $\alpha \Rightarrow \beta$ has the modal interpretation that in every *reachable* world where α holds, β holds as well. That is, α constructively implies β if *necessarily* α classically implies β. The accessibility relation is required to be reflexive and transitive, but not necessarily symmetric. In the terminology of modal semantics, these are the S_4 models [Kri63].

Definition 3.1 *Let $\mathfrak{M} = \langle \mathbf{U}, P, R \rangle$ be a realizability model. Let S be a set of labelled assumptions. A* valuation of S on $\mathbf{U}$ *is a function v from the labels of S to evidence such that*

$$x{:}\,\alpha \in S \quad \textit{implies} \quad v(x) \in P(\alpha)$$

If v is a valuation of S and $a \in P(\alpha)$ is a piece of evidence, then $[x := a]v$ is the valuation defined as follows:

$$[x := a]v(y) = \begin{cases} v(y), & \textit{if } y \not\equiv x; \\ a, & \textit{if } y \equiv x. \end{cases}$$

Definition 3.2 *A* Kripke model *for positive constructive propositional logic is a triple* $\mathfrak{M} = \langle \mathbf{W}, \preceq, \nu \rangle$*, where* $\mathbf{W}$ *is a set of* worlds*,* $\preceq$ *is a reflexive, transitive binary relation on* $\mathbf{W}$ *called* reachability*, and* ν *is a function, called* satisfaction*, from* $\mathbf{W}$ *to valuations on atomic propositional symbols, which is closed under reachability, i.e., if* $\mathfrak{v} \preceq \mathfrak{w}$*, and if* $\nu_{\mathfrak{v}}(\alpha)$*, then* $\nu_{\mathfrak{w}}(\alpha)$*.*

Notice that we write $\nu_{\mathfrak{w}}(\alpha)$ rather than $\nu(\mathfrak{w})(\alpha)$.

We extend the satisfaction function from atomic formulae to arbitrary formulae as follows:

Definition 3.3 *Let* $\mathfrak{M} = \langle \mathbf{W}, \preceq, \nu \rangle$ *be a Kripke model. The relation* $\mathfrak{w} \Vdash \alpha$ *is defined inductively for* $\mathfrak{w} \in \mathbf{W}$ *and propositional formulae* α *as follows:*

- *If* α *is an atomic propositional symbol, then* $\mathfrak{w} \Vdash \alpha$ *if and only if* $\nu_{\mathfrak{w}}(\alpha)$*;*
- $\mathfrak{w} \Vdash \alpha \wedge \beta$ *if and only if* $\mathfrak{w} \Vdash \alpha$ *and* $\mathfrak{w} \Vdash \beta$*;*
- $\mathfrak{w} \Vdash \alpha \vee \beta$ *if and only if* $\mathfrak{w} \Vdash \alpha$ *or* $\mathfrak{w} \Vdash \beta$*;*
- $\mathfrak{w} \Vdash \alpha \Rightarrow \beta$ *if and only if for every world* $\mathfrak{v}$ *such that* $\mathfrak{w} \preceq \mathfrak{v}$*, and such that* $\mathfrak{v} \Vdash \alpha$*, we also have* $\mathfrak{v} \Vdash \beta$*.*

In addition, we have the following abbreviated forms:

- $\mathfrak{M} \models_{\mathrm{K}} \alpha$ *if and only if* $\mathfrak{w} \Vdash \alpha$ *for every world* $\mathfrak{w}$ *in* $\mathbf{W}$*;*
- $\models_{\mathrm{K}} \alpha$ *if and only if* $\mathfrak{M} \models_{\mathrm{K}} \alpha$ *for every Kripke model* $\mathfrak{M}$*; and*
- $\Gamma \models_{\mathrm{K}} \alpha$ *if and only if* $\mathfrak{M} \models_{\mathrm{K}} \alpha$ *whenever* $\mathfrak{M} \models_{\mathrm{K}} \gamma$ *for every* $\gamma \in \Gamma$*.*
- $\mathsf{Th}_{\mathrm{K}}(\mathfrak{M}) = \{\alpha \mid \mathfrak{M} \models_{\mathrm{K}} \alpha\}$*.* $\mathsf{Th}_{\mathrm{K}}(\mathfrak{M})$ *is called the* theory *of* $\mathfrak{M}$*.*

Although our results concern only the positive propositional calculus, it may be helpful in understanding Kripke models to consider examples of formulae with negation as well. The semantic rule for negation is $\mathfrak{w} \Vdash \neg\alpha$ if and only if there is no $\mathfrak{v}$ such that $\mathfrak{w} \preceq \mathfrak{v}$ and $\mathfrak{v} \Vdash \alpha$. In this treatment, $\neg\alpha$ is equivalent to $\alpha \Rightarrow \bot$, where $\bot$ is false in every world of every model.

The most natural intuitive explanation of a Kripke model is that each world $\mathfrak{w}$ represents a conceivable state of constructive knowledge, and $\mathfrak{w} \models_{\mathrm{K}} \alpha$ means that α is *knowable* in world $\mathfrak{w}$, and $\mathfrak{w} \preceq \mathfrak{w}'$ means that it is possible to be in state $\mathfrak{w}$ at one moment, then in state $\mathfrak{w}'$ at some future moment. The closure of ν under $\preceq$ means that knowability of atomic formulae is never lost, and a simple induction shows that $\models_{\mathrm{K}}$ is similarly closed, so no ability to know the truth of a formula is ever lost. Notice that $\wedge$ and $\vee$ behave classically at each world, but $\alpha \Rightarrow \beta$ asserts that, in all possible futures in which α is knowable, β is knowable as well. Beth models are similar to Kripke models, except that $\alpha \vee \beta$ may hold in a world $\mathfrak{w}$ where neither α nor β holds, as long as every possible future from $\mathfrak{w}$ leads eventually to a world where one of α and β holds. This variation has important technical advantages for some semantic studies. In particular, the theory (set of true formulae) of each world is no longer a prime set. This, together the inclusion of *fallible* models, namely those with possibly inconsistent theories at some worlds, has opened the way for a constructive treatment of completeness by Veldman, de Swart, Friedman and others. See [Tv88] for a discussion. These advantages of Beth semantics are not directly relevant to our analysis. We seek more explicit semantic representation of the concept of construction within the semantics itself. It remains an intriguing question, however, how the fallible Beth approach might be captured in Läuchli semantics. A first attempt at a solution is considered in [LO84].

Many interesting examples of Kripke models may be pictured as labelled trees, where the nodes denote worlds, the reachability relation is the transitive closure of the directed edges, and the labels describe which atomic propositional symbols hold at the associated worlds. To understand Kripke semantics better, consider the Kripke models pictured in Figures 7–10. Each one is the countermodel to one or more interesting formulae that are classically, but not constructively, true.

Without loss of generality in the description of constructive propositional theories, Kripke models may be restricted to be finitely branching (or even binary), well-founded forests. Every theory of a Kripke model that is *disjunctively closed,* (i.e., whenever $\alpha \vee \beta$ is in the theory either α or β is in the theory) is the theory of a finitely branching, well-founded, rooted tree. Theories that are not disjunctively closed require multiple roots with α true at one root and β true at another. In Kripke models with a unique root (i.e., a minimal element under $\preceq$) the root world may be thought of as the present time. A formula holds in a uniquely rooted model if and only if it holds at the root. The restriction of Kripke models to trees does not affect the valid formulae (α such that $\models_{\mathrm{K}} \alpha$) nor the valid consequences (pairs $\langle \Gamma, \alpha \rangle$ such that $\Gamma \models_{\mathrm{K}} \alpha$), even though it restricts the theories of individual models.

Notice that branching is essential to Kripke models. Linearly ordered Kripke models satisfy $\neg\alpha \vee \neg\neg\alpha$, $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha)$, $(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$, $((\alpha \Rightarrow \beta) \Rightarrow \gamma) \Rightarrow ((\beta \Rightarrow \alpha) \Rightarrow \gamma) \Rightarrow \gamma$. The nonlinear model in Figure 8 falsifies all of these formulae. Also, the use of arbitrarily long chains is crucial. For example, Kripke models with chains of length at most 1 satisfy $\alpha \vee (\alpha \Rightarrow (\beta \vee \neg\beta))$, $\alpha \vee (\alpha \Rightarrow (\beta \vee (\beta \Rightarrow \alpha)))$. The model in Figure 9 uses a chain of length 2 to falsify these formulae. Finally, arbitrarily deep nesting of branching is essential. For example, every model with only one level of branching satisfies

$$\neg\alpha \vee \neg\neg\alpha \vee (\alpha \Rightarrow (\neg\beta \vee \neg\neg\beta)),$$

and

$$(\alpha \Rightarrow \gamma) \vee ((\alpha \Rightarrow \gamma \Rightarrow \gamma) \vee (\alpha \Rightarrow ((\beta \Rightarrow \gamma) \vee ((\beta \Rightarrow \gamma) \Rightarrow \gamma)))).$$

But, the model in Figure 10 uses nested branching to falsify these formulae.

It is sometimes crucial to distinguish two worlds $\mathfrak{v}$ and $\mathfrak{w}$ in a Kripke model, even though the same set of atomic formulae holds for each, that is to say $\{\alpha\colon \nu_{\mathfrak{v}}(\alpha)\} = \{\alpha\colon \nu_{\mathfrak{w}}(\alpha)\}$), because the different possible futures may make nonatomic formulae behave differently on the two. For example, the simplest model falsifying $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha)$ is illustrated in Figure 11. In this model, the same atomic propositional formulae are true at $\mathfrak{v}$ and at $\mathfrak{w}$, but, if the

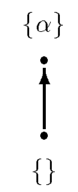

Figure 7: Kripke countermodel for
$\alpha \vee \neg\alpha$,
$\alpha \vee (\alpha \Rightarrow \beta)$,
$\neg\neg\alpha \Rightarrow \alpha$,
$((\alpha \Rightarrow \beta) \Rightarrow \alpha) \Rightarrow \alpha$

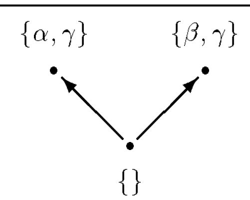

Figure 8: Kripke countermodel for
$\neg\alpha \vee \neg\neg\alpha$, $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha)$
$(\alpha \Rightarrow \beta) \vee (\beta \Rightarrow \alpha)$, $((\alpha \Rightarrow \beta) \Rightarrow \gamma) \Rightarrow ((\beta \Rightarrow \alpha) \Rightarrow \gamma) \Rightarrow \gamma$

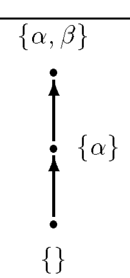

Figure 9: Kripke countermodel for
$\alpha \vee (\alpha \Rightarrow (\beta \vee \neg\beta))$
$\alpha \vee (\alpha \Rightarrow (\beta \vee (\beta \Rightarrow \alpha)))$

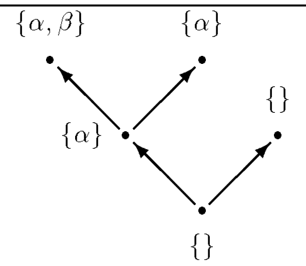

Figure 10: Kripke countermodel for
$\neg\alpha \vee \neg\neg\alpha \vee (\alpha \Rightarrow (\neg\beta \vee \neg\neg\beta))$
$(\alpha \Rightarrow \gamma) \vee ((\alpha \Rightarrow \gamma) \Rightarrow \gamma) \vee (\alpha \Rightarrow ((\beta \Rightarrow \gamma) \vee ((\beta \Rightarrow \gamma) \Rightarrow \gamma)))$

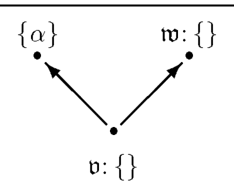

Figure 11: A Kripke countermodel for $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha)$

two worlds are identified, the resulting model satisfies $(\alpha \Rightarrow \beta) \vee ((\alpha \Rightarrow \beta) \Rightarrow \alpha$. The explanation for this apparent paradox is that in world $\mathfrak{v}$ it is possible to come to know α, while this is not possible in world $\mathfrak{w}$.

The temporal-epistemic reading of Kripke models is unsatisfying as a foundation for constructive logic, since there is no explicit treatment of constructions as objects. While Kripke models are very convenient and useful technical tools, particularly for demonstrating that certain formulae are not constructive theorems, a proof of completeness with respect to Kripke models does *not* demonstrate fullness of constructive logic, because we have no reason to believe that *every* Kripke model represents a conceivable constructive reality. It may well be that the fundamental nature of constructions restricts the possible temporal developments of knowledge in a significant way. But, the completeness of the Heyting calculus for Kripke semantics depends critically on the use of rather complex reachability structures on worlds. A temporal-epistemic reading of this ordering seems to serve pedagogical aim at best. It is far from a convincing explanation of the alleged constructive nature of this semantics.

It is not even clear that soundness for Kripke semantics supports faithfulness. The absence of an explicit representation of constructions raises doubt whether the models contain sufficient information to determine truth and falsehood. For example, suppose that in every world for which α holds, β holds as well, but the knowledge of that fact is not constructive. Is it reasonable to say that $\alpha \Rightarrow \beta$ holds constructively?

Dummet [Dum77] has written a much more thorough critique of Kripke and Beth models as interpretations of constructive logic, and has shown why they do not qualify, in his view, as explanations of constructive meaning. The Läuchli realizability models that we construct in Section 6 will turn out to correspond in a very simple way to certain restricted Kripke models, but not all Kripke models will correspond to Läuchli realizability models. So, the proof of completeness with respect to Läuchli models is *prima facie* stronger than completeness with respect to Kripke models, and leads in fact to a reasonable demonstration of fullness.

4 Realizability Semantics

Realizability semantics is based on the intuition that a formula is constructively valid precisely if it is *realized* by some construction that demonstrates its validity. We associate with each propositional formula a class of objects including all conceivable pieces of evidence that might be advanced in support of the formula—even some objects that do *not* constitute constructive proof, along with those that do. The distinction between constructively valid and invalid evidence seems to be better made as a judgement applied *a posteriori*, rather than as a restriction applied *a priori.*

We will consider three different notions of realizability semantics. The first is defined informally in this section, to correspond as closely as possible to a plausible intuition about constructions. The other two are defined formally in subsequent sections to approximate the informal concept. In each of the three, the classes of evidence are essentially the same—all those objects of appropriate type to support a given formula. Each semantics identifies a different subclass of evidence for a formula, called the *realizers,* and defines a valid formula as one that has a realizer in every model. The three semantics, in spite of their different definitions of realizers, turn out to have the same theory.

The following two definitions, in essence a reformulation of the so-called BHK, or Brouwer-Heyting-Kolmogorov interpretation [Tv88] of the logical connectives, are intended as almost precise intuitive ones, not as formal mathematical definitions.

Definition 4.1 (Informal) *A* realizability structure **R** *is a collection of classes $P(\alpha)$ of* pieces of evidence *for each proposition α, such that*

- *For each atomic formula α, $P(\alpha)$ contains all conceivable constructive evidence for α.*
- *$P(\alpha \wedge \beta)$ is the class of all pairs containing a member of $P(\alpha)$ marked in some way, and a member of $P(\beta)$, marked in a different and distinguishable way.*
- *$P(\alpha \vee \beta)$ is the class containing each member of $P(\alpha)$ marked in some way, and each member of $P(\beta)$ marked in a different and distinguishable way.*

- $P(\alpha \Rightarrow \beta)$ *is the class of all* not necessarily uniform or effective rules of correspondence *mapping* $P(\alpha)$ *to* $P(\beta)$.

It is important in Definition 4.1 that $P(\alpha)$ may contain evidence that *does not* constitute constructive proof of α. Although such insufficient evidence will never be taken as legitimate proof, it will have an indirect effect because a construction demonstrating $\alpha \Rightarrow \beta$ is a function that must operate on all conceivable evidence for α, not just that evidence that constitutes constructive proof.

It is tempting to mark each member of $P(\alpha \vee \beta)$ and components of members of $P(\alpha \wedge \beta)$ with a token indicating only which of the propositions denoted by α or β is supported, not whether it is on the left or the right of the $\vee$ or $\wedge$. When α and β denote the same proposition, the marking could be ambiguous. The difference between unambiguous markings based on position in a formula, and possibly ambiguous ones based only on the proposition represented by a given subformula, does not seem to affect any of the results of the semantic analysis in this paper. It would be interesting, nonetheless, to see a careful development of the ambiguous-marking approach.

Several plausible restrictions on the classes of evidence allowed in realizability structures do not seem to matter, and we have aimed for generality as much as possible. For example, our definition does not view evidence as either characterizing, or as being characterized by, the proposition for which it is evidence. Thus, a given piece of evidence is allowed to support different propositions, so $P(\alpha) \cap P(\beta)$ may be nonempty although α and β are not equivalent. Even the extreme restriction that $P(\alpha) \cap P(\beta) = \emptyset$ whenever $\alpha \not\equiv \beta$ does not affect the propositional theories that can be generated by realizability structures.

As long as the classes $P(\alpha)$ are allowed to have arbitrarily large finite sizes, other restrictions on their sizes are irrelevant to the formal propositional theory of constructive logic. For example, we may require them to be all nonempty, all finite, or all infinite, without losing any potential theories. Thus, our fullness results are valid even if there exists potential evidence for every proposition.

Now, we define *satisfaction* of a proposition by a realizability structure. This basic type of relation is central to our semantic analysis, and it appears in several variations, all based on the general form $\mathbf{R} \models a\colon \alpha$. This form asserts that the evidence represented by a constitutes a valid construction verifying the proposition represented by α, in the conceivable state of the world represented by $\mathbf{R}$. In this paper, propositions are always represented by propositional formulae, but the ways in which each of the other two concepts is represented varies, as well as the technical criteria for validity of a construction. A construction may be represented variously by a proof formula, or by a functional associated with the construction. A conceivable state of the world may be represented by a realizability structure, or by any of several types of formal model.

Definition 4.2 (Informal) *Let* $\mathbf{R}$ *be a realizability structure, and* $P(\alpha)$ *one of its classes. When* a *is a* uniformly constructible *member of* $P(\alpha)$*, we write* $\mathbf{R} \models a\colon \alpha$*, and we say that* a realizes α in $\mathbf{R}$.

In addition, we use the following abbreviated forms:

- $\mathbf{R} \models \alpha$ *(read* $\mathbf{R}$ satisfies α*) if and only if there exists an a such that* $\mathbf{R} \models a\colon\alpha$.
- $\models \alpha$ *if and only if* $\mathbf{R} \models \alpha$ *for every realizability structure* $\mathbf{R}$.
- $\Gamma \models \alpha$ *if and only if, whenever* $\mathbf{R}$ *satisfies every formula in* Γ*, then* $\mathbf{R}$ *also satisfies* α.

There are two sources of informality in Definition 4.2. First is the inherited informality from Definition 4.1, where the mechanism for constructing evidence for a composite formula from evidence for its subformulae is given by intuitive description. Then, the intuitive use of the phrase "uniformly constructible" to restrict the realizer a adds a second, and more profound, informality.

The first informality can be removed by formalizing evidence classes in set theory, say in CZF or IZF, or in a sufficiently strong constructive type theory.

Definition 4.3 *A* realizability model *is a triple* $\mathfrak{R} = \langle \mathbf{U}, P, R \rangle$ *satisfying the following conditions.*

- $\mathbf{U}$ *is a set.*
- P *maps atomic propositional formulae to subsets of* $\mathbf{U}$*. We extend* P *to all positive propositional formulae inductively as follows:*
  - $P(\alpha \wedge \beta) = P(\alpha) \times P(\beta)$.
  - $P(\alpha \vee \beta) = (\{0\} \times P(\alpha)) \bigcup (\{1\} \times P(\beta))$.
  - $P(\alpha \Rightarrow \beta) = P(\beta)^{P(\alpha)}$.

  An element of $P(\alpha)$ *is called* evidence *for* α.
- *For each formula* α*,* $R(\alpha) \subseteq P(\alpha)$*. An element of* $R(\alpha)$ *is called a* realizer *for* α.
- $\mathfrak{R} \models_{\mathrm{R}} a\colon\alpha$ *if and only if* $a \in R(\alpha)$
- $\mathfrak{R} \models_{\mathrm{R}} \alpha$ *if and only if there exists an a such that* $\mathfrak{R} \models_{\mathrm{R}} a\colon\alpha$

Note that our formalization of realizability models involves several commitments that are not a part of the informal definition of realizability structures. First, the intuitive classes $P(A)$ of objects providing evidence for atomic formulae are modeled as subsets of a formal set $\mathbf{U}$. Second, the distinguished pairs of $P(A \wedge B)$ are represented by *ordered pairs*, although no order was required for the distinguishing principle in $P(A \wedge B)$. Third, the marked objects of $P(A \vee B)$ are represented by using the particular marks 0 and 1. Finally, the rules of $P(A \Rightarrow B)$ are represented by deterministic *functions in extension*. In each application of realizability models, we must inspect the effects of these changes to insure that the formal properties of the models reflect the true properties of intuitive realizability structures, rather than artifacts of the formal representation techniques.

Resolving the informality in the reference to "uniformly constructible" elements is much more difficult. One reasonable candidate for the class of constructible functions is the *total computable* functions, which is not even recursively enumerable[Odi89]. It is not clear which is the right class of functions to use among those having a satisfactory effective formal definition. So, instead of attempting a direct and precise formal characterization of uniform constructibility, we leave the definition of realizers arbitrary, and investigate different classes of models with different realizers. Given a model $\mathfrak{R}$, notice how the class of formulae α such that $\mathfrak{R} \models_{\mathrm{R}} \alpha$ in Definition 4.2 depends on which members of $P(\alpha)$ are accepted as realizers. In the extreme case, where every piece of evidence is a realizer, we get classical logic.

Definition 4.4 *A realizability model* $\mathfrak{R} = \langle \mathbf{U}, P, R\ \rangle$ *is a* classical *realizability model if and only if* $R\ (\alpha) = P(\alpha)$ *for all formulae* α.

$\Gamma \models_{\mathrm{C}} \alpha$ *if and only if* $\mathfrak{R} \models_{\mathrm{R}} \alpha$ *for every classical realizability model* $\mathfrak{R}$ *such that* $\mathfrak{R} \models_{\mathrm{R}} \gamma$ *for every* $\gamma \in \Gamma$.

In classical realizability models, $\emptyset$ simulates *falsehood,* and each nonempty set simulates *truth.* In particular, $R\ (\alpha \Rightarrow \beta) = R\ (\beta)^{R\ (\alpha)} = \emptyset$ if and only if $\alpha\ \neq \emptyset$ and $\beta\ = \emptyset$. It is easy to show that $\Gamma \models_{\mathrm{C}} \alpha$ if and only if α follows from Γ in classical logic.

If we expand the set of allowable realizers, we may satisfy more formulae, and if we contract it we may satisfy less. So, variations in the class of acceptable realizers have monotonic effects on the relations $\mathfrak{R} \models_{\mathrm{R}} \alpha$ and $\models_{\mathrm{R}} \alpha$. So, we can study the informal concept of uniformly constructible realizers indirectly through formal approximations. That is, we find *two* different approximate criteria for uniform constructibility—a necessary but not sufficient one, and a sufficient but not necessary one. These two formal definitions of realizability bracket uniform constructibility above and below. But, the propositional theories of the two bracketing formalisms turn out to be the same, and therefore they are the same as the intuitive and informal theory of constructive propositional logic.

First, we consider the realizers that are defined by formal constructive proof formulae (the λ-realizers), and the relation $\models_\lambda$ where $\mathfrak{R} \models_\lambda \alpha$ whenever there is a λ-realizer in $P(\alpha)$. It will be obvious from the definition that every λ-realizer is uniformly constructible, and therefore $\models_\lambda \subseteq \models$. So, a formal proof of soundness for $\models_\lambda$ supports the faithfulness of formal constructive logic. Next, we consider a class of *invariant* realizers proposed by Läuchli [Läu70] and the associated relation $\models_{\mathrm{L}}$. Then we argue that every uniformly constructible realizer is necessarily an invariant realizer, and therefore $\models \subseteq \models_{\mathrm{L}}$. So, a formal proof of completeness for $\models_{\mathrm{L}}$ supports the fullness of constructive logic.

In effect, we establish the chain of implications

$$\begin{aligned} \Gamma \vdash \alpha \quad & \text{implies} \quad \Gamma \models_\lambda \alpha && (4.5)\\ & \text{implies} \quad \Gamma \models \alpha && (4.6)\\ & \text{implies} \quad \Gamma \models_{\mathrm{L}} \alpha && (4.7)\\ & \text{implies} \quad \Gamma \vdash \alpha && (4.8) \end{aligned}$$

Implications 4.5 and 4.8 are formal theorems (Theorems 5.6 and 10.12 respectively). Implications 4.6 and 4.7 are informal claims that we justify with detailed intuitive arguments (Propositions 5.4 and 6.12 respectively), but which cannot be proven formally. This strategy lets us show that the formal proof system of Definition 2.1 captures precisely the power of constructive reasoning for positive propositional logic, in spite of the lack of a precise formal and effective characterization of constructions.

5 Lambda Realizability

In this section, we develop the relation $\models_\lambda$, mentioned in Section 4, so as to prove the soundness of $\vdash$ with respect to $\models_\lambda$. The idea is to show that each proof formula $a\colon\alpha$ (cf. Definition 2.1) names a specific piece of evidence in $P(\alpha)$, and then to let these particular named pieces of evidence be the realizers that determine the satisfaction relation $\models_\lambda$. This structure makes the justification of implication 4.6 especially transparent since it is intuitively clear that each named piece of evidence is constructible.

Recall the definition of a valuation from a set of labels to evidence from Section 3, Definition 3.1. The evidence denoted by a proof formula is defined by:

Definition 5.1 *Let $S \vdash a\colon\alpha$, and let v be a valuation of S. The evidence $\mathcal{F}_v(a)$* denoted by a with v *is defined by induction as follows:*

- *if x is a label, then $\mathcal{F}_v(x)$ is $v(x)$;*
- *$\mathcal{F}_v(\langle a, b\rangle)$ is the ordered pair of $\mathcal{F}_v(a)$ and $\mathcal{F}_v(b)$;*
- *$\mathcal{F}_v(\pi_1 a)$ is b when $\mathcal{F}_v(a)$ is the ordered pair of b and c;*
- *$\mathcal{F}_v(\pi_2 a)$ is c when $\mathcal{F}_v(a)$ is the ordered pair of b and c;*
- *$\mathcal{F}_v(\langle 0, a\rangle)$ is the ordered pair of 0 and $\mathcal{F}_v(a)$;*
- *$\mathcal{F}_v(\langle 1, a\rangle)$ is the ordered pair of 1 and $\mathcal{F}_v(a)$;*
- *$\mathcal{F}_v(\chi abc)$ is $\mathcal{F}_v(b)(d)$ when $\mathcal{F}_v(a)$ is the ordered pair of 0 and d, and $\mathcal{F}_v(c)(e)$ when $\mathcal{F}_v(a)$ is the ordered pair of 1 and e;*
- *$\mathcal{F}_v(ba)$ is $\mathcal{F}_v(b)(\mathcal{F}_v(a))$;*
- *$\mathcal{F}_v(\lambda x\colon\alpha\ .\ b)$ is the function mapping each $a \in P(\alpha)$ to $\mathcal{F}_{[x:=a]v}(b)$.*

A piece of evidence b is a lambda proof over v *if there is a proof formula a such that $b = \mathcal{F}_v(a)$.*

When a is a closed proof formula—one with no assumptions—no valuation is required, and we write simply $\mathcal{F}(a)$, and call this a *closed lambda proof.* Läuchli calls the pieces of evidence $\mathcal{F}(a)$ definable by closed proof formulae the *simple functions.* A *lambda model* is just a realizability model, with some specified objects that are postulated to be lambda proofs, and therefore realizers.

Definition 5.2 *A* lambda model *is a triple* $\langle \mathbf{U}, P, v\rangle$*, where* $\mathbf{U}$*,* P *are the first two components of a realizability model, and* v *is a valuation on* $\mathbf{U}$*.*

Lambda semantics is defined by taking the realizers to be the lambda proofs, so a formula α is true precisely when there is a lambda proof in $P(\alpha)$.

Definition 5.3 *If* $\mathfrak{M} = \langle \mathbf{U}, P, v\rangle$ *is a lambda model, then* $\mathfrak{M} \models_\lambda a{:}\,\alpha$ *if and only if* $a \in P(\alpha)$ *is a lambda proof over* v*. As in Definition 3.3, we have the following abbreviated forms:*

- $\mathfrak{M} \models_\lambda \alpha$ *if and only if there is an* a *such that* $\mathfrak{M} \models_\lambda a{:}\,\alpha$*;*
- $\models_\lambda \alpha$ *if and only if* $\mathfrak{M} \models_\lambda \alpha$ *for every lambda model* $\mathfrak{M}$*;*
- $\Gamma \models_\lambda \alpha$ *if and only if* $\mathfrak{M} \models_\lambda \alpha$ *whenever* $\mathfrak{M} \models_\lambda \gamma$ *for every* $\gamma \in \Gamma$*.*

If $\Gamma \models_\lambda \alpha$*, then* α *is a* lambda consequence *of* Γ*.*

There is a very strong and reliable intuitive consensus that all lambda-definable objects are uniformly constructible. Normalization procedures, and implementations of functional programming languages, are practical demonstrations of this constructibility. Because it allows only some of the uniformly constructible objects—the lambda-definable ones—to be used as proofs, lambda semantics is at least as stringent as intuitive realizability semantics. We are now able to establish implication 4.6.

Proposition 5.4 *For all formulae* α *and sets of formulae* Γ*, if* $\Gamma \models_\lambda \alpha$ *then* $\Gamma \models \alpha$*.*

Justification: Each lambda proof is defined by a particular proof formula a. We simply reinterpret a within an arbitrary realizability structure satisfying all members of Γ. Each assumption may be interpreted as referring to some unknown uniformly constructible object in an appropriate class $P(\alpha)$ for some $\alpha \in \Gamma$. The primitive operations creating and projecting from pairs are clearly uniformly constructible, as is the conditional operation χ. The two forms of construction in proofs are application of a function to an argument (ba) and explicit definition of a function by a term $(\lambda x{:}\,\alpha\ .\ b)$. Both of these forms clearly preserve uniform constructibility.

$\square$ Proposition 5.4

Proposition 5.4 argues, in effect, that lambda-definability is a sufficient condition for uniform constructibility.

Now, we can show the soundness of provability with respect to lambda-realizability semantics (i.e., implication 4.5) in a formal theorem.

Theorem 5.5 *For all labelled formulae* $a{:}\,\alpha$ *and sets of labelled formulae* $\{x_0{:}\,\gamma_0, \ldots, x_k{:}\,\gamma_k\}$*, if* $x_0{:}\,\gamma_0, \ldots, x_k{:}\,\gamma_k \vdash a{:}\,\alpha$*, then for all lambda models* $\mathfrak{M} = \langle \mathbf{U}, P, v\rangle$ *with* $v(x_0) \in P(\gamma_0), \ldots v(x_k) \in P(\gamma_k)$*, we have* $\mathfrak{M} \models_\lambda \mathcal{F}_v(a){:}\,\alpha$*.*

Proof: Straightforward.

$\square$ Theorem 5.5

Theorem 5.6 (Soundness of $\vdash$ for $\models_\lambda$) *If* $\Gamma \vdash \alpha$, *then* $\Gamma \models_\lambda \alpha$.

Proof: Assume $\Gamma \vdash \alpha$. Let a be a proof formula such that $\Gamma \vdash a{:}\,\alpha$. Then there exist labels $x_0, \ldots, x_k$ and formulae $\gamma_0, \ldots, \gamma_k$ in Γ such that

$$x_0{:}\,\gamma_0, \ldots, x_k{:}\,\gamma_k \vdash a{:}\,\alpha$$

Let $\mathfrak{M} = \langle \mathbf{U}, P, v \rangle$ be a lambda model such that $\mathfrak{M} \models_\lambda \gamma$ for all $\gamma \in \Gamma$, and let $g_0, \ldots, g_k$ be proof formulae such that $\mathfrak{M} \models_\lambda \mathcal{F}_v(g_i){:}\,\gamma_i$ for each $i \leq k$.

Let a' be the proof formula obtained by replacing every free occurrence of a label x_i by the proof formula g_i. As usual, renaming of bound labels may be necessary to prevent capture of free labels. The free labels that occur in a' are precisely the free labels it inherited from the g_i's. As $\mathcal{F}_v$ is defined on the g_i's, the valuation v must be defined on the free labels that occur in the g_i's, $\mathcal{F}_v$ is defined on a', and therefore $\mathfrak{M} \models_\lambda \mathcal{F}_v(a'){:}\,\alpha$ as required.

$\square$ Theorem 5.6

Along with Proposition 5.4, Theorem 5.6 shows that the conventional systems for constructive propositional logic are faithful to the intuitive realizability semantics. Notice that Theorem 5.6 provides not only *theorem* soundness, i.e., $\vdash \alpha$ implies $\models_\lambda \alpha$, but in fact provides *deductive* soundness, i.e., $\Gamma \vdash \alpha$ implies $\Gamma \models_\lambda \alpha$.

The converse of Theorem 5.6, i.e., the completeness of $\vdash$ for $\models_\lambda$, is also true and not difficult to prove. We omit the proof here, as it will be a trivial consequence of our proof of the completeness of $\vdash$ for $\models_{\mathrm{L}}$. Since the converse to Proposition 5.4 is not intuitively apparent, we cannot infer the fullness of $\vdash$ with respect to $\models$ from the completeness of $\vdash$ with respect to $\models_\lambda$. The technical completeness with respect to lambda-realizability semantics is not satisfying, because it begs the question whether the lambda-definable functions include sufficiently many constructions.

In order to prove an intuitively satisfying completeness result, we must choose tractable *necessary conditions* for uniform constructibility. Most work on realizability follows Kleene [Kle45, Kle59, KV65] in concentrating on the *computability* of functions, using recursive function theory. Rose showed that Kleene's computability-based realizability does not support completeness for the Heyting Calculus—the formula

$$((\neg\neg\gamma \Rightarrow \gamma) \Rightarrow (\neg\neg\gamma \vee \neg\gamma)) \Rightarrow (\neg\neg\gamma \vee \neg\gamma)$$

where

$$\gamma = (\neg\alpha \vee \neg\beta)$$

has a Kleene realizer [Ros53], but is not provable in Heyting Calculus. Läuchli showed [Läu70] how completeness follows from uniformity alone, without considering computability.

6 Läuchli's Realizability Semantics

Logical demonstrations are linguistic constructions, manipulating the *names* of objects rather than the objects themselves. Since the atomic formulae have no predetermined logical meanings, the assignment of names to constructions for atomic formulae is essentially arbitrary, except that a single name may not be assigned to two different objects. The symbols $\wedge$, $\vee$, and $\Rightarrow$, however, have definite logical meanings, so the names for their constructions have logical content, and the assignment of names to constructions is derived from the assignments associated with atomic formulae.

Given the arbitrary nature of the assignment of names to constructions of atomic formulae, we would expect that different minds would make different assignments. Indeed, a single mind might make different assignments at different times. This renders the task of communicating (or even remembering, which might viewed as the special case of communicating with oneself) constructions for atomic formulae by purely logical methods impossible. For composite formulae, however, there is hope. For example, the term $(\lambda x{:}\, \alpha\ .\ x)$ reliably names a specific construction of $\alpha \Rightarrow \alpha$. We are confident that we will interpret this term the same way next year as we do today, and that when we communicate this term to others, they will interpret it in the same way as we do.

Every constructive proof formula should, like $(\lambda x{:}\, \alpha\ .\ x)$, denote a *uniform* piece of evidence, in the sense that it reliably names a construction independently of the assignment of names to evidence for atomic formulae. Läuchli's insight is that constructive proof systems are complete for a semantics ($\models_{\mathrm{L}}$) based on uniformity alone. This lets us define semantics independent of any specific notation for constructions.

Two functions m_1 and m_2 (for *meaning*) from a fixed set of names to a fixed set of pieces of evidence are *type-consistent* if and only if for every name x and formula α, $m_1(x) \in P(\alpha)$ if and only if $m_2(x) \in P(\alpha)$. In a realizability model, we can lift m_1 and m_2 to act consistently on names of realizers of arbitrary type. A name x is (m_1, m_2)-*uniform* if $m_1(x) = m_2(x)$.

Now we need formal *necessary conditions* for a piece of evidence to be a construction, and those necessary conditions will be taken to define the *realizers* for Läuchli models. In the special case where m_1 and m_2 are *bijections* such conditions are simple to define. In this case, $\pi = m_2 \circ m_1^{-1}$, is a permutation on the set of evidence. Moreover, a name x is (m_1, m_2)-uniform if and only if $m_1(x)$ is *invariant* under π. Thus, we can avoid the need to formalize the relation between names and evidence, and instead analyze abstract permutations of the evidence. Our formal definition is based on extending this case to handle more than two meanings (thereby giving rise to more permutations). Notice that if an object is invariant under two permutations, it is necessarily invariant under their compositions and inverses. Therefore, there is no advantage in considering arbitrary sets of permutations: it suffices to consider just those sets that form groups. Since we only require necessary conditions for a piece of evidence to be a construction, we ignore meaning functions that are not bijections (they would lead to more stringent conditions), but a thorough analysis of those functions might be interesting.

We do not claim that all invariant evidence is constructive, nor do we claim that invariance provides an adequate formalization of the notion of uniformity. It suffices for our purposes that we can justify the informal claim that every uniform construction is an invariant piece of evidence, and that we can prove formally that a formula has invariant evidence in every Läuchli model if and only if it is a theorem of the Heyting calculus.

Definition 6.1 *A* Läuchli-model *is a triple* $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$*, where*

- $\mathbf{U}$ *is a set*
- P *maps atomic propositional formulae to subsets of* $\mathbf{U}$
- Π *is a group of permutations of* $\mathbf{U}$*, setwise stabilizing* $P(\alpha)$ *for each atomic formula* α

A Läuchli-model is Abelian *if the group* Π *is.*

Having defined our models, our next goal is to define what it means for a formula to hold in a model. To this end, we lift the permutations of Π from $\mathbf{U}$ to $P(\alpha)$ for arbitrary formula α. We then define $\mathfrak{M} \models_{\mathrm{L}} \alpha$ if and only if there is evidence $a \in P(\alpha)$ that is invariant under (fixed by every element of) Π. The invariant pieces of evidence are the realizers for Läuchli models.

Given a permutation π and an atomic formula α, we define π_α to be the restriction of π to $P(\alpha)$. This is guaranteed to be a permutation by the *setwise stabilizing* clause of Definition 6.1. Assume that we have defined the action of $\pi \in \Pi$ on $P(\alpha)$ and $P(\beta)$. Clearly, the action of π on an element of $P(\alpha \wedge \beta)$ ought to be to permute the components of each distinguished pair independently, according to the action that has already been defined on $P(\alpha)$ and $P(\beta)$. Similarly, the action of π on $P(\alpha \vee \beta)$ ought to be to permute each marked member of $P(\alpha)$ according to its action on $P(\alpha)$, leaving the mark unchanged, and analogously for marked elements of $P(\beta)$.

The permutation π of $P(\alpha \Rightarrow \beta)$ must map $f \in P(\alpha \Rightarrow \beta)$ to a new function $\pi_{\alpha\Rightarrow\beta} f$. Since f operates from and to *unpermuted* evidence, and $\pi_{\alpha\Rightarrow\beta} f$ operates from and to *permuted* evidence, $\pi_{\alpha\Rightarrow\beta} f$ should have the same action on permuted evidence that f has on unpermuted evidence (see Figure 12). From this diagram, it is clear that $\pi_{\alpha\Rightarrow\beta} f$ must satisfy $\pi_{\alpha\Rightarrow\beta} f a = \pi_\beta(f(\pi_\alpha^{-1} a))$, so $\pi_{\alpha\Rightarrow\beta} f = \pi_\beta \circ f \circ \pi_\alpha^{-1}$. We see now that, e.g., the identity function on each class $P(\alpha)$ is invariant under all permutations, as well as the function in $P(\alpha \wedge (\alpha \Rightarrow \beta) \Rightarrow \beta)$ that applies the $P(\alpha \Rightarrow \beta)$ component of its input to the $P(\alpha)$ component.

Definition 6.2 *If* $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$ *is a Läuchli model, then we define the permutations* π_α *of* $P(\alpha)$ *for every* $\pi \in \Pi$ *and formula* α *as follows:*

- *If* α *is an atomic formula, and* $a \in P(\alpha)$*, then* $\pi_\alpha a = \pi a$*;*
- *if* $\langle a, b \rangle \in P(\alpha \wedge \beta)$*, then* $\pi_{\alpha\wedge\beta}\langle a, b \rangle = \langle \pi_\alpha a, \pi_\beta b \rangle$*;*
- *if* $\langle 0, a \rangle \in P(\alpha \vee \beta)$*, then* $\pi_{\alpha\vee\beta}\langle 0, a \rangle = \langle 0, \pi_\alpha a \rangle$*;*

$$
\begin{tikzcd}
\alpha & \beta \\
a_1 \arrow[r, "c"] \arrow[d, "f_\alpha"'] & b_1 \arrow[d, "f_\beta"'] \\
a_2 \arrow[r, "f_{\alpha \Rightarrow \beta} c"] & b_2
\end{tikzcd}
$$

Figure 12: Diagram of the permuted function $\pi_{\alpha \Rightarrow \beta} f$

- *if* $\langle 1, b \rangle \in P(\alpha \vee \beta)$, *then* $\pi_{\alpha \vee \beta} \langle 1, b \rangle = \langle 1, \pi_\beta b \rangle$
- *if* $b \in P(\alpha \Rightarrow \beta)$, *then* $\pi_{\alpha \Rightarrow \beta} b = \pi_\beta \circ b \circ \pi_\alpha^{-1}$

Finally, there are several special types of Läuchli models that we consider in the sequel.

Definition 6.3 *Let* $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$ *be a Läuchli model.*
$\mathfrak{M}$ *is* finite *if and only if* $\mathbf{U}$ *is finite.*
$\mathfrak{M}$ *is* well-ordered *if and only if* $P(\alpha)$ *has a well-ordering for every propositional formula* α.

Notice that the finiteness of $\mathbf{U}$ implies the finiteness of $P(\alpha)$ for all α, but a well-ordering of $\mathbf{U}$ does *not* imply that all $P(\alpha)$s are well-ordered.

Definition 6.4 *If* Π *is the set of all permutations setwise stabilizing* $P(\alpha)$ *for each atomic formula* α, *then* $\langle \mathbf{U}, P, \Pi \rangle$ *is the* full *Läuchli model for* $\langle \mathbf{U}, P \rangle$.

Full models are natural sorts of maximally permuted models. Läuchli implicitly uses a sort of minimally permuted model, in which all permutations are generated by a single π—in this setting he reduces all group-theoretic reasoning to number-theoretic reasoning about the lengths of cycles in π.

Definition 6.5 *If* Π *is a cyclic group of permutations (i.e.* $\Pi = \{\pi^i \mid i \geq 0\}$ *for some permutation* π*), then* $\langle \mathbf{U}, P, \Pi \rangle$ *is a* cyclic *Läuchli model for* $\langle \mathbf{U}, P \rangle$.

We are now ready to define what it means for a formula to hold in a Läuchli model.

Definition 6.6 *If* $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$ *is a Läuchli model and* α *is a formula, then* $\mathfrak{M} \models_{\mathrm{L}} a\colon\alpha$ *if and only if* $a \in P(\alpha)$*, and* $\pi_\alpha(a) = a$ *for every* $\pi \in \Pi$ *(i.e.,* a *is invariant under* Π*).*

Equivalently, we can define a realizability model $\mathfrak{M}' = \langle \mathbf{U}, P, R \rangle$*, where* $R\ (\alpha) = \{a \in P(\alpha) \mid \pi_\alpha a = a \text{ for all } \pi \in \Pi\}$*. Then* $\mathfrak{M} \models_{\mathrm{L}} a\colon\alpha$ *if and only if* $\mathfrak{M}' \models_{\mathrm{R}} a\colon\alpha$*.*

As usual, we have the following abbreviated forms:

- $\mathfrak{M} \models_{\mathrm{L}} \alpha$ *if there exists* $a \in P(\alpha)$ *such that* $\mathfrak{M} \models_{\mathrm{L}} a\colon\alpha$*.*
- $\models_{\mathrm{L}} \alpha$ *if and only if* $\mathfrak{M} \models_{\mathrm{L}} \alpha$ *for every Läuchli model* $\mathfrak{M}$*.*
- $\Gamma \models_{\mathrm{L}} \alpha$ *if and only if* $\mathfrak{M} \models_{\mathrm{L}} \alpha$ *whenever* $\mathfrak{M} \models_{\mathrm{L}} \gamma$ *for every* $\gamma \in \Gamma$*.*
- $\mathsf{Th}_{\mathrm{L}}(\mathfrak{M} = \{\alpha \mid \mathfrak{M} \models_{\mathrm{L}} \alpha\}$*.* $\mathsf{Th}_{\mathrm{L}}(\mathfrak{M})$ *is called the* theory *of* $\mathfrak{M}$*.*

The hierarchies of permutations $\{\pi_\alpha\}$ defined by Läuchli models are examples of the *logical relations* [Sta85] used to study definability in the lambda calculus. Läuchli's results [Läu70] are stated only for full models, but his proofs construct only cyclic models. For generality and smoothness of reasoning we decided to cover both extremes, and everything in between.

Clearly, larger groups make for fewer theorems.

Theorem 6.7 *Let* $\mathfrak{M}_1 = \langle \mathbf{U}, P, \Pi_1 \rangle$ *and* $\mathfrak{M}_2 = \langle \mathbf{U}, P, \Pi_2 \rangle$ *be Läuchli models. If* $\Pi_1 \subseteq \Pi_2$*, then* $\mathfrak{M}_2 \models_{\mathrm{L}} \alpha$ *implies that* $\mathfrak{M}_1 \models_{\mathrm{L}} \alpha$*.*

If $\mathfrak{M} \models_{\mathrm{L}} \alpha$ *for all full models* $\mathfrak{M}$*, then* $\models_{\mathrm{L}} \alpha$*.*

Proof: Trivial.

Theorem 6.7 □

Now, we argue that Läuchli semantics is more liberal than intuitive realizability semantics—that it allows at least as many inferences. Since this claim is not expressed formally, we must appeal to intuition in the argument. In order to illuminate more precisely the intuitive assumptions that are required to support our arguments, we consider three propositions: first that $\models \alpha$ implies $\models_{\mathrm{L}} \alpha$, then that $\Gamma \models \alpha$ implies $\Gamma \models_{\mathrm{L}} \alpha$ for finite sets Γ, and finally that $\Gamma \models \alpha$ implies $\Gamma \models_{\mathrm{L}} \alpha$ for all Γ. The first proposition depends only on basic intuitions about uniformity of constructions, the second requires the deduction property for $\models$, and the third requires compactness.

Proposition 6.8 *For all formulae* α*,* $\models \alpha$ *implies* $\models_{\mathrm{L}} \alpha$*.*

Justification: It suffices to show, for every Läuchli model $\mathfrak{M} = \langle \mathbf{U}, P_{\mathfrak{M}}, \Pi \rangle$, that there exists a realizability structure $\mathbf{R}$ with evidence classes $P_{\mathbf{R}}(\alpha)$ such that every uniformly constructible object in $P_{\mathbf{R}}(\alpha)$ corresponds to an invariant member of $P_{\mathfrak{M}}(\alpha)$. We have already argued that permutation invariance is a necessary condition for uniformity. We must now argue that the formal set-theoretic definition of $P_{\mathfrak{M}}(\alpha)$ in a Läuchli model gives a sufficiently accurate representation of $P_{\mathbf{R}}(\alpha)$ in a corresponding realizability structure.

The basis case where α is atomic is straightforward, if we accept that every formal set corresponds to some intuitive class. $P_{\mathfrak{M}}(\alpha \wedge \beta)$ is the set of *ordered pairs* of elements from $P_{\mathfrak{M}}(\alpha)$ and $P_{\mathfrak{M}}(\beta)$. The informal definition of $P_{\mathbf{R}}(\alpha \wedge \beta)$ required only that the elements of each pair be marked distinguishably, but not necessarily ordered. It is clear that the ordering is ignored in Definition 6.6, so the added specificity of ordering, rather than distinguishing, has no impact on our results. Similarly, the specific choice of 0 and 1 as the marks for the members of $P_{\mathfrak{M}}(\alpha \vee \beta)$ is not used in Definition 6.6—any other marks would have the same effect.

Next, we have gone from the functions-as-rules of $P_{\mathbf{R}}(\alpha \Rightarrow \beta)$ to the functions-in-extension of $P_{\mathfrak{M}}(\alpha \Rightarrow \beta)$. Notice, however, that we did not require the rules constituting evidence in Definition 4.1 to be uniform or effective. So, every extension corresponds to at least one rule—the rule that looks each input up in a graph of the extension. Certainly, every rule has an extension. Different rules for the same extension might be permuted to different extensions, since there is no reason to require any sort of uniformity of the permutations themselves. But limiting attention to extensions *reduces* our freedom in choosing a permutation, which can only *increase* the occurrence of invariants and the set of theorems.

Proposition 6.8 $\square$

Proposition 6.8 above claims that all valid formulae according to a certain constructive intuition are Läuchli valid. In order to support a claim that valid *inferences* are confirmed by Läuchli semantics, we need the *deduction property* to connect inferences and theorems.

Definition 6.9 *A logical consequence relation $\models_{\mathrm{Q}}$ between sets of and formulae has the* deduction property *if and only if, for all formulae α, β, and for all sets of formulae Γ*

$$\Gamma, \alpha \models_{\mathrm{Q}} \beta \textit{ implies } \Gamma \models_{\mathrm{Q}} \alpha \Rightarrow \beta$$

Proposition 6.10 *$\models$—the logical consequence relation of the intuitive realizability semantics of Definition 4.2—satisfies the deduction property.*

Discussion: We can find no detailed rigorous argument in support of the deduction property for $\models$. $\Gamma, \alpha \models \beta$ requires only that every structure satisfying $\Gamma \cup \{\alpha\}$ also satisfies β. This is enough to guarantee a uniform construction mapping each realizability structure $\mathbf{R}$ with uniform constructions for $\Gamma \cup \{\alpha\}$ to a uniform construction for β in $\mathbf{R}$. But, a uniform construction for $\alpha \Rightarrow \beta$ must also map *nonconstructive* evidence a for α to evidence for β. It is very plausible that such mapping can be achieved by considering a modified structure $\mathbf{R}'$, similar to $\mathbf{R}$, but with a corresponding piece of evidence a' arranged to be uniformly constructible. More definite support for the constructive validity of this technique requires a deeper analysis of the properties of uniform constructions than we have achieved here. Perhaps for some constructivists, the deduction property is an article of faith.

Proposition 6.10 $\square$

Proposition 6.11 *For all finite sets* Γ *of formulae, and for all formulae* α

$$\Gamma \models \alpha \textit{ implies } \Gamma \models_{\mathrm{L}} \alpha$$

Justification: Let $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$, and assume $\Gamma \models \alpha$. By Proposition 6.10 applied n times, $\models \gamma_1 \Rightarrow \cdots \gamma_n \Rightarrow \alpha$. By Proposition 6.8, $\models_{\mathrm{L}} \gamma_1 \Rightarrow \cdots \gamma_n \Rightarrow \alpha$. By the validity of *modus ponens* for $\models_{\mathrm{L}}$ (easy to prove), $\Gamma \models_{\mathrm{L}} \alpha$.

Proposition 6.11 □

Proposition 6.12 *For all sets* Γ *of formulae, and for all formulae* α

$$\Gamma \models \alpha \textit{ implies } \Gamma \models_{\mathrm{L}} \alpha$$

Justification: Direct from Propositions 6.11. Even if compactness fails to hold, we must accept Proposition 6.12 for possibly infinite sets of formulae if we believe that the deduction property (Definition 6.9) holds for an extended infinitary or second-order language in which, for each (possibly infinite) set Γ of formulae, there is a formula expressing the conjunction of all the formulae in Γ.

Proposition 6.12 □

The converses to Propositions 6.8–6.12 are not intuitively apparent, but will follow from a later formal proof that $\models_{\mathrm{L}}$ is equivalent to $\models_{\lambda}$.

It is not really crucial to our project of establishing the fullness of constructive proofs with respect to intuitive realizability semantics to have soundness with respect to Läuchli semantics. But, such a soundness result simplifies life, and is easy to prove anyway.

Theorem 6.13 *Let* α *be a formula,* $\{x_0 : \gamma_0, \ldots, x_k : \gamma_k\}$ *be a set of labelled formulae, and* a *be a proof formula. If* $\{x_0 : \gamma_0, \ldots, x_k : \gamma_k\} \vdash a : \alpha$*, then for all Läuchli models* $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$ *and valuations* v *with*

$$v(x_0) \in P(\gamma_0), \ldots v(x_k) \in P(\gamma_k),$$

we have $\mathfrak{M} \models_{\mathrm{L}} \mathcal{F}_v(a) : \alpha$.

Proof: Elementary induction on the structure of the closure of a. It is easy to see that each of the inductive steps in Definition 5.1 of $\mathcal{F}_v(a)$ preserves invariance.

Theorem 6.13 □

Corollary 6.14 (Soundness of $\vdash$ for $\models_{\mathrm{L}}$) *If* $\Gamma \vdash \alpha$*, then* $\Gamma \models_{\mathrm{L}} \alpha$.

Notice that it is absolutely crucial to modeling constructive logic that we require invariant functions to map all objects in their domains, not just the invariant ones. If we redefine P so that $P(\alpha \Rightarrow \beta)$ is the set of invariant functions from $P(\alpha)$ to $P(\beta)$—that is, we filter out variants as we perform the inductive construction of the $P(\alpha)$s, rather than globally after the construction—then we get classical logic again. Similar considerations hold for the intuitive realizability semantics. So, it is crucial to constructive realizability semantics that unacceptable evidence be included in models, in the form of not-uniformly-constructible pieces of evidence, which are to be ruled out as realizers. That is, the inputs to constructions are treated as black boxes that are postulated to be constructions, rather than as explicit constructions in a known system of notation.

The observations above show that, according to realizability semantics, it is not proper to read *"α constructively implies β"* as *"if there is an actual constructive realizer for α, then there is also such a realizer for β."* Rather, it means the same as *"there is a uniform construction which, given arbitrary evidence for α, produces evidence for β that is constructive relative to the presumed construction of α."* So, from our point of view, the implication in constructive logic has a subjunctive, and not simply conditional, quality. And, constructive logic is especially conservative, in that it demands explicit constructions of conclusions, but does not assume that hypotheses are constructed in any particular notation.

In order to manipulate and analyze Läuchli models more conveniently, we develop alternate characterizations of validity in Läuchli models. In particular, $\wedge$ and $\vee$ behave classically, and the validity of $\Rightarrow$ in one model depends on models whose permutation groups are subgroups of the first.

Theorem 6.15 *Let $\mathfrak{L}$ be a Läuchli model.*

- $\mathfrak{L} \models_{\mathrm{L}} \alpha \wedge \beta$ *if and only if* $\mathfrak{L} \models_{\mathrm{L}} \alpha$ *and* $\mathfrak{L} \models_{\mathrm{L}} \beta$.
- $\mathfrak{L} \models_{\mathrm{L}} \alpha \vee \beta$ *if and only if* $\mathfrak{L} \models_{\mathrm{L}} \alpha$ *or* $\mathfrak{L} \models_{\mathrm{L}} \beta$.

Proof: Elementary.

Theorem 6.15 □

Definition 6.16 *Let $\mathfrak{L} = \langle \mathbf{U}, P, \Pi \rangle$ be a Läuchli model, $a \in P(\alpha)$ for some propositional formula $\alpha \in \mathsf{PF}$.*

The orbit *of a in a subgroup $\Pi' \subseteq \Pi$ is $\{\pi(a) \mid \pi \in \Pi'\}$.*

The stabilizer *of a in Π is $\mathbf{stab}_\Pi(a) = \{\pi \in \Pi \mid \pi_\alpha(a) = a\}$.*

Similarly, if $A \subseteq \bigcup\{P(\alpha) \mid \alpha \in \mathsf{PF}\}$, then the pointwise stabilizer of A in Π *is $\mathbf{stab}_\Pi(A) = \{\pi \in \Pi \mid (\forall \alpha \in \mathsf{PF}, a \in A \cap P(\alpha))\pi_\alpha(a) = a\}$.*

If $A \subseteq \bigcup\{P(\alpha) \mid \alpha \in \mathsf{PF}\}$, then the setwise stabilizer of A in Π *is $\mathbf{set\text{-}stab}_\Pi(A) = \{\pi \in \Pi \mid (\forall \alpha \in \mathsf{PF}, a \in A \cap P(\alpha))\pi_\alpha(a) \in A \cap P(\alpha)\}$.*

When α above is atomic, we get the conventional definitions of pointwise and setwise stabilizers. It is easy to prove by induction on the structure of α that the extended concepts are well-defined.

Theorem 6.17 *Let $\mathfrak{L} = \langle \mathbf{U}, P, \Pi \rangle$ be a Läuchli model. Of the three propositions below, (1)$\Rightarrow$(2)$\Leftrightarrow$(3). Furthermore, if $\mathfrak{L}$ is well-ordered, then (3)$\Rightarrow$(1), so all three are equivalent.*

1. $\mathfrak{L} \models_{\mathrm{L}} \alpha \Rightarrow \beta$
2. *For all subgroups $\Pi' \subseteq \Pi$ such that $\langle \mathbf{U}, P, \Pi' \rangle \models_{\mathrm{L}} \alpha$, $\langle \mathbf{U}, P, \Pi' \rangle \models_{\mathrm{L}} \beta$*
3. *For all $a \in P(\alpha)$, $\langle \mathbf{U}, P, \mathbf{stab}_\Pi(a) \rangle \models_{\mathrm{L}} \beta$*

Proof: The proof that (1)⇒(2)⇔(3) is elementary.

For (3)⇒(1), assume that there is a well-ordering of $\bigcup\{P(\alpha) \mid \alpha \in \mathsf{PF}\}$, and that (3) holds. We construct a function $f \in P(\alpha \Rightarrow \beta)$, such that f is invariant under Π.

First, let g be a function from $P(\alpha)$ to $P(\beta)$ such that for all $a \in P(\alpha)$, $\mathbf{stab}_\Pi(a)$ fixes $g(a)$ as well. Such a function is guaranteed by (3). Now, let $f(a) = \pi_\beta^{-1}(g(\pi_\alpha(a)))$, where $\pi_\alpha(a)$ is the least element (under the given well-ordering) in the orbit of a (i.e., in $\{\pi_\alpha(a) \mid \pi \in \Pi\}$).

To see that f is invariant, consider an arbitrary $\pi \in \Pi$, $a \in P(\alpha)$.

$$\pi_{\alpha\Rightarrow\beta}(f)(a) = \pi_\beta^{-1}(f(\pi_\alpha(a))) = \pi_\beta^{-1}(\sigma_\beta^{-1}(g(\sigma_\alpha(\pi_\alpha(a))))) = \\ (\sigma_\beta \circ \pi_\beta)^{-1}(g((\sigma_\alpha \circ \pi_\alpha)(a))) = f(a),$$

where $(\sigma_\alpha \circ \pi_\beta)(a)$ is the least element in the orbit of a, which is the same as the orbit of $\pi_\alpha(a)$.

Theorem 6.17 □

Theorem 6.18 *Let $\mathfrak{L} = \langle \mathbf{U}, P, \Pi \rangle$ be a Läuchli model. Then (1)⇒(2) below.*

1. $\mathfrak{L} \models_{\mathrm{L}} \alpha \Rightarrow \beta$
2.
   - $\mathfrak{L} \models_{\mathrm{L}} \alpha$ *implies* $\mathfrak{L} \models_{\mathrm{L}} \beta$, *and*
   - *for all a such that* $\mathbf{stab}_\Pi(a) \neq \Pi$, $\langle \mathbf{U}, P, \mathbf{stab}_\Pi(a) \rangle \models_{\mathrm{L}} \alpha \Rightarrow \beta$

If $\mathfrak{L}$ is well-ordered, then (1) and (2) are equivalent.

Proof: (1)⇒(2) is direct, and the rest is an elementary application of Theorem 6.17.

Theorem 6.18 □

7 Converting Läuchli Models to Kripke Models.

In this section, we show that for every well-ordered Läuchli Model $\mathfrak{L}$, there is a logically or *elementarily equivalent* rooted Kripke model $\mathfrak{K}$. First we make this notion of equivalence precise:

Definition 7.1 *Let $\mathfrak{A}$ and $\mathfrak{B}$ be models for propositional logic. We say the models are* **elementarily equivalent**, *and write $\mathfrak{A} \equiv \mathfrak{B}$ if they have the same theories, that is to say, if precisely the same set of propositions is true in each.*

To construct a Kripke model from a Läuchli model, we let the subgroups of the permutation group be Kripke worlds.

Definition 7.2 *Let $\mathfrak{L} = \langle \mathbf{U}, P, \Pi \rangle$ be a Läuchli model. $\mathcal{K}(\mathfrak{L}) = \langle \mathbf{W}, \preceq, \nu \rangle$, where $\mathbf{W}$ is the set of all subgroups of Π, $\preceq$ is the subgroup relation, and $\nu_{\mathfrak{w}}(\alpha)$ if $P(\alpha)$ contains an object invariant under the permutations in $\mathfrak{w}$.*

Lemma 7.3 *If $\mathfrak{L}$ is a Läuchli model, then $\mathcal{K}(\mathfrak{L})$ is a rooted Kripke model.*

Proof: To prove that they are Kripke models, we need only show that ν is closed upward under $\preceq$, which follows directly from the fact that a point invariant in a permutation group is invariant in all of its subgroups. Rootedness is elementary, since Π is the root.

Theorem 7.4 *Let $\mathfrak{L} = \langle \mathbf{U}, P, \Pi \rangle$ be a well-ordered Läuchli model. Then* $\mathsf{Th}_{\mathrm{L}}(\mathfrak{L}) = \mathsf{Th}_{\mathrm{K}}(\mathcal{K}(\mathfrak{L}))$.

Proof: By induction on the structure of a formula α, using Theorems 6.15 and 6.17, α holds in each Läuchli model $\langle \mathbf{U}, P, \Pi' \rangle$ iff it holds at the world Π' of $\mathcal{K}(\mathfrak{L})$.

Theorem 7.4 □

8 A Technique for Constructing Kripke and Läuchli Models

The essence of a completeness proof is a procedure which, given a candidate sequent $\Gamma \vdash \alpha$, Γ, constructs either a proof of α from Γ or a countermodel satisfying Γ but not α. If the procedure is effective, and always halts, and the properties of the procedure can be proved constructively, then we get a constructive proof that *every sequent is either provable, or false.* We can prove such a result for the propositional calculus, because it is decidable. For more powerful languages, the procedure is forced to compute infinitely when it fails to find a proof, and it constructs an infinite countermodel by a limiting process. In such a case, the best we can hope for is a constructive proof that *if a sequent is true then it is provable,* or the even weaker contrapositive *if a sequent is not provable, then it is false.*

The proof/countermodel procedure above requires a basic model constructor to produce countermodels for a sequent $\Gamma \vdash \alpha$ from countermodels for other sequents that arise in the search for a proof. It is easier to understand such a constructor for Kripke models, before developing an analogous constructor for Läuchli models.

Tree-like Kripke models are naturally conceived as constructed using several operations: one, *sconing* that attaches to a given model a new least node, another, a *joining* or *gluing* operation that combines several given models above a new root. The latter operation can be thought of as sconing of each model followed by identification of root nodes. The result of such operations depends, not only on the models to be combined, but also on the settings of atomic formulae at the new root. In order for the operation to make sense, the settings at the new root must be consistent with those at the roots of the original models.

Definition 8.1 *Let $\mathfrak{K} = \langle \mathbf{W}, \preceq, \nu \rangle$, $\mathfrak{K}_1 = \langle \mathbf{W}_1, \preceq_1, \nu_1 \rangle$, and $\mathfrak{K}_2 = \langle \mathbf{W}_2, \preceq_2, \nu_2 \rangle$ be Kripke models, and let Γ be a set of atomic formulae such that*

$$\mathfrak{K} \models_{\mathrm{K}} \gamma, \mathfrak{K}_1 \models_{\mathrm{K}} \gamma \text{ and } \mathfrak{K}_2 \models_{\mathrm{K}} \gamma$$

for all $\gamma \in \Gamma$. We define the Γ**-scone** $\mathfrak{K}^{\Gamma}$ *of $\mathfrak{K}$ to be the model $\langle \mathbf{W}^{\mathbf{o}}, \preceq^{o}, \nu^{\Gamma} \rangle$, where*

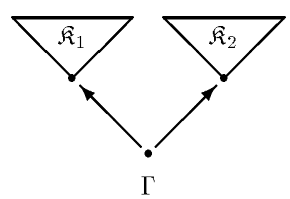

Figure 13: The joined Kripke model $\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2$

- $W^o = \{\mathfrak{r}_o\} \cup W$, *with* $\mathfrak{r}_o \notin W$
- $\mathfrak{r}_o \preceq^o \mathfrak{w}$ *for every world* $\mathfrak{w}$ *in* W, *and* $\preceq^o\restriction_W = \preceq$.
- $\nu^\Gamma_{\mathfrak{r}_o}(\gamma)$ *if and only if* $\gamma \in \Gamma$, *and for every other world* $\mathfrak{w}$, $\nu^\Gamma_{\mathfrak{w}}(\alpha) \Leftrightarrow \nu_{\mathfrak{w}}(\alpha)$.

We define the Γ-join $\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2$ *to be the model* $\langle \mathbf{W}^*, \preceq^*, \nu^* \rangle$, *where*

- $\mathbf{W}^* = (\{1\} \times \mathbf{W}_1) \bigcup (\{2\} \times \mathbf{W}_2) \bigcup \{\mathfrak{w}_0\}$,
  for some $\mathfrak{w}_0 \notin (\{1\} \times \mathbf{W}_1) \bigcup (\{2\} \times \mathbf{W}_2)$.
- $\preceq^*$ *is the least relation satisfying*
  - $\mathfrak{w}_0 \preceq^* \mathfrak{w}$ *for all* $\mathfrak{w} \in \mathbf{W}$
  - $\mathfrak{w} \preceq_i \mathfrak{v}$ *implies* $\langle i, \mathfrak{w} \rangle \preceq^* \langle i, \mathfrak{v} \rangle$
- ν^* *is defined by*
  - $\nu^*_{\mathfrak{w}_0}(\alpha)$ *if and only if* $\alpha \in \Gamma$
  - $\nu^*_{\langle i, \mathfrak{w} \rangle}(\alpha)$ *if and only if* $\nu_{i\mathfrak{w}}(\alpha)$

$\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2$ *is shown pictorially in Figure 13.*

We will not require a precise characterization of the theory of joined and sconed models here. The construction, and its connections with realizability and the slash operation are discussed at length in Smorynski's [Smo73], in [Tv88]. An equivalent operation on Heyting algebras was defined by Freyd (who first called it *sconing*) in the 1970's, who extended it to arbitrary categories (see [FS90]). Freyd's scone is also discussed and applied in [LS86, SS82, MS93] .

We will be interested here in the correspondence between joins of Kripke models (and of Läuchli models), treated in this section, and the use of joins in the completeness proof of Section 10. The following partial characterization is included just to exercise the intuition about joins.

Definition 8.2 *A set of formulae* Γ *is* deductively closed *if and only if* $\Gamma \vdash \alpha$ *implies that* $\alpha \in \Gamma$.

Γ *is* disjunctively closed *(or* prime *if and only if* $(\alpha \vee \beta) \in \Gamma$ *implies that either* $\alpha \in \Gamma$ *or* $\beta \in \Gamma$.

Observe that $\mathsf{Th}_{\mathrm{K}}(\mathfrak{K})$ is deductively and disjunctively closed, for every rooted Kripke model $\mathfrak{K}$.

Theorem 8.3 *Let $\mathfrak{K}_1$ and $\mathfrak{K}_2$ be Kripke models satisfying all formulae in Γ.*

- $\mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2) \subseteq \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1) \cap \mathsf{Th}(\mathfrak{K}_2)$
- *For each atomic formula α, $\alpha \in \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2)$ if and only if $\alpha \in \Gamma$.*
- $\mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2)$ *is a function of* $\mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1)$ *and* $\mathsf{Th}(\mathfrak{K}_2)$ *and* Γ.

Proof: Elementary induction on the structure of formulae.

Theorem 8.3 □

In fact, of those deductively and disjunctively closed sets of formulae that are subsets of the intersection of the theories of Kripke models $\mathfrak{K}_1$, and $\mathfrak{K}_2$, and that contain exactly the atomics in Γ, the theory of the join is a maximal one.

Theorem 8.4 *Let $\mathfrak{K}_1$ and $\mathfrak{K}_2$ be Kripke models satisfying all formulae in Γ. Let Ψ be a deductively and disjunctively closed set of formulae, such that* $\mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2) \subseteq \Psi \subseteq \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1) \cap \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_2)$ *and* $\{\alpha \mid \alpha \in \Psi$ *and* α *is atomic*$\} = \Gamma$. *Then* $\Psi = \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2)$.

Proof: By elementary induction on the structure of a formula α, if $\alpha \in \Psi$, then
$\alpha \in \mathsf{Th}_{\mathrm{K}}(\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2)$.

Theorem 8.4 □

It is easy to see that the theory of the join of Kripke models, although it is *maximal* in the class described above, is not *maximum*. Consider joining two one-world Kripke models, each verifying α and β, but no other atomic formula. Take the $\emptyset$-join. The theory of the join contains $\alpha \Rightarrow \beta$, but not $(\alpha \Rightarrow \beta) \Rightarrow \alpha$. There is another maximal deductively and disjunctively closed set containing $(\alpha \Rightarrow \beta) \Rightarrow \alpha$ but not $\alpha \Rightarrow \beta$, particularly the theory of the ternary $\emptyset$-join that adds a third one-world Kripke model with α true but β false.

Theorem 8.4 □

In a completeness proof, the join operation may be used to combine a countermodel $\mathfrak{K}_1 \not\models_{\mathrm{K}} \alpha_1 \Rightarrow \beta_1$ and a countermodel $\mathfrak{K}_2 \not\models_{\mathrm{K}} \alpha_2 \Rightarrow \beta_2$ into a simultaneous countermodel for both implications: $\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2 \not\models_{\mathrm{K}} \alpha_1 \Rightarrow \alpha_2$ and $\mathfrak{K}_1 \oplus_\Gamma \mathfrak{K}_2 \not\models_{\mathrm{K}} \alpha_1 \Rightarrow \alpha_2$.

9 Separating, Sconing and Gluing

We now define some useful geometric constructions on Läuchli models. We will use the following conventions.

9.1 Separated Group Actions

Let G be an abelian group. We will write $H < G$ to mean H is a subgroup of G, and G/H to mean the set of all *cosets* of H in G, that is to say, the set

$$\{gH : g \in G\} \qquad \text{where} \qquad gH = \{gx : x \in H\}.$$

We will write $[g]_H$ as alternative notation for gH. Recall that G/H can be viewed as the set of equivalence classes of G modulo the equivalence relation

$$x \sim y \Leftrightarrow x^{-1}y \in H.$$

A *group action* of a group G on a set U is a triple (G, U, φ) where φ is a map

$$\varphi : G \times U \;\rightarrow\; U$$

satisfying $\varphi(e, x) = x$ and $\varphi(g, \varphi(h, x)) = \varphi(gh, x)$, where $g, h \in G$, e is the identity in G and $x \in U$.

The value $\varphi(g, x)$ is called the action of g on x and is often written gx when the mapping φ is clear from context. Any group action of G on U can be identified with a group of permutations of U by identifying each g in G with the permutation $x \mapsto gx$ For the rest of this section we will only be considering group actions given by a group of permutations on a set, so we will just denote such actions by *pairs* (G, U) where U is a set and G a group of permutations of U. The group of all permutations of a set U will be denoted S_U.

Definition 9.1 *We say a group action (Π, U) is* **separated** *if for every subgroup H of Π there is an element of U fixed by H and its subgroups, but by no other subgroup of Π. We say the Läuchli model $\mathfrak{L} = \langle U, P, \Pi\rangle$ is* **separated** *if $\langle \Pi, U\rangle$ is.*

Definition 9.2 *Let $\mathfrak{L} = \langle U, P, \Pi\rangle$ be a Läuchli model. Define, for each atomic formula α*

$$[\![\alpha]\!] = \{H < \Pi : Fix(H) \cap P(\alpha) \neq \varnothing\}$$

Let $\mathcal{K}(\Pi)$ be the subgroup lattice of Π ordered by reverse inclusion. A subset of $\mathcal{K}(\Pi)$ is **open** *if it is upwards closed. Observe that each $[\![\alpha]\!]$ is open. For $\mathcal{O}$ an open set, define $SQ(\Pi, \mathcal{O}) = \{\Pi/H : H < \Pi$ and $H \in \mathcal{O}\}$. When the group Π is clear from context we will simply write $SQ(\mathcal{O})$. The* **separated universe** *of Π is $\bigcup SQ(\Pi, \mathcal{K}(\Pi))$. The separated* **closure** *of Π is the group action $\langle \Pi, \bigcup SQ(\Pi, \mathcal{K}(\Pi))\rangle$ with $\pi([x]_H) = [\pi x]_H$.*

The separated closure $Sep(\mathfrak{L})$ of the Läuchli model $\mathfrak{L}$ is the Läuchli model $\langle \bigcup SQ(\Pi, \mathcal{K}(\Pi)), \widehat{P}, \Pi^o\rangle$ where $\widehat{P}(\alpha) = \bigcup SQ(\Pi, [\![\alpha]\!])$.

A group action of the form $(\Pi, \bigcup SQ(\mathcal{O}))$ just described is sometimes referred to as a *group acting on its right cosets* in the literature.

Lemma 9.3 *Let G and H be subgroups of a group Π acting on its cosets as described above. Then some member of Π/G is fixed by every member of H if and only of $H < G$. Thus, if $\mathcal{O}$ is an open set of subgroups of Π and $H < \Pi$ then every member of H fixes a member of $\bigcup SQ(\Pi, \mathcal{O})$ if and only if $H \in \mathcal{O}$.*

Proof: Suppose $h([x]_G) = [x]_G$. Then $[hx]_G = [x]_G$, whence $h \in G$. If every $h \in H$ fixes $[x]_G$ then $H < G$. By definition of open, if every such h fixes some $[x]_G$ in $\bigcup SQ(\Pi, \mathcal{O})$ for some $G \in \mathcal{O}$ then $H \in \mathcal{O}$.

9.3 $\square$

An easy consequence of this fact is the following lemma, which asserts that every Läuchli model has a separated equivalent.

Lemma 9.4 *If $\mathfrak{L}$is abelian, then $Sep(\mathfrak{L})$ is separated, and* $\mathsf{Th}(Sep(\mathfrak{L})) = \mathsf{Th}(\mathfrak{L})$.

Proof: Let $\mathfrak{L} = \langle U, P, \Pi \rangle$. By 7.4 the theories of $Sep(\mathfrak{L})$ and $\mathfrak{L}$ are the same as those of their associated Kripke models $\mathcal{K}(Sep(\mathfrak{L}))$ and $\mathcal{K}(\mathfrak{L})$. The underlying partial order (subgroup lattices) of the two models are the same, so it suffices to show that the same atomic formulas are true at the same corresponding worlds, i.e. that for each subgroup H of Π the intersection of H with $P(\alpha)$ is nonempty just when its intersection with $\widehat{P}(\alpha)$ is. But this is true by definition of $\widehat{P}(\alpha)$.

9.4 $\square$

In this section, we make repeated use of the equivalence of a Läuchli model $\mathfrak{L}$ with its associated Kripke model $\mathcal{K}(\mathfrak{L})$, established in theorem 7.4.

As shall be seen shortly, the *group action* yielded by separated closure of the group of a model is of greater interest than the separated closure of $\mathfrak{L}$ itself. Using the separated universe $\bigcup SQ(\Pi, \mathcal{K}(\Pi))$ we can exert quite a bit of control on the structure of the resulting Läuchli model by selectively picking out the subgroups H whose quotients Π/H we wish to place in $P(\alpha)$.

9.2 Sconing Läuchli Models

In this section we will define a Läuchli-model analogue of the sconing and gluing operations defined for Kripke models above. In particular, given Läuchli models $\mathfrak{L}_1, \mathfrak{L}_2$ and an atomic subtheory Γ of $\mathsf{Th}(\mathfrak{L}_1) \cap \mathsf{Th}(\mathfrak{L}_2)$ we will define models $\mathfrak{L}^\Gamma$ and $\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2$ elementarily equivalent to the Kripke models $\mathcal{K}(\mathfrak{L})^\Gamma$ and $\mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$ respectively. For the remainder of this section it will be convenient to describe the atomic assignment of a Kripke model $\mathfrak{K} = \langle K, \leq, \nu \rangle$ as a *function* $\mathfrak{A}_\mathfrak{K}$ from nodes to sets of atomic formulas, instead of as a predicate ν as in former sections. That is to say, we will write $\mathfrak{A}_\mathfrak{K}(\mathfrak{w}) = \mathfrak{S}$ for some set of formulas S as another way of saying $\nu_\mathfrak{w}(\alpha) \Leftrightarrow \alpha \in S$.

Definition 9.5 *Let $\mathfrak{K}_1 = \langle K_1, \leq_1, \mathfrak{A}_1 \rangle, \mathfrak{K}_2 = \langle K_2, \leq_2, \mathfrak{A}_2 \rangle$ be Kripke models. A* **logomorphism** *between $\mathfrak{K}_1$ and $\mathfrak{K}_2$ is a binary relation $\mathfrak{R}$ on $W_1 \times W_2$ (the worlds of the models) s.t.*

1. *$\mathfrak{R}$ covers K_1 and K_2 (it relates every member of K_1 to something in K_2 and vice versa)*
2. *If $\mathfrak{w}_1 \mathfrak{R} \mathfrak{w}_2$, then the atomic formulae forced by $\mathfrak{w}_1$ are precisely the same as those forced by $\mathfrak{w}_2$.*

3. *If the theory of* $\mathfrak{w}_1$ *in* K_1 *is consistent, i.e.* **not** *the set* PF *of all propositions, then, if* $\mathfrak{w}_1 \Re \mathfrak{w}_2$ *and* $\mathfrak{w}_1 <_1 \mathfrak{w}'_1$ *there is a corresponding* $\mathfrak{w}'_2$ *s.t.* $\mathfrak{w}'_1 \Re \mathfrak{w}'_2$ *and* $\mathfrak{w}_2 <_2 \mathfrak{w}'_2$ *and conversely for consistent worlds* $\mathfrak{w}_2$ *in* K_2.

We will say two Kripke models are **logomorphic** *if there is a logomorphism between them.*[2]

Lemma 9.6 *Let* $\mathfrak{K}_1, \mathfrak{K}_2$ *be Kripke models, and let* $\Re$ *be a logomorphism between them. Then, whenever* $\mathfrak{w}_1 \Re \mathfrak{w}_2$, $\mathfrak{w}_1$ *and* $\mathfrak{w}_2$ *force precisely the same formulae. A fortiori,* K_1 *and* K_2 *are elementarily equivalent.*

Proof: by induction on formula structure.

Basis: it holds for the atomics by definition of logomorphism.

Induction: the only interesting case is implication. Suppose $\alpha \Rightarrow \beta$ is forced at $\mathfrak{w}_1$ in $\mathfrak{K}_1$, where $\mathfrak{w}_1 \Re \mathfrak{w}_2$ and $\mathfrak{w}_2 < \mathfrak{w}'_2$. If α is forced at $\mathfrak{w}'_2$, then there is a $\mathfrak{w}'_1$ with $\mathfrak{w}_1 < \mathfrak{w}'_1$ and $\mathfrak{w}'_1 \Re \mathfrak{w}'_2$. By the induction hypothesis, α is forced at $\mathfrak{w}'_1$. Since αimpliesβ is forced at $\mathfrak{w}_1$, β is forced at $\mathfrak{w}'_1$. By the induction hypothesis again β is forced at $\mathfrak{w}'_2$.

9.6 □

Now we define the scone and join of Läuchli models.

The scone of a Läuchli model is achieved first by *lifting* and then slightly modifying the atomic assignments of the separated closure of the result.

Definition 9.7 *Let* $\mathfrak{L} = U, P, \Pi$ *be a Läuchli model. The* **lifted** *model* $\mathfrak{L}^* = \langle U^*, P^*, \Pi^* \rangle$ *is given by the following data:*

1. $U^* = \{0,1\} \times U$.
2. $P^*(\alpha) = \{0,1\} \times P(\alpha)$
3. $\Pi^* = \langle \sigma, \overline{\Pi} \rangle$, *that is to say, the subgroup of* $S_{(U^*)}$ *generated by* σ *and* $\overline{\Pi} = \{\bar{\pi} : \pi \in \Pi\}$, *where*
   - $\sigma(i,x) = (1-i, x)$ *for* $i = 0, 1$.
   - $\bar{\pi}(i,x) = (i, \pi x)$.

The group in the lifted model is the result of augmenting the original group by a single permutation σ of order 2 which guarantees that no $P(\alpha)$ has a global fixed point (that is to say a point fixed by every permutation in Π^*). σ commutes with all $\bar{\pi}$, so all elements of the augmented group are of the form $\sigma\bar{\pi}$ or just $\bar{\pi}$, for some π in Π. If Π is abelian, so is Π^* Notice that the subgroup lattice of the augmented group consists of two planes, a lower one consisting of σ-**groups**: subgroups containing elements of the form $\sigma\bar{\pi}$ and an upper one consisting of subgroups of the embedded copy $\overline{\Pi}$ of Π. Now we transform this

[2] A logomorphism $\Re$ induces a bijection between the up-ideal (Alexandroff) topologies of the two models, associating with each open $U \subset K_1$, the set $\{z \in K_2 | \exists_x \in K_1\ x \Re z\} \subset K_2$, and conversely. This in turn induces an open geometric morphism between the corresponding presheaf topoi, as the categorically minded reader can check. Lemma 9.6 follows readily.

into the structure we are seeking by taking the separated closure as underlying group action, but carefully redefining atomic assignments so as to produce a model whose associated Kripke model is equivalent to $\mathcal{K}(\mathfrak{L})$ in the sense of lemma 9.6.

Definition 9.8 *Let $\mathfrak{L} = \langle U, P, \Pi\rangle$, be a Läuchli model, Γ an atomic subtheory of $\mathsf{Th}(\mathfrak{L})$, and $\mathfrak{L}^* = \langle U^*, P^*, \Pi^*\rangle$ the lifted model. For each atomic proposition α we have, as before,*

$$[\![\alpha]\!] = \{H < \Pi\colon Fix(H) \cap P(\alpha) \neq \varnothing\}$$

*and we define its image $\overline{[\![\alpha]\!]}$ to be the set of corresponding subgroups $\overline{H}$ in $\overline{\Pi}$, and its *-image $[\![\alpha]\!]^*$ to be the set of subgroups*

$$\{H < \Pi^*\colon \overline{H} \cap \overline{\Pi} \in \overline{[\![\alpha]\!]} \text{ and } H \neq \Pi^*\}.$$

That is to say, $[\![\alpha]\!]^*$ consists of all **proper** *subgroups of Π^* whose image under the map*

$$X \mapsto X \cap \overline{\Pi}$$

lies in (the copy in $\overline{\Pi}$ of) the original $[\![\alpha]\!]$.

The **Γ-scone** *$\mathfrak{L}^\Gamma$ of $\mathfrak{L}$ is the model $\langle \bigcup SQ(\mathcal{K}(\Pi^*)), P_s, \Pi^*\rangle$ with group action the separated closure of Π^* and where*

$$P_s(\alpha) = \begin{cases} \bigcup SQ([\![\alpha]\!]^*) & \text{if } \alpha \notin \Gamma \\ \bigcup SQ(\mathcal{K}(\Pi^*)) & \text{otherwise.} \end{cases}$$

Theorem 9.9 *Let $\mathfrak{L} = \langle U, P, \Pi\rangle$ be an abelian Läuchli model, $\mathcal{K}(\mathfrak{L})$ its associated Kripke model. Then the scones of each are elementarily equivalent, i.e.*

$$\mathsf{Th}_{\mathrm{L}}(\mathfrak{L}^\Gamma) \equiv \mathsf{Th}_{\mathrm{K}}(\mathcal{K}(\mathfrak{L})^\Gamma).$$

Proof: We show that there is a logomorphism between $\mathcal{K}(\mathfrak{L}^\Gamma)$ and $\mathcal{K}(\mathfrak{L})^\Gamma$ satisfying the premiss of lemma 9.6. Let K_1 be the underlying partial order of $\mathcal{K}(\mathfrak{L}^\Gamma)$, namely the lattice of subgroups of Π^*, and K_2 the underlying partial order of $\mathcal{K}(\mathfrak{L})^\Gamma)$.

Observe that K_1 is the union of the poset $\{H < \langle\sigma, \overline{\Pi}\rangle\colon \sigma \in H\}$, of σ-groups and the lattice of σ-free groups $\mathcal{K}(\overline{\Pi})$. Thus every subgroup W in K_1 is of the form $\overline{\mathrm{H}}$ or $\langle\sigma, \overline{\mathrm{H}}\rangle$ for some $H \in \mathcal{K}(\mathfrak{L})$.

Also observe that K_2 is the union $\{\mathfrak{r}_0\} \cup \mathcal{K}(\Pi)$ (see def. 7.2).

Now Define the relation $\Re$ on $K_1 \times K_2$ by

$$\Re = \{(\overline{\mathrm{H}}, H)\colon H < \Pi\} \cup \{(\langle\sigma, \overline{\mathrm{H}}\rangle, H)\colon H < \Pi \text{ and } H \neq \Pi\} \cup \{(\langle\sigma, \overline{\Pi}\rangle, \mathfrak{r}_0)\}$$

By the preceding remarks $\Re$ covers $K_1 \times K_2$. By the definition of the atomic assignment P for the Γ-scone, $\Re$-related worlds have the same atomic assignment. The reader can check that the third condition of definition 9.5 is also satisfied. Thus, by 9.6, $\Re$-related nodes have the same theory, and the two models are elementarily equivalent.

Theorem 9.9 $\square$

9.3 Gluing

Definition 9.10 *Let $\mathfrak{L}_1 = \langle U_1, P_1, \Pi_1\rangle$ and $\mathfrak{L}_2 = \langle U_2, P_2, \Pi_2\rangle$ be Läuchli models, and $U_1 \oplus U_2$ their disjoint union $\{1\} \times U_1 \cup \{2\} \times U_2$. We then define the following groups of permutations of $U_1 \oplus U_2$*

- $\overline{\Pi}_i = \{\bar{\pi}: \pi \in \Pi_i\}$ *where* $\bar{\pi}(i,x) = (i, \pi x)$ *and* $\bar{\pi}(2-i,x) = (2-i,x)$ *if* $\bar{\pi} \in \overline{\Pi}_i \quad (i = 1,2)$
- $\overline{\Pi}_1\overline{\Pi}_2 = \{\bar{\pi}_1\bar{\pi}_2: \bar{\pi}_1 \in \overline{\Pi}_1$ *and* $\bar{\pi}_2 \in \overline{\Pi}_2\}$

Define a subgroup W of $\overline{\Pi}_1\overline{\Pi}_2$ to be **pure** *if $W < \overline{\Pi}_i$ for $i = 1$ or 2. Call a subgroup W* **clean** *if W is of the form $\overline{\Pi}_1 H$ with $H < \overline{\Pi}_2$ or $H\overline{\Pi}_2$ with $H < \overline{\Pi}_1$, and unclean otherwise.*

We are now ready to define the Γ-join (or gluing of $\mathfrak{L}_1$ and $\mathfrak{L}_2$ along Γ). The notation is from the preceding definition.

Definition 9.11 *Let $\mathfrak{L}_1$ and $\mathfrak{L}_2$ be two Läuchli models and let Γ be an atomic subtheory of $\mathsf{Th}_{\mathrm{L}}(\mathfrak{L}_1) \cap \mathsf{Th}_{\mathrm{L}}(\mathfrak{L}_2)$. Let $\mathfrak{L}_1^\Gamma = \langle U_1, P_1, \Pi_1\rangle$ and $\mathfrak{L}_2^\Gamma = \langle U_2, P_2, \Pi_2\rangle$ be the Γ-scones, respectively, of the two models. For each propositional letter α let*

$$[\![\alpha]\!]^* = \{\overline{H\Pi_2} < \overline{\Pi}_1\overline{\Pi}_2 : H \in [\![\alpha]\!]_1\} \cup \{\overline{\Pi}_1\overline{H} < \overline{\Pi}_1\overline{\Pi}_2 : H \in [\![\alpha]\!]_2\}$$

where

$$[\![\alpha]\!]_i = \{H < \Pi_i : Fix(H) \cap P_i(\alpha) \neq \varnothing\}$$

and let

$$\Upsilon = \{W < \overline{\Pi}_1\overline{\Pi}_2 : W \text{ unclean }\}$$

Then $\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2 = \langle\widehat{\mathfrak{U}}, \widehat{P}, \widehat{\Pi}\rangle$ where

1. $\widehat{\mathfrak{U}} = \bigcup SQ(\mathcal{K}(\overline{\Pi}_1\overline{\Pi}_2))$,
2. $\widehat{\Pi} = (\overline{\Pi}_1\overline{\Pi}_2)$ *and*
3. $\widehat{P}(\alpha) = \bigcup SQ([\![\alpha]\!]^* \cup \Upsilon)$.

In particular, $\langle\widehat{\mathfrak{U}}, \widehat{\Pi}\rangle$ is the separated closure of $\overline{\Pi}_1\overline{\Pi}_2$.

Observe that since the constituent models $\mathfrak{L}_i$ of the join were already sconed, the atomic theory of the join will be Γ.

Lemma 9.12 *Let $\mathfrak{L}_1, \mathfrak{L}_2, \widehat{\Pi}$ and $\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2$ be as in the preceding definition, with $\mathfrak{L}_1, \mathfrak{L}_2$ abelian, and with $H < \Pi_1$ $G < \Pi_2$ and Υ the poset of unclean subgroups of Π^*. Then*

1. $\mathsf{Th}(\overline{H\Pi_2}) = \mathsf{Th}(H)$ *and* $\mathsf{Th}(\overline{\Pi}_1\bar{G}) = \mathsf{Th}(G)$.
2. *For* $W \in \Upsilon, \mathsf{Th}(W) = \mathsf{PF}$, the set of all propositions.
3. $Fix(\widehat{\Pi}) \cap \widehat{P}(\alpha) \neq \varnothing \Leftrightarrow \alpha \in \Gamma$.

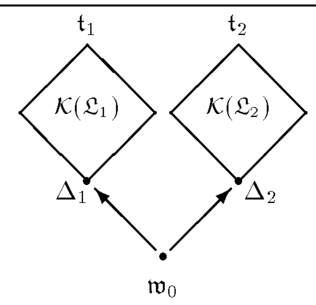

Figure 14: The Kripke model $\mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$

Proof: Easy and left to the reader.

Lemma 9.12 □

We now come to the main result of the section.

Theorem 9.13 *Let $\mathfrak{L}_1$ and $\mathfrak{L}_2$ be well-ordered abelian Läuchli models, and Γ be an atomic subtheory of both models. Then*

$$\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2 \equiv \mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$$

that is to say, the theory of their join over Γ is the same as that of the Kripke-join of their corresponding Kripke models.

Proof: Suppose $\mathfrak{L}_1 = \langle V_1, Q_1, \Delta_1 \rangle$, $\mathfrak{L}_2 = \langle V_2, Q_2, \Delta_2 \rangle$, and that the corresponding sconed models are ${\mathfrak{L}_1}^\Gamma = \langle U_1, P_1, \Pi_1 \rangle$ and ${\mathfrak{L}_2}^\Gamma = \langle U_2, P_2, \Pi_2 \rangle$, where the Π_i are generated by the Δ_i and a permutation of order 2, as in definitions 9.7 and 9.8.

Observe that the embedded sub-models $\mathcal{K}(\mathfrak{L}_1)$ and $\mathcal{K}(\mathfrak{L}_2)$ of the join $\mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$ have root nodes Δ_1 and Δ_2 and top nodes $\mathfrak{t}_i$ corresponding to the trivial subgroups of the permutation groups Δ_i, as in figure 14.

We now construct a logomorphism $\Re$ between $\mathcal{K}(\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2)$ and the Kripke join $\mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$. Using the terminology from definition 9.11, we put

$$\Re = \{(\widehat{\Pi}, \mathfrak{w}_0)\} \cup \{(\overline{\mathrm{H}\Pi}_2, H) \colon H < \Delta_1\} \cup \{(\overline{\Pi}_1 \overline{\mathrm{H}}, H) \colon H < \Delta_2\} \cup \Upsilon \times \{\mathfrak{t}_1, \mathfrak{t}_2\}$$

Now observe that every subgroup in $\mathcal{K}(\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2)$ is either in Υ or of the form $\overline{\mathrm{H}\Pi}_2$ for $H < \Delta_1$ or $\overline{\Pi}_1 \overline{\mathrm{H}}$ for some $H < \Delta_2$. Thus, by lemma 9.12, $\Re$ is a logomorphism. Hence,

$$\mathcal{K}(\mathfrak{L}_1 \oplus_\Gamma \mathfrak{L}_2) \equiv \mathcal{K}(\mathfrak{L}_1) \oplus_\Gamma \mathcal{K}(\mathfrak{L}_2)$$

Theorem 9.13 □

10 Completeness With Respect to Läuchli Semantics

The key result about Läuchli models is the completeness of the standard formal systems of constructive proof. Läuchli proved completeness for the first-order predicate calculus [Läu70], but his proof takes an unnecessary digression through Kripke models and number theory, and obscures the constructive content of the result. We would like to see a proof of completeness that constructs a formal constructive proof directly from an arbitrary invariant function of appropriate type, but more study of the structure of all invariants of a given type seems to be required for such a construction. Instead, we take Fitting's proof of completeness with respect to Kripke semantics [Fit69], and adapt it to an explicit direct construction of Läuchli models. We try to clear up some confusing ambiguities in Fitting's presentation along the way, and we use the redundant rule ($\Rightarrow$RW) to simplify the termination criterion for his procedure, in particular avoiding the need for loop detection.

Before proceeding, we remark that we can obtain completeless immediately just from the closure of Läuchli semantics under sconing and gluing, and the fact that e.g. the Jaskowsky sequence of Kripke models, definable from one-world models just be use of the scone and join, is complete for the Heyting calculus. The Jaskowsky sequence, as well as the relevance of scone-and-join closure for completeness of a class of models, is discussed at length in Smorynski's [Smo73].

The basic idea behind Fitting's proof is the same as that behind proofs of completeness for the classical propositional calculus—define a procedure using rules of inference to transform sequents into forms where they are either directly derivable, or directly refutable by a countermodel.

Definition 10.1 *A Läuchli model $\mathfrak{M}$ is a* countermodel *for a sequent $\Gamma \vdash \Psi$ if and only if $\mathfrak{M} \models_{\mathrm{L}} \alpha$ for every $\alpha \in \Gamma$ and $\mathfrak{M} \not\models_{\mathrm{L}} \beta$ for every $\beta \in \Psi$.*

In the case of constructive sequents, *directly derivable* means that the basis rule (B) of Theorem 2.8 applies, in which case we say that the sequent is *closed.*

Definition 10.2 *A sequent $\Gamma \vdash \Psi$ is* closed *if and only if $\Gamma \cap \Psi \neq \emptyset$, it is* open *otherwise.*

In order to refute a sequent $\Gamma \vdash \Psi$ by a countermodel, we must generate enough information to know how to treat the atomic formulae, in order to make formulae in Γ true and those in Ψ false. As with classical logic, it is easy to do so when a sequent is *saturated* in such a way that every nonatomic formula is supported by one or both of its principal subformulae. For constructive logic, the ($\Rightarrow$RS) rule complicates the analysis, because in the process of following rules backward from a desired conclusion toward saturated sequents, this rule throws away formulae from the right-hand side of the sequent, potentially cancelling progress toward saturation. Applications of ($\Rightarrow$RS) are *choice points*, in the sense that we must choose the correct implication to process in the right-hand side of a sequent, in order to arrive at a proof or a set of saturated sequents. Because of this complication, Fitting considers *semisaturated* sequents, resulting from all possible backward applications of the choice-free rules, and the *associated* results of single backwards applications of ($\Rightarrow$RS) to semisaturated sets as well. Our addition of the redundant rule ($\Rightarrow$RW) simplifies the analysis considerably, by tightening the correspondence between semisaturation and the results of applying all choiceless rules.

Definition 10.3 *A sequent* $\Gamma \vdash \Psi$ *is* semisaturated *if and only if, for all formulae* α *and* β

1. $\alpha \wedge \beta \in \Gamma$ *implies* $\alpha \in \Gamma$ *and* $\beta \in \Gamma$*;*
2. $\alpha \wedge \beta \in \Psi$ *implies* $\alpha \in \Psi$ *or* $\beta \in \Psi$*;*
3. $\alpha \vee \beta \in \Gamma$ *implies* $\alpha \in \Gamma$ *or* $\beta \in \Gamma$*;*
4. $\alpha \vee \beta \in \Psi$ *implies* $\alpha \in \Psi$ *and* $\beta \in \Psi$*;*
5. $\alpha \Rightarrow \beta \in \Gamma$ *implies* $\alpha \in \Psi$ *or* $\beta \in \Gamma$*;*
6. $\alpha \Rightarrow \beta \in \Psi$ *implies* $\beta \in \Psi$*.*

A sequent $\Gamma \vdash \Psi$ *is* saturated *if and only if it is semisaturated and, in addition,*

7. $\alpha \Rightarrow \beta \in \Psi$ *implies* $\alpha \in \Gamma$*.*

Definition 10.4 *A sequent* $\Gamma' \vdash \Psi'$ *is a* semisaturation *of* $\Gamma \vdash \Psi$ *if and only if* $\Gamma \subseteq \Gamma'$, $\Psi \subseteq \Psi'$, $\Gamma' \vdash \Psi'$ *is semisaturated, and for all* Γ'' *and* Ψ'' *with* $\Gamma \subseteq \Gamma'' \subseteq \Gamma'$ *and* $\Psi \subseteq \Psi'' \subseteq \Psi'$ *such that* $\Gamma'' \vdash \Psi''$ *is semisaturated,* $\Gamma'' = \Gamma'$ *and* $\Psi'' = \Psi'$*. That is, a semisaturation is a (not usually unique) minimal extension of a sequent to semisaturated form.*

Definition 10.5 *Let* $\Gamma \vdash \alpha \Rightarrow \beta, \Psi$ *be semisaturated, with* $\alpha \notin \Gamma$*. Then* $\Gamma, \alpha \vdash \beta$ *is an* associate *of* $\Gamma \vdash \alpha \Rightarrow \beta, \Psi$*.*

The final goal of our sequent transforming procedure is a saturated sequent, which must be either closed and derivable, or open and refutable by a countermodel. The countermodel for an open saturated sequent is an especially simple Läuchli model with trivial subgroup structure, corresponding to a one-world Kripke model, or equivalently to a classical model.

Lemma 10.6 *If* $\Gamma \vdash \Psi$ *is closed, then it is derivable.*

Proof: Trivial from the basis rule (B) of Theorem 2.8.

Lemma 10.6 $\square$

Lemma 10.7 *If* $\Gamma \vdash \Psi$ *is saturated and open, then there is a countermodel* $\mathfrak{M}$ *for* $\Gamma \vdash \Psi$*.*

Proof: Let $\mathbf{U} = \{0,1,2\}$, $P(\alpha) = \{0,1,2\}$ for all atomic formulae $\alpha \in \Gamma$, $P(\alpha) = \{1,2\}$ for all atomic formulae $\alpha \notin \Gamma$, and let $\pi(0) = 0$, $\pi(1) = 2$, $\pi(2) = 1$. Let $\mathfrak{M} = \langle \mathbf{U}, p, \{\pi, e\}\rangle$. By an elementary induction on the structure of γ, using saturation at each step, $\gamma \in \Gamma$ implies $\mathfrak{M} \models_{\mathrm{L}} \gamma$, and $\gamma \in \Psi$ implies $\mathfrak{M} \not\models_{\mathrm{L}} \gamma$.

Lemma 10.7 □

Given a sequent $\Gamma \vdash \Psi$ that is not semisaturated, we must consider all possible ways of semisaturating it, i.e., all possible ways of applying the sequent rules other than ($\Rightarrow$RS) to obtain it. If all semisaturations are derivable, then we get a proof of $\Gamma \vdash \Psi$; if any one of them has a countermodel it serves as a countermodel for $\Gamma \vdash \Psi$ as well.

Lemma 10.8 *If all semisaturations of $\Gamma \vdash \Psi$ are derivable, then $\Gamma \vdash \Psi$ is derivable.*

Proof: The semisaturations of $\Gamma \vdash \Psi$ are precisely the results of backwards derivation from $\Gamma \vdash \Psi$ using all of the rules of Theorem 2.8 except ($\Rightarrow$RS).

Lemma 10.8 □

Lemma 10.9 *If $\Gamma' \vdash \Psi'$ is a semisaturation of $\Gamma \vdash \Psi$, and $\mathfrak{M}$ is a countermodel for $\Gamma' \vdash \Psi'$, then $\mathfrak{M}$ is a countermodel for $\Gamma \vdash \Psi$.*

Proof: Direct, since $\Gamma \subseteq \Gamma'$ and $\Psi \subseteq \Psi'$.

Lemma 10.9 □

Given a semisaturated but not saturated sequent $\Gamma \vdash \Psi$, we must consider all of its associates. In this case, if some associate is derivable, we get a proof of $\Gamma \vdash \Psi$, but if *all* associates have countermodels, we can combine them into a countermodel for $\Gamma \vdash \Psi$.

Lemma 10.10 *If some associate of $\Gamma \vdash \Psi$ is derivable, then $\Gamma \vdash \Psi$ is derivable.*

Proof: Direct, by the ($\Rightarrow$RS) rule.

Lemma 10.10 □

Lemma 10.11 *If $\Gamma \vdash \Psi$ is open and semisaturated, and if every associate $\Gamma' \vdash \Psi'$ has a finite countermodel, then $\Gamma \vdash \Psi$ has a finite countermodel.*

Proof: Let $\Psi_0 = \{\alpha_1 \Rightarrow \beta_1, \ldots, \alpha_n \Rightarrow \beta_n\} = \{\alpha \Rightarrow \beta \in \Psi \mid \alpha \notin \Gamma\}$ The associates of $\Gamma \vdash \Psi$ are precisely the sequents $\Gamma, \alpha \vdash \beta$, where $\alpha \Rightarrow \beta \in \Psi_0$. Let $\Gamma_{\mathrm{at}} = \{\gamma \in \Gamma \mid \gamma \text{ is atomic}\}$

Let $\mathfrak{M}_i$ be a finite countermodel for $\Gamma, \alpha_i \vdash \beta_i$. Let $\mathfrak{M} = \oplus_{\Gamma_{\mathrm{at}}} \mathfrak{M}_1, \ldots, \mathfrak{M}_n$ (see Definition 9.11). By theorem 9.13 and induction on the structure of a formula γ, $\gamma \in \Gamma$ implies $\mathfrak{M} \models_{\mathrm{L}} \gamma$, and $\gamma \in \Psi$ implies $\mathfrak{M} \not\models_{\mathrm{L}} \gamma$.

Basis: If γ is an atomic element of Γ, then $\mathfrak{M} \models_{\mathrm{L}} \gamma$ by definition of $\mathfrak{M}$. If γ is an atomic element of Ψ, then $\gamma \notin \Gamma$ by the openness of $\Gamma \vdash \Psi$, and so $\mathfrak{M} \not\models_{\mathrm{L}} \gamma$ by definition.

Induction: The cases ($\vee$L), ($\vee$R), ($\wedge$L), and ($\wedge$R) follow by arguments similar to Lemma 10.7. The two interesting cases are ($\Rightarrow$L) and ($\Rightarrow$R).

($\Rightarrow$L) Let $\gamma \equiv \alpha \Rightarrow \beta \in \Gamma$. Now, either $\alpha \in \Psi$ or $\beta \in \Gamma$ by semisaturation.
If $\alpha \in \Psi$, then $\mathfrak{M} \not\models_{\mathrm{L}} \alpha$ by our inductive hypothesis. Furthermore, each $\mathfrak{M}_i \models_{\mathrm{L}} \alpha \Rightarrow \beta$, and so $\mathfrak{M} \models_{\mathrm{L}} \alpha \Rightarrow \beta$ by Theorem 6.18. If $\beta \in \Gamma$, then $\mathfrak{M} \models_{\mathrm{L}} \beta$ by our inductive hypothesis, and so $\mathfrak{M} \models_{\mathrm{L}} \alpha \Rightarrow \beta$ trivially.

($\Rightarrow$R) Let $\gamma \equiv \alpha \Rightarrow \beta \in \Psi$. We have $\beta \in \Psi$ by semisaturation. If $\alpha \in \Gamma$, then $\mathfrak{M} \models_{\mathrm{L}} \alpha$ and $\mathfrak{M} \not\models_{\mathrm{L}} \beta$, so $\mathfrak{M} \not\models_{\mathrm{L}} \alpha \Rightarrow \beta$. If $\alpha \notin \Gamma$, then α is α_i and β is β_i for some $\alpha_i \Rightarrow \beta_i \in \Psi_0$. By definition of $\mathfrak{M}_i$, $\mathfrak{M}_i \not\models_{\mathrm{L}} \alpha \Rightarrow \beta$, so by the properties if the join (theorem 9.13), $\mathfrak{M} \not\models_{\mathrm{L}} \alpha \Rightarrow \beta$.

Lemma 10.11 □

We are now in a position to establish the Completeness Theorem:

Theorem 10.12 (Completeness Theorem) *For all finite sets of propositional formulae Γ, Ψ with $\Psi \neq \emptyset$, either $\Gamma \vdash \Psi$ or there is a finite Läuchli model $\mathfrak{M}$ such that $\mathfrak{M} \models_{\mathrm{L}} \alpha$ for all $\alpha \in \Gamma$ and $\mathfrak{M} \not\models_{\mathrm{L}} \beta$ for all $\beta \in \Psi$.*

Proof. Let Σ be the set of all subformulae of formulae in $\Gamma \bigcup \Psi$. Our proof is by induction on $|\Sigma - \Gamma|$.

Basis: If $|\Sigma - \Gamma| = 0$, then $\Gamma = \Sigma$, and so $\Psi \subseteq \Gamma$. Therefore $\Gamma \vdash \Psi$ is closed, and by Lemma 10.6, derivable.

Induction: There are three possibilities, according to whether $\Gamma \vdash \Psi$ is saturated, not saturated but semisaturated, or not semisaturated.

$\Gamma \vdash \Psi$ is saturated. By Lemmas 10.6 and 10.7, $\Gamma \vdash \Psi$ is either closed and derivable, or open and has a countermodel.

$\Gamma \vdash \Psi$ is semisaturated, but not saturated. Consider its associates,

$$\Gamma_0 \vdash \Psi_0, \ldots, \Gamma_n \vdash \Psi_n.$$

There are only finitely many such associates (since Ψ is finite), and there is at least one (since $\Gamma \vdash \Psi$ is not saturated). We have $\Gamma \subset \Gamma_i$ for each i (in fact, $|\Gamma_i| = |\Gamma| + 1$), because each associate adds precisely one new formula to the left-hand side of a sequent.

The set Σ_i of subformulae of $\Gamma_i \bigcup \Psi_i$ is a subset of Σ, and we have $|\Sigma_i - \Gamma_i| < |\Sigma - \Gamma|$, and the inductive hypothesis applies to each $\Gamma_i \vdash \Psi_i$. If some $\Gamma_i \vdash \Psi_i$ is derivable, then by Lemma 10.10, $\Gamma \vdash \Psi$ is derivable. If no $\Gamma_i \vdash \Psi_i$ is derivable, then by our inductive hypothesis there is a finite countermodel to each $\Gamma_i \vdash \Psi_i$, and so by Lemma 10.11 there is a finite countermodel to $\Gamma \vdash \Psi$.

$\Gamma \vdash \Psi$ is not semisaturated. If $\Gamma \vdash \Psi$ is not semisaturated, let $\Gamma_0 \vdash \Psi_0$, ..., $\Gamma_m \vdash \Psi_m$ be its (finitely many, but at least one) semisaturations. $\Gamma \subseteq \Gamma_i$, and $\Sigma_i \subseteq \Sigma$, where Σ_i is the set of subformulae of formulae in $\Gamma_i \bigcup \Psi_i$. So, one of the previous cases applies to each sequent $\Gamma_i \vdash \Psi_i$. Now, either all of the sequents $\Gamma_i \vdash \Psi_i$ are derivable, in which case by Lemma 10.8, $\Gamma \vdash \Psi$ is derivable; or some $\Gamma_i \vdash \Psi_i$ is refutable, in which case Lemma 10.9 yields a countermodel for $\Gamma \vdash \Psi$, as desired.

Theorem 10.12 □

In effect the proof of the completeness theorem provides a procedure that takes as input a finite sequent $\Gamma \vdash \Psi$, and builds a tree of attempted derivations of $\Gamma \vdash \Psi$. The search for a derivation works backwards, by taking sequents from which $\Gamma \vdash \Psi$ has been derived, and adding to the beginning of a derivation more saturated sequents from which the later ones may be derived. In some cases, the search succeeds, by reducing all sequents to closed form, producing a formal derivation of $\Gamma \vdash \Psi$, which, by Theorem 2.8 may be converted to a proof formula a such that $\Gamma \vdash a{:}\psi$ for some $\psi \in \Psi$. When the search fails to find a derivation, it reduces all hypothetical sequents to open saturated form. It is trivial, by Lemma 10.7 to provide countermodels for these open saturated sequents, and Lemmas 10.9 and 10.11 show how to construct a countermodel to $\Gamma \vdash \Psi$ from the countermodels of the open saturated sequents, and the structure of the search for a derivation. In the process we have shown that validity for the Heyting calculus is decidable[3].

11 Negation

Most conventional treatments of the propositional calculus, whether classical or constructive, include the unary propositional operator *negation*, which we write as $\neg$. Technically, it is not particularly hard to provide realizability semantics for conventional constructive negation, using Läuchli models. The results look rather peculiar intuitively, however, and we do not consider them to be a satisfactory explanation of negation.

Realizability semantics for negation seem clearest when presented through the reduction of negation to implication and falsehood. We introduce a special atomic propositional symbol $\perp$ to denote *falsehood*, in the sense of *absurdity*, *inconsistency*, or *contradiction*. Then, let $\neg\alpha$ be an abbreviation for $\alpha \Rightarrow \perp$.

[3] A well known result. In fact Statman has shown validity is PSPACE-complete, see e.g. [Ner91, Fit69, Tv88]

Now, we need only extend each Läuchli model $\mathfrak{M} = \langle \mathbf{U}, P, \Pi \rangle$ to include the new propositional symbol, by providing values for $P(\bot)$.

The obvious candidate for $P(\bot)$ is the empty set, since we never intend to consider alleged constructions of inconsistent falsehoods. It is not even clear what it should mean for something to constitute evidence for falsehood. But, $P(\bot) = \emptyset$ does *not* yield the conventional theory of negation formalized in the Heyting Calculus. Notice that the empty interpretation of falsehood implies that, whenever $P(\alpha) \neq \emptyset$, then $P(\neg\alpha) = \emptyset$, since there can be no function from a nonempty set to an empty one. Since there is a *unique* function from $\emptyset$ to any other set, that function cannot be permuted to anything else, and so is invariant. So, $\mathfrak{M} \models_{\mathrm{L}} \neg\alpha$ if and only if $P(\alpha) = \emptyset$. When $P(\alpha) \neq \emptyset$, then $P(\neg\alpha) = \emptyset$, so $\mathfrak{M} \models_{\mathrm{L}} \neg\neg\alpha$. Thus, we cannot construct a model with empty interpretation of falsehood that simultaneously invalidates $\neg\alpha$ and $\neg\neg\alpha$, nor one that invalidates $(\neg\alpha) \vee (\neg\neg\alpha)$.

A technically correct solution is to allow each model to set $P(\bot)$ arbitrarily, as long as there is an invariant function from $P(\bot)$ to $P(\alpha)$ for each atomic formula α. It is easy to see, by induction on formula structure, that this property must hold for nonatomics as well. This solution seems no better than a thinly disguised introduction of $\bot \Rightarrow \alpha$ as an axiom schema, and we do not find any useful *semantic* intuition about falsehood there. A variant that is not quite so transparent an encoding of the axiom schema above is to require that $P(\bot) \subseteq P(\alpha)$ for each atomic formula α, so that the required invariant function is simply a restricted identity function. While the latter solution does not beg the question of *why* falsehood should imply everything so transparently as the former, it does not appear to answer the question either. Since the containment of $P(\bot)$ as a subset does not pull up inductively to evidence for nonatomic formulae (the presence of an invariant function does, but on nonatomic domains the function is no longer a restricted identity), this solution makes a counterintuitive distinction between atomic formulae and nonatomics.

We do not know of an intuitively satisfying semantic treatment of conventional negation in the Heyting Calculus. We suspect that the formal treatment of negation in the Heyting Calculus is, in fact, *not* well-founded on the sort of constructive intuition that is captured by Läuchli semantics. We hope to see investigation of other notions of negation with better intuitive explanations.

12 Realizability Semantics and Completeness for the Predicate Calculus

All of the realizability concepts discussed here for the propositional calculus extend in well-known ways to the predicate calculus. Läuchli's proof of completeness [Läu70] is, in fact, for the predicate calculus. Rather than present in detail the definitions and theorems appropriate to realizability semantics for the predicate calculus, we merely sketch the ways in which the constructive quality of the completeness proof deteriorates. First, validity for the predicate calculus is not decidable, so there is no hope of proving that either $\Gamma \vdash \Psi$ or there is a countermodel $\mathfrak{M}$ with $\mathfrak{M} \models_{\mathrm{L}} \Gamma$ and $\mathfrak{M} \not\models_{\mathrm{L}} \Psi$. The best approximation that we can hope for is that $\Gamma \models_{\mathrm{L}} \Psi$ implies $\Gamma \vdash \Psi$ (completeness) and the converse (soundness). The weaker implicative form is just as good as the stronger disjunctive form as a justification of the appropriateness of a formal calculus to a semantic analysis, so this is not a serious problem to the intuition.

Now, consider the procedure that searches for a derivation of $\Gamma \vdash \Psi$, and produces a countermodel when it fails to find a derivation. In the propositional calculus, that procedure always terminates, and produces either a derivation or a *finite* countermodel. In the predicate calculus, the procedure may succeed, or fail after a finite time, or it may fail by searching infinitely. In the second case, there is no derivation, and we may extract from the trace of the procedure an infinite model. It is not difficult to make the definition of this infinite model perfectly constructive, but the proof that the model satisfies Γ but not Ψ is not constructive. The problem arises in the attempt to generalize the proof of Lemma 10.11, in which we joined together the countermodels for all the associates of a sequent $\Gamma \vdash \Psi$ to get a countermodel for $\Gamma \vdash \Psi$. In order to guarantee that the joined model satisfies an implication $(\alpha \Rightarrow \beta) \in \Gamma$, we appealed to Lemma 6.18 and thence to Lemma 6.17 to construct an invariant function f in $P(\alpha \Rightarrow \beta)$ from a function g in the same domain mapping each member of $P(\alpha)$ to a member of $P(\beta)$ that is at least as stable under permutation. In the construction of f, Lemma 6.17 must pick a canonical element from each orbit in $P(\alpha)$, since the value of f on a determines the value of f on everything in the orbit of a. Such a choice of canonical element was justified in the propositional calculus proof, because all models constructed were finite, and therefore well-ordered. When joining infinite models, as we do in the model construction for an underivable sequent in the predicate calculus, there is no apparent way to choose the canonical orbit elements.

So, Läuchli's proof of completeness for the predicate calculus, and all variations that we know currently, use the axiom of choice, which is repugnant to a constructive intuition. We expect that there is a variation of the technical definition of Läuchli model that allows a constructive proof of completeness—such an improvement is roughly analogous to the fallible Kripke, and fallible Beth-style approaches of Veldman, Friedman, de Swart and other (see e.g. [Tv88, Tro77, Lip92, Lip]) Our reason for optimism is that the objects in Läuchli models are consciously intended to include more than the truly constructible objects. Therefore, it should be possible to replace the proper constructive concept of function used in the definition of $P(\alpha \Rightarrow \beta)$ by some sort of protofunctional concept, satisfying an axiom of choice. A constructive Läuchli-style semantics based on a relational interpretation of $P(\alpha \Rightarrow \beta)$ is sketched in [LO84]). For any new proposal along these lines, it will be necessary to re-examine the intuitive argument that formal realizability semantics is more liberal than intuitive semantics under the new definition of $P(\alpha)$, of course.

13 Categorical Models

There has been important work by Makkai and Harnik in recent years to clarify the deeper mathematical structure of Läuchli semantics via a categorical analysis

[HM92, HM93a, HM93b]. Blute and Scott have developed a categorical Läuchli semantics (and a strong completeness theorem) for Linear Logic [BS96].

We defer a categorical treatment of the constructions described above to a subsequent paper, as our main points here do not require it. However a few comparative remarks are in order. In [HM92] the Läuchli completeness theorem is strengthened to the following, (where connectionally closed means cartesian closed with binary coproducts).

Theorem 13.1 (Makkai-Harnik) *Let $\mathbb{C}$ be a countable free connectionally closed category. For each pair of objects A, B of $\mathbb{C}$ there is a representation $F_{A,B}: \mathbb{C} \longrightarrow \mathsf{Set}^{\mathbb{Z}}$ which is weakly full with respect to A, B, i.e., if $\mathbb{C}(A, B) = \emptyset$ then $\mathsf{Set}^{\mathbb{Z}}(FA, FB) = \emptyset$.*

In case $\mathbb{C}$ has the disjunction property, the authors exhibit a weakly full representation from $\mathbb{C}$ into $\mathsf{Set}^{\mathbb{Z}}$.

There are a number of significant differences between our treatment and theirs. Given a pair A, B where A is a set of premisses (or a conjunction thereof) and B a formula, Kripke countermodels $K_{A,B}$ are built via tableaux from the pair A, B. We then build a Läuchli model $L(K_{A,B})$ over a group $G_{A,B}$ whose structure is determined by A, B, via gluing and sconing, the map L commuting with these operations on both kinds of models. This operation induces a representation $F_{A,B}$ from the proof-theoretic category (the free connectionally closed category) $\mathbb{C}$ into $\mathsf{Set}^{G_{A,B}}$.

A fundamental gain in formulating Läuchli completeness in terms of a more general class of groups, instead of the single group $\mathbb{Z}$ is to expose the algebraic structure implicit in Läuchli's coding of Kripke structure into $\mathbb{Z}$. We also show the canonical role played by the modular groups used in this coding: they provide separated group actions that permit gluing of Läuchli models.

Nonetheless, a full clarification of the role played by gluing as an operator on Kripke and Läuchli models would seem to require a categorical treatment of more components of the structure of the completeness argument. Our constructions seem to suggest that the proper place for formalizing the "Läuchli representation theorem" along the lines developed in this paper is some subcategory of the category of topoi Set^{G}, with varying G. These questions will be taken up in a sequel to this paper.

14 Acknowledgements

This paper has been through an inordinate number of drafts since 1986. Early versions of the ideas in this paper were shaken down in a seminar by Michael O'Donnell at the University of Chicago. Frank Corley helped particularly by taking notes on the seminar, and participating in discussions. Around 1988–1989, Michael O'Donnell worked with Stuart Kurtz and John Mitchell on an early draft, which led to the outline published in [KMO92]. It has become impossible to determine precisely which parts of the paper are from the work with Kurtz and Mitchell, and which are later additions and revisions. Kurtz

and Mitchell have kindly given permission to use the draft of 1989, which was partly their work, as a basis for the current paper.

References

[Bee80] M. J. Beeson. *Foundations of Constructive Mathematics.* Springer-Verlag, 1980.

[Bet59] E. W. Beth. *The Foundations of Mathematics, A Study in the Philosophy of Science.* Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, 1959.

[BS96] R. Blute and P. J. Scott. Linear läuchli semantics. *Annals of Pure and Applied Logic*, 962, 1996.

[Dum77] M. A. E. Dummett. *Elements of Intuitionism.* Oxford University Press, 1977.

[Fef79] S. Feferman. Constructive theories of functions and classes. In M. Boffa, D. van Dalen, and K. McAloon, editors, *Logic Colloquium 78*, pages 159–224. North-Holland, 1979. Proceedings of the colloquium held in Mons, August 1978.

[Fit69] M. C. Fitting. *Intuitionistic Logic, Model Theory, and Forcing.* Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, London, 1969.

[Fit74] F. B. Fitch. *Elements of Combinatory Logic.* Yale University Press, 1974.

[FS90] P. J. Freyd and A. Scedrov. *Categories, Allegories.* North-Holland, Amsterdam, 1990.

[GLT89] J.-Y. Girard, Y. Lafont, and P. Taylor. *Proofs and Types.* Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1989.

[Hey30] A. Heyting. Die formalen Regeln der intuitionistischen Logik. *Sitzungsberichte der Preussischen Academie der Wissenschaften, Physikalisch-Matematische Klasse*, pages 42–56, 1930.

[HM92] V. Harnik and M. Makkai. Lambek's categorical proof theory and läuchli's abstract realizability. *The Journal of Symbolic Logic*, 57(1):200–230, 1992.

[HM93a] V. Harnik and M. Makkai. The fibrational formulation of intuitionistic predicate logic i: Completeness according to gödel, kripke and läuchli, part i. *Notre Dame Journal of Formal Logic*, 34(3):334–377, 1993.

[HM93b] V. Harnik and M. Makkai. The fibrational formulation of intuitionistic predicate logic i: Completeness according to gödel, kripke and läuchli, part ii. *Notre Dame Journal of Formal Logic*, 34(4):471–498, 1993.

[How80] W. A. Howard. The formulae-as-types notion of construction. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, pages 479–490. Academic Press, 1980.

[Kle45] S. C. Kleene. On the interpretation of intuitionistic number theory. *The Journal of Symbolic Logic*, 10(4):109–124, December 1945.

[Kle59] S. C. Kleene. Realizability. In A. Heyting, editor, *Constructivity in Mathematics*, pages 285–289. North-Holland, Amsterdam, 1959. Proceedings of the Colloquium Held in Amsterdam, August 26–31, 1957.

[KMO92] Stuart A. Kurtz, John Mitchell, and Michael J. O'Donnell. Connecting formal semantics to constructive intuitions. In *Proceedings of a Symposium on Constructivity in Computer Science, Trinity University, San Antonio TX, June 1991*, volume 613 of *Lecture Notes in Computer Science*, pages 1–21. Springer-Verlag, 1992.

[Kri63] S. A. Kripke. Semantical analysis of modal logic I: Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

[Kri65] S. A. Kripke. Semantical analysis of intuitionistic logic, I. In J. N. Crossley and M. A. E. Dummett, editors, *Formal Systems and Recursive Functions*, pages 92–130. North-Holland, Amsterdam, 1965. Proceedings of the Eighth Logic Colloquium, Oxford, July 1963.

[KV65] S. C. Kleene and R. E. Vesley. *The Foundations of Intuitionistic Mathematics, Especially in Relation to Recursive Functions.* Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, London, 1965.

[Läu70] H. Läuchli. An abstract notion of realizability for which intuitionistic predicate calculus is complete. In A. Kino, J. Myhill, and R. E. Vesley, editors, *Intuitionism and Proof Theory*, Studies in Logic and the Foundations of Mathematics, pages 277–234. North-Holland, Amsterdam, London, 1970. Proceedings of the Conference on Intuitionism and Proof Theory, Buffalo, New York, August 1968.

[Lip] J. Lipton. Type theory and inhabitation: A constructive completeness theorem for martin-löf type theory. Submitted to the *Annals of Pure and Applied Logic*.

[Lip92] James Lipton. Constructive kripke semantics and realizability. In Y. Moschovakis, editor, *Logic for Computer Science*, MSRI workshop proceedings. Springer-Verlag, 1992.

[LO84] James Lipton and Michael J. O'Donnell. Intuitive counterexamples for constructive fallacies. In Igor Prívara, Branislav Rovan, and Peter Ružička, editors, *Mathematical Foundations of Computer Science 1994 — 19th International Symposium, MFCS '94, Košice, Slovakia, August 1994 — Proceedings*, volume 841 of *Lecture Notes in Computer Science*, pages 87–111. Springer-Verlag, 1984.

[LS86] J. Lambek and P. J. Scott. *Introduction to Higher Order Categorical Logic*. Cambridge University Press, 1986.

[MS93] J. Mitchell and A. Scedrov. Notes on sconing and relators. In *CSL '92: 1992 Computer Science Logic Proceedings*, Lecture Notes in Computer Science. Springer, 1993.

[Ner91] A. Nerode. Some lecures on intuitionistic logic i. In *Logical Foundations of Computer Science, 1988 CIME Summer School, Montecatini, Italy*, Lecture Notes in Computer Science. Springer-Verlag, 1991.

[Odi89] P. Odifreddi. *Classical recursion Theory*. North-Holland, 1989.

[Pei85] C. S. Peirce. On the algebra of logic: A contribution to the philosophy of notation. *American Journal of Mathematics*, 8:180–202, 1885.

[Pra65] D. Prawitz. *Natural Deduction*. Almqvist & Wiksell, Stockholm, 1965.

[Ros53] G. F. Rose. Propositional calculus and realizability. *Transactions of the American Mathematical Society*, 75:1–19, July–September 1953.

[Smo73] C. A. Smorynski. Applications of kripke models. In *Mathematical Investigation of Intuitionistic Arithmetic and Analysis*, volume 344 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1973.

[SS82] A. Scedrov and P. Scott. A note on the friedman slash and the freyd covers. In A. Troesltra and D. van Dalen, editors, *The L. E. J. Brouwer Centenary Symposium*. North-Holland, 1982.

[Sta85] R. Statman. Logical relations and the typed lambda calculus. *Information and Control*, 65:85–97, 1985.

[Ste72] S. Stenlund. *Combinators, λ-terms, and Proof Theory*. D. Riedel, Dordrecht-Holland, 1972.

[Tro77] A. S. Troelstra. *Choice Sequences: a chapter of intuitionistic mathematics*. Oxford Logic Guides. Clarendon Press, Oxford, 1977.

[Tv88] A. S. Troelstra and D. van Dalen. *Constructivism in Mathematics: an Introduction, volume II*. Studies in Logic and the Foundations of Mathematics. North-Holland, 1988.

[vO90] J. van Oosten. *Exercises in Realizability*. PhD thesis, Univ. of Amsterdam, 1990.