WILFRIED SIEG
JOHN BYRNES

Normal Natural Deduction Proofs (in classical logic)*

Abstract. Natural deduction (for short: nd-) calculi have not been used systematically as a basis for automated theorem proving in classical logic. To remove objective obstacles to their use we describe (1) a method that allows to give *semantic proofs of normal form theorems for nd-calculi* and (2) a framework that allows to *search directly for normal nd-proofs*. Thus, one can try to answer the question: How do we bridge the gap between claims and assumptions in heuristically motivated ways? This informal question motivates the formulation of *intercalation calculi*. Ic-calculi are the technical underpinnings for (1) and (2), and our paper focuses on their detailed presentation and meta-mathematical investigation in the case of classical predicate logic. As a central theme emerges the connection between restricted forms of nd-proofs and (strategies for) proof search: normal forms are not obtained by removing local "detours", but rather by constructing proofs that directly reflect proof-strategic considerations. That theme warrants further investigation.

1. Proof search

Natural deduction calculi have been available since the mid-thirties and reflect "as accurately as possible the actual logical reasoning involved in mathematical proofs".[1] They capture the logical structure of arguments, in part, by incorporating *inferences from* and *to* complex formulas with characteristic principal connectives. The rules for the "proper" logical connectives, $\wedge$, $\vee$, $\rightarrow$, $\forall$, and $\exists$ are consequently divided into "Elimination", i.e., *proper* E-, and "Introduction", i.e., *proper* I-rules. Rules for negation do not fit fully into this schematic approach, in particular not, if they are formulated in the standard (Gentzen-Prawitz) mould. We use instead a very symmetric formulation: the first rule for negation, $\perp_c$, is the distinctive rule of classical

* The work reported here continues the metamathematical investigations basic for the Carnegie Mellon Proof Tutor, see [Sieg and Scheines]; sections 1–4 are corrected and much improved versions of [Sieg 1992] and [Sieg 1994], whereas sections 5 and 6 expand that work and provide the theoretical basis for automated proof search in predicate logic.

We are grateful for the continued, lively discussions in the weekly meetings of the "CPT-group", i.e., with J. Hughes, M. Ravaglia, and R. Scheines. We also appreciate the comments of R. Dyckhoff, A. M. Conti, and two anonymous referees which helped us to improve an earlier version of this paper.

[1] Gentzen in his "Investigations into logical deduction", [Gentzen], p. 74.

Studia Logica **60**: 67–106, 1998.
©1998 *Kluwer Academic Publishers. Printed in the Netherlands.*

logic and is needed, for example, to prove the law of excluded middle and Peirce's law;

$$\dfrac{\begin{array}{cc}[\neg\psi] & [\neg\psi] \\ \vdots & \vdots \\ \varphi & \neg\varphi\end{array}}{\psi}$$

the second rule, $\perp_i$, captures the form of indirect argumentation admitted also intuitionistically and used, most classically, in the Pythagorean proof of the irrationality of $\sqrt{2}$.

$$\dfrac{\begin{array}{cc}[\psi] & [\psi] \\ \vdots & \vdots \\ \varphi & \neg\varphi\end{array}}{\neg\psi}$$

We consider the rules for negation as both E- and I-rules, but not as *proper* E- or I-rules.

Generally, E-rules specify how components of assumed or established complex formulas can be used in an argument; I-rules provide conditions under which complex formulas can be inferred from already established components. This leads directly to the formulation of very intuitive strategies. Technically, the strategies exploit that the structure of nd-proofs can be made to depend on the *syntactic context* provided by assumptions and conclusions: the nd-calculi share, as Prawitz [1965] discovered, important metamathematical properties with sequent calculi. For the statement of the first of these properties recall that the premise of an E-rule with the characteristic connective is called the *major premise*; a proof is called *p-normal*[2], when no formula occurrence in the proof is the conclusion of a proper I-rule or $\perp_c$ and also the major premise of a proper E-rule. To be quite accurate, we have to exclude *segments* of formula occurrences, such that the first formula in the segment is the conclusion of a proper I-rule or $\perp_c$ and the last formula the major premise of a proper E-rule. Here and below we make use of terminology used by Prawitz — with just one exception, we use 'branch' for his 'thread'. Note also that we have not yet defined 'normal'. In order to obtain a definition matching that of Prawitz, we first define the *adjacency condition*: the major premise of a $\perp$-rule must not be inferred by a $\perp$-rule. A *normal* proof, then, is p-normal and satisfies the adjacency condition.

The first central property, the Normalization Theorem, was established by Prawitz for a restricted language[3]: (by a sequence of special "reductions")

[2] This term is not related to the term "p-normal" used by Troelstra and van Dalen.

[3] Without $\exists$ and $\vee$; $\perp_c$ was applicable only to get atomic conclusions.

any proof of G from α in the nd-calculus can be transformed into a normal proof leading from α to G, where α is a sequence of formulas[4]. The second central property for the nd-calculus concerns the logical complexity of formulas in proofs: normal proofs Σ leading from α to G have a (modified) *subformula property*; i.e., every formula occurring in Σ is (the negation of) either a subformula of G or of an element in α. This is a consequence of the third central property, a structural feature of paths in (the tree presentation of) normal nd-proofs: every path contains a uniquely determined E-part and I-part, consisting only of segments that are major premises of proper E-rules, respectively premises of proper I-rules; these two parts are separated by the minimum segment that is the premise of an I-rule.

Despite the naturalness of nd-calculi, the part of proof theory that deals with them has hardly influenced developments in automated theorem proving. For that the proof theoretic tradition rooted in Herbrand's work and Gentzen's work on sequent calculi has been more important. The keywords here are *resolution*, *tableaux*, and *logic programming*. From a purely logical point of view this is prima facie peculiar: it is after all the subformula property of special kinds of derivations[5] that makes resolution and related techniques possible; normal derivations in natural deduction calculi, as we just noticed, have that very property with the minor addition mentioned. *Why is it then that nd-calculi have not been exploited for automated proof search?* The answer to this general question seems, in part, to lie in answers to three crucial questions: (1) How can one specify through a calculus *only* normal proofs? (2) How can one construct a search space that allows the formulation of strategies for finding such proofs? and (3) How can one prove the termination of search strategies?

In the case of sequent calculi the analogues to these questions have direct answers: use calculi without the cut rule; invert systematically their rules; prove their completeness! In this rough description of the theoretical background for automated deduction based on sequent calculi the syntactic normalization or cut-elimination procedure is not mentioned, since the semantic completeness proof for the cut-free part is fundamental, not Gentzen's cut-elimination procedure. Indeed, algorithms for finding *cut-free* derivations are refinements of strategies used in that completeness proof. Such strate-

[4] Prawitz's proof for the intuitionistic calculus can be extended to the full classical case with the negation rules formulated in the symmetric way as above; that was established by Byrnes. The strong normalization theorem for the full language (with restricted $\vee$- and $\exists$-inferences) was proved by Statman [1974].

[5] Derivations in Herbrand's calculus and derivations in the sequent calculus without cut have the (full) *subformula property*: they contain only subformulas of their endformula, respectively endsequent.

gies realize the heuristic idea of *searching for semantic counterexamples* and yield trees Σ such that *either* one of Σ's branches allows the definition of a counterexample to "α has G as a logical consequence" *or* Σ constitutes a cut-free derivation of the sequent $\neg\alpha, G$.[6] In the case of nd-calculi normal proofs are also sufficient to obtain all logical consequences from given assumptions. However, this fact has not been established directly: its proof combines the completeness theorem for the calculus with the normalization theorem. In order to obtain a direct proof of the fact and an answer to (1), *intercalation calculi* are introduced. They provide frameworks for answering (2), and completeness proofs for these calculi answer (3).

The broad problem is this: How can we derive a conclusion or goal G from assumptions $\phi_1, \ldots, \phi_n$? or, more vividly: How can we close the gap between G and the $\phi_1, \ldots, \phi_n$ via logical rules? This question is at the heart of spanning search spaces via ic-calculi: their basic rules are reformulations of those for Gentzen's nd-calculi, but it is the preservation of inferential information and the restricted way in which the rules are used to close the gap (and thus to build up derivations) that is distinctive. The ic-calculi provide the underpinning for specifying informal approaches to proof search: their rules are used to construct a search space that contains all possible ways of closing the gap between assumptions and G via the ic-rules. In this space we search for a gap-closing subspace that determines, in turn, a unique normal or p-normal nd-proof from the assumptions to G. If the search fails, the search space contains enough information to yield a semantic counterexample. This sketch of the completeness proof for ic-calculi shows the family resemblance to completeness proofs for the sequent calculus without cut. The difference can be put sharply as follows: *In the case of the sequent calculus, one tries to find a semantic counterexample and, if that search fails, one actually has found a proof*[7]*; in the case of ic-calculi, one tries to find a proof and, if that search fails, one has a counterexample.* Let us turn to the rigorous metamathematical discussion.

We will discuss at first only classical sentential logic with the connectives $\neg$, $\wedge$, $\vee$, $\rightarrow$; however, the considerations will then be extended to predicate logic and can be used to treat non-classical logics, for example, intuitionistic

[6] $\neg\alpha$ consists of the negations of the formulas in α.

[7] A sequent proof is far from reflecting the structure of ordinary arguments. Thus, we have here and in the case of resolution based procedures the non-trivial problem of finding associated nd-proofs. Cf. Shanin e.a., but also Andrews and Pfenning. The issue is also addressed in implementations of, e.g., NUPRL and ISABELLE. Bledsoe's way of using nd-methods is not systematic in the logical setting. Cf. our remark at the end of section 3 and also note 18.

logic.[8] The ic-rules operate on triples of the form $\alpha;\beta?G$. α is the sequence of *available assumptions*; G is the current *goal*; β is a sequence of formulas obtained by $\wedge$-elimination and $\rightarrow$-elimination from elements in α[9]. To facilitate the description of rules and parts of search trees let us agree on some conventions. Lowercase Greek letters $\alpha, \beta, \gamma, \ldots$ range over finite sequences of formulas; as syntactic variables over formulas we use $\phi, \psi, \chi, \ldots$, and also G and H; $\Pi, \Sigma, \mathrm{T}, \ldots$ range over trees. $\phi \in \alpha$ expresses that ϕ is an element of the sequence α; α, β or $\alpha\beta$ is short for the concatenation $\alpha * \beta$ of the sequences α and β; α, ϕ stands for the sequence $\alpha * \langle\phi\rangle$, where $\langle\phi\rangle$ is the sequence with ϕ as its only element. There are three kinds of ic-rules: those corresponding to the proper E-rules for $\wedge$, $\vee$, $\rightarrow$; those corresponding to the proper I-rules for $\wedge$, $\vee$, $\rightarrow$; finally, the rules for negation. Let us list the rules of the first kind, i.e., $\downarrow$-rules.

$\wedge_i\downarrow$: $\alpha;\beta?G,\ \phi_1 \wedge \phi_2 \in \alpha\beta \Longrightarrow \alpha;\beta,\phi_i?G$ for i=1 or 2

$\vee\downarrow$: $\alpha;\beta?G,\ \phi_1 \vee \phi_2 \in \alpha\beta \Longrightarrow \alpha,\phi_1;\beta?G$ AND $\alpha,\phi_2;\beta?G$

$\rightarrow\downarrow$: $\alpha;\beta?G,\ \phi_1 \rightarrow \phi_2 \in \alpha\beta \Longrightarrow \alpha;\beta?\phi_1$ AND $\alpha;\beta,\phi_2?G$

Now we formulate the rules that correspond to inverted proper I-rules, i.e., $\uparrow$-rules.

$\wedge\uparrow$: $\alpha;\beta?\phi_1 \wedge \phi_2 \Longrightarrow \alpha;\beta?\phi_1$ AND $\alpha;\beta?\phi_2$

$\vee_i\uparrow$: $\alpha;\beta?\phi_1 \vee \phi_2 \Longrightarrow \alpha;\beta?\phi_i$ for $i = 1$ or 2

$\rightarrow\uparrow$: $\alpha;\beta?\phi_1 \rightarrow \phi_2 \Longrightarrow \alpha,\phi_1;\beta?\phi_2$

Finally, we come to the rules for negation:

$\bot_c(\mathcal{F})$: $\alpha;\beta?G, \varphi \in \mathcal{F}(\alpha,\neg G) \Longrightarrow \alpha,\neg G;\beta?\varphi$ AND $\alpha,\neg G;\beta?\neg\varphi$

$\bot_i(\mathcal{F})$: $\alpha;\beta?\neg G, \varphi \in \mathcal{F}(\alpha, G) \Longrightarrow \alpha, G;\beta?\varphi$ AND $\alpha, G;\beta?\neg\varphi$

$\mathcal{F}(\gamma)$ is obtained as follows. Let F_γ consist of all proper subformulas of formulas in γ and of all negations occurring in γ. $\mathcal{F}(\gamma)$ then consists of all unnegated formulas in F_γ and the unnegated part ψ of all negations $\neg\psi$ in

[8] That was done for sentential logic by Cittadini in his M.S. thesis written in May 1991; see [Cittadini 1992]. The case of intuitionistic predicate logic and other non-classical logics will be considered in a joint paper with Cittadini, "Normal Natural Deduction Proofs (in non-classical logics)".

[9] The reason for this separation is that some important syntactic constructions will refer only to the available assumptions; for example, concerning the indirect rules and, later on in predicate logic, concerning the analogue of $\forall$-introduction.

F_γ. $\mathcal{F}(\gamma)$ is obviously finite; that is crucial for the finiteness of the search space. Operations $\mathcal{O}$ leading to smaller and yet sufficient classes can be specified; cf. the end of section 3. The different calculi we are considering are distinguished through the operation $\mathcal{O}$, and we denote a particular calculus by $\mathrm{IC}_0(\mathcal{O})$, or simply $\mathrm{IC}(\mathcal{O})$ — as long as it is clear that we are dealing with sentential logic; the corresponding systems for first order logic will later be denoted by $\mathrm{IC}_1(\mathcal{O})$.

REMARKS. (1) Intuitionistic versions of ic-calculi are obtained by using the rule *ex falso quodlibet* $\alpha;\beta?G,\ \varphi \in \mathcal{O}(\alpha) \Longrightarrow \alpha;\beta?\varphi$ and $\alpha;\beta?\neg\varphi$ instead of $\perp_c(\mathcal{O})$. For the classical system $\mathrm{IC}(\mathcal{F})$, the rule $\rightarrow\downarrow$ can be weakened to $\alpha;\beta?G,\ \phi_1 \rightarrow \phi_2 \in \alpha\beta,\ \phi_1 \in \alpha\beta \Longrightarrow \alpha;\beta,\phi_2?G$. But this formulation, as Cittadini noticed, is too weak for intuitionistic logic (and unnatural for proof search even in the classical case).

(2) We formulated the ic-rules as Post-productions, but they can also be represented in the standard way with appropriate side conditions; however, the natural application of these rules is "bottom-up". Here are three reformulations:

$$\rightarrow\downarrow\colon \frac{\alpha;\beta?\phi_1 \quad \alpha;\beta,\phi_2?G}{\alpha;\beta?G} \qquad \text{with } \phi_1 \rightarrow \phi_2 \in \alpha\beta$$

$$\rightarrow\uparrow\colon \frac{\alpha,\phi_1;\beta?\phi_2}{\alpha;\beta?\phi_1 \rightarrow \phi_2}$$

$$\perp_i(\mathcal{F})\colon \frac{\alpha,G;\beta?\varphi \quad \alpha,G;\beta?\neg\varphi}{\alpha;\beta?\neg G} \qquad \text{just in case } \varphi \in \mathcal{F}(\alpha,G)$$

Because of this correspondence we call the consequent(s) of a Post-production, *premise(s)* of the appropriate rule. This reformulation brings out the restrictive character of the $\downarrow$-rules: the principal formula of a $\downarrow$-inference must already be in $\alpha\beta$.

Next we turn to the construction of the search or problem space, using these rules; indeed, we shall interleaf the nodes of a tree-like arrangement of questions with "rule nodes" that provide information on the rule that is connecting the questions.

2. The problem space for sentential logic

As an example of how the ic-rules are used to build up the search space for a question $\alpha?G$, let us show the search tree for the question $?P \vee \neg P$. It is partially presented in Diagrams 1, 1.A, and 1.B of the Appendix. We start

out by applying the three possible ic-rules to obtain new questions, namely, $?P$ or $?\neg P$ or, proceeding indirectly, $\neg(P \vee \neg P)?\varphi$ and $\neg(P \vee \neg P)?\neg\varphi$ with each element φ of $\mathcal{F}(\neg(P \vee \neg P))$. Let us pursue the leftmost branch in the tree. To answer $?P$ we have to use $\perp_{\mathbf{c}}$ and, because of the restriction on the choice of contradictory pairs, we have only to ask $\neg P?P$ and $\neg P?\neg P$. In the first case only $\perp_{\mathbf{c}}$ could be applied, but would lead to the question we just analyzed. Thus we *close this branch* with **N**. In the second case the gap between assumption and goal is obviously closed, so we top this branch with **Y**. No rule is applicable to the question $?\neg P$; so that branch is closed with **N** as well. The other parts of the tree are constructed in a similar manner. Each application of $\perp_{\mathbf{c}}$ ($\perp_{\mathbf{i}}$) is labeled "$\perp_{\mathbf{c}}, \phi$" ("$\perp_{\mathbf{i}}, \phi$"), where ϕ is the minor premise of the rule application. The subtree in diagram 1.A is not full, but at the numbered nodes 1 through 4 the resulting trees do not help in closing the gap. In contrast, the subtree in diagram 1.B is of interest, and we discuss it below.

The composition of Diagrams 1, 1.A, and 1.B contains enough information for the extraction of derivations in a variety of styles of natural deduction. For our calculus we can easily obtain corresponding derivations; namely, first:

$$\dfrac{\dfrac{\dfrac{\dfrac{[\neg P]}{P \vee \neg P} \quad [\neg(P \vee \neg P)]}{P}}{P \vee \neg P} \quad [\neg(P \vee \neg P)]}{P \vee \neg P}$$

(Here we use square brackets to indicate cancelation of an assumption.) The second derivation is "dual" to this one with the roles of P and $\neg P$ interchanged. Finally, the derivation that emerges from Diagram 1.B:

$$\dfrac{\dfrac{\dfrac{[\neg P]}{P \vee \neg P} \quad [\neg(P \vee \neg P)]}{P} \quad \dfrac{\dfrac{[P]}{P \vee \neg P} \quad [\neg(P \vee \neg P)]}{\neg P}}{P \vee \neg P}$$

The proof represented in the second diagram above is p-normal, but it is *not* a normal proof, as the major premise $\neg P$ of the last inference with rule $\perp_{\mathbf{c}}$ has been obtained by $\perp_{\mathbf{i}}$. (Natural normalization steps reduce this derivation either to the first derivation or its dual.)

The full search or ic-tree is specified inductively by applying ic-rules to the initial question or to the "non-terminal" leaves of an already obtained partial search tree — in all possible ways, unless the application of a rule

leads to a question that is not *new for* the branch determined by the appropriate leaf ($\alpha;\beta?G$ is the *same question as* $\alpha^*;\beta^*?G$ just in case the sets of formulas in the sequences $\alpha\beta$ and $\alpha^*\beta^*$ are identical.) In either case one addresses questions of the form $\alpha;\beta?G$ at a particular node:

if G is an element of $\alpha\beta$, then close the branch determined by the current question node with **Y**;

if G is not an element of $\alpha\beta$ and every applicable rule leads to a question that is not new for the branch determined by the current question node then close with **N**;

if G is not an element of $\alpha\beta$ and some applicable rule leads to a new question, then extend the tree at the current question node for all such rules by appropriate rule and question nodes (with a fixed ordering of rules)[10].

For any implementation of a proof search procedure it is crucial to decide quickly, whether a particular rule will lead, at the current question node, to a new question or not. A first easy step is to impose local side conditions on the $\downarrow$-rules that prevent the application of a rule, in case it does lead to the same question; this can be done, for example, as follows:

$\wedge_i\downarrow$: $\alpha;\beta?G,\phi_1\wedge\phi_2\in\alpha\beta,\phi_i\notin\alpha\beta\Longrightarrow\alpha;\beta,\phi_i?G$ for $i=1$ or 2

$\vee\downarrow$: $\alpha;\beta?G,\phi_1\vee\phi_2\in\alpha\beta,\phi_1\notin\alpha\beta,\phi_2\notin\alpha\beta\Longrightarrow\alpha,\phi_1;\beta?G$ AND $\alpha,\phi_2;\beta?G$

$\rightarrow\downarrow$: $\alpha;\beta?G,\phi_1\rightarrow\phi_2\in\alpha\beta,\phi_2\notin\alpha\beta,\phi_1\neq G\Longrightarrow\alpha;\beta?\phi_1$ AND $\alpha;\beta,\phi_2?G$

Indeed, these local side conditions are now taken as part of the ic-rules. A second, more intricate step involves a careful analysis of the conditions under which repeated questions can occur. This allows us to avoid checking for repetitions in many instances. A third step would restrict the application of the indirect rules: $\perp_c$ is never applied to negated formulas. Thus, to a given question node only one $\perp$-rule is applied. We do not pursue such issues in any systematic way, as we are intending to present only the broad theoretical framework for proof search via ic-calculi; there will be some additional remarks at the end of sections 3 and 4.

The ic-tree is constructed in the above general way for questions $\alpha?G$; its branches determine sequences of *subquestions* for $\alpha?G$. Due to the finiteness of $\mathcal{F}$ and the form of the rules, only finitely many different subquestions for $\alpha?G$ can be formulated. This together with the requirement not to repeat

[10] For example we could use the order $\wedge_1\downarrow$, $\wedge_2\downarrow$, $\rightarrow\downarrow$, $\vee\downarrow$, $\wedge\uparrow$, $\rightarrow\uparrow$, $\vee_1\uparrow$, $\vee_2\uparrow$, $\perp_i$, $\perp_c$. We also need to order multiple applications of each rule, say by the order in which the formulas to which it is applied appear in $\alpha\beta$.

questions on a branch yields the *Proposition*: The ic-trees for questions $\alpha?G$ are finite, and their branches are closed with either **Y** or **N**. This assignment to questions at leaves of an ic-tree can be extended to all questions in the tree and determines a unique value for the original question $\alpha?G$; the value of a question $\alpha^*;\beta^*?G^*$ is indicated by $[\![\alpha^*;\beta^*?G^*]\!]$. In the remainder of this section we will show: if **Y** is assigned to the root of the ic-tree, then there is a p-normal proof leading from the assumptions to the goal of the question. In the next section this fact will be complemented by a second fact: if **N** is assigned to the root of the ic-tree, then there is not only no p-normal proof, but no proof at all; i.e., the ic-tree contains enough information to show that the inference from α to G is semantically invalid. We will also show that a certain restricted calculus $\mathrm{IC}_0(\mathcal{I})$ is still complete; nd-proofs obtained from "derivations" in that calculus are actually normal.

We saw through the $P \vee \neg P$ example, how an nd-proof can be read off from a properly chosen partial ic-tree whose root evaluates to **Y**. To formulate the underlying general fact properly we define first the notion of an *ic-derivation.*

DEFINITION. An *ic-derivation for the question* $\alpha;\beta?G$ is a subtree T of the ic-tree Σ for $\alpha;\beta?G$ satisfying: (i) $\alpha;\beta?G$ is the root of T, (ii) all branches of T are **Y**-closed branches of Σ, and (iii) every question node in T (that is not a leaf) is followed by exactly one rule node (to obtain the next question(s)).

One can easily extract ic-derivations from ic-trees that evaluate to **Y**. Let Σ be the ic-tree for $\alpha?G$ and assume that $[\![\alpha?G]\!] = \mathbf{Y}$. We can determine from Σ a canonical **Y**-subtree T as $f(hg(\Sigma))$, where $hg(\Sigma)$ is the height of Σ and f a function defined recursively as follows:

$$
\begin{aligned}
f(0) &= \alpha?G \\
f(2n+1) &= \begin{cases} \epsilon_1(f(2n)) & \text{if some branch of } f(2n) \text{ can be extended} \\ f(2n) & \text{otherwise} \end{cases} \\
f(2n+2) &= \begin{cases} \epsilon_2(f(2n)) & \text{if } f(2n+1) \neq f(2n) \\ f(2n) & \text{otherwise} \end{cases}
\end{aligned}
$$

f extends the open branches of a partial ic-derivation by their "left-most **Y**-expansions" in Σ. More explicitly, the open branches of $f(2n)$ are open branches of Σ and are consequently expanded by ic-rules; at least one of these rules must have a (pair of) premise(s) evaluating to **Y**; ϵ_1 chooses the left-most such rule application in each case, and ϵ_2 expands the tree by the appropriate question node(s). The main point is that from an ic-derivation

we can construct uniquely an nd-proof[11]; indeed, that proof is p-normal.

PROPOSITION. *For any Σ, α, β, G: if Σ is an ic-derivation for $\alpha;\beta?G$, then there is a uniquely determined p-normal nd-proof Π_Σ leading from $\alpha\beta$ to G.*

PROOF. (by induction on the height of Σ). If $hg(\Sigma) = 1$, the ic-derivation simply consists of the question $\alpha;\beta?G$ with $G \in \alpha\beta$, as Σ evaluates to $\mathbf{Y}$. Π_Σ is the nd-proof consisting of the node G. — If $hg(\Sigma) > 1$, distinguish cases as to the ic-rule that is applied to $\alpha;\beta?G$ in Σ. The induction hypothesis asserts: for any ic-derivation T with $hg(\mathrm{T}) < hg(\Sigma)$ there is a uniquely determined p-normal nd-proof Π_T answering the question at the root of T.

$\wedge_i\downarrow$: The immediate subderivation T_i of Σ has root $\alpha;\beta,\phi_i?G$; by induction hypothesis there is a uniquely determined p-normal nd-proof Π_{T_i} leading from assumptions in $\alpha\beta,\phi_i$ to G. If Π_{T_i} contains occurrences of ϕ_i as open assumptions, then replace those occurrences by $\dfrac{\phi_1 \wedge \phi_2}{\phi_i}$. The resulting p-normal proof of G from $\alpha\beta$ is the associated nd-proof Π_Σ.

$\vee\downarrow$: The immediate subderivations T_i of Σ have roots $\alpha,\phi_i;\beta?G$ for $i =$ 1 or 2; by induction hypothesis there are uniquely determined p-normal nd-proofs Π_{T_i} leading from α,ϕ_i,β to G. The associated p-normal nd-proof Π_Σ of G from $\alpha\beta$ is:

$$\dfrac{\begin{array}{ccc} & [\phi_1] & [\phi_2] \\ & \vdots & \vdots \\ \phi_1 \vee \phi_2 & G & G \end{array}}{G}$$

This construction is proper, as $\vee\downarrow$ has as its major premise an element of $\alpha\beta$, and G is the endformula of Π_{T_i}.

$\rightarrow\downarrow$: The immediate subderivations T_1 and T_2 of Σ have roots $\alpha;\beta?\phi_1$ and $\alpha;\beta,\phi_2?G$; by induction hypothesis there are uniquely determined p-normal nd-proofs Π_{T_1} and Π_{T_2} leading from $\alpha\beta$ to ϕ_1, respectively from $\alpha\beta,\phi_2$ to G. Use Π_{T_1} and the fact that $\phi_1 \rightarrow \phi_2 \in \alpha\beta$ to construct a p-normal proof Π of ϕ_2 from assumptions in $\alpha\beta$.

$$\dfrac{\begin{array}{cc} \vdots & \\ \phi_1 & \phi_1 \rightarrow \phi_2 \end{array}}{\phi_2}$$

[11] An analogous procedure for the sequent calculus is outlined roughly by Prawitz (p. 91); however, note that no "choices" have to be made in our procedure.

If Π_{T_2} contains any occurrences of ϕ_2 as open assumptions, then replace those occurrences by Π. This construction yields the p-normal proof Π_Σ of G from assumptions in $\alpha\beta$.

$\wedge\uparrow$: The immediate subderivations T_i of Σ have roots $\alpha;\beta?\phi_i$, for $i =$ 1 or 2, and G is $(\phi_1 \wedge \phi_2)$; by induction hypothesis there are uniquely determined p-normal nd-proofs Π_{T_i} leading from $\alpha\beta$ to ϕ_i. The nd-proof Π_Σ is obtained by joining Π_{T_1} and Π_{T_2} via $\wedge$-introduction.

$\vee_i\uparrow$: The immediate subderivation T_i of Σ has root $\alpha;\beta?\phi_i$ and G is $(\phi_1 \vee \phi_2)$; by induction hypothesis there is a uniquely determined p-normal nd-proof Π_{T_i} leading from $\alpha\beta$ to ϕ_i. The p-normal nd-proof Π_Σ is obtained by $\vee$-introduction.

$\rightarrow\uparrow$: The immediate subderivation T of Σ has root $\alpha, \phi_1; \beta?\phi_2$ and G is $(\phi_1 \rightarrow \phi_2)$; by induction hypothesis there is a uniquely determined p-normal nd-proof Π_T leading from α, ϕ_1, β to ϕ_2. The nd-proof Π_Σ is obtained by $\rightarrow$-introduction with ϕ_1 and ϕ_2.
Finally, we treat the rules for negation.

$\perp_i$: The immediate subderivation T [$T_\neg$] of Σ has root $\alpha, \psi; \beta?[\neg]\varphi$, where G is $\neg\psi$ and $\varphi \in \mathcal{F}$; by induction hypothesis there are uniquely determined nd-proofs Π_T and $\Pi_{T_\neg}$ leading from α, ψ, β to φ, respectively $\neg\varphi$. The nd-proof Π_Σ is obtained by applying $\perp_i$ to infer G. — The classical rule $\perp_c$ is treated in the same way as $\perp_i$. ■

The nd-proof Π_Σ uses exactly the same rules as Σ. (One parenthetical remark is appropriate here: the structural similarity between ic-derivations and nd-proofs is even more apparent, when the latter are represented graphically by Fitch-diagrams[12]. The ic-derivations can then be viewed as prescriptions for constructing isomorphic Fitch-diagrams.) Joining the proposition and the earlier observation concerning the extraction of ic-derivations from ic-trees we have:

PROOF EXTRACTION THEOREM. *For any α and G: If the ic-tree Σ for $\alpha?G$ evaluates to* **Y**, *then a p-normal nd-proof of G from assumptions in α can be found.*

[12] Prawitz (1965, p. 98–99) asserts that already Jaśkowski introduced this representation in the late twenties. In any event, for computer implementation Fitch-diagrams are convenient for the representation of nd-proofs: they reflect dependencies as graphically as trees do, but are easier to put on a screen and avoid the duplication of parts of proofs necessary in tree representations.

It is extremely easy to obtain the interpolation theorem (and other metamathematical results); the argument is a modification of that for the proof extraction theorem.

INTERPOLATION THEOREM. *For any α, G: if G is a logical consequence of α, then there is an interpolating formula ϕ together with p-normal nd-proofs Π_ϕ and $\Pi_{\phi,G}$, such that Π_ϕ leads from α to ϕ, and $\Pi_{\phi,G}$ leads from ϕ to G.*

The theorem follows from the next proposition, when observing (with the counterexample extraction theorem established in the next section) that — on account of the fact that G is a logical consequence from α — the ic-tree for the question $\alpha?G$ evaluates to **Y** and thus contains an ic-derivation answering the question $\alpha?G$.

PROPOSITION. *For any Σ, α, β, G: if Σ is an ic-derivation for $\alpha;\beta?G$, then there is a uniquely determined interpolant ϕ, an nd-proof Π_ϕ leading from $\alpha\beta$ to ϕ, and an nd-proof $\Pi_{\phi,G}$ leading from ϕ to G. Furthermore, Π_ϕ and $\Pi_{\phi,G}$ are p-normal.*

3. Normal form theorems for sentential logic

By the evaluation of ic-trees we know that a question $\alpha?G$ obtains the value **Y** or **N**. In case the value is **Y** we can determine an associated p-normal proof. In case the question has value **N**, we have as an immediate consequence: "The search failed!" But that only means that the particular possibilities of building up derivations — as reflected in the construction of the ic-tree — do not lead to a proof establishing G from assumptions in α. We will do better: a *specially* selected branch in the ic-tree can be used to define a semantic counterexample to the inference from α to G.

COUNTEREXAMPLE EXTRACTION THEOREM. *For any α and G : If the ic-tree Π for $\alpha?G$ evaluates to* **N**, *then it contains a canonical refutation branch* P *that determines a valuation* **v** *with* $\mathbf{v} \models \phi$ *for all* $\phi \in \alpha$ *and* $\mathbf{v} \not\models G$. *(That is,* **v** *is a counterexample to the inference from α to G.)*

Clearly, if the question $\alpha?G$ evaluates to **N**, so does one of the questions $\alpha, G^-?\phi$ and $\alpha, G^-?\neg\phi$ for each $\phi \in \mathcal{F}(\alpha, G^-)$, where we define

$$\phi^- = \begin{cases} \psi & \text{if } \phi = \neg\psi \\ \neg\phi & \text{otherwise} \end{cases} \qquad \text{and} \qquad \phi^+ = \begin{cases} \psi & \text{if } \phi = \neg\neg\psi \\ \phi & \text{otherwise} \end{cases}$$

It will be quite direct to see that the following construction leads to a branch P through Π if $\mathcal{F}(\alpha, G^-)$ is non-empty. If this set is empty, α, G^- consists

only of sentential letters. The valuation v, defined for sentential letters P by $\mathbf{v} \models P$ iff P occurs in α, G^-, provides a counterexample. If $\mathcal{F}(\alpha, G^-)$ is not empty, we need a more sophisticated argument and, naturally, some auxiliary definitions.

The finite set $\mathcal{F}(\alpha, G^-)$ for the negation rules can be enumerated (without repetition) by $\langle H_i \rangle_{i \in I}$, where $I = \{i \mid 1 \leq i \leq n\}$. Let $H_0 = G$. Define:

$$\kappa(\gamma, i) = \begin{cases} \mu k.(i < k \leq n \wedge H_k \notin \gamma \wedge \neg H_k \notin \gamma) & \text{if there is such an } H_k \\ 0 & \text{otherwise} \end{cases}$$

The sequence of nodes of $\mathrm{P}^* = \mathrm{P}^*(0), \ldots$ is defined as follows:

$$\begin{array}{rcll} \alpha_0 & = & \alpha & \\ \lambda_0 & = & 0 & \\ \lambda_{m+1} & = & \kappa(\alpha_m, \lambda_m) & \\ G_m & = & \begin{cases} H_{\lambda_m} & \text{if } [\![\alpha_m ? H_{\lambda_m}]\!] = \mathrm{N} \\ \neg H_{\lambda_m} & \text{otherwise} \end{cases} & \\ \alpha_{m+1} & = & \alpha_m, G_m^- & \\ \mathrm{P}^*(2m) & = & \alpha_m ? G_m & \\ \mathrm{P}^*(2m+1) & = & \begin{cases} \bot_{\mathbf{i}}, H_{\lambda_{m+1}} & \text{if } G_m \text{ is a negation} \\ \bot_{\mathbf{c}}, H_{\lambda_{m+1}} & \text{otherwise} \end{cases} & \end{array}$$

Let ν be the smallest m with $\lambda_{m+1} = 0$. Define P to be P^* restricted to $\{m \mid m \leq 2\nu\}$. P is the initial segment of some branch in the search tree; we call the leftmost such branch the canonical refutation branch. Let us illustrate and clarify this construction through Diagram 2 in the Appendix: At each step in selecting the next question node of the canonical branch P one or the other indicated possibility of proceeding must obtain (as long as the set of assumptions can be properly extended), because not both conclusions of the appropriate $\bot$-rule with the contradictory pair H_k and $\neg H_k$ can be evaluated as **Y**. (In case both are evaluated as N, we choose the leftmost.) The top node of P is $\alpha_\nu ? G_\nu$. Let $A = \{\phi \mid \phi \in \alpha_\nu, G_\nu^-\}$. The set A has important syntactic closure properties and this can be exploited to define a valuation that will serve as a model for α, G^-, i.e., a counterexample to $\alpha \models G$. We establish first the closure properties.

CLOSURE LEMMA. *For all formulas ψ :*

(i) $\psi \in A \Longrightarrow \psi^- \notin A$;

(ii) *ψ is a subformula of an element in $A \Longrightarrow \psi^+ \in A$ or $\psi^- \in A$;*

(iii) *ψ is $\neg\neg\phi_1$, $\neg\neg\phi_1 \in A \Longrightarrow \phi_1 \in A$;*

(iv) ψ *is* $(\phi_1 \wedge \phi_2)$, $(\phi_1 \wedge \phi_2) \in A \Longrightarrow \phi_1^+ \in A$ *and* $\phi_2^+ \in A$;
ψ *is* $\neg(\phi_1 \wedge \phi_2)$, $\neg(\phi_1 \wedge \phi_2) \in A \Longrightarrow \phi_1^- \in A$ *or* $\phi_2^- \in A$;

(v) ψ *is* $(\phi_1 \vee \phi_2)$, $(\phi_1 \vee \phi_2) \in A \Longrightarrow \phi_1^+ \in A$ *or* $\phi_2^+ \in A$;
ψ *is* $\neg(\phi_1 \vee \phi_2)$, $\neg(\phi_1 \vee \phi_2) \in A \Longrightarrow \phi_1^- \in A$ *and* $\phi_2^- \in A$;

(vi) ψ *is* $(\phi_1 \to \phi_2)$, $(\phi_1 \to \phi_2) \in A \Longrightarrow \phi_1^- \in A$ *or* $\phi_2^+ \in A$;
ψ *is* $\neg(\phi_1 \to \phi_2)$, $\neg(\phi_1 \to \phi_2) \in A \Longrightarrow \phi_1^+ \in A$ *and* $\phi_2^- \in A$.

PROOF. We assume for simplicity that $G_\nu^- \notin \alpha_\nu$; thus no question node which has α_ν, G_ν^- on the left-hand side will repeat a question in P. (If $G_\nu^- \in \alpha_\nu$ and $\alpha_\nu, G_\nu^- ?G^*$ repeats a question $\alpha_m ?G^*$ on P, then the arguments below are carried out for that earlier question.) (i) Let $\psi \in A$. If ψ is not a negation and $\neg\psi \in A$ then the following subtree is in the search space:

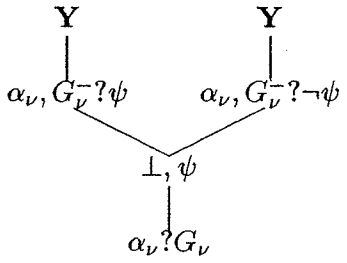

Thus $[\![\alpha_\nu ?G_\nu]\!] = \mathbf{Y}$, contradicting the construction of P. If ψ is a negation, the argument proceeds similarly.

For (ii), let $\phi \in A$ and ψ a subformula of ϕ. Assume as case 1, that $\psi^+ = \psi$. If $\psi = \phi$, we are clearly done; so suppose ψ is a proper subformula of ϕ. Then either ψ or ψ^- is an element of $\mathcal{F}(\alpha_\nu, G_\nu^-) = \mathcal{F}(\alpha, G^-)$. If $\psi \in \alpha, G^-$ or $\psi^- \in \alpha, G^-$, we are done. Otherwise some G_m, $m > 0$, is one of $\psi, \neg\psi, \psi^-, \neg(\psi^-)$. G_m^- is one of $\psi^-, (\neg\psi)^-, \psi^{--}, (\neg(\psi^-))^-$. These are, respectively, $\psi^-, \psi, \psi, \psi^-$, as we are supposing $\psi^+ = \psi$. Thus, G_m^- is either ψ^+ or ψ^-, and as $G_m^- \in A$, either $\psi^+ \in A$ or $\psi^- \in A$.

Assume as case 2 now that $\psi^+ \neq \psi$, i.e., $\psi = \neg\neg\chi$ for some χ, and $\psi^+ = \chi$. Then $\neg\chi$ is an element of $\mathcal{F}(\alpha_\nu, G_\nu^-) = \mathcal{F}(\alpha, G^-)$. If neither $\neg\chi$ nor χ is in α, then some G_m, $m > 0$, is either $\neg\chi$ or $\neg\neg\chi$. G_m^- is then either χ or $\neg\chi$; the former is ψ^+, the latter ψ^-. Thus, as before, $\psi^+ \in A$ or $\psi^- \in A$.

For (iii) assume $\neg\neg\phi_1 \in A$. By (ii) either $(\neg\neg\phi_1)^+ \in A$ or $(\neg\neg\phi_1)^- \in A$; by (i) the latter case cannot arise. Thus $(\neg\neg\phi_1)^+ \in A$; $(\neg\neg\phi_1)^+$ is ϕ_1 and we can conclude that $\phi_1 \in A$.

The arguments for the remaining items are similar. We present the argument only for (iv). First let $(\phi_1 \wedge \phi_2) \in A$ and assume $\phi_1^+ \notin A$ (the case $\phi_2^+ \notin A$ is symmetric); by (ii) $\phi_1^- \in A$. If ϕ_1 is not a negation, then $\phi_1^- = \neg\phi_1$, and the following subtree is in the search space:

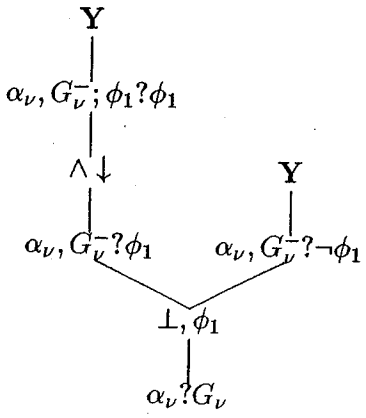

so $[\![\alpha_\nu?G_\nu]\!] = \mathbf{Y}$ which contradicts the construction of P. If ϕ_1 is a negation, then we have a symmetric tree which again yields a contradiction.

Here the application of $\wedge\downarrow$ is crucial. To establish the second part of (iv) $\wedge\uparrow$ is used analogously. So assume that $\neg(\phi_1 \wedge \phi_2) \in A$, $\phi_1^- \notin A$, and $\phi_2^- \notin A$. By (ii) we have that $\phi_1^+ \in A$ and $\phi_2^+ \in A$. For simplicity's sake let us first consider the case that $\phi_i^+ = \phi_i$. Then we have:

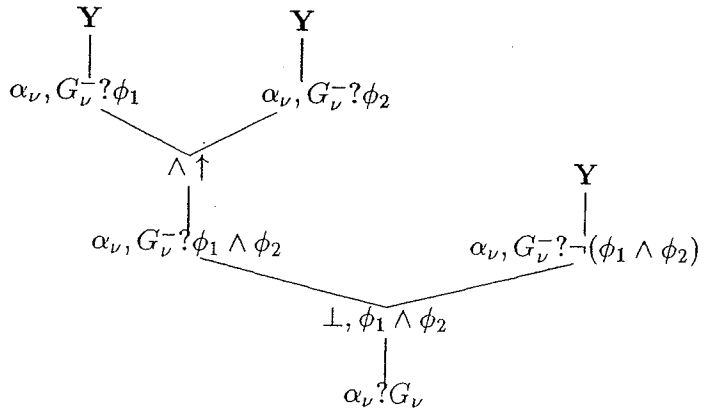

If ϕ_1 is $\neg\neg\chi_1$ and ϕ_2^+ is ϕ_2, then the left branch over the question node $\wedge\uparrow$ has to be replaced by

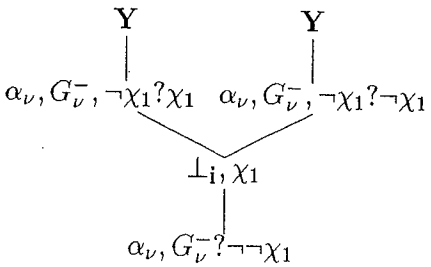

The remaining cases ϕ_1^+ is ϕ_1, but ϕ_2 is $\neg\neg\chi_2$, and ϕ_i is $\neg\neg\chi_i$ are treated similarly. ■

Now define a valuation by $\mathbf{v} \models P$ *iff* $P \in A$. Using this valuation and the closure lemma we can prove the proposition: for every $\phi \in A$, $\mathbf{v} \models \phi$. Hence $\mathbf{v}$ is a model for α, G^-; this concludes the proof of the theorem concerning the extraction of counterexamples. Putting these considerations together, we have a completeness theorem for classical sentential logic in the following form:

COMPLETENESS THEOREM. *The ic-tree for the question $\alpha?G$ allows us to determine either a p-normal proof of G from α or a branch that provides a counterexample to the inference from α to G.*

This yields, a *semantic proof of the p-normal form theorem* for the natural deduction calculus.

P-NORMAL FORM THEOREM. *If G can be proved from assumptions in α, then there is a p-normal proof of G from α.*

As a matter of fact, the proof establishes more, as the nd-proofs obtainable from ic-derivations are a proper subclass of p-normal derivations; for example, the following derivations

$$\dfrac{\dfrac{[\phi_1] \quad \phi_1 \to \phi_2}{\phi_2}}{\phi_1 \to \phi_2}$$

and

$$\dfrac{\dfrac{\phi_1 \wedge \phi_2}{\phi_1} \quad \dfrac{\phi_1 \wedge \phi_2}{\phi_2}}{\phi_1 \wedge \phi_2}$$

are p-normal, but not obtainable from an ic-derivation. Notice that these derivations are actually normal and that one can construct such derivations of arbitrary length. As a matter of fact, the "normal" form can be further restricted. But before considering such additional restrictions we would like to re-emphasize one absolutely central point: the normality of the nd-proofs obtained from ic-derivations is a direct consequence of (the very intuitive strategy for constructing nd-proofs that underlies) the generation of ic-trees for particular questions. That intuitive strategy consists of trying to close the gap between assumptions and conclusion "from above" (by elimination rules) and "from below" (by inverted introduction rules); if neither works,

one proceeds indirectly. Thus, a $\downarrow$-rule can only be applied to assumptions or to formulas that have been inferred by $\downarrow$-rule applications; similarly, the conclusion of $\bot_{\mathbf{c}}$ cannot be the major premise of a proper elimination rule.

Further restrictions on "normal" forms are obtained by restricting the generation of ic-trees; we discuss this here only for modifications of the $\bot$-rules. This will lead to normal nd-proofs. In the above discussion we considered $\bot_{\mathbf{c}}$ essentially as an I-rule for complex non-negated formulas; to a formula thus introduced no E-rule can be applied. Why not consider also negated formulas and disallow subsequent applications of $\bot_{\mathbf{c}}$ (now viewed as an E-rule)? That excludes then in particular nd-proofs of the form

$$\dfrac{\dfrac{[\neg\psi]}{\vdots}\quad\dfrac{\begin{array}{cc}[\phi] & [\phi]\\ \vdots & \vdots\\ \chi & \neg\chi\end{array}}{\neg\phi}}{\phi}\qquad\Bigg/\ \psi$$

Indeed, this is just a special case: in a normal proof, no major premise of a $\bot$-rule is the conclusion of a $\bot$-rule. That the ic-calculus can be restricted in such a way as to provide only normal nd-proofs (without loss of completeness) will be a consequence of the subsequent considerations.

For proof search it is important that ic-trees be pruned — without losing completeness. That can be achieved by restricting the formulas with which contradictory pairs are formed; one can do this through four successively more restrictive versions of the operation $\mathcal{F}(\gamma)$, namely, $\mathcal{N}(\gamma)$, $\mathcal{P}(\gamma)$, $\mathcal{S}(\gamma)$, and $\mathcal{I}(\gamma)$. N_γ (P_γ, S_γ) consists of all negations that occur as (positive, strictly positive)[13] subformulas in γ; I_γ contains exactly the elements of γ that are negations. $\mathcal{N}(\gamma)$ ($\mathcal{P}(\gamma)$, $\mathcal{S}(\gamma)$, $\mathcal{I}(\gamma)$) consists then of the formulas ψ with $\neg\psi$ in N_γ (P_γ, S_γ, I_γ). The $\bot$-rules for these operations are now formulated, except for $\mathcal{I}$, as indicated earlier; the $\bot(\mathcal{I})$-rules are given as follows:

$\bot_{\mathrm{c}}(\mathcal{I})$: $\alpha;\beta?G,\varphi\in\mathcal{I}(\alpha\beta,\neg G)\Longrightarrow\alpha,\neg G;\beta?\varphi$ AND $\alpha,\neg G;\beta?\neg\varphi$

$\bot_{\mathrm{i}}(\mathcal{I})$ is given in a similar way. Clearly, these rules can be reformulated as

$\bot_{\mathrm{c}}(\mathcal{I})$: $\alpha;\beta?G,\varphi\in\mathcal{I}(\alpha\beta,\neg G)\Longrightarrow\alpha,\neg G;\beta?\varphi$

and similarly for $\bot_{\mathrm{i}}(\mathcal{I})$; this brings out most clearly that $\neg\varphi$ is "immediately available".[14]

[13] These notions are defined in the Appendix.

[14] A trivial modification is now needed in the proof extraction lemma.

To establish completeness for each of the resulting variations of the ic-calculus, it suffices to show that the restricted ic-trees (built up by the $\downarrow$- and $\uparrow$-rules, and the restricted $\perp$-rules) allow the extraction of a counterexample in case $[\![\alpha?G]\!] = \mathbf{N}$. The construction of a canonical refutation branch involves now not only the $\perp$-rules, but possibly all the other rules. In defining such a branch one has to make sure that the appropriate version of the closure lemma can be established. From this fact for $\mathrm{IC}(\mathcal{I})$ we can infer the normal form theorem below: the adjacency condition is obviously satisfied in this case, as the major premise $\alpha, G^{-}; \beta?\neg\varphi$ of the $\perp$-rules has an immediate **Y**-answer.

NORMAL FORM THEOREM. *If G can be proved from assumptions in α, then there is a normal proof of G from α.*

REMARK. Before extending our considerations to full predicate logic, let us return to some general remarks we made in section 1. There we emphasized the role of the ic-calculus as a technical tool in the search for nd-proofs. The rules are directly modeled after the I-, E-, and $\perp$-rules of the classical natural deduction calculus (with a special treatment of classical negation). However, due to the way in which assumptions are indicated and $\downarrow$-rules are represented, there is also a certain resemblance with the sequent calculus.[15]

Two distinctive features of the ic-calculus were already mentioned in note 9 and remark 2 at the end of section 1. Here we note some additional (and obvious) differences with the sequent calculus: (i) the ic-calculus always has exactly one formula on the right-hand side; (ii) every formula on the left-hand side of a conclusion appears on the left-hand side of the premise(s); (iii) redundant formulas may not be inferred on the left-hand side; (iv) the negation rules have been altered. To put it briefly and informally: the ic-calculus is a special form of natural deduction, where the goal is never left out of sight!

It has been suggested that the sequent calculus could be used as well as the ic-calculus in the search for nd-proofs, i.e., one would proceed in two steps:

(i) search for a proof in, say, Gentzen's **LK** (or alternatively in a tableau system, which can be viewed as a notational variant of **LK**);

(ii) translate the resulting proof into an nd-proof.

[15] If it were just for the first feature, we would have essentially the formulation of NK as given in [Gentzen 1936], p. 512–515.

In (i), **LK** may be restricted so as to make for more efficient search, and, in addition, allow an easier translation to **NK** or provide more natural **NK** proofs. In (ii), the **LK** proof itself may be manipulated before translating to achieve the additional goals just mentioned. (We referred to work along these lines in note 7, in particular that of Shanin e.a.) From this point of view one might look at the technical aspects of our paper as imposing particular restrictions on **LK** which make search more efficient and translation into **NK** trivial.

However, this view is rather forced: it brushes aside not only all differing "details" of the calculi, but also the strategic use of the ic-calculus for building up an *appropriate* search space. The search space should be appropriate for our main goal, i.e., it should allow us, from the very beginning, to focus on the question of *finding* "natural" **NK** proofs by using "natural" search strategies.[16] The most obvious of these strategies is to work backward from the goal formula and forward from derived lines, both in a restricted and goal-directed manner, i.e., to perform sequences of intercalation steps. We try to find simple representations of the states of the search and of the transition steps taking us from one state to the next. The representation of the final state of a successful search must encode a "proof" that can be directly viewed as an **NK** proof. This difference in the strategic use of the ic-calculus comes out clearly in the completeness proof for the calculus presented in this section (and was emphatically stated in section 1).

4. Normal form theorems for predicate logic

The metamathematical considerations for sentential logic can be extended to predicate logic. To that end we use the following formulation of the E- and I-rules for the quantifiers; note that writing ϕt assumes that t is free for x in ϕx or, alternatively, that some bound variables in ϕx have been renamed. For $\forall$ we have the rules:

$$\frac{(\forall x)\phi x}{\phi t}\ \forall\mathrm{E} \qquad\qquad \frac{\phi y}{(\forall x)\phi x}\ \forall\mathrm{I}$$

Applications of the I-rule must satisfy the restriction that y does not have a free occurrence in any assumption on which the derivation of ϕy depends. —

[16] Indeed, this opens interesting questions for proof theoretic study, e.g., how is the form of nd-proofs related to the strategies used in their search?

For $\exists$ we have the rules:

$$\dfrac{(\exists x)\phi x \qquad \begin{array}{c}[\phi y]\\ \vdots \\ \eta\end{array}}{\eta}\ \exists\mathrm{E} \qquad\qquad \dfrac{\phi t}{(\exists x)\phi x}\ \exists\mathrm{I}$$

with the usual restriction on the E-rule, namely, y must not have free occurrences in η or $(\exists x)\phi x$ nor in any assumption (other than ϕy) on which the proof of (the upper occurrence of) η depends.

To build up ic-trees one applies now also quantifier rules "to close the gap between assumptions and conclusion" in the ic-format. In the formulation of the ic-rules $\mathcal{T}(\gamma)$ denotes the finite set of terms occurring in the formulas of γ.[17]

$\forall\downarrow$: $\alpha;\beta?G$, $(\forall x)\phi x \in \alpha\beta$, $t \in \mathcal{T}(\alpha\beta, G) \Longrightarrow \alpha;\beta,\phi t?G$

$\exists\downarrow$: $\alpha;\beta?G$, $(\exists x)\phi x \in \alpha\beta$, y is new for α, $(\exists x)\phi x, G \Longrightarrow \alpha,\phi y;\beta?G$

$\forall\uparrow$: $\alpha;\beta?(\forall x)\phi x$, y is new for α, $(\forall x)\phi x \Longrightarrow \alpha;\beta?\phi y$

$\exists\uparrow$: $\alpha;\beta?(\exists x)\phi x$, $t \in \mathcal{T}(\alpha\beta, (\exists x)\phi x) \Longrightarrow \alpha;\beta?\phi t$

Ic-trees are specified inductively: if $\alpha^*;\beta^*?G^*$ is an open question, all possibilities of intercalating formulas are considered as in the case of sentential logic. Let us just remark that for applications of the $\perp$-rules we are considering as (proper) subformulas of quantified formulas all instances with the finitely many terms in $\mathcal{T}$. The resulting calculus is denoted by $\mathrm{IC}_1(\mathcal{O})$, depending on the set of formulas admitted for the $\perp$-rules. Branches are closed with **Y** and **N** under the same conditions as before. In general, however, ic-trees will not be finite. Thus, at every stage of construction there may be an open question at some leaf; to evaluate finite *partial* ic-trees Σ a third value **U** is assigned to such a leaf. Given the valuation v_Σ, the value of the question at Σ's root is determined by recursion on Σ following Kleene's scheme [p. 334] for three-valued logic: If N is a leaf of Σ, $[\![N]\!]_\Sigma = \mathrm{v}_\Sigma(N)$, and in case N is the unique successor of M, $[\![N]\!]_\Sigma = [\![M]\!]_\Sigma$. In case N is at a conjunctive branching,

$$[\![N]\!]_\Sigma = \begin{cases} \mathbf{Y} & \text{if for all immediate predecessors } M \text{ of } N\text{: } [\![M]\!]_\Sigma = \mathbf{Y} \\ \mathbf{N} & \text{if for some immediate predecessor } M \text{ of } N\text{: } [\![M]\!]_\Sigma = \mathbf{N} \\ \mathbf{U} & \text{otherwise} \end{cases}$$

[17] As in the propositional calculus, we add restrictions to the $\downarrow$-rules which prune the search space. In the case of $\exists\downarrow$, the restriction is that there is no t such that $\phi t \in \alpha\beta$; in the case of $\forall\downarrow$, we require that $\phi t \notin \alpha\beta$.

and in case N is at a disjunctive branching,

$$[\![N]\!]_\Sigma = \begin{cases} \mathbf{N} & \text{if for all immediate predecessors } M \text{ of } N\text{: } [\![M]\!]_\Sigma = \mathbf{N} \\ \mathbf{Y} & \text{if for some immediate predecessor } M \text{ of } N\text{: } [\![M]\!]_\Sigma = \mathbf{Y} \\ \mathbf{U} & \text{otherwise} \end{cases}$$

The full ic-tree Σ for $\alpha?G$ is defined in stages as follows: Σ_0 is $\alpha?G$; Σ_{n+1} is Σ_n if $[\![\alpha?G]\!]_{\Sigma_n}$ is either $\mathbf{Y}$ or $\mathbf{N}$, otherwise Σ_{n+1} is obtained from Σ_n by expanding each open branch by all applicable rules. Three possibilities can arise: (1) for some $n \in N$, $[\![\alpha?G]\!]_{\Sigma_n} = \mathbf{Y}$, (2) for some $n \in N$, $[\![\alpha?G]\!]_{\Sigma_n} = \mathbf{N}$, and (3) for all $n \in N$, $[\![\alpha?G]\!]_{\Sigma_n} = \mathbf{U}$. In the first case a p-normal derivation can be associated with a subtree of Σ_n — by selecting an ic-derivation and by proving (inductively) that each ic-derivation determines a unique p-normal derivation of G from elements in α. In the second case we can construct a finite canonical refutation branch as in sentential logic and define from it a counterexample. The third case, whose treatment is clearly crucial to complete this sketch of the completeness proof, requires additional considerations.

COUNTEREXAMPLE EXTRACTION THEOREM. *For any α and G: if the ic-tree Σ for $\alpha?G$ is such that for every natural number n $[\![\alpha?G]\!]_{\Sigma_n} = \mathbf{U}$, then Σ contains an infinite refutation branch* P *that determines a structure $\mathcal{M}$ with $\mathcal{M} \models \phi$, for all $\phi \in \alpha$, and $\mathcal{M} \models \neg G$. Thus, $\mathcal{M}$ is a counterexample to the inference from α to G.*

The extraction of a counterexample from an infinite ic-tree requires some circumspection: Instead of constructing a refutation branch directly, we determine first a particular infinite subtree Σ^* of the ic-tree Σ and then apply König's Lemma to this *canonical refutation tree*. The reason for having to cut down the ic-tree Σ to the canonical refutation tree Σ^* is this: Refutation branches have to satisfy suitable closure conditions, and it is trivial to construct infinite branches of Σ that don't. So we define Σ^* in such a way that all of its infinite branches satisfy the closure conditions. The pertinent considerations extend those for sentential logic with variations on familiar Henkin and "fair" tableau constructions; thus we emphasize only the crucial points.

The construction of Σ^* (as a subtree of the ic-tree Σ) for the question $\alpha?G$ proceeds in two waves: The first aims for "sub-maximization" with respect to a given finite set of formulas, whereas the second introduces new subformulas by witnessing — through instances with new variables — existential and negated universal formulas that occur on the l.h.s. of ?. We start out the construction of the binary tree Σ^* (using conventions and definitions

from the sentential case) with the first wave for the enumeration of the formulas in $\mathcal{F}(\alpha, G^-)$ as in the sentential logical case with $\Sigma^*(0) = \alpha?G$. For $m \geq 0$, we extend each open branch of $\Sigma^*(2m)$ (i.e., its leaf evaluates to **U**) with a rule node of the form $\bot, \phi$

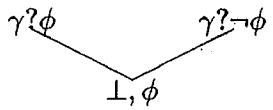

if both questions $\gamma?\phi$ and $\gamma?\neg\phi$ evaluate to **U**; if only one of them evaluates to **U**, then the branch is extended at just that question. One of these cases must hold, because the rule node $\bot, \phi$ has value **U**. (Clearly, as before, ϕ is the first element in the given enumeration that extends γ properly.) After finitely many steps this construction cannot be continued. However, at least one branch in the tree constructed so far has to be open for extensions by rules other than the $\bot$-rules, as for all $n \in N$ $[\![\alpha?G]\!]_{\Sigma_n} = \mathbf{U}$. In sentential logic, we saw, that cannot happen; the resulting set of formulas A is deductively closed in the sense of the earlier Closure Lemma. Here, some of the A's associated with leaves cannot satisfy the closure conditions $(\exists x)\phi x \in A \Longrightarrow \phi^+ t \in A$ for some term t, and $\neg(\forall x)\phi x \in A \Longrightarrow \phi^- t \in A$ for some term t. In the first case the rule $\exists\downarrow$ is applicable with a canonically chosen new variable; in the second case we are able to extend the branch in the following way using also a canonically chosen new variable:

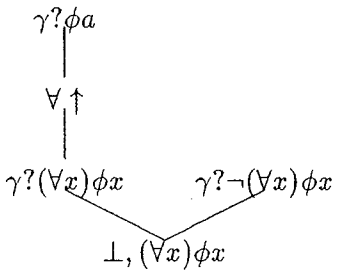

The right extension closes with **Y**, whereas the left one remains open. This brings us to the second wave: We apply $\exists\downarrow$ in all needed cases and then perform the above analysis on those $\neg(\forall x)\phi x$ for which no negated instance is available. The first wave can be repeated now for an extended set of formulas and so on, obviously! We obtain in this way an infinite subtree Σ^* of the ic-tree; König's Lemma applied to Σ^* yields an infinite branch P. Define $A_{\mathrm{P}} = \{\psi \mid \psi$ occurs on the l.h.s. of ? in some question on P$\}$; this set has all the appropriate closure properties needed to serve as the basis for the counterexample definition. Let $\mathcal{T}(A_{\mathrm{P}})$ consist of all terms that occur in some formula of A_{P}.

CLOSURE LEMMA. *For all formulas ψ:*

(i) $\psi \in A_P \Longrightarrow \psi^- \notin A_P$;

(ii) *ψ is a subformula of an element in* $A_P \Longrightarrow \psi^+ \in A_P$ *or* $\psi^- \in A_P$;

(iii) *ψ is* $\neg\neg\phi_1$, $\neg\neg\phi_1 \in A_P \Longrightarrow \phi_1 \in A_P$;

(iv) *ψ is* $(\phi_1 \wedge \phi_2)$, $(\phi_1 \wedge \phi_2) \in A_P \Longrightarrow \phi_1^+ \in A_P$ *and* $\phi_2^+ \in A_P$;
ψ is $\neg(\phi_1 \wedge \phi_2)$, $\neg(\phi_1 \wedge \phi_2) \in A_P \Longrightarrow \phi_1^- \in A_P$ *or* $\phi_2^- \in A_P$;

(v) *ψ is* $(\phi_1 \vee \phi_2)$, $(\phi_1 \vee \phi_2) \in A_P \Longrightarrow \phi_1^+ \in A_P$ *or* $\phi_2^+ \in A_P$;
ψ is $\neg(\phi_1 \vee \phi_2)$, $\neg(\phi_1 \vee \phi_2) \in A_P \Longrightarrow \phi_1^- \in A_P$ *and* $\phi_2^- \in A_P$;

(vi) *ψ is* $(\phi_1 \rightarrow \phi_2)$, $(\phi_1 \rightarrow \phi_2) \in A_P \Longrightarrow \phi_1^- \in A_P$ *or* $\phi_2^+ \in A_P$;
ψ is $\neg(\phi_1 \rightarrow \phi_2)$, $\neg(\phi_1 \rightarrow \phi_2) \in A_P \Longrightarrow \phi_1^+ \in A_P$ *and* $\phi_2^- \in A_P$

(vii) *ψ is* $(\exists x)\phi x$, $(\exists x)\phi x \in A_P \Longrightarrow \phi^+ t \in A_P$ *for some term* $t \in \mathcal{T}(A_P)$;
ψ is $\neg(\exists x)\phi x$, $\neg(\exists x)\phi x \in A_P \Longrightarrow \phi^- t \in A_P$ *for all terms* $t \in \mathcal{T}(A_P)$;

(viii) *ψ is* $(\forall x)\phi x$, $(\forall x)\phi x \in A_P \Longrightarrow \phi^+ t \in A_P$ *for all terms* $t \in \mathcal{T}(A_P)$;
ψ is $\neg(\forall x)\phi x$, $\neg(\forall x)\phi x \in A_P \Longrightarrow \phi^- t \in A_P$ *for some term* $t \in \mathcal{T}(A_P)$.

The definition of a structure $\mathcal{M}$ from A_P is now standard, and we obtain a completeness theorem for classical predicate logic in the form:

COMPLETENESS THEOREM. *The ic-tree for the question $\alpha?G$ determines either a p-normal nd-proof of G from α or a branch that provides a counterexample $\mathcal{M}$ to the inference from α to G.*

So we have a semantic argument for the p-normalizability of nd-proofs; and from ic-derivations we can construct not only p-normal nd-proofs, but also as in the case of sentential logic interpolants to obtain the interpolation theorem.

P-NORMAL FORM THEOREM. *If G can be proved from assumptions in α, then there is a p-normal nd-proof of G from α.*

The $\perp$-rules can be restricted to smaller classes of formulas; that provides then, as in the case of sentential logic, the argument for the normal form theorem.

NORMAL FORM THEOREM. *If G can be proved from assumptions in α, then there is a normal nd-proof of G from α.*

If we were just concerned with establishing normal form theorems, we could end our paper right here. However, we want to provide the broad theoretical basis for proof search in first order logic. That requires additional work,

namely, to find the basis for a natural extension of the search algorithm for sentential logic, as implemented in the Carnegie Mellon Proof Tutor.[18]

5. Skolem-Herbrand expansion

For the search algorithm the language of predicate logic is expanded by *new free variables* and *Skolem and Herbrand* functions as done, for example, in Fitting's book. It is in this expansion that quantifiers are eliminated during the search in a "canonical" way. To direct the search we use heuristics employed for sentential logic together with two novel features, namely an appropriately narrow concept of "strictly positive canonical subformula" and a unification algorithm for quantified formulas, see [Sieg and Kauffmann]. We will come back to these issues at the end of our paper, briefly. Here we focus on the description of the search space, i.e., the generation of appropriate ic-trees.

$\mathcal{L}_1$ is the language of IC_1; the *Skolem-Herbrand expansion* $\mathrm{IC}_{\mathrm{SH}}$ has as its underlying language an expansion $\mathcal{L}_{\mathrm{SH}}$ of $\mathcal{L}_1$; $\mathcal{L}_1$ is fixed here to have just the set $X = \{x, x_0, x_1, \ldots\}$ as its set of variables. $\mathcal{L}_{\mathrm{SH}}$ has in addition a set Y of bound variables $\{y, y_0, y_1, \ldots\}$, a set Z of parameters $\{z, z_0, z_1, \ldots\}$, and a set F of function symbols $\{f, f_0, f_1, \ldots\}$. The sets X, Y, and Z are all disjoint; F contains infinitely many function symbols for each arity n, n a natural number; the 0-ary symbols are constants. Terms and formulas of $\mathcal{L}_{\mathrm{SH}}$ are inductively generated as usual. Let us just note that we call a given variable or function symbol *new* for a given tree (or set) if that symbol does not occur in the tree (or set). For a sequence of formulas γ, $\mathcal{T}_{\mathrm{SH}}(\gamma)$ is the set of terms in γ; $\mathcal{T}_1(\gamma)$ contains exactly those terms in γ which are terms in $\mathcal{L}_1$; $\mathcal{T}^*(\gamma) = \mathcal{T}_{\mathrm{SH}}(\gamma) - \mathcal{T}_1(\gamma)$. Similarly we use $\mathcal{L}^* = \mathcal{L}_{\mathrm{SH}} - \mathcal{L}_1$. Finally, the set of parameters of γ is given by $\mathrm{FV}(\gamma) = \mathcal{T}_{\mathrm{SH}}(\gamma) \cap Z$. The $\mathrm{IC}_{\mathrm{SH}}$ calculus is obtained from IC_1 by replacing the quantifier rules with those which appear below:[19]

[18] Quite sophisticated strategies are involved in the algorithm underlying the Proof Tutor that searches automatically for nd-proofs in classical sentential logic; that program was developed by Richard Scheines and Wilfried Sieg with assistance from Jonathan Pressler and Chris Walton. Presently we are redesigning it in collaboration with Jesse Hughes, Mark Ravaglia, Richard Scheines, and Frank Wimberly, and we have extended the search algorithm to predicate logic along the lines sketched here. — Similarly motivated programs have been developed by Jeff Pelletier and Fred Portoraro; cf. their papers in this volume.

[19] We can, of course, make the same kind of restrictions as before concerning inference of repeated formulas; cf. note 17. We can also improve efficiency by restricting ourselves to canonically chosen function symbols (one for each formula up to renaming of variables) and by taking as parameters for the term only the "relevant" variables — both strategies are discussed and analyzed for tableaux by Baaz and Fermüller, 1995.

$\forall\downarrow$: $\alpha;\beta?G, (\forall x)\phi x \in \alpha\beta \Longrightarrow \alpha;\beta,\phi z?G$ for some new z

$\exists\downarrow$: $\alpha;\beta?G, (\exists x)\phi x \in \alpha\beta, \overline{z} = \mathrm{FV}(\alpha,(\exists x)\phi x, G) \Longrightarrow \alpha,\phi f(\overline{z});\beta?G$ for some new f

$\forall\uparrow$: $\alpha;\beta?(\forall x)\phi x, \overline{z} = \mathrm{FV}(\alpha,(\forall x)\phi x) \Longrightarrow \alpha;\beta?\phi f(\overline{z})$ for some new f

$\exists\uparrow$: $\alpha;\beta?(\exists x)\phi x \Longrightarrow \alpha;\beta?\phi z$ for some new z

Parameters and function symbols are new relative to (partial) $\mathrm{ic}_{\mathrm{SH}}$-trees. Such trees are built up in the most straightforward way by using the rules of $\mathrm{IC}_{\mathrm{SH}}$; what is not as straightforward is the formulation of appropriate closure conditions. I.e., branches will be closed with **Y**, **N**, and **U** under roughly the same conditions as before, but now we consider also "partial" yes-answers $\mathbf{Y}_\sigma$ relative to a "unifying substitution" σ. The reason is simple, as the question "$\alpha;\beta?G$" is now asking "Is G unifiable with an element in $\alpha\beta$?" In case we find unifying substitutions, we close the branch with a sequence of $\mathbf{Y}_\sigma$'s and, in case other rules can be applied, also with **U**. In the last case, all other options of intercalating formulas are used to expand the partial $\mathrm{ic}_{\mathrm{SH}}$-tree.

Three points have to be taken up: (1) appropriate unification, i.e., a substitution concept generalized to formulas; (2) evaluation of partial $\mathrm{ic}_{\mathrm{SH}}$-trees that uses the unification information properly; (3) extraction of nd-proofs from $\mathrm{ic}_{\mathrm{SH}}$-trees. The last issue and normal form theorems will be addressed in the next section. (2) will be quite naturally resolved, as soon as (1) is properly set up.

DEFINITION. A *term assignment* is a mapping σ from Z to the terms of $\mathcal{L}_{\mathrm{SH}}$ such that $\sup(\sigma) = \{z \mid \sigma(z) \neq z\}$ is finite. If $\sup(\sigma) = \{z_0,\ldots,z_n\}$ then σ can be represented by $\langle\sigma(z_0)/z_0,\ldots,\sigma(z_n)/z_n\rangle$; $\langle\rangle = \mathrm{id}_Z$.

Substitutions, based on term assignments, will include a canonical renaming of bound variables. For that we have to consider "modifications" of term assignments $\sigma^{(t_0/w_0,\ldots,t_n/w_n)}$ for variables $w_0,\ldots,w_n \in X \cup Z$ and terms $t_0,\ldots,t_n \in \mathcal{L}_{\mathrm{SH}}$. The modification is given for $w \in X \cup Z$ by

$$\sigma^{(t_0/w_0,\ldots,t_n/w_n)}(w) = \begin{cases} t_i & \text{if } w = w_i \text{ for some } i \leq n \\ \sigma(w) & \text{otherwise} \end{cases}$$

Note that this will not be, in general, a term assignment, as variables from X may appear in the support. For a given modified term assignment σ we define a family of (i,σ)-substitutions on $X \cup Z$ as follows: $\sigma_i[x] = \sigma(x)$ and $\sigma_i[z] = \sigma(z)$ in the base case; σ_i distributes over function and relation

symbols, but also over the sentential connectives. For quantified formulas $(Qw)\phi$, where Q is $\forall$ or $\exists$,

$$\sigma_i[(Qw)\phi] = (Qy_i)\sigma_{i+1}^{(y_i/w)}[\phi].$$

For a given term assignment σ, we write σ for σ_0 and call any (i,σ)-substitution simply a substitution.

Notice that $\sigma[(\forall x_1)(\exists x_2)Px_1x_2]$ is $(\forall y_0)(\exists y_1)Py_0y_1$; applying σ to $(\forall x_7)(\exists x_3)Px_7x_3$ yields the same result; i.e., σ literally identifies formulas that are identical only up to renaming of bound variables. It is the canonical renaming of bound variables that allows us to extend unifiability from terms to formulas: Two formulas ϕ and ψ in $\mathcal{L}_{\mathrm{SH}}$ are called *unifiable* iff there is a term assignment σ, such that $\sigma[\phi] = \sigma[\psi]$.[20] Let us look at some examples on how to prove statements in the expanded calculus and motivate the additional technical steps we have to take.

EXAMPLE 1. $Pa \vdash (\exists x)Px$

$$\begin{array}{c} \mathbf{Y}_\sigma \\ | \\ Pa?Pz_1 \\ | \\ \exists \uparrow \\ | \\ Pa?(\exists x)Px \end{array}$$

σ is the substitution $\langle a/z_1\rangle$; as example 3 will show, closing with $\mathbf{Y}_\sigma$ will not always guarantee success of the proof search.

EXAMPLE 2. $(\forall x_0)Px_0, (\forall x_1)Qx_1 \vdash (\forall x_2)Px_2 wedge (\forall x_3)Qx_3$

$$\begin{array}{ccc} \mathbf{Y}_\emptyset & & \mathbf{Y}_\emptyset \\ | & & | \\ (\forall x_0)Px_0, (\forall x_1)Qx_1?(\forall x_2)Px_2 & & (\forall x_0)Px_0, (\forall x_1)Qx_1?(\forall x_3)Qx_3 \\ & \searrow \quad \swarrow & \\ & \wedge \uparrow & \\ & | & \\ & (\forall x_0)Px_0, (\forall x_1)Qx_1?(\forall x_2)Px_2 \wedge (\forall x_3)Qx_3 & \end{array}$$

This transforms directly into an nd-proof, as we assume the general renaming rule; cf. beginning of section 6.

[20] Standard unification algorithms can be easily adapted to provide a most general idempotent unifier; cf. [Sieg and Kauffmann].

EXAMPLE 3. Closing every branch with a substitution is not enough to guarantee that a derivation has been found. For example,

$$(\exists x_1)(\exists x_2)Px_1x_2, (\exists x)Qx \not\vdash (\exists x_1)(\exists x_2)(Px_1x_2 \wedge Qx_1).$$

Consider the following partial $\mathrm{ic}_{\mathrm{SH}}$-tree (where α abbreviates the appropriate sequence of assumptions):

$$\begin{array}{c}
\mathbf{Y}_{\sigma} \\
\alpha, (\exists x_2)Pf_1x_2, Pf_1f_2?Pz_1z_2 \\
\exists\downarrow \qquad\qquad\qquad \mathbf{Y}_{\eta} \\
\alpha, (\exists x_2)Pf_1x_2?Pz_1z_2 \qquad\qquad \alpha, Qf_3?Qz_1 \\
\exists\downarrow \qquad\qquad\qquad \exists\downarrow \\
\alpha?Pz_1z_2 \qquad\qquad\qquad \alpha?Qz_1 \\
\wedge\uparrow \\
(\exists x_1)(\exists x_2)Px_1x_2, (\exists x)Qx?Pz_1z_2 \wedge Qz_1 \\
\exists\uparrow \\
(\exists x_1)(\exists x_2)Px_1x_2, (\exists x)Qx?(\exists x_2)(Pz_1x_2 \wedge Qz_1) \\
\exists\uparrow \\
(\exists x_1)(\exists x_2)Px_1x_2, (\exists x)Qx?(\exists x_1)(\exists x_2)(Px_1x_2 \wedge Qx_1)
\end{array}$$

Here, $\sigma = \langle f_1/z_1, f_2/z_2\rangle$ and $\eta = \langle f_3/z_1\rangle$. σ unifies Pf_1f_2 with Pz_1z_2 and η unifies Qf_3 with Qz_1, but applying both of these means that Pf_1f_2 and Qf_3 will be the premises to $\wedge\uparrow$ with a conclusion which should unify with $Pz_1z_2 \wedge Qz_1$. The only possible conclusion from these premises is $Pf_1f_2 \wedge Qf_3$, which does not unify with $Pz_1z_2 \wedge Qz_1$. Thus the failure of this tree to be a derivation is not determined at the leaves, but rather when the unification information is passed down the tree. Note also that in applying $\exists\downarrow$ to obtain Qf_3, we introduce a function symbol which is new to the entire tree, not just to the branch below the rule application. If we required only that these names were new to the given branch, we could have instantiated $(\exists x)Qx$ with Qf_1. Doing so would have allowed us to use $\eta' = \langle f_1/z_1\rangle$. The tree would then be a "derivation".

EXAMPLE 4. $(\forall x_1)(\exists x_2)Px_1x_2, (\exists x)Qx \vdash (\exists x_1)(\exists x_2)(Px_1x_2 \wedge Qx_1)$ The derivation is as follows:

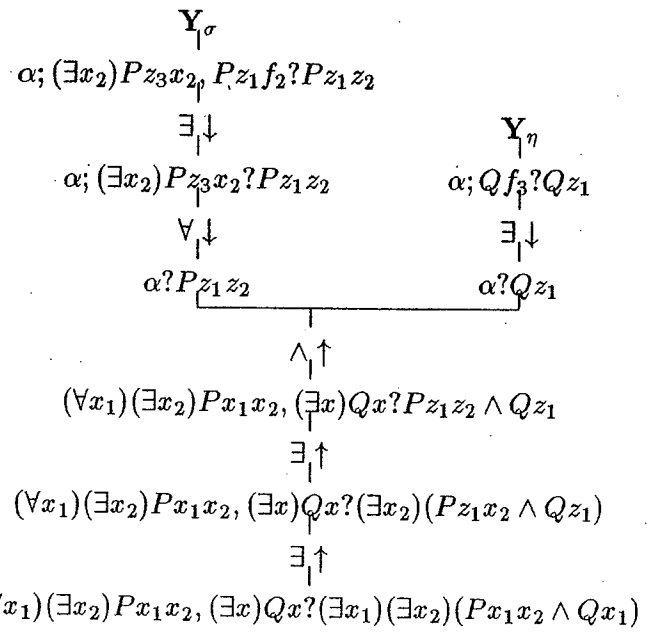

Here σ is $\langle z_3/z_1, f_2/z_2\rangle$, and η is $\langle f_3/z_1\rangle$. In this tree, when passing down the unifier information, we can "merge" σ and η to a substitution $\langle f_3/z_1,\ f_3/z_3,\ f_2/z_2\rangle$. This substitution will work for the part of the tree where the two branches merge.

The following considerations serve to make explicit the mechanism underlying the "merging" mentioned in example 4. For that purpose we review first some standard definitions as found, for example, in [Snyder]. We indicate the composition of two substitutions σ and ρ by $\sigma\rho$, with $\sigma\rho(w) = \sigma(\rho(w))$. A substitution σ is *idempotent* just in case $\sigma\sigma = \sigma$. Finally, $\rho \leq \sigma$ ("ρ is more general than σ" or, perhaps better, "ρ is less specific than σ") if and only if there is a substitution η, such that $\eta\rho = \sigma$. For idempotent substitutions we note that $\rho \leq \sigma$ implies $\sigma\rho = \sigma$. Finally, we come to our crucial definitions.

DEFINITION. (i) For substitutions σ_1 and σ_2, $\sigma = \sigma_1 \dot{\vee} \sigma_2$ is the least (with respect to $\leq$) substitution σ such that $\sigma_1 \leq \sigma$ and $\sigma_2 \leq \sigma$. σ is called the *join* of σ_1 and σ_2. Note that the join is not always defined, and that sometimes multiple substitutions may be joins for a given pair of substitutions (in which case we simply pick one). (ii) Substitutions σ and ρ are called *consistent* exactly when $\sigma \dot{\vee} \rho$ is defined.

Let us consider two examples. If $\sigma = \langle a/z_1, b/z_2\rangle$ and $\eta = \langle c/z_1\rangle$, then $\sigma[Pz_1z_2 \wedge Qz_3] = Pab \wedge Qz_3$ and $\eta[Pz_1z_2 \wedge Qz_3] = Pcz_2 \wedge Qz_3$, whereas

$\sigma\dot{\vee}\eta$ is undefined. Now let $\sigma = \langle z_3/z_1, b/z_2\rangle$ and $\eta = \langle c/z_1\rangle$, then $\sigma[Pz_1z_2 \wedge Qz_3] = Pz_3b \wedge Qz_3$ and $\eta[Pz_1z_2 \wedge Qz_3] = Pcz_2 \wedge Qz_3$. In this case $\sigma\dot{\vee}\eta = \langle c/z_1, b/z_2, c/z_3\rangle$ and $\sigma\dot{\vee}\eta[Pz_1z_2 \wedge Qz_3] = Pcb \wedge Qc$.

To summarize our discussion (through examples): When asking a question in the calculi described prior to $\mathrm{IC}_{\mathrm{SH}}$, closing a branch with $\mathbf{Y}$ guaranteed success along that branch, and succeeding for the whole proof required only that we succeed on sufficiently many branches to build a derivation. The above examples illustrate that closing off a branch with a unifier does not guarantee success on that branch, as it may cause the new free variables occurring in the branch to be instantiated with terms which are not consistent with the rest of the tree. Thus, a unifier gives us success only modulo its compatability with unifiers from other branches of the tree. Moreover, at any stage there may be multiple possible unifiers, any of which may or may not succeed further down the tree.

In order to keep track of all of these possibilities we modify the valuation function accordingly. We will introduce the value $\mathbf{Y}_\sigma$ for every idempotent substitution σ. Roughly speaking, a node N will be given the value $\mathbf{Y}_\sigma$ if applying σ to the subtree rooted at N will result in a tree having value $\mathbf{Y}$. To do this rigorously, we first introduce a means for "joining" *values*, so that the value for the whole tree can be determined — when the leaves have values, namely, sets of $\mathbf{Y}_\sigma$'s.

DEFINITION. Let $A = \{\mathbf{Y}_{\rho_1}, \ldots, \mathbf{Y}_{\rho_m}\}$, $B = \{\mathbf{Y}_{\sigma_1}, \ldots, \mathbf{Y}_{\sigma_n}\}$. $A\dot{\vee}B = \{\mathbf{Y}_{\rho_i\dot{\vee}\sigma_j} \mid 1 \leq i \leq m, 1 \leq j \leq n,$ and ρ_i and σ_j are consistent$\}$.

The earlier evaluation function $[\![N]\!]_\Sigma$ has to be modified, as sets of values are assigned to nodes. Let Σ be a partial $\mathrm{ic}_{\mathrm{SH}}$-tree and v_Σ the valuation for the leaves of Σ. In case N is a leaf, say $N = \alpha;\beta?G$, $[\![N]\!]_\Sigma = \{Y_\sigma \mid \sigma[G] \in \sigma[\alpha\beta]$ and σ idempotent$\}$. In case N is the unique successor of M, $[\![N]\!]_\Sigma = [\![M]\!]_\Sigma$; in case N is at a conjunctive branching, i.e., a rule node for a two-premise rule with M_1 and M_2 above N, $[\![N]\!]_\Sigma = [\![M_1]\!]_\Sigma\dot{\vee}[\![M_2]\!]_\Sigma$; finally, in case N is at a disjunctive branching, i.e., a question node with rule nodes $M_1, \ldots, M_k$ above N, $[\![N]\!]_\Sigma = \bigcup_{1\leq i\leq k}[\![M_i]\!]_\Sigma$. If Σ has (question node) N as its root, we set $[\![\Sigma]\!] = [\![N]\!]_\Sigma$.

6. Correctness of proof search (in the SH-expansion)

We consider the SH-expansion "just" as a convenient technical tool for automated proof search; thus, we ask basic questions $\alpha?G$ only when the elements of α and G are formulas in $\mathcal{L}_1$. Clearly, if we find an $\mathrm{ic}_{\mathrm{SH}}$-derivation for such a question, we want to associate with it an nd-proof in $\mathcal{L}_1$ of G from assumptions in α. This is immediate with p-normal or normal derivations, as

soon as we know how to transform partial $\mathrm{ic}_{\mathrm{SH}}$-trees into ic_1-trees. For this purpose, we define a *canonical renaming function* as follows:

DEFINITION. Let σ be a substitution and Π an $\mathrm{ic}_{\mathrm{SH}}$-tree. Let $\{t_1, \ldots, t_n\} = \mathcal{T}^*(\sigma[\Pi])$. Let $x_1, \ldots, x_n \in X$ be new for Π. R_σ^Π is the *tree-renaming generated by* Π *and* σ.

$$R_\sigma^\Pi[t] = \begin{cases} x_i & \text{if } \sigma[t] = t_i, 1 \leq i \leq n \\ \sigma[t] & \text{otherwise} \end{cases}$$

R_σ^Π distributes over relation symbols, sentential connectives, and quantifiers. It is applied to a question node by applying it to every formula at the node, and to a tree by applying it to every question node in the tree (as well as to the formulas displayed at $\perp$-rule nodes). We abbreviate $R_\sigma^\Pi[\Pi]$ by $R_\sigma[\Pi]$.

For such renamings we will show that they associate with partial $\mathrm{ic}_{\mathrm{SH}}$-trees partial ic_1-trees. Then we will show that an $\mathrm{ic}_{\mathrm{SH}}$-derivation exists for a given $\mathcal{L}_1$-question $\alpha?G$ if and only if an ic_1-derivation exists for $\alpha?G$. We add to our formulation of the natural deduction calculus a rule that allows for renaming of bound variables:

(R) $\dfrac{\phi}{\phi^*}$ if $\sigma[\phi^*] = \sigma[\phi]$ for every term-assignment σ.

We need to make the corresponding adjustment to IC_1, i.e., we close a branch with leaf node $\alpha;\beta?G$ with Y whenever $G \in \alpha\beta$ up to renaming of bound variables.

LOCAL CORRECTNESS LEMMA. *Let* T *be an* $\mathrm{ic}_{\mathrm{SH}}$*-tree for* $\alpha_0?G_0 - \alpha_0, G_0$ *in* $\mathcal{L}_1$. *Then for any substitution* σ, $R_\sigma[\mathrm{T}]$ *is an* ic_1*-tree for* $\alpha_0?G_0$.

PROOF. Fix $R = R_\sigma^{\mathrm{T}}$; by the definition of R, $R[\mathrm{T}]$ has root $\alpha_0?G_0$ (since the support of R is contained in $\mathcal{L}^*$). We show now by induction on the height of Π^* that for every subtree Π^* of T, $R[\Pi^*]$ is an ic_1-tree. Let $\alpha;\beta?G$ be the root of Π and let $\alpha';\beta'?G' = R[\alpha;\beta?G]$. If $\alpha;\beta?G$ is a leaf node then $R[\alpha;\beta?G]$ is a leaf in $\mathcal{L}_1$ and hence an ic_1-tree. For the inductive step, it will be enough to show: for each rule node r immediately above $\alpha;\beta?G$, the tree Π consisting of $\alpha;\beta?G$ and the subtree of Π^* with root r above it is mapped by R to an ic_1-tree. The tree Π is represented by

$$\frac{\Sigma_0 \quad \Sigma_1}{\alpha;\beta?G}$$

where we allow Σ_1 to be empty in case r has only one premise. $R[\Sigma_0]$ and $R[\Sigma_1]$ are ic_1-trees by the induction hypothesis. Let $\Sigma_0' = R[\Sigma_0]$ and

$\Sigma_1' = R[\Sigma_1]$. We proceed by case according to the rule r. Consider $\wedge\uparrow$; here Π is

$$\frac{\Sigma_0\left\{\begin{array}{c}\vdots\\ \alpha;\beta?\phi_1\end{array}\right. \quad \Sigma_1\left\{\begin{array}{c}\vdots\\ \alpha;\beta?\phi_2\end{array}\right.}{\alpha;\beta?\phi_1\wedge\phi_2}$$

If $R[\alpha;\beta?\phi_i] = \alpha';\beta'?\phi_i'$ for each $i = 1, 2$, then $R[\alpha;\beta?\phi_1\wedge\phi_2] = \alpha';\beta'?\phi_1'\wedge\phi_2'$ and $R[\Pi]$ is

$$\frac{\Sigma_0'\left\{\begin{array}{c}\vdots\\ \alpha';\beta'?\phi_1'\end{array}\right. \quad \Sigma_1'\left\{\begin{array}{c}\vdots\\ \alpha';\beta'?\phi_2'\end{array}\right.}{\alpha';\beta'?\phi_1'\wedge\phi_2'}$$

this is, by $\wedge\uparrow$, an ic_1-tree. All other propositional rules follow in a similar way. Now let us consider the quantifier rules. In the case of $\forall\downarrow$, Π is

$$\frac{\Sigma_0\left\{\begin{array}{c}\vdots\\ \alpha;\beta,\phi z?G\end{array}\right.}{\alpha;\beta?G}$$

where $(\forall x)\phi x \in \alpha\beta$. Let $t = R[z]$. If $\sigma[z] \in \mathcal{L}_1$, then $t = \sigma[z] \in \mathcal{L}_1$, by the definition of R; if not, then $t = x$ for some appropriate $x \in \mathcal{L}_1$. In either case, $t \in \mathcal{L}_1$ and $R[\Sigma_0] = \Sigma_0'$ is an ic_1-tree for $\alpha';\beta',\phi't?G'$. $R[\Pi]$ is then

$$\frac{\Sigma_0'\left\{\begin{array}{c}\vdots\\ \alpha';\beta',\phi't?G'\end{array}\right.}{\alpha';\beta'?G'}$$

where $(\forall x)\phi' x \in \alpha'\beta'$. By $\forall\downarrow$, $R[\Pi]$ is an ic_1-tree. The case for $\exists\uparrow$ is similar. Finally, consider $\exists\downarrow$; here Π is

$$\frac{\Sigma_0\left\{\begin{array}{c}\vdots\\ \alpha,\phi f\overline{z};\beta?G\end{array}\right.}{\alpha;\beta?G}$$

where $(\exists x)\phi x \in \alpha\beta$ and $f\overline{z} \notin \mathcal{T}_{\mathrm{SH}}(\alpha\beta, G)$. There is no $z \in \mathcal{T}_{\mathrm{SH}}(\alpha\beta, G)$ such that $\sigma[z] = \sigma[f\overline{z}]$, because any such z would be an argument of f.[21] Thus, $R[f\overline{z}] \notin \mathcal{T}_{\mathrm{SH}}(\alpha\beta, G)$. Say $R[f\overline{z}] = x$; $R[\Pi]$ is now

$$\frac{\Sigma_0'\left\{\begin{array}{c}\vdots\\ \alpha',\phi' x;\beta'?G'\end{array}\right.}{\alpha';\beta'?G'}$$

[21] Here one uses the fact that, given any term $f(z_1,\ldots,z_n)$, there is no substitution σ such that $\sigma[f(z_1,\ldots,z_n)] = \sigma[z_i]$ for any $1 \le i \le n$.

$R[\Pi]$ is an ic_1-tree by $\exists\downarrow$, since $(\exists x)\phi' x \in \alpha'\beta'$, and x is new to α', G'. The case for $\forall\uparrow$ is similar. ∎

Having associated with partial $\mathrm{ic}_{\mathrm{SH}}$-trees partial ic_1-trees, we show now that the "association" preserves global correctness in the sense of the following theorem:

VALUATION THEOREM. *Let* T *be an* $\mathrm{ic}_{\mathrm{SH}}$*-tree,* σ *a substitution.*

(i) *If* $\mathbf{Y}_\sigma \in [\![\mathrm{T}]\!]$, *then* $[\![R_\sigma[\mathrm{T}]]\!] = \mathbf{Y}$.

(ii) *If* $[\![\mathrm{T}]\!] = \emptyset$, *then* $[\![R_\sigma[\mathrm{T}]]\!] = \mathbf{U}$ *or* $\mathbf{N}$.

PROOF. Let T and σ be given as above; we use R as an abbreviation for R^{T}_σ. To establish the theorem it suffices to show by induction on Π that for every subtree Π of T (rooted at a question node):

$$\mathbf{Y}_\rho \in [\![\Pi]\!]_{\mathrm{T}} \text{ for some } \rho \leq \sigma \text{ if and only if } [\![R[\Pi]]\!]_{R[\mathrm{T}]} = \mathbf{Y}.$$

Note, that if $\mathbf{Y}_\rho \in [\![\Pi]\!]_{\mathrm{T}}$ for some $\rho \leq \sigma$, then $R[\Pi]$ is an ic_1-tree by the local correctness lemma. Let $N = \alpha;\beta?G$ be the root of Π. For the base case assume that N is a leaf node. If $\mathbf{Y}_\rho \in [\![\Pi]\!]_{\mathrm{T}}$ for some $\rho \leq \sigma$, then choose $\phi \in \alpha\beta$ such that $\rho[\phi] = \rho[G]$. $\sigma[\phi] = \sigma\rho[\phi] = \sigma[\rho[\phi]] = \sigma[\rho[G]] = \sigma\rho[G] = \sigma[G]$. Thus $R[G] \in R[\alpha\beta]$ up to renaming of bound variables, so $[\![R[\Pi]]\!] = \mathbf{Y}$. Conversely, assume $[\![R[\Pi]]\!] = \mathbf{Y}$. Then $R[G] \in R[\alpha\beta]$ up to renaming of bound variables, so $\mathbf{Y}_\rho \in [\![\alpha;\beta?G]\!]$.

In the inductive step, N is not a leaf node. For the $\Rightarrow$-direction, assume $\mathbf{Y}_\rho \in [\![N]\!]_{\mathrm{T}}$ for some $\rho \leq \sigma$. Let M be any node immediately above N such that $\mathbf{Y}_\rho \in [\![M]\!]$. (Such an M exists by the definition of $[\![N]\!]$.) If M has a single premise M_0, then $\mathbf{Y}_\rho \in [\![M_0]\!]$ by the definition of $[\![M]\!]$. By inductive hypothesis, $[\![R[M_0]]\!] = \mathbf{Y}$. Then $[\![R[M]]\!] = \mathbf{Y}$ and hence $[\![R[N]]\!] = \mathbf{Y}$ as well. If M has two premises, M_0 and M_1, then by the definition of $[\![M]\!]$, there are ρ_0 and ρ_1 such that $\mathbf{Y}_{\rho_0} \in [\![M_0]\!]$, $\mathbf{Y}_{\rho_1} \in [\![M_1]\!]$, and $\rho = \rho_0 \dot{\vee} \rho_1$. Since $\rho_0 \leq \rho \leq \sigma$, we have $\rho_0 \leq \sigma$ and, thus, $[\![R[M_0]]\!] = \mathbf{Y}$. Similarly we have $[\![R[M_1]]\!] = \mathbf{Y}$. But then $[\![R[M]]\!] = \mathbf{Y}$ and hence $[\![R[N]]\!] = \mathbf{Y}$.
For the $\Leftarrow$-direction, assume $[\![R[N]]\!] = \mathbf{Y}$, and choose rule node M immediately above N, such that $[\![R[M]]\!] = \mathbf{Y}$. If M has only one premise M_0, then $[\![R[M_0]]\!] = \mathbf{Y}$. By inductive hypothesis, $\mathbf{Y}_\rho \in [\![M_0]\!]_{\mathrm{T}}$ for some $\rho \leq \sigma$. $[\![N]\!] = [\![M]\!] = [\![M_0]\!]$, so we are done. If M has two premises M_0 and M_1, then $[\![R[M_0]]\!] = \mathbf{Y}$ and $[\![R[M_1]]\!] = \mathbf{Y}$ by the definition of $[\![R[M]]\!]$. By the inductive hypothesis, choose (for $i = 0, 1$) ρ_i such that $\mathbf{Y}_{\rho_i} \in [\![M_i]\!]$ and $\rho_i \leq \sigma$. $\mathbf{Y}_{\rho_0 \dot{\vee} \rho_1} \in [\![M]\!] \subseteq [\![N]\!]$. Since $\rho_0 \leq \sigma$ and $\rho_1 \leq \sigma$, $\rho_0 \dot{\vee} \rho_1 \leq \sigma$; so we are done. ∎

The following corollary to this theorem establishes the usefulness of the $\mathrm{IC_{SH}}$ calculus.

CORRECTNESS FOR $\mathrm{IC_{SH}}$. *Let α, G be in $\mathcal{L}_1$; then there is an $\mathrm{ic_{SH}}$-derivation for $\alpha?G$ if and only if there is an $\mathrm{ic_1}$-derivation for $\alpha?G$.*

Indeed, the nd-proofs that are then associated with $\mathrm{ic_1}$-derivations are, depending on the operation chosen in the $\perp$-rules, either p-normal or normal. The SH-expansion is thus a tool that provides correctly nd-proofs. In the introductory remarks to the previous section we mentioned that the SH-expansion is to be used for proof search, indeed, proof search that extends in a most natural way the strategic considerations for sentential logic — implemented in the Carnegie Mellon Proof Tutor. Those strategic considerations are described in [Sieg and Scheines]; here we review just the coarse structure of the (very efficient) search procedure. The search for an answer, i.e., an ic-derivation, to the question $\alpha;\beta?G$ involves three distinct components: (i) use of $\downarrow$-rules, (ii) use of $\uparrow$-rules, (iii) use of $\perp$-rules (with a limited set of contradictory pairs of formulas). It is step (i) that is central and taken in a *goal-directed way*. If the question

(*) Is G a strictly positive subformula of a formula in $\alpha\beta$?

has an affirmative answer, this step provides sequences of $\downarrow$-rule applications that *extract* G from strictly positive occurrences of G in elements of $\alpha\beta$. The connecting formula sequences consist of the major premises of the $\downarrow$-rules and require, in general, answers to new questions, namely, those raised in the minor premises of the rule applications.

It is for the appropriate generalization of this *extraction strategy* that the SH-expansion is absolutely critical. Recall that the question in sentential logic "Is G an element of $\alpha\beta$?" is generalized in predicate logic to "Is G unifiable with an element of $\alpha\beta$?". The goal-directedness of applications of $\downarrow$-rules is now obtained by generalizing the question (*) above to

Is G unifiable with a strictly positive canonical subformula of a formula in $\alpha\beta$?

A subformula is considered to be a *canonical* one, if quantifiers are instantiated by terms that match the $\downarrow$-quantifier rules of $\mathrm{IC_{SH}}$, i.e., those terms would be used, if the formula were "extractable" by $\downarrow$-rules. — This natural extension of the sentential logical search satisfies three important desiderata: (i) logical truths of sentential logic, e.g., instances of the law of excluded middle with complex formulas of $\mathcal{L}_1$, are recognized without appealing to quantificational rules; (ii) the selection of terms for $\forall\downarrow$ and $\exists\uparrow$ is delayed;

(iii) extractability is the central feature of the search. The details of our approach to automated proof search will be presented in a later publication (together with a discussion of benchmark examples).

7. So what?

This work is to address, ultimately, the question of finding proofs in mathematics with logical *and* mathematical understanding. If one looks at Georg Polya's writings on mathematical reasoning and heuristics one realizes quickly that his most general strategies for argumentation are simple logical ones. Clearly, logical formality per se does not facilitate the finding of proofs. Logic within a natural deduction framework does help, however, to bridge the gap between assumptions and conclusions by suggesting *very rough* structures for arguments, i.e. *logical structures* that depend solely on the syntactic form of assumptions and conclusions. This role of logic, though modest, is the starting-point for moving up to subject-specific considerations that support a theorem.

Proofs provide explanations of what they prove by putting their conclusions in a context that shows them to be correct. The deductive organization of parts of mathematics is *the* classical methodology for specifying such contexts. This methodology has two well-known aspects: the formulation of principles, i.e. axioms, and the reasoning from such principles; the latter is mediated through logical inferences and subject-specific lemmata. Heuristic considerations and "leading mathematical ideas" for particular parts of mathematics have to be found and properly articulated. Saunders MacLane (1934) suggested to include in the scope of logic such a *structure-theory of proofs*: this extension of the traditional role of logic and, in particular, of proof theory interacts directly and, we are convinced, fruitfully with a sophisticated, automated search for humanly intelligible proofs.

Appendix

In this appendix we give first a definition used at the end of section 3; then two diagrams are drawn that complement the text of sections 2 and 3.

Positive and *strictly positive subformulas* of a given formula are defined by induction; indeed, for the first concept one defines simultaneously, when ϕ is a *positive subformula* of $\psi[\phi \in \boldsymbol{P}(\psi)]$ or ϕ is a *negative subformula* of $\psi[\phi \in \boldsymbol{N}(\psi)]$, namely by the rules,

(i) ϕ is $\psi \Longrightarrow \phi \in \boldsymbol{P}(\psi)$

(ii) (a) ψ is $\neg\psi_1$, $\phi \in \boldsymbol{N}(\psi_1) \Longrightarrow \phi \in \boldsymbol{P}(\psi)$
(b) ψ is $\neg\psi_1$, $\phi \in \boldsymbol{P}(\psi_1) \Longrightarrow \phi \in \boldsymbol{N}(\psi)$

(iii) (a) ψ is $\psi_1 \wedge \psi_2$, $\phi \in \boldsymbol{P}(\psi_1) \cup \boldsymbol{P}(\psi_2) \Longrightarrow \phi \in \boldsymbol{P}(\psi)$
(b) ψ is $\psi_1 \wedge \psi_2$, $\phi \in \boldsymbol{N}(\psi_1) \cup \boldsymbol{N}(\psi_2) \Longrightarrow \phi \in \boldsymbol{N}(\psi)$

(iv) (a) ψ is $\psi_1 \vee \psi_2$, $\phi \in \boldsymbol{P}(\psi_1) \cup \boldsymbol{P}(\psi_2) \Longrightarrow \phi \in \boldsymbol{P}(\psi)$
(b) ψ is $\psi_1 \vee \psi_2$, $\phi \in \boldsymbol{N}(\psi_1) \cup \boldsymbol{N}(\psi_2) \Longrightarrow \phi \in \boldsymbol{N}(\psi)$

(v) (a) ψ is $\psi_1 \rightarrow \psi_2$, $\phi \in P(\psi_2) \Longrightarrow \phi \in \boldsymbol{P}(\psi)$
(b) ψ is $\psi_1 \rightarrow \psi_2$, $\phi \in \boldsymbol{N}(\psi_1) \Longrightarrow \phi \in \boldsymbol{P}(\psi)$
(c) ψ is $\psi_1 \rightarrow \psi_2$, $\phi \in \boldsymbol{P}(\psi_1) \cup \boldsymbol{N}(\psi_2) \Longrightarrow \phi \in \boldsymbol{N}(\psi)$

Finally, ϕ is a *strictly positive subformula* of $\psi[\phi \in \mathcal{S}(\psi)]$ if and only if it can be obtained by just the rules (i), (iii)(a), (iv)(a), and (v)(a).

Diagram 1 illustrates the construction of an ic-tree in an interesting case, namely the proof of *tertium non datur*. Diagram 2 illustrates the construction of canonical refutation branches discussed in section 3.

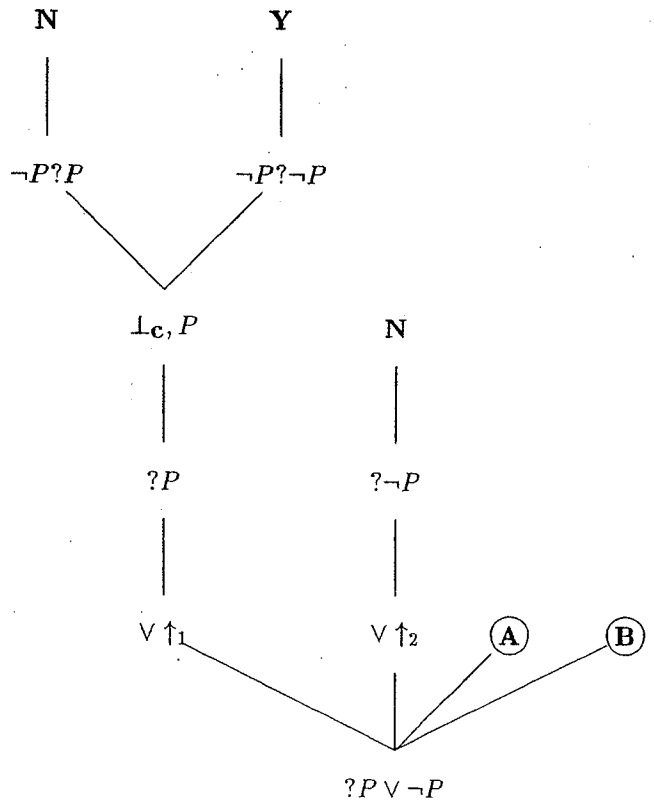

Diagram 1.

Diagrams 1.A and 1.B expand nodes A and B, respectively.

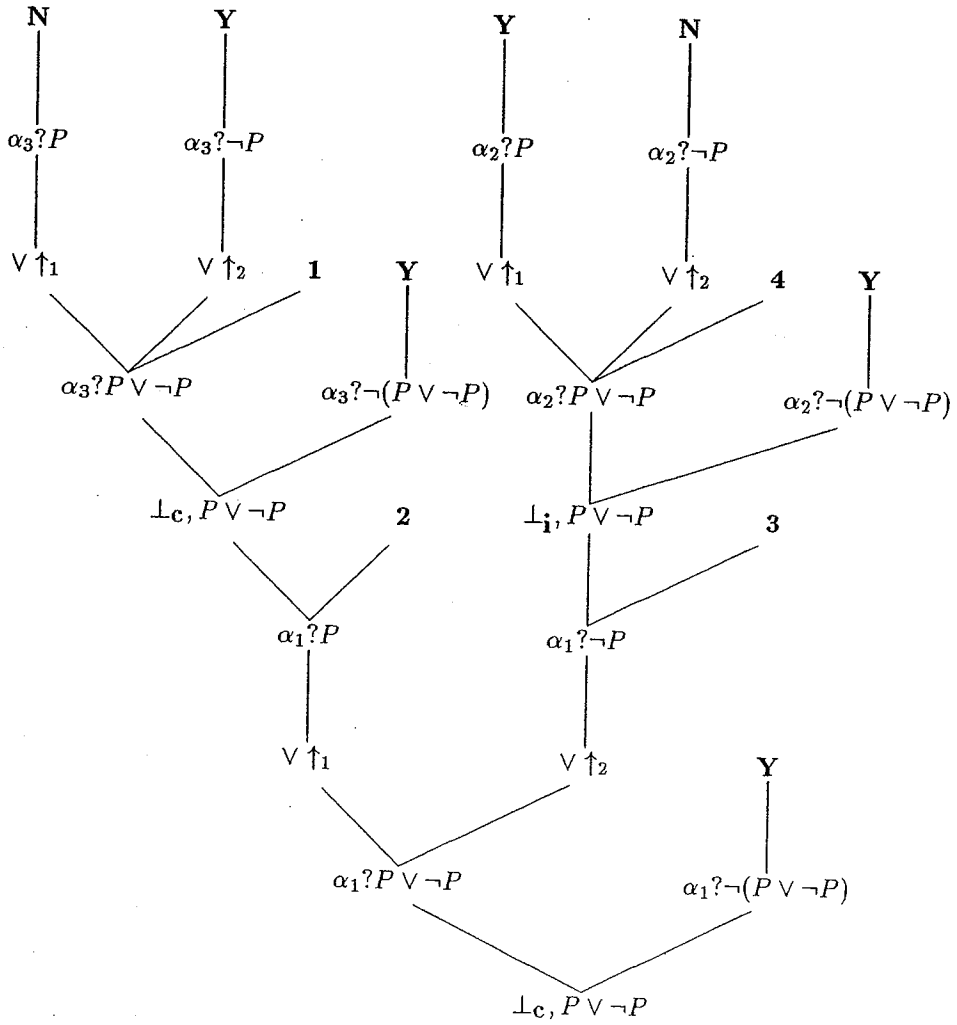

$\alpha_1 = \neg(P \vee \neg P)$
$\alpha_2 = \neg(P \vee \neg P), P$
$\alpha_3 = \neg(P \vee \neg P), \neg P$

Diagram 1.A.

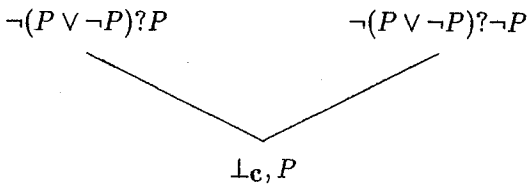

Diagram 1.B.

The question nodes above are expanded in exactly the same way as the corresponding questions in diagram 1.A.

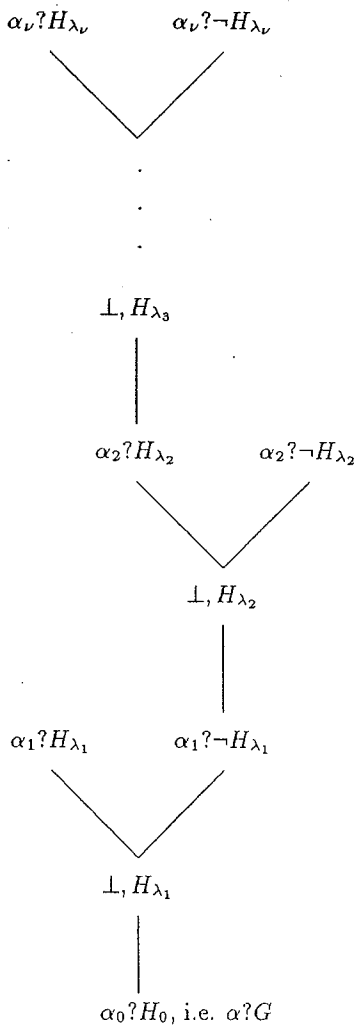

Diagram 2.

References

[1] ANDREWS, P., 'Transforming matings into natural deduction proofs', in *5th Conference on Automated Deduction*, 281–292, New York, Berlin, Springer-Verlag, 1980.

[2] BAAZ M., and C. G. FERMÜLLER, 'Non-elementary speedups between different versions of tableaux', in Baumgartner et al., editors, *Theorem Proving with Analytic Tableaux and Related Methods*, 217–230. Springer, 1995.

[3] BLEDSOE, W., 'Non-resolution theorem proving', *Artificial Intelligence* 9, 1977, 1–35.

[4] BLEDSOE, W., 'The UT natural-deduction prover', Technical Report, Departments of Mathematics and Computer Science, University of Texas, April 1983.

[5] CITTADINI, S., 'Intercalation calculus for intuitionistic propositional logic', Carnegie Mellon Technical Report PHIL-29, Philosophy, Methodology, and Logic, 1992.

[6] CONSTABLE, R. et al, *Implementing Mathematics with the Nuprl Proof Development System*, Prentice-Hall, Englewood Cliffs, NJ, 1986.

[7] FITCH, F. B., *Symbolic Logic: An Introduction*, Ronald, 1952.

[8] FITTING, M., *First-Order Logic and Automated Theorem Proving*, Springer-Verlag, New York, Berlin, 1990.

[9] GENTZEN, G., 'Untersuchungen über das logische Schließen i, ii', *Math. Zeitschrift* 39, 1934, 176–210, 1935, 405–431. English translation in [11].

[10] GENTZEN, G., 'Die Widerspruchsfreiheit der reinen Zahlentheorie', *Mathematische Annalen* 112, 1936, 493–565. English translation in [11].

[11] GENTZEN, G., *The Collected Papers of Gerhard Gentzen*, North-Holland Publishing Company, Amsterdam, edited by M. E. Szabo, 1969.

[12] HERBRAND, J., *Logical Writings*, Cambridge, Harvard University, edited by W. Goldfarb, 1971.

[13] JAŚKOWSKI, S., 'On the rules of suppositions in formal logic' *Studia Logica* (1), 1934.

[14] KLEENE, S. C., 1952, *Introduction to Metamathematics*, Wolters-Noordhoff Publishing, Groningen.

[15] MACLANE, S., *Abgekürzte Beweise im Logikkalkül*, PhD thesis, University of Göttingen, 1934.

[16] MACLANE, S., 'A logical analysis of mathematical structure', *The Monist*, 1935, 118–130. The paper was read to the American Mathematical Society on December 28, 1933.

[17] NEVINS, A. J., 'A human oriented logic for automatic theorem proving', *J. ACM* 21, 1974, 606–621.

[18] PAULSON, L., *Logic and Computation: Interactive Proof with Cambridge LCF*, Cambridge University Press, 1987.

[19] PELLETIER, J., 'Automated natural deduction in Thinker', this volume, 3–43.

[20] PFENNING, F., *Proof Transformations in Higher-Order Logic*, PhD thesis, Carnegie Mellon University, 1987.

[21] PORTORARO, F., 'Strategic construction of Fitch-style proofs', this volume, 45–66.

[22] PRAWITZ, D., *Natural Deduction: A Proof-Theoretical Study*, Almqvist & Wiskell, Stockholm, 1965.

[23] SHANIN, N. A. et al., 'An algorithm for a machine search of a natural logical deduction in a propositional calculus', in Siekmann and Wrightson, editors, *Automation of Reasoning*, vol. 1, 424–483, Springer-Verlag, New York, Berlin. Reprinted from Izdat. Nauka, Moscow, 1965.

[24] SIEG, W., *Mechanisms and Search: Aspects of Proof Theory*, Associazione Italiana di Logica e sue Applicazioni, Padova, 1992.

[25] SIEG, W., 'Intercalation calculi for classical logic', Carnegie Mellon Technical Report PHIL-46, Philosophy, Methodology, and Logic, 1994.

[26] SIEG, W., and B. KAUFFMANN, 'Unification for quantified formulae', Carnegie Mellon Technical Report PHIL-44, Philosophy, Methodology, and Logic, 1993.

[27] SIEG, W., and R. SCHEINES, 'Searching for proofs (in sentential logic)', in L. Burkholder, editor, *Philosophy and the Computer*, 137–159, Westview Press, Boulder, San Francisco, Oxford, 1992.

[28] SNYDER, W., *A Proof Theory for General Unification*, Birkhäuser, Boston, Basel, Berlin, 1991.

[29] STÅLMARCK, G., 'Normalization theorems for full first order classical natural deduction', *J. Symbolic Logic* 56, 1991, 129–149.

[30] STATMAN, R., *Structural Complexity of Proofs*, PhD thesis, Stanford, 1974.

[31] TROELSTRA, A. S., and D. VAN DALEN, *Constructivism in Mathematics: An Introduction*, North Holland, Amsterdam, 1988.

WILFRIED SIEG and JOHN BYRNES
Department of Philosophy
Carnegie Mellon University
Pittsburgh, PA 15213-3891