Epsilon Terms in Intuitionistic Sequent Calculus

Giselle Reis
Inria & LIX/École Polytechnique, France

Bruno Woltzenlogel Paleo
Australian National University, Canberra, Australia

Abstract

Skolemization is unsound in intuitionistic logic in the sense that a Skolemization $sk(F)$ of a formula F may be derivable in the intuitionistic sequent calculus **LJ** while F itself is not. This paper defines a transformation T_ε that differs from Skolemization only by its use of ε-terms instead of Skolem terms; and shows that, for a simple locally restricted sequent calculus $\mathbf{LJ}^\star$, this transformation is sound: if $T_\varepsilon(F)$ is derivable in $\mathbf{LJ}^\star$, then so is F.

1 Introduction

It is well-known that there are formulas whose Skolemizations are derivable in the intuitionistic sequent calculus **LJ** while the formulas themselves are not. Consequently, there exists no immediate method of de-Skolemization, i.e. a method to eliminate Skolem terms from intuitionistic proofs by introducing quantifiers without obtaining just classical proofs. The usual reaction to this fact is to conclude that Skolemization is intrinsically unsound in intuitionistic logic and, consequently, must be either avoided or modified in sophisticated ways [4, 9]. These approaches assume (quite naturally) that provability in **LJ** correctly captures validity in intuitionistic logic even in the presence of Skolem terms.

This paper explores a different approach that regards **LJ** as an unsound calculus for reasoning about formulas containing Skolem terms. From this perspective,

The authors would like to thank: Christian Retoré for organizing the workshop on *Hilbert's Epsilon and Tau in Logic, Informatics and Linguistics*; the anonymous reviewers of this workshop for providing useful feedback on an extended abstract of this paper; and Sergei Soloviev, for pointing out an interesting related work of Grigori Mints.

the reason why underivable formulas become derivable in **LJ** after they have been Skolemized is due to **LJ**'s inference rules being too permissive: they fail to recognize the special status of Skolem terms and allow them to be used in ways that should be forbidden. Therefore, the interesting question is not how to modify Skolemization in order to obtain an intuitionistically sound Skolemization-like transformation w.r.t. to **LJ**, but how to modify and restrict **LJ** so that Skolemization is sound w.r.t. the restricted calculus.

The main contribution of this paper is the design of a restricted sequent calculus $\mathbf{LJ}^\star$ for which *epsilonization* is sound: if $T_\varepsilon(S)$ (the epsilonization of the sequent S) is derivable in $\mathbf{LJ}^\star$, then so is S. In particular, we define a method of de-epsilonization of intuitionistic proofs transforming intuitionistic proofs with ε-terms into ordinary intuitionistic proofs. The transformation T_ε differs from Skolemization mainly in its use of Hilbert's ε-terms instead of Skolem terms. But in contrast to Hilbert's traditional ε-calculus, where all quantifiers are eliminated, T_ε eliminates only strong quantifiers. Skolem terms can be regarded as abbreviations of Hilbert's ε-terms [1]; conversely, ε-terms can be regarded as more informative Skolem terms. $\mathbf{LJ}^\star$ restricts the use of ε-terms in the instantiations performed by weak quantifier rules. The restrictions are local and purely syntactic; they use the extra information available in ε-terms but not in Skolem terms.

2 LJ and Epsilonization

We assume the reader is familiarized with the language of first-order logic. The rules of **LJ** are depicted in Figures 1 and 2. $\forall_l$ and $\exists_r$ are called *weak quantifier rules*, while $\forall_r$ and $\exists_l$ are called *strong quantifier rules*. $\forall$-quantifiers of positive polarity and $\exists$-quantifiers of negative polarity are called *strong quantifiers*.

$$\frac{\Gamma_1, A \vdash F \quad \Gamma_2, B \vdash F}{\Gamma_1, \Gamma_2, A \vee B \vdash F} \vee_l \qquad \frac{\Gamma \vdash A}{\Gamma \vdash A \vee B} \vee_r^1 \qquad \frac{\Gamma \vdash B}{\Gamma \vdash A \vee B} \vee_r^2 \qquad \frac{\Gamma \vdash A}{\Gamma, \neg A \vdash} \neg_l \qquad \frac{\Gamma, A \vdash}{\Gamma \vdash \neg A} \neg_r$$

$$\frac{\Gamma, A, B \vdash F}{\Gamma, A \wedge B \vdash F} \wedge_l \qquad \frac{\Gamma_1 \vdash A \quad \Gamma_2 \vdash B}{\Gamma_1, \Gamma_2 \vdash A \wedge B} \wedge_r \qquad \frac{\Gamma_1 \vdash A \quad \Gamma_2, B \vdash F}{\Gamma_1, \Gamma_2, A \to B \vdash F} \to_l \qquad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \to B} \to_r$$

$$\frac{}{A \vdash A} a \ (A \text{ is atomic}) \qquad \frac{\Gamma \vdash F}{\Gamma, A \vdash F} w_l \qquad \frac{\Gamma \vdash}{\Gamma \vdash A} w_r \qquad \frac{\Gamma, A, A \vdash F}{\Gamma, A \vdash F} c_l \qquad \frac{\Gamma_1 \vdash A \quad \Gamma_2, A \vdash F}{\Gamma_1, \Gamma_2 \vdash F} cut$$

Figure 1: Propositional and Structural Rules for **LJ**

$$\frac{\Gamma, A[t] \vdash F}{\Gamma, \forall x.A[x] \vdash F}\ \forall_l \qquad \frac{\Gamma \vdash A[\alpha]}{\Gamma \vdash \forall x.A[x]}\ \forall_r \qquad \frac{\Gamma, A[\alpha] \vdash F}{\Gamma, \exists x.A[x] \vdash F}\ \exists_l \qquad \frac{\Gamma \vdash A[t]}{\Gamma \vdash \exists x.A[x]}\ \exists_r$$

where:

- α must satisfy the eigenvariable condition.

Figure 2: Quantifier Rules for **LJ**

Skolemization is a transformation that removes all strong quantifiers from first-order formulas and replaces the variables they quantify by *Skolem terms*. There are various Skolemization methods, which may differ in the proof complexity of the transformed formula [5]. To see that Skolemization does not preserve derivability in the sequent calculus **LJ**, consider the formula $\neg\forall x.P(x) \rightarrow \exists y.\neg P(y)$, in which the $\forall$ quantifier is strong (note that it would be introduced by a $\forall_r$ inference in a sequent calculus proof). While it is clear that $\nvdash_{\mathbf{LJ}} \neg\forall x.P(x) \rightarrow \exists y.\neg P(y)$, the proof below shows that its Skolemization $\neg P(s) \rightarrow \exists y.\neg P(y)$ (where s is a skolem constant) is derivable:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{}{Ps \vdash Ps}}{Ps, \neg Ps \vdash}\ \neg_l}{\neg Ps \vdash \neg Ps}\ \neg_r}{\neg Ps \vdash \exists y.\neg Py}\ \exists_r}{\vdash \neg Ps \rightarrow \exists y.\neg Py}\ \rightarrow_r$$

In this example, the use of s on the weak quantifier rule could be avoided if we had more information about it. In order to obtain more informative terms, we choose to use ε-terms instead of Skolem terms for replacing the strongly quantified variables of a formula.

ε-terms[1] are formed with two binders: ε and τ. The intended meaning of ε-terms is delimited by the following *epsilon axioms*:

$$\exists x.A[x] \rightarrow A[\varepsilon_x A[x]] \quad \text{and} \quad A[\tau_x A[x]] \rightarrow \forall x.A[x]$$

In classical logic, the following equivalences hold, and hence τ is definable using ε:

$$A[\tau_x A[x]] \leftrightarrow \forall x.A[x] \leftrightarrow \neg\exists x.\neg A[x] \leftrightarrow \neg\neg A[\varepsilon_x \neg A[x]] \leftrightarrow A[\varepsilon_x \neg A[x]]$$

[1]We assume that the usual inductively defined *terms* of first-order logic are extended to include ε-terms. Hence, in general, a term may or may not contain ε-binders. ε-terms, on the other hand, are assumed to have ε or τ binders as their outermost symbols. Therefore, every ε-term is a term, but not every term is an ε-term.

In intuitionistic logic, however, the equivalences above do not hold. Therefore, both binders are needed. Epsilonization is analogous to Skolemization, but it uses ε-terms instead of Skolem terms.

Definition 1 (Epsilonization). *An* epsilonization $T_\varepsilon(F)$ *of a formula F is defined inductively on the structure of F using two functions T_ε^+ and T_ε^-. On the definitions below, $p \in \{+,-\}$ and $\bar{p}$ is $+$ if $p = -$ and $-$ if $p = +$.*

$$
\begin{array}{rcl}
T_\varepsilon(F) & = & T_\varepsilon^+(F) \\
T_\varepsilon^p(A) & = & A \text{ if } A \text{ is atomic.} \\
T_\varepsilon^p(\neg A) & = & \neg T_\varepsilon^{\bar{p}}(A) \\
T_\varepsilon^p(A \wedge B) & = & T_\varepsilon^p(A) \wedge T_\varepsilon^p(B) \\
T_\varepsilon^p(A \vee B) & = & T_\varepsilon^p(A) \vee T_\varepsilon^p(B) \\
T_\varepsilon^p(A \to B) & = & T_\varepsilon^{\bar{p}}(A) \to T_\varepsilon^p(B) \\
T_\varepsilon^+(\exists x.A) & = & \exists x.T_\varepsilon^+(A) \\
T_\varepsilon^+(\forall x.A) & = & A'\{x \mapsto \tau_x A'\} \text{ for } A' = T_\varepsilon^+(A) \\
T_\varepsilon^-(\forall x.A) & = & \forall x.T_\varepsilon^-(A) \\
T_\varepsilon^-(\exists x.A) & = & A'\{x \mapsto \varepsilon_x A'\} \text{ for } A' = T_\varepsilon^-(A)
\end{array}
$$

Definition 2 (Epsilonization of sequents). *The* epsilonization $T_\varepsilon(S)$ *of a sequent S of the form $A_1, \ldots, A_n \vdash B_1, \ldots, B_m$ is a sequent of the form $T_\varepsilon^-(A_1), \ldots, T_\varepsilon^-(A_n) \vdash T_\varepsilon^+(B_1), \ldots, T_\varepsilon^+(B_m)$.*

In Skolemization one needs to explicitly keep track of weakly quantified variables in order to add them as arguments of the Skolem function. In epsilonization such book-keeping is not needed. Since the whole formula will be a sub-expression of the ε-term, the weakly quantified variables will occur naturally in the term. In contrast to what is done in Hilbert's ε-calculus [1], the epsilonization procedure defined here does not eliminate the weak quantifiers; therefore ε-terms may contain quantified formulas. Like in the standard ε-calculus, innermost strong quantifiers are removed first. Using this strategy, strong quantifiers will never occur inside an ε-term. Instead, it will contain *nested* ε-terms corresponding to the variables that were bound by those strong quantifiers. An ε-term that is not nested inside another ε-term is a *top-level* ε-term.

Example 1. *Consider the formula $\forall x.\exists y.\exists z.P(x,y,z)$. In a negative context, its epsilonization would be:*

$$\forall x.P(x, \varepsilon_y P(x, y, \varepsilon_z P(x, y, z)), \varepsilon_z P(x, \varepsilon_y P(x, y, \varepsilon_z P(x, y, z)), z))$$

As desired, the weakly quantified variable x naturally occurs inside the ε-terms for y and z. The weak quantifier $\forall x$ remained. The innermost strong quantifier $\exists z$ within the scope of the strong quantifier $\exists y$ resulted in an ε-term for y (i.e. $\varepsilon_y P(x, y, \varepsilon_z P(x, y, z))$) containing a nested ε-term for z (i.e. $\varepsilon_z P(x, y, z)$) as a sub-term. The ε-terms $\varepsilon_y P(x, y, \varepsilon_z P(x, y, z))$ and $\varepsilon_z P(x, \varepsilon_y P(x, y, \varepsilon_z P(x, y, z)), z)$ are top-level ε-terms in the formula above. Comparing the epsilonization with a Skolemization of the same formula, such as $\forall x.P(x, sk_y(x), sk_z(x))$, the Skolem terms $sk_y(x)$ and $sk_z(x)$ can be seen as abbreviations for the two top-level ε-terms.

The treatment of strong quantifiers from inside out is compatible to our principal aim: the epsilonization of proofs. Since this procedure (presented in Definition 5) traverses the proof from the axioms to the end-sequent, innermost quantifiers are treated first. The motivation for removing strong quantifier inferences from proofs is due to the CERES method for intuitionistic logic [6, 10], a cut-elimination procedure based on the resolution calculus. To apply this method, the proof must not contain strong quantifier inferences on end-sequent ancestors. This is easily accomplished for classical logic via Skolemization (as we can eventually de-Skolemize the constructed cut-free proof), but it is not straightforward for intuitionistic proofs.

3 LJ*: a restricted LJ

We now define **LJ***, a version of **LJ** with restricted weak quantifier rules, which uses information available in the ε-terms to decide if they can be used on the instantiation of weak quantifiers. In what follows we will use ν to denote any of the ε-binders ε or τ, and $\rightsquigarrow$ as a rewriting relation.

Definition 3. *A term t is* accessible *in a formula F iff:*

- *for any top-level ε-term $\nu_x G$ in t it is the case that $F[\nu_x G \rightsquigarrow x]$ is a sub-formula of G; or*
- *t contains a nested ε-term $\nu_y H$ such that $\nu_y H$ is accessible in F and $t[\nu_y H \rightsquigarrow y]$ is accessible in $F[\nu_y H \rightsquigarrow y]$.*

The recursion in Definition 3 is necessary for coping with arbitrarily nested ε-terms.

Example 2. *Consider the formula F below:*

$$P(w, \varepsilon_y P(w, y, \varepsilon_z P(w, y, z)), \varepsilon_z P(w, \varepsilon_y P(w, y, \varepsilon_z P(w, y, z)), z))$$

Let t_1 be the term $\varepsilon_y P(w, y, \varepsilon_z P(w, y, z))$. The term t_1 is accessible in F, because $F[t_1 \rightsquigarrow y]$ (which is equal to $P(w, y, \varepsilon_z P(w, y, z))$) is a sub-formula of G_1 (where

G_1 is, in accordance with Definition 3, $P(w, y, \varepsilon_z P(w, y, z)))$. Let t_2 be the term $\varepsilon_z P(w, \varepsilon_y P(w, y, \varepsilon_z P(w, y, z)), z)$. The term t_2 is accessible in F, because t_1 is accessible in F and $t_2[t_1 \rightsquigarrow y]$ (which is $\varepsilon_z P(w, y, z)$) is accessible in $F[t_1 \rightsquigarrow y]$ (which is $P(w, y, \varepsilon_z P(w, y, z)))$, since $F[t_1 \rightsquigarrow y][t_2[t_1 \rightsquigarrow y] \rightsquigarrow z]$ (which is $P(w, y, z)$) is a sub-formula of G_2 (where G_2 is the formula under the scope of the ε-binder in $t_2[t_1 \rightsquigarrow y]$: $P(w, y, z)$).

Definition 4. *A term t is* accessible *in a sequent S iff all top-level ε-terms in t are accessible in some formula occurring in S.*

When thinking about bottom-up proof search, a term is accessible only after the strong quantifier inference introducing its corresponding eigenvariable in a regular **LJ** proof is applied. This means that, at this point, the term (or the eigenvariable) is already available for use in a weak quantifier inference. Take our previous unprovable sequent: $\vdash \neg\forall x.Px \to \exists x.\neg Px$. As shown before, its Skolemization is provable in **LJ** because the skolem term used for $\forall x.Px$ is available to be used in $\exists x.\neg Px$. The epsilonization of this sequent is: $\vdash \neg P(\tau_x Px) \to \exists x.\neg Px$. The fact that Px is a sub-formula of $\neg Px$ informs us that the strong quantifier was within the scope of the negation, and therefore a negation inference would have to be applied in order to make the ε-term accessible before it could be used in a weak quantifier inference. Therefore, as desired, the epsilonized sequent is not provable in **LJ***.

Additionally, the ε-terms used in this calculus will contain labels. The purpose of these labels is two-fold.

Firstly, they will restrict the shape of the proofs in **LJ*** in order to make de-epsilonization possible. Without the restriction, the removal of ε-terms and re-introduction of strong quantifiers could generate incorrect **LJ** proofs that violate the eigenvariable condition. Take, for example, the following proof of the epsilonization of $\neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash$:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\overline{P(\tau_x\neg Px) \vdash P(\tau_x\neg Px)} \quad \overline{P(\tau_x\neg Px) \vdash P(\tau_x\neg Px)}}{P(\tau_x\neg Px), P(\tau_x\neg Px) \vdash P(\tau_x\neg Px) \wedge P(\tau_x\neg Px)}\wedge_r}{P(\tau_x\neg Px), P(\tau_x\neg Px), \neg(P(\tau_x\neg Px) \wedge P(\tau_x\neg Px)) \vdash}\neg_l}{P(\tau_x\neg Px), \neg(P(\tau_x\neg Px) \wedge P(\tau_x\neg Px)) \vdash \neg P(\tau_x\neg Px)}\neg_r}{P(\tau_x\neg Px), \neg\neg P(\tau_x\neg Px), \neg(P(\tau_x\neg Px) \wedge P(\tau_x\neg Px)) \vdash}\neg_l}{P(\tau_x\neg Px), \neg\neg P(\tau_x\neg Px), \forall y.\neg(P(\tau_x\neg Px) \wedge Py) \vdash}\forall_l}{P(\tau_x\neg Px), \neg\neg P(\tau_x\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\forall_l}{\neg\neg P(\tau_x\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash \neg P(\tau_x\neg Px)}\neg_r}{\neg\neg P(\tau_x\neg Px), \neg\neg P(\tau_x\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\neg_l}{\neg\neg P(\tau_x\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}c_l
$$

When de-epsilonizing, two strong quantifiers need to be introduced in this proof; both of them between $\neg_l$ and $\neg_r$ inferences: one in the second/third level and the other in the sixth/seventh, bottom-up. The proof with the strong quantifiers is:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P(\alpha)\vdash P(\alpha)} \quad \dfrac{}{P(\beta)\vdash P(\beta)}}{P(\alpha),P(\beta)\vdash P(\alpha)\wedge P(\beta)}\;\wedge_r}{P(\alpha),P(\beta),\neg(P(\alpha)\wedge P(\beta))\vdash}\;\neg_l}{P(\alpha),\neg(P(\alpha)\wedge P(\beta))\vdash \neg P(\beta)}\;\neg_r}{P(\alpha),\neg(P(\alpha)\wedge P(\beta))\vdash \forall x.\neg P(x)}\;\forall_r *}{P(\alpha),\neg\forall x.\neg P(x),\neg(P(\alpha)\wedge P(\beta))\vdash}\;\neg_l}{P(\alpha),\neg\forall x.\neg P(x),\forall y.\neg(P(\alpha)\wedge Py)\vdash}\;\forall_l}{P(\alpha),\neg\forall x.\neg P(x),\forall z.\forall y.\neg(Pz\wedge Py)\vdash}\;\forall_l}{\neg\forall x.\neg P(x),\forall z.\forall y.\neg(Pz\wedge Py)\vdash\neg P(\alpha)}\;\neg_r}{\neg\forall x.\neg P(x),\forall z.\forall y.\neg(Pz\wedge Py)\vdash\forall x.\neg P(x)}\;\forall_r}{\neg\forall x.\neg P(x),\neg\forall x.\neg P(x),\forall z.\forall y.\neg(Pz\wedge Py)\vdash}\;\neg_l}{\neg\forall x.\neg P(x),\forall z.\forall y.\neg(Pz\wedge Py)\vdash}\;c_l
$$

Note that the top-most $\forall_r *$ inference violates the eigenvariable condition. In fact, as this rule is applied after (above) both weak quantifiers, a violation is unavoidable. The only way of de-epsilonizing the proof into a valid **LJ** proof would be to perform more complex operations, such as re-ordering of inferences. Instead of pursuing a more complicated de-epsilonization procedure, we restrict proof search in $\mathbf{LJ}^\star$ by using labels and avoiding the construction of such proofs in the first place. The restriction still preserves completeness.

Secondly, the labels will make epsilonization of **LJ** proofs an injective function. If labels were not used, the two following derivations would map to the same one:

$$
\dfrac{\dfrac{\dfrac{P\alpha, P\beta\vdash}{P\alpha,\exists x.Px\vdash}\;\exists_l}{\exists x.Px,\exists x.Px\vdash}\;\exists_l}{\exists x.Px\vdash}\;c_l
\qquad
\dfrac{\dfrac{P\alpha, P\alpha\vdash}{P\alpha\vdash}\;c_l}{\exists x.Px\vdash}\;\exists_l
$$

$$
\curlyvee
$$

$$
\dfrac{P(\varepsilon_x Px), P(\varepsilon_x Px)\vdash}{P(\varepsilon_x Px)\vdash}\;c_l
$$

Figure 3 shows the inferences of $\mathbf{LJ}^\star$ that are different than those of **LJ**, all others remain the same. The labels in ε-terms can be variables or constants. When epsilonizing a formula according to Definition 1, each ε-term receives a different label variable. When using $\mathbf{LJ}^\star$ for proof search, the following conditions must be enforced:

- On the initial rule, the corresponding ε-terms in the antecedent and consequent must have the same labels, and these must be constants.
- On the weak quantifier rules, the term t used for the substitution must be accessible and, *additionally*, its ε-subterms must have constants as labels. If this is not the case, the label variables of the ε-terms in t that occur in accessible positions in the conclusion sequent are substituted by a new (fresh) constant.
- Upon contracting a formula with ε-terms that have a variable label, there are two cases:
  - For accessible ε-terms, the same variable is used in the contracted occurrences in the premise.
  - For inaccessible ε-terms, new variable labels are created to be used in the contracted occurrences in the premise[2].

  If the label is a constant, then it was already used by a weak quantifier inference below contraction, which means the term is accessible. In this case, the constant label is simply copied to the contracted occurrences in the premise.

$$\frac{\Gamma, A[t] \vdash F}{\Gamma, \forall x.A[x] \vdash F}\ \forall_l' \qquad \frac{\Gamma \vdash A[t]}{\Gamma \vdash \exists x.A[x]}\ \exists_r' \qquad \frac{}{A[\nu_x^l\ F] \vdash A[\nu_x^l\ F]}\ a$$

where:

- the term t must be accessible in the conclusion sequent (*accessibility condition*).
- accessible occurrences of t or any of its ε-subterms in Γ and F must have a constant as a label (*label condition*).
- l is a constant in a (*initial condition*).

Figure 3: Rules for $\mathbf{LJ}^\star$

One might wonder about the (im)possibility to devise a simpler treatment of labels or stronger restrictions on contracted formulas in order to avoid the problems shown before. An immediate thought would be to use always constant labels and force contraction to create two different labels on the premises. But this restriction is too strong and would render the calculus incomplete, if it were adopted. The sequent

[2]This means that proofs in $\mathbf{LJ}^\star$ might contract formulas with different labels in its ε-terms.

$\exists x.Px \vdash \exists x.(Px \wedge Px)$ is an example. Its epsilonization is $P(\varepsilon_x Px) \vdash \exists x.(Px \wedge Px)$ and a proof in $\mathbf{LJ}^\star$ is shown below:

$$\dfrac{\dfrac{\dfrac{\overline{P(\varepsilon_x Px) \vdash P(\varepsilon_x Px)} \quad \overline{P(\varepsilon_x Px) \vdash P(\varepsilon_x Px)}}{P(\varepsilon_x Px), P(\varepsilon_x Px) \vdash P(\varepsilon_x Px) \wedge P(\varepsilon_x Px)}\,\wedge_r}{P(\varepsilon_x Px), P(\varepsilon_x Px) \vdash \exists x.(Px \wedge Px)}\,\exists_r}{P(\varepsilon_x Px) \vdash \exists x.(Px \wedge Px)}\,c_l$$

If different labels were used when contracting, the sequent would not be provable.

Another simpler potential solution would be to restrict contraction to formulas that only have accessible ε-terms. Unfortunately, this does not work in the general case. Consider the sequent $\neg(\forall x.Px \vee \neg\forall x.Px) \vdash$, whose epsilonization is $\neg(P(\tau_x Px) \vee \neg\forall x.Px) \vdash$. The term $\tau_x Px$ is obviously not accessible, thus should contraction on this formula not be allowed, the sequent would not be provable (whereas the original sequent is intuitionistically valid).

Theorem 1 (Soundness)**.** *For an ε-free formula F, if $\vdash_{\boldsymbol{LJ}^\star} F$ then $\vdash_{\boldsymbol{LJ}} F$.*

Proof. Let ψ' be an $\mathbf{LJ}^\star$-proof of F. Then an $\mathbf{LJ}$-proof ψ of F can be constructed simply by replacing $\forall_l'$ and $\exists_r'$ inferences by, respectively, $\forall_l$ and $\exists_r$ inferences. Since F is ε-free, the rules a and c_l are the same as those in $\mathbf{LJ}$. $\square$

$\mathbf{LJ}^\star$ is also sound relative to $\mathbf{LJ}$ for formulas with ε-terms (i.e., if $\vdash_{\mathbf{LJ}^\star} T_\varepsilon(F)$ then $\vdash_{\mathbf{LJ}} T_\varepsilon(F)$). We simply need to ignore the labels when transforming the proof.

Theorem 2 (Completeness)**.** *For an ε-free formula F, if $\vdash_{\boldsymbol{LJ}} F$ then $\vdash_{\boldsymbol{LJ}^\star} F$.*

Proof. Since F is ε-free and $\vdash_{\mathbf{LJ}} F$, there is an ε-free $\mathbf{LJ}$-proof ψ of F. An $\mathbf{LJ}^\star$-proof ψ' of F can be constructed simply by replacing all $\forall_l$ and $\exists_r$ inferences by, respectively, $\forall_l'$ and $\exists_r'$ inferences. No accessibility or label violation occurs, because no term in ψ' contains ε-terms. Also, the conditions for the inferences a and c_l are not violated, for the same reason. $\square$

The *epsilonization* of a proof removes the strong quantifier inferences that operate on ancestors of formulas occurring in the end-sequent and replaces the corresponding eigenvariables by ε-terms.

Definition 5 (Epsilonization of proofs)**.** *Let ψ be an $\boldsymbol{LJ}^\star$ proof of an ε-free sequent S. We define $T_\varepsilon(\psi)$, an $\boldsymbol{LJ}^\star$ proof of $T_\varepsilon(S)$, inductively on the inference rules.*

Base case: *ψ consists of only one axiom. Then $T_\varepsilon(\psi) = \psi$.*

Step case: *ψ ends with an inference ρ, as in the following cases.*

- *ρ is $\forall_r$ or $\exists_l$ applied to an end-sequent ancestor.*

  Let $(Qx)F$ be the main formula, ψ' be the proof of ρ's premise and α the eigenvariable used to instantiate the strongly quantified variable x. By induction hypothesis, $T_\varepsilon(\psi')$ is well defined. Then $T_\varepsilon(\psi)$ is $T_\varepsilon(\psi')\{\alpha \mapsto \nu^l_x\ F\}$, where ν is ε if Q is $\exists$ and τ if Q is $\forall$, and l is a fresh constant label.

  Note that strong quantifiers that go to cut-formulas are not replaced.

- *ρ is $\forall_l$ or $\exists_r$ applied to an end-sequent ancestor.*

  Let $(Qx)Fx$ be the main formula, Ft the auxiliary formula and ψ' the proof of ρ's premise. By induction hypothesis, $T_\varepsilon(\psi')$ is well defined. Then $T_\varepsilon(\psi)$ is $T_\varepsilon(\psi')$ plus the inference $\forall_l$ or $\exists_r$ (depending whether Q is $\forall$ or $\exists$) which introduces the quantifier and replaces t by x in F, including the occurrences of t inside copies of F occurring in ε-terms. The variable x used may not be bound.

- *ρ is c_l applied to an end-sequent ancestor, and the formulas contracted contain ε-terms with labels.*

  Let ψ' be the proof of ρ's premise. By induction hypothesis, $T_\varepsilon(\psi')$ is well defined. Then $T_\varepsilon(\psi)$ is $T_\varepsilon(\psi')$ plus the contraction, where its main formula will have new variables as labels.

- *ρ is another inference. Then $T_\varepsilon(\psi) = \psi$.*

Observe that, apart from possibly different labels, contraction will always operate on equal terms, since weak quantifiers also operate on formulas inside ε-terms:

$$\dfrac{\dfrac{\dfrac{\begin{array}{c}\psi\\ P(a,\alpha), P(b,\beta) \vdash\end{array}}{\exists y.P(a,y), \exists y.P(b,y) \vdash}\ \exists_l \times 2}{\forall x.\exists y.P(x,y), \forall x.\exists y.P(x,y) \vdash}\ \forall_l \times 2}{\forall x.\exists y.P(x,y) \vdash}\ c_l \quad \rightsquigarrow \quad \dfrac{\dfrac{\begin{array}{c}T_\varepsilon(\psi)\\ P(a,\varepsilon^{l_1}_y.P(a,y)), P(b,\varepsilon^{l_2}_y.P(b,y)) \vdash\end{array}}{\forall x.P(x,\varepsilon^{l_1}_y.P(x,y)), \forall x.P(x,\varepsilon^{l_2}_y.P(x,y)) \vdash}\ \forall_l \times 2}{\forall x.P(x,\varepsilon^{l}_y.P(x,y)) \vdash}\ c_l$$

Lemma 1. *If an **LJ*****-proof ψ has end-sequent S, then $T_\varepsilon(\psi)$ has end-sequent $T_\varepsilon(S)$ (modulo renaming of labels).*

Proof. By induction on the structure of ψ and by Definition 5. □

Lemma 2. *If ψ is an **LJ*****-proof of S, then no weak quantifier inference in $T_\varepsilon(\psi)$ violates the accessibility condition.*

Proof. First of all, note that the order in which the inferences are applied in $T_\varepsilon(\psi)$ is the same as in ψ, with the only difference being that strong quantifier inferences were removed.

Let $Qx.F$ be a strong quantified formula in S, α the eigenvariable used for this strong quantifier in ψ and $Q'x.G$ a weak quantifier in S which is instantiated in ψ with a term containing α. Since ψ is a correct proof, the weak quantifier inference ρ_w on $Q'x.G$ occurs after (above) the strong quantifier inference ρ_s on $Qx.F$.

Now consider the proof $T_\varepsilon(\psi)$. Given Definition 5, at the point where ρ_s was applied, the formula $Qx.F$ will have the shape $F'[\nu^l_x\ F']$, where F' is possibly F without strong quantifiers. Since F' is a sub-formula of F', the ε-term is already accessible. All inferences above this point will either decompose (the outer-most) F' into more sub-formulas or keep it unchanged. In this way, the ε-term $\nu^l_x\ F'$ will remain accessible. As ρ_w is applied after (above) the considered point, the accessibility relation will not be violated. □

Lemma 3. *If ψ is an **LJ***-proof of S, then no weak quantifier inference in $T_\varepsilon(\psi)$ violates the label condition.*

Proof. By Definition 5, the eigenvariables in a proof are always replaced by ε-terms with constant labels. Since a weak quantifier that uses an eigenvariable α occurs above the strong quantifier that introduced such variable, the label condition will hold in the epsilonized proof. □

Lemma 4. *If ψ is an **LJ***-proof of S, then no axiom inference in $T_\varepsilon(\psi)$ violates the initial condition.*

Proof. Trivial by Definition 5 and by the fact that there are no inferences operating above axioms. □

Theorem 3. *If $\vdash_{\boldsymbol{LJ}^\star} S$, then $\vdash_{\boldsymbol{LJ}^\star} T_\varepsilon(S)$.*

Proof. Let ψ be an **LJ***-proof of S. Then, by Lemmas 1, 2, 3 and 4 $T_\varepsilon(\psi)$ is a correct **LJ***-proof of $T_\varepsilon(S)$. □

De-epsilonization of proofs, denoted by T_ε^{-1}, replaces ε-terms by eigenvariables and introduces strong quantifier inferences in appropriate places. To detect the appropriate places, the following definition is helpful.

Intuitively, the de-epsilonization procedure will traverse a proof ψ in a top-down manner, re-applying the inference rules from ψ. As this is done, the sequents will contain formulas of the form $A[\nu_x B[x]]$ where $\nu_x B[x]$ is a top-level ε-term, for increasingly more complex A. Thus, $\nu_x B[x]$ is initially accessible in the formula and in the sequent, while A is a subformula of B. Replacing the ε-term by an eigen-variable and introducing a strong quantifier inference for this eigen-variable becomes possible when A becomes exactly equal to B, in which case the ε-term is said to be *ready*. However, to avoid violations of the eigen-variable condition, it is still necessary to postpone the introduction of strong quantifier inferences as much as possible. That is why the de-epsilonization procedure seeks to introduce them just before they become inaccessible. However, introducing them earlier may be necessary if a contraction operates on occurrences of $\nu_x B[x]$ with different labels.

Definition 6 (De-epsilonization of proofs)**.** *Let F be an ε-free formula and ψ an* ***LJ**** *proof of $T_\varepsilon(F)$. The de-epsilonization $T_\varepsilon^{-1}(\psi)$ is constructed inductively on the inference rules.*

Base case: *ψ consists of only one axiom. Then $T_\varepsilon^{-1}(\psi) = \psi$.*

Step case: *ψ ends with an inference ρ. By the induction hypothesis, the de-epsilonization of ρ's premises: $T_\varepsilon^{-1}(\psi_1)$ and (for the case of binary inferences) $T_\varepsilon^{-1}(\psi_2)$ are well-defined. Then $T_\varepsilon^{-1}(\psi)$ is defined according to the possible cases for ρ:*

- *ρ is a weakening.*

  Then $T_\varepsilon^{-1}(\psi)$ is simply $T_\varepsilon^{-1}(\psi_1)$ followed by the same weakening.

- *ρ is a cut.*

  Then $T_\varepsilon^{-1}(\psi)$ is the proof obtained by applying the same cut on $T_\varepsilon^{-1}(\psi_1)$ and $T_\varepsilon^{-1}(\psi_2)$.

- *ρ is a contraction on a formula F.*

  If F contains no ε-terms, then $T_\varepsilon^{-1}(\psi)$ is defined as $T_\varepsilon^{-1}(\psi_1)$ followed by the contraction. Otherwise, if the contracted formulas contain ε-terms $\nu_x^{l_1}\ G$ and $\nu_x^{l_2}\ G$, then $T_\varepsilon^{-1}(\psi)$ depends on the following cases for l_1 and l_2:

  - *The labels l_1 and l_2 are equal, regardless whether they are variables or constants. In this case, $T_\varepsilon^{-1}(\psi)$ is defined as $T_\varepsilon^{-1}(\psi')$ followed by the same contraction.*
  - *The labels l_1 and l_2 are two different constants*[3]*. In this case, F and G are the same, then $T_\varepsilon^{-1}(\psi)$ is defined as $T_\varepsilon^{-1}(\psi_1)\{\nu_x^{l_1}\ Fx \mapsto \alpha\}\{\nu_x^{l_2}\ Fx \mapsto \beta\}$ followed by two strong quantifier inferences ($\forall_r$ if ν is τ and $\exists_l$ if ν is ε) and a contraction on the quantified formulas.*

[3]This case occurs for epsilonized proofs, but not in proofs obtained by proof search in **LJ***.

The case where l_1 *and* l_2 *are two different variables does not occur for one of two reasons: (1) if* ψ *was obtained via proof search in* $\boldsymbol{LJ}^\star$*, then contraction of formulas with accessible* ε*-terms copies the variables to the premise, and thus they will be instantiated with the same constant at a later step; or (2) if* ψ *was obtained via epsilonization of a* **LJ** *proof, the labels will be constants.*

- ρ *is a logical inference.*

  If ρ *operates on* ε*-free formulas or all top-level* ε*-terms in* ρ*'s auxiliary formulas are still accessible in the conclusion, then* $T_\varepsilon^{-1}(\psi)$ *is defined as* ρ *applied to the de-epsilonization of its premise(s).*

  Otherwise, while there exists a top-level ε*-term* $\nu_x^l\ F$ *that would no longer be accessible in* ρ*'s conclusion, we add the appropriate strong quantifier and apply the replacement* $\{\nu_x^l.F \mapsto \alpha\}$ *to the proof with a fresh variable* α *as well as the replacements* $\{\nu_x^l\ F' \mapsto \alpha\}$ *(with the same variable* α*) for any* F' *that differs from* F *only in the presence of nested* ε*-terms. Finally, when there are no more* ε*-terms that would become inaccessible,* $T_\varepsilon^{-1}(\psi)$ *becomes* ρ *applied to the proof resulting from this iterative quantifier reintroduction procedure.*

If after this process the end-sequent still contains ε*-terms, then additional strong quantifier inferences are added accordingly.*

An example illustrating the need for replacing nested terms and for the while loop in the last case of Definition 6 is available in Section 4.3.

We can now prove soundness of the epsilonization method.

Lemma 5. *If* ψ *is an* $\boldsymbol{LJ}^\star$*-proof of an end-sequent* $T_\varepsilon(S)$*, then the end-sequent of* $T_\varepsilon^{-1}(\psi)$ *is* S*.*

Proof. In Definition 6, all ε-terms from ψ are replaced by eigenvariables and strong quantifier rules are applied, so that eventually formulas of the form $A[\varepsilon_x A[x]]$ (or $A[\tau_x A[x]]$) in $T_\varepsilon(S)$ are replaced by $\exists x.A[x]$ (or, respectively, $\forall x.A[x]$), innermost subformulas first. Notice that, at the time of the introduction of the strong quantifier, the outer formulas and those bound by the ε-term are indeed the same, since possibly nested ε-terms correspond to innermost quantifiers which will have been already introduced above in the proof. $\square$

Lemma 6. *If* ψ *is an* $\boldsymbol{LJ}^\star$*-proof of an end-sequent* $T_\varepsilon(S)$*, then there is an* $\boldsymbol{LJ}^\star$*-proof* ψ' *of* S *obtainable from* $T_\varepsilon^{-1}(\psi)$ *by reductive cut-elimination.*

Proof. The key point is to show that any violation of the eigenvariable condition in $T_\varepsilon^{-1}(\psi)$ can be removed by reductive cut-elimination. Assume that there is a strong quantifier inference ρ in $T_\varepsilon^{-1}(\psi)$ that violates the eigenvariable condition. This means that $T_\varepsilon^{-1}(\psi)$ has one of the following forms near ρ:

$$\dfrac{\begin{array}{c}\vdots\\ \hline\hline \Gamma, A[\alpha] \vdash B[\alpha]\end{array}}{\begin{array}{c}\Gamma, \exists x.A[x] \vdash B[\alpha]\\ \hline\hline \vdots\end{array}}\ \rho : \exists_l$$

$$\dfrac{\begin{array}{c}\vdots\\ \hline\hline \Gamma, B[\alpha], A[\alpha] \vdash C\end{array}}{\begin{array}{c}\Gamma, B[\alpha], \exists x.A[x] \vdash C\\ \hline\hline \vdots\end{array}}\ \rho : \exists_l$$

$$\dfrac{\begin{array}{c}\vdots\\ \hline\hline \Gamma, B[\alpha] \vdash A[\alpha]\end{array}}{\begin{array}{c}\Gamma, B[\alpha] \vdash \forall x.A[x]\\ \hline\hline \vdots\end{array}}\ \rho : \forall_r$$

For each of the cases above, there are four potential subcases. We show below that three of them cannot occur, because they would lead to contradictions, whereas the fourth can be fixed by reductive cut-elimination:

- $B[\alpha]$ propagates down to the end-sequent: α would then occur in the end-sequent of $T_\varepsilon^{-1}(\psi)$, but this would contradict Lemma 5.

- $B[\alpha]$ propagates down to a strong quantifier inference ρ' which has eigenvariable α: this case cannot occur, because ψ would then violate the label condition, thus contradicting the assumption that ψ is a correct $\mathbf{LJ}^\star$-proof.

- $B[\alpha]$ propagates down to a weak quantifier inference ρ' with an auxiliary formula $D[t[\alpha]]$: then ρ' would have auxiliary formula $D[t[\varepsilon_x B'[x]]]$ in ψ. If B were a proper super-formula of B', the term $t[\varepsilon_x B'[x]]$ would not be accessible and ρ' would be violating the accessibility condition. If B were equal to B', then ρ' would be occurring below ρ, which contradicts the fact that, in Definition 6, strong quantifier inferences such as ρ are introduced as low as possible. Indeed, notice that as the weak quantifier inference occurs in the proof with ep-terms, it will be applied during de-epsilonization in the same place, while the strong quantifier is only added when absolutely necessary (i.e., the term is no longer accessible or at the end-sequent).

- $B[\alpha]$ propagates down to a cut: in this case, the eigenvariable violation can be removed by shifting the cut upward, using Gentzen's reductive cut-elimination method.

□

Theorem 4 (Soundness of Epsilonization)**.** *If* $\vdash_{\boldsymbol{LJ}^\star} T_\varepsilon(S)$*, then* $\vdash_{\boldsymbol{LJ}} S$*.*

Proof. Let ψ be an $\mathbf{LJ}^\star$-proof of $T_\varepsilon(S)$. Then, by Lemmas 5 and 6, $T_\varepsilon^{-1}(\psi)$ is a correct $\mathbf{LJ}^\star$-proof of S. By Theorem 1, $\vdash_{\mathbf{LJ}} S$. □

4 Examples

This section presents a set of examples that help understand the epsilonization and de-epsilonization of proofs. Each example demonstrates the need for some aspect of the definitions.

4.1 Labels, Contractions and Inaccessible ε-terms

This section illustrates the need for different labels when contracting formulas with inaccessible ε-terms. We start with an end-sequent already considered before:

$$\neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash$$

whose epsilonization is

$$\neg\neg P(\tau_x \neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash$$

We have seen that, had labels not been used, the later sequent would admit a proof whose de-epsilonization would generate a proof with eigenvariable violations. Taking the labels into account, the proof found by proof search in $\mathbf{LJ}^\star$ is the following:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\overline{P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px)} \quad \overline{P(\tau^{l_2}_x.\neg Px) \vdash P(\tau^{l_2}_x.\neg Px)}}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px) \wedge P(\tau^{l_2}_x.\neg Px)}\ \wedge_r}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px), \neg(P(\tau^{l_1}_x.\neg Px) \wedge P(\tau^{l_2}_x.\neg Px)) \vdash}\ \neg_l}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash}\ \forall'_l}{P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash \neg P(\tau^{l_2}_x.\neg Px)}\ \neg_r}{P(\tau^{l_1}_x.\neg Px), \neg\neg P(\tau^{l_2}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash}\ \neg_l}{P(\tau^{l_1}_x.\neg Px), \neg\neg P(\tau^{l_2}_x.\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\ \forall'_l}{\neg\neg P(\tau^{l_2}_x.\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash \neg P(\tau^{l_1}_x.\neg Px)}\ \neg_r}{\neg\neg P(\tau^{l_1}_x.\neg Px), \neg\neg P(\tau^{l_2}_x.\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\ \neg_l}{\neg\neg P(\tau^{l}_x.\neg Px), \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\ c_l
$$

Note how the ε-term labelled with l_2 is not available for the weak quantifier $\forall y$. Observe also how the two labels of the contracted formulas need to be different. Had they been the same, we would be able to obtain the same proof as before, which de-epsilonizes to an incorrect proof.

The de-epsilonization procedure constructs, in a top-down manner, the same proof up to this point:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\overline{P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px)} \quad \overline{P(\tau^{l_2}_x.\neg Px) \vdash P(\tau^{l_2}_x.\neg Px)}}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px) \wedge P(\tau^{l_2}_x.\neg Px)}\ \wedge_r}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px), \neg(P(\tau^{l_1}_x.\neg Px) \wedge P(\tau^{l_2}_x.\neg Px)) \vdash}\ \neg_l}{P(\tau^{l_2}_x.\neg Px), P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash}\ \forall'_l}{P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash \neg P(\tau^{l_2}_x.\neg Px)}\ \neg_r
$$

If the next inference, $\neg_l$, were applied, the ε-term $\tau^{l_2}_x.\neg Px$ would no longer be accessible. Therefore, it is time to introduce a strong quantifier. Since the ε-term is bound by τ, the de-epsilonization procedure introduces a $\forall_r$ inference and replaces $\tau^{l_2}_x.\neg Px$ by a new fresh variable α.

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\overline{P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px)} \quad \overline{P\alpha \vdash P\alpha}}{P\alpha, P(\tau^{l_1}_x.\neg Px) \vdash P(\tau^{l_1}_x.\neg Px) \wedge P\alpha}\ \wedge_r}{P\alpha, P(\tau^{l_1}_x.\neg Px), \neg(P(\tau^{l_1}_x.\neg Px) \wedge P\alpha) \vdash}\ \neg_l}{P\alpha, P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash}\ \forall'_l}{P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash \neg P\alpha}\ \neg_r}{P(\tau^{l_1}_x.\neg Px), \forall y.\neg(P(\tau^{l_1}_x.\neg Px) \wedge Py) \vdash \forall x.\neg Px}\ \forall_r
$$

The re-construction of the proof is continued until the next point where a strong quantifier is needed, for the same reason as before. The same procedure is followed, now replacing $\tau_x^{l_1}.\neg Px$ by a new variable β. The final result is the following valid $\mathbf{LJ}^\star$ (and also $\mathbf{LJ}$) proof:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P\beta \vdash P\beta} \quad \dfrac{}{P\alpha \vdash P\alpha}}{P\alpha, P\beta \vdash P\beta \wedge P\alpha}\,\wedge_r}{P\alpha, P\beta, \neg(P\beta \wedge P\alpha) \vdash}\,\neg_l}{P\alpha, P\beta, \forall y.\neg(P\beta \wedge Py) \vdash}\,\forall'_l}{P\beta, \forall y.\neg(P\beta \wedge Py) \vdash \neg P\alpha}\,\neg_r}{P\beta, \forall y.\neg(P\beta \wedge Py) \vdash \forall x.\neg Px}\,\forall_r}{P\beta, \neg\forall x.\neg Px, \forall y.\neg(P\beta \wedge Py) \vdash}\,\neg_l}{P\beta, \neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\,\forall'_l}{\neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash \neg P\beta}\,\neg_r}{\neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash \forall x.\neg Px}\,\forall_r}{\neg\forall x.\neg Px, \neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\,\neg_l}{\neg\forall x.\neg Px, \forall z.\forall y.\neg(Pz \wedge Py) \vdash}\,c_l
$$

4.2 Contraction with Distinct Labels

The example in this section illustrates the need for allowing contraction on formulas with different constant labels. Consider the following $\mathbf{LJ}^\star$ proof of an ε-free end-sequent:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P\alpha \vdash P\alpha}}{P\alpha \vdash \exists x.Px}\,\exists_r \quad \dfrac{\dfrac{}{P\beta \vdash P\beta}}{P\beta \vdash \exists x.Px}\,\exists_r}{P\alpha, P\beta \vdash \exists x.Px \wedge \exists x.Px}\,\wedge_r}{P\alpha, \exists x.Px \vdash \exists x.Px \wedge \exists x.Px}\,\exists_l}{\exists x.Px, \exists x.Px \vdash \exists x.Px \wedge \exists x.Px}\,\exists_l}{\exists x.Px \vdash \exists x.Px \wedge \exists x.Px}\,c_l
$$

Following the epsilonization procedure, the proof obtained is:

$$\cfrac{\cfrac{\cfrac{}{P(\varepsilon^a_x.Px) \vdash P(\varepsilon^a_x.Px)}}{P(\varepsilon^a_x.Px) \vdash \exists x.Px}\ \exists_r \quad \cfrac{\cfrac{}{P(\varepsilon^b_x.Px) \vdash P(\varepsilon^b_x.Px)}}{P(\varepsilon^b_x.Px) \vdash \exists x.Px}\ \exists_r}{\cfrac{P(\varepsilon^a_x.Px), P(\varepsilon^b_x.Px) \vdash \exists x.Px \wedge \exists x.Px}{P(\varepsilon^l_x.Px) \vdash \exists x.Px \wedge \exists x.Px}\ c_l}\ \wedge_r$$

Note how contraction must allow the two ε-terms to have different constant labels. The label in the conclusion can be arbitrary. Such flexibility makes it possible to map the epsilonized proof to the exact intention of the original proof, which was to use two different eigenvariables for the strong quantifier. Interestingly, this situation only occurs if a proof is epsilonized. Had we searched for a proof of the same end-sequent in **LJ***, only one "eigenvariable" would have been used.

4.3 Nested ε-terms

When sequents contain blocks of strong quantifiers or strong quantifiers inside the scope of other strong quantifiers, epsilonization results in sequents with nested ε-terms. In this section, we look at an example of this kind.

Let F be $\forall x.(\exists y.\exists z.P(x,y,z) \to \exists w.\exists v.\exists q.P(q,v,w))$. Then $T_\varepsilon(F)$ is:

$$P(t_x, t_y, t_z) \to \exists w.\exists v.\exists q.P(q,v,w)$$

where:

- $t_z = \delta_z(\delta_x, \delta_y(\delta_x))$
- $t_y = \delta_y(\delta_x)$
- $t_x = \delta_x$
- $\delta_x = \gamma_x[\delta_y(x), \delta_z(x, \delta_y(x))]$
- $\delta_y(x) = \gamma_y[\delta_z(x,y)](x)$
- $\delta_z(x,y) = \gamma_z(x,y) = \varepsilon_z P(x,y,z)$
- $\gamma_y[t](x) = \varepsilon_y P(x,y,t)$
- $\gamma_x[t_1,t_2] = \tau_x(P(x,t_1,t_2) \to \exists w.\exists v.\exists q.P(q,v,w))$

Let ψ be the following **LJ***-proof of $T_\varepsilon(F)$:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{}{P(t_x,t_y,t_z) \vdash P(t_x,t_y,t_z)}}{P(t_x,t_y,t_z) \vdash \exists q.P(q,t_y,t_z)}\ \exists_r}{P(t_x,t_y,t_z) \vdash \exists v.\exists q.P(q,v,t_z)}\ \exists_r}{P(t_x,t_y,t_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r}{\vdash P(t_x,t_y,t_z) \to \exists w.\exists v.\exists q.P(q,v,w)}\ \to_r$$

During the top-down construction of $T_\varepsilon^{-1}(\psi)$, initially the three $\exists_r$ inferences are simply reapplied:

$$\cfrac{\cfrac{\cfrac{\overline{P(t_x,t_y,t_z) \vdash P(t_x,t_y,t_z)}}{P(t_x,t_y,t_z) \vdash \exists q.P(q,t_y,t_z)}\ \exists_r}{P(t_x,t_y,t_z) \vdash \exists v.\exists q.P(q,v,t_z)}\ \exists_r}{P(t_x,t_y,t_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r$$

At this point, t_z is ready, but applying $\rightarrow_R$ would make it inaccessible. Therefore, it is time to introduce the strong quantifier for z:

$$\cfrac{\cfrac{\cfrac{\cfrac{\overline{P(t'_x[\alpha_z],t'_y[\alpha_z],\alpha_z) \vdash P(t'_x[\alpha_z],t'_y[\alpha_z],\alpha_z)}}{P(t'_x[\alpha_z],t'_y[\alpha_z],\alpha_z) \vdash \exists q.P(q,t'_y[\alpha_z],\alpha_z)}\ \exists_r}{P(t'_x[\alpha_z],t'_y[\alpha_z],\alpha_z) \vdash \exists v.\exists q.P(q,v,\alpha_z)}\ \exists_r}{P(t'_x[\alpha_z],t'_y[\alpha_z],\alpha_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r}{\exists z.P(t'_x[z],t'_y[z],z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l$$

where $t'_y[z] = t_y\{\delta_z(.,.) \mapsto z\}$ and $t'_x[z] = t_x\{\delta_z(.,.) \mapsto z\}$. Note that the replacement of terms of the general form $\delta_z(.,.)$ by z de-epsilonizes occurrences of $\varepsilon_z P(.,.,z)$ nested inside t_x and t_y. This illustrates the need for substituting not only top-level ε-terms, but also nested ε-terms in the last case of Definition 6.

Now $t'_y[z]$ becomes ready and it would not be accessible anymore after application of $\rightarrow_R$. Therefore, it is time to introduce the strong quantifier inference for y:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\overline{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash P(t''_x[\alpha_y],\alpha_y,\alpha_z)}}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists q.P(q,\alpha_y,\alpha_z)}\ \exists_r}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists v.\exists q.P(q,v,\alpha_z)}\ \exists_r}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r}{\exists z.P(t''_x[\alpha_y],\alpha_y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l}{\exists y.\exists z.P(t''_x[y],y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l$$

$t''_x[y]$ is not ready yet, and the $\rightarrow_r$ inference rule can be applied:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash P(t''_x[\alpha_y],\alpha_y,\alpha_z)}}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists q.P(q,\alpha_y,\alpha_z)}\ \exists_r}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists v.\exists q.P(q,v,\alpha_z)}\ \exists_r}{P(t''_x[\alpha_y],\alpha_y,\alpha_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r}{\exists z.P(t''_x[\alpha_y],\alpha_y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l}{\exists y.\exists z.P(t''_x[y],y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l}{\vdash \exists y.\exists z.P(t''_x[y],y,z) \rightarrow \exists w.\exists v.\exists q.P(q,v,w)}\ \rightarrow_r
$$

The fact that we had to introduce two strong quantifier inferences, de-epsilonizing two different ε-terms, before being able to reapply a logical inference rule illustrates the need for a while loop in the last case of Definition 6.

Finally, a $\forall_r$ inference rule has to be applied, because t_x is now ready in the end sequent and there are no further inferences from ψ to be reapplied:

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P(\alpha_x,\alpha_y,\alpha_z) \vdash P(\alpha_x,\alpha_y,\alpha_z)}}{P(\alpha_x,\alpha_y,\alpha_z) \vdash \exists q.P(q,\alpha_y,\alpha_z)}\ \exists_r}{P(\alpha_x,\alpha_y,\alpha_z) \vdash \exists v.\exists q.P(q,v,\alpha_z)}\ \exists_r}{P(\alpha_x,\alpha_y,\alpha_z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_r}{\exists z.P(\alpha_x,\alpha_y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l}{\exists y.\exists z.P(\alpha_x,y,z) \vdash \exists w.\exists v.\exists q.P(q,v,w)}\ \exists_l}{\vdash \exists y.\exists z.P(\alpha_x,y,z) \rightarrow \exists w.\exists v.\exists q.P(q,v,w)}\ \rightarrow_r}{\vdash \forall x.(\exists y.\exists z.P(x,y,z) \rightarrow \exists w.\exists v.\exists q.P(q,v,w))}\ \forall_r
$$

5 Related Work

Other methods for a Skolemization-like procedure for intuitionistic logic have been investigated. This has been the topic of a series of papers by Baaz and Iemhoff that study the use of an existence predicate, introduced by Scott [11], for Skolemization. They start by defining *eSkolemization* [2], a process for removing strong existential quantifiers in intuitionistic logic. In the same paper there is a semantical proof of completeness of eSkolemization and later on they provide a proof-theoretical proof [4]. In [3] the authors extend the method for strong universal quantifiers, but the solution is more ad-hoc, as it requires the addition of a pre-order to the logic and introduces weak quantifiers. Roughly, the eSkolemization method replaces strong occurrences of $\forall\bar{x}.A\bar{x}$ by $E(f(\bar{x})) \rightarrow A(f(\bar{x}))$ and strong occurrences of $\exists\bar{x}.A\bar{x}$ by

$E(f(\bar{x})) \wedge A(f(\bar{x}))$, where f is a new function symbol. In contrast to our approach, which only adds to the language the ε and τ operators, eSkolemization requires extending the language with infinitely many symbols, including a predicate. Moreover, the treatment of existential and universal quantifiers is not uniform whereas in our method those treatments are naturally dual. The calculus LJE, presented in [4], contains different rules for the quantifiers which add the existence predicate to the premises. Therefore, it does not have the sub-formula property. Also, the rules for $\forall_l$ and $\exists_r$ are binary, adding yet another complexity for proof search. It is worth noting that $\mathbf{LJ}^{\star}$ presents none of these issues.

The approach that comes closer to what is presented here is that of Mints [7, 8]. Although the precise relation is not easy to pinpoint and describe, it is straightforward to note important differences. Firstly, whereas Mints is concerned with the extension of **LJ** by an epsilonization rule in the calculus (which acts only at whole formulas), we consider epsilonization as a pre-processing step, acting deeply on all strongly quantified subformulas in the end-sequent. In Mints' calculus, the epsilonization rule is essentially a strong quantifier rule that instantiates the variable by an ε-term instead of an eigen-variable. In contrast, $\mathbf{LJ}^{\star}$-proofs of epsilonized end-sequents contain no inferences that act as strong quantifier inferences in disguise. It was a significant challenge, and one of the main distinguishing contributions of this paper, to discover that ε-terms are informative enough to tell where strong quantifier inferences need to be introduced when de-epsilonizing. Mints also describes a condition for the correctness of proofs, requiring that all sequents are *intelligent*[4]. The definition of *intelligence* is related to the definition of *accessibility* presented here. However, the definition of *intelligence* is not local: to decide whether a sequent S is intelligent in a proof ψ, it may be necessary to look at every sequent S' occurring below S in ψ. This is undesirable in the context of bottom-up proof search, because the whole derivation may have to be traversed and checked in order to decide if an inference is allowed. The definition of *accessibility*, on the other hand, is local: to decide if a weak quantifier inference is allowed, only its conclusion sequent needs to be checked. Furthermore, while Mints [7, 8] restricts all inference rules (by requiring that all sequents be *intelligent*), in the $\mathbf{LJ}^{\star}$ calculus presented here, only the weak quantifier rules need to be restricted. Therefore, the restrictions described here are weaker. Another difference is that Mints [7, 8] considers only the ε binder, whereas here τ is also taken into account.

Decades later, Mints [9] proposed a new calculus where he dropped the global intelligibility condition and adopted binary weak quantifier rules (thus following

[4]Mints used the adjective осмысленный in the Russian original [7]. This was translated as *intelligent* in [8]. In [12], Soloviev uses the better translation *meaningful*.

the trend of [4]) whose left premises require proving that the instantiating term is *defined*. While the notion of *defined* is arguably more local than the notion of *intelligent*, it requires proof search and is semantically inspired. Moreover, the definition of *defined* is incomplete because it is defined only for top-level ε-terms. It is not clear what should be done, for example, when the instantiating term is not a top-level ε-term but contains an ε-term as a sub-term. Furthermore, in Mints' new calculus, epsilonization is still treated as an inference rule, not as a pre-processing step.

6 Conclusion

We have shown that, whereas Skolemization is unsound for **LJ** (as is well-known), the new epsilonization transformation defined here is sound for the restricted calculus $\mathbf{LJ}^\star$ proposed. Although the definitions and proofs are technically complex, the underlying idea is conceptually very simple. The unsoundness of Skolemization for **LJ** is essentially due to violations of the eigenvariable condition, which happen implicitly and unnoticed, because Skolemization replaces eigenvariables by Skolem terms. In the case of epsilonization, on the other hand, ε-terms are informative enough to allow us to know where strong quantifier inferences introducing their corresponding eigenvariables would be located if the sequent had not been epsilonized. This information allows us to restrict the weak quantifier rules in $\mathbf{LJ}^\star$ that use ε-terms, so that they only occur above those implicit strong quantifier inferences' locations. Consequently, as desired, de-epsilonizing $\mathbf{LJ}^\star$ proofs never results in violations of the eigenvariable condition.

The approach presented here distinguishes itself from related work primarily by being the only purely syntactic, deterministic (not requiring additional proof search) and local restriction of the intuitionistic sequent calculus where a Skolemization-like pre-processing transformation is sound.

References

[1] J. Avigad and R. Zach. The Epsilon Calculus. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2013 edition, 2013.

[2] M. Baaz and R. Iemhoff. The Skolemization of existential quantifiers in intuitionistic logic. *Annals of Pure and Applied Logic*, 142(1–3):269 – 295, 2006.

[3] M. Baaz and R. Iemhoff. On Skolemization in Constructive Theories. *The Journal of Symbolic Logic*, 73(3):pp. 969–998, 2008.

[4] M. Baaz and R. Iemhoff. Eskolemization in Intuitionistic Logic. *J. Log. Comput.*, 21(4):625–638, 2011.

[5] M. Baaz and A. Leitsch. On Skolemization and Proof Complexity. *Fundam. Inform.*, 20(4):353–379, 1994.

[6] A. Leitsch, G. Reis, and B. W. Paleo. Towards CERes in intuitionistic logic. In P. Cégielski and A. Durand, editors, *CSL*, volume 16 of *LIPIcs*, pages 485–499. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

[7] G. Mints. Heyting Predicate Calculus with Epsilon Symbol. *Zapiski Nauchnykh Seminarov Leningradskogo Otdeleniya Matematicheskogo Instituta im. V. A. Steklova AN SSSR*, 40:110–118, 1974.

[8] G. Mints. Heyting Predicate Calculus with Epsilon Symbol. *Journal of Soviet Mathematics*, 8(3):317–323, 1977.

[9] G. Mints. Intuitionistic Existential Instantiation and Epsilon Symbol. *CoRR*, abs/1208.0861, 2012.

[10] G. Reis. *Cut-elimination by resolution in intuitionistic logic*. PhD thesis, Vienna University of Technology, July 2014.

[11] D. Scott. Identity and existence in intuitionistic logic. In M. Fourman, C. Mulvey, and D. Scott, editors, *Applications of Sheaves*, volume 753 of *Lecture Notes in Mathematics*, pages 660–696. Springer Berlin Heidelberg, 1979.

[12] S. Soloviev. Studies of Hilbert's ε-operator in the USSR. *Journal of Logic and its Applications*, 2016.

Received 23 October 2015