A New Indefinite Semantics for Hilbert's epsilon

Claus-Peter Wirth

FR Informatik, Saarland Univ., D–66041 Saarbrücken, Germany
cp@ags.uni-sb.de

Abstract. After reviewing the literature on semantics of Hilbert's epsilon symbol, we present a new one that is similar to the referential interpretation of indefinite articles in natural languages.

1 Motivation

In [10] we have studied the combination of mathematical induction in the style of Fermat's *descente infinie* with state-of-the-art logical deduction into a formal system in which a working mathematician can straightforwardly develop his proofs. This system's soundness proof required a notion of reduction like the one of Def. 5.19 below that "*preserves solutions*". By this we mean the following: All solutions that transform a proof attempt (to which a proposition has been reduced) into a closed proof (i.e. the "closing substitutions" for the free variables) are also solutions of the original proposition, just as a proof in Prolog computes answers of a query proposition. The *liberalized* δ-rules as found in [4] do not satisfy this notion. The addition of our choice-conditions—which can be understood as a new semantics for Hilbert's ε-symbol that mirrors the referential interpretation of indefinite articles in natural languages—finally turned out to be the only way to repair this defect of the liberalized δ-rules. As Hilbert's ε is of universal interest and applicability, we suppose that our new semantics will turn out to be useful in many other areas where logic is applied as a tool for description and reasoning. The paper organizes as follows: Sect. 2 introduces the ε. In Sect. 3 we show the problems with indefinite and committed choice and review the literature on the ε's semantics w.r.t. adequacy and Hilbert's intentions. Sect. 4 introduces our new approach to the ε's semantics informally. A formal discussion follows in Sect. 5, where proofs are omitted but can be found in [10].

2 What is Hilbert's ε?

Just like Bertrand Russell's ι-symbol (cf. e.g. [7, Vol. I]), Hilbert's ε-symbol is a binder that takes a variable x and a formula A and produces the term $\varepsilon x.\ A$.

Example 2.1 (ι) *(Buggy!)*
Let III and IV denote Henry III and Henry IV, resp., and consider Father to be a predicate for which $\mathsf{Father}(\mathsf{III}, \mathsf{IV})$ holds, i.e. "Henry III is father of Henry IV". Now, "*the* father of Henry IV" can be denoted by $\iota x.\ \mathsf{Father}(x, \mathsf{IV})$, and because this is nobody but Henry III, i.e. $\mathsf{III} = \iota x.\ \mathsf{Father}(x, \mathsf{IV})$, we know that

$\mathsf{Father}(\iota x.\ \mathsf{Father}(x, \mathsf{IV}), \mathsf{IV})$ holds. Similarly,

$$\mathsf{Father}(\iota x.\ \mathsf{Father}(x, \mathsf{Adam}), \mathsf{Adam}), \tag{Ex. 2.1.1}$$

and thus $\exists y.\ \mathsf{Father}(y, \mathsf{Adam})$, but, oops! Adam and Eve do not have any fathers. If you think differently on this, consider the following problem that occurs when somebody has two fathers:

$$\mathsf{Father}(\mathsf{HolyGhost}, \mathsf{Jesus}) \quad \wedge \quad \mathsf{Father}(\mathsf{Joseph}, \mathsf{Jesus}). \tag{Ex. 2.1.2}$$

Then the Holy Ghost is *the* father of Jesus and Joseph is *the* father of Jesus, i.e.

$$\mathsf{HolyGhost} = \iota x.\, \mathsf{Father}(x, \mathsf{Jesus}) \ \wedge\ \mathsf{Joseph} = \iota x.\, \mathsf{Father}(x, \mathsf{Jesus}) \tag{Ex. 2.1.3}$$

which implies something the Pope may not accept, namely that $\mathsf{HolyGhost} = \mathsf{Joseph}$, and he anathematized Henry IV in the year 1076.

Thus, in order to be able to write down $\iota x.\ A$ without further consideration, following the (well justified) standard in classical non-modal logics that any term uniquely denotes something that exists, we have to treat $\iota x.\ A$ as an uninterpreted term about which we only know

$$\exists! x.\ A \ \Rightarrow\ A\{x \mapsto \iota x.\ A\}$$

or in different notation $(\exists! x.\ (A(x))) \ \Rightarrow\ A(\iota x.\ (A(x)))$. Now the problems of Ex. 2.1 disappear because (Ex. 2.1.1) and (Ex. 2.1.3) are not valid. The price we have to pay for this is that—roughly speaking— $\iota x.\ A$ is of no use unless the unique existence $\exists! x.\ A$ can be derived.

Compared to this, the ε is more useful than the ι because it comes with the stronger axiom

$$\exists x.\ A \ \Rightarrow\ A\{x \mapsto \varepsilon x.\ A\} \tag{ε_0}$$

As the basic methodology of David Hilbert's formal program is to treat all symbols as meaningless, he does not give us any semantics but only the axiom (ε_0). Although no meaning is required, it furthers the understanding, and therefore Paul Bernays writes [7, Vol. II, p. 12] in the fundamental book on the work of David Hilbert and his group on the foundations of mathematics that

> $\varepsilon x.\ A$ ... "is an object of the universe for which—according to the semantical translation of the formula (ε_0)—*the predicate A holds, provided that A holds for any object of the universe.*" (our translation)

Example 2.2 (ε instead of ι) *(continuing Ex. 2.1)*
We still have $\mathsf{III} = \varepsilon x.\ \mathsf{Father}(x, \mathsf{IV})$ and $\mathsf{Father}(\varepsilon x.\ \mathsf{Father}(x, \mathsf{IV}), \mathsf{IV})$. But, from the contrapositive of (ε_0) and $\neg\mathsf{Father}(\varepsilon x.\ \mathsf{Father}(x, \mathsf{Adam}), \mathsf{Adam})$, we now can conclude that $\neg\exists y.\ \mathsf{Father}(y, \mathsf{Adam})$.

David Hilbert did not need any semantics or precise intention for the ε-symbol because it was introduced merely as a formal syntactic device to facilitate proof-theoretic investigations, motivated by the possibility to get rid of the existential and universal quantifiers via

$$\exists x.\ A \quad \Leftrightarrow \quad A\{x \mapsto \varepsilon x.\ A\} \tag{ε_1}$$

and

$$\forall x.\ A \quad \Leftrightarrow \quad A\{x \mapsto \varepsilon x.\ \neg A\} \tag{ε_2}$$

Cf. [9] for an excellent modern treatment of the subject.

Wilhelm Ackermann, Paul Bernays, and David Hilbert finally succeeded [7, Vol. II] in giving a proof of the fact that the ε (just as the ι [7, Vol. I]) is a conservative extension in the sense that any formal proof of an ε-free formula can be transformed into a formal proof that does not use the ε at all ("*ε-elimination theorems*"). Generally, this is not the case anymore if we go beyond first-order logic, e.g. by considering finite structures only, cf. [3] for references.

While the historical and technical research on ε-elimination is still going on, this is not the subject of this paper. Moreover, it is my opinion that in our days we should be less interested in Hilbert's formal program and the consistency of mathematics than in the powerful use of logic in creative processes. And, instead of the tedious syntactical proof transformations that easily lose their usefulness and elegance within their technical complexity and that—more importantly—can only refer to an already existing logic, we need *semantical* means for finding new logics and new applications. And the question that then arises here—esp. for classical logic—is: What would be the proper semantics for Hilbert's ε?

3 Is Indefinite Choice the Proper Semantics?

Just as the ι-symbol is generally taken to be the referential interpretation of the *definite* article in natural languages, it is our opinion that the ε-symbol should be that of the *indefinite* one.

Example 3.1 (ε instead of ι) *(continuing Ex. 2.1)*
It may well be the case that

$$\mathsf{HolyGhost} = \varepsilon x.\ \mathsf{Father}(x, \mathsf{Jesus}) \quad \wedge \quad \mathsf{Joseph} = \varepsilon x.\ \mathsf{Father}(x, \mathsf{Jesus})$$

i.e. that "The Holy Ghost is $\underline{a}$ father of Jesus and Joseph is $\underline{a}$ father of Jesus." But this does not bring the Pope into trouble because we do not know whether any father of Jesus is equal to any father of Jesus. This will become clearer when reconsidering this example in Ex. 4.4.

3.1 Committed Choice

Closely connected to indefinite choice (also called "indeterminism" or "don't care nondeterminism") is the notion of "*committed choice*". E.g., when you have a new telephone, you *don't care* which number you get, but once the provider has chosen a number for your telephone, you want the provider to *commit to its choice*, i.e. not to change your phone number between two incoming calls.

Example 3.2 (Committed Choice) *(Buggy!)*
Suppose we want to prove $\exists x.\ (x \neq x)$.
According to (ε_1) this reduces to $\varepsilon x.\ (x \neq x) \neq \varepsilon x.\ (x \neq x)$.
Since there is no solution to $x \neq x$ we can replace $\varepsilon x.\ (x \neq x)$ with anything.
Thus, the above reduces to $\mathsf{0} \neq \varepsilon x.\ (x \neq x)$,
and then, by exactly the same argumentation, to $\mathsf{0} \neq \mathsf{1}$,
which is valid.

Thus we have proved our original formula $\exists x.\ (x \neq x)$, which, however, happens to be invalid. What went wrong? Of course, we have to commit to our choice for the ε-term introduced for the elimination of the existential quantifier.

The following lengthy example shows that the elimination of $\forall$- and $\exists$-quantifiers with the help of ε-terms may be more difficult than one might think. The problem is that some ε-terms may become "subordinate" to others. An ε-term $\varepsilon v.\ B$ (or more generally a binder on v) is *superordinate* to an (occurrence of an) ε-term $\varepsilon x.\ A$ if $\varepsilon x.\ A$ is a subterm of B and an occurrence of the variable v in $\varepsilon x.\ A$ is free in B. An (occurrence of an) ε-term a is *subordinate* to an ε-term b if b is superordinate to a. These subordinate ε-terms [7, Vol. II, p. 24] are responsible for the difficulty to prove the theorems on ε-elimination constructively.

Example 3.3 (Subordinate ε-terms)
Consider the formula $\forall x.\ \exists y.\ \forall z.\ \mathsf{P}(x,y,z)$. Let us apply (ε_1) and (ε_2) in order to completely remove the three quantifiers. The resulting term does not depend on the order in which we do this, but it is quite deep; in general n nested quantifiers result in an ε-nesting depth of 2^n-1. Moreover, huge ε-terms occur up to n-times with commitment to their choice. Therefore, we should carefully identify the terms that have multiple occurrence. The result is $\mathsf{P}(x_1, y_2, z_4)$, where

$z_4 = \varepsilon z.\ \neg\mathsf{P}(x_1, y_2, z)$
$y_2 = \varepsilon y.\ \mathsf{P}(x_1, y, z_3(y))$ with $z_3(\underline{y}) = \varepsilon z.\ \neg\mathsf{P}(x_1, \underline{y}, z)$
$x_1 = \varepsilon x.\ \neg\mathsf{P}(x, y_1(x), z_2(x))$ with $z_2(\underline{x}) = \varepsilon z.\ \neg\mathsf{P}(\underline{x}, y_1(\underline{x}), z)$
and $y_1(\underline{x}) = \varepsilon y.\ \mathsf{P}(\underline{x}, y, z_1(\underline{x}, y))$ with $z_1(\underline{x}, \underline{y}) = \varepsilon z.\ \neg\mathsf{P}(\underline{x}, \underline{y}, z)$.

Firstly note that the equations for z_3, z_2, y_1, z_1 are a little problematic because the underlined variables $\underline{x}$ and $\underline{y}$ occurring on the right-hand sides are only seemingly free but actually bound by the next ε to the left, to which the closest ε-term thus becomes subordinate. E.g., the ε-term $z_3(y)$ is subordinate to the ε-term y_2. Secondly, the top ε-binders on the right-hand sides are exactly those that require a commitment to their choice. This means that each of z_1, z_2, z_3, z_4 and of y_1, y_2 may be chosen differently without affecting soundness of the equivalence transformation. Note that the variables are strictly nested into each other. Thus we must choose in the order of z_1, y_1, z_2, x_1, z_3, y_2, z_4. Moreover, for z_3, z_2, y_1, z_1 we actually have to choose a function instead of a simple value. In Hilbert's view, however, there are no objects at all, but only terms, where x_1 reads

$$\varepsilon x.\ \mathsf{P}(x, \varepsilon y.\ \mathsf{P}(x, y, \varepsilon z_1.\ \mathsf{P}(x, y, z_1)), \varepsilon z_2.\ \mathsf{P}(x, \varepsilon y.\ \mathsf{P}(x, y, \varepsilon z_1.\ \mathsf{P}(x, y, z_1)), z_2))$$

and y_2 and z_4 take several lines to write them down.

For $\forall x.\ \forall y.\ \forall z.\ \mathsf{P}(x,y,z)$ instead of $\forall x.\ \exists y.\ \forall z.\ \mathsf{P}(x,y,z)$, we get the same exponential explosion as in Ex. 3.3 when we eliminate the ε-terms completely using (ε_2). The only difference is that we get two more occurrences of '$\neg$'. But when we have quantifiers of the same kind like '$\exists$' or '$\forall$', we had better choose them in parallel, e.g., for $\forall x.\ \forall y.\ \forall z.\ \mathsf{P}(x,y,z)$ we take $v_1 := \varepsilon v.\ \neg\mathsf{P}(1^{\mathrm{st}}(v), 2^{\mathrm{nd}}(v), 3^{\mathrm{rd}}(v))$, and then $x_1 = 1^{\mathrm{st}}(v_1)$, $y_2 = 2^{\mathrm{nd}}(v_1)$, $z_4 = 3^{\mathrm{rd}}(v_1)$.

3.2 Free Existential and Weak Free Universal Variables

Suppose we want to prove the existential property $\exists x.\ A$. While in old-fashioned inference systems the γ-rules (according to Smullyan's classification) require us to choose a witnessing term t immediately when removing the quantifier, more modern inference systems like the ones in [4] enable us to delay the choice of the term t (which is crucial for the success of the later proof) until the state of the proof attempt provides information that is sufficient to make a successful decision. This delay is achieved with a special kind of variable, sometimes called "dummy", "free variable", or "meta variable". We will call these variables *free existential variables* and write them like $x^{\exists}$. Now, $\exists x.\ A$ is first reduced to $A\{x \mapsto x^{\exists}\}$ and later in the proof we can globally substitute $x^{\exists}$ with an appropriate term.

In theorem proving, cf. [10] e.g., the explosion of Ex. 3.3 is reduced by not removing the quantifiers below ε-binders and by replacing existentially quantified variables with free existential variables. For the case of Ex. 3.3, this yields $z_4 = \varepsilon z.\ \neg\mathsf{P}(x_1, y_1^{\exists}, z)$ and $x_1 = \varepsilon x.\ \neg\exists y.\ \forall z.\ \mathsf{P}(x, y, z)$. Thus, in general, the binder nesting does not become deeper than $\frac{1}{4}(n+1)^2$. Moreover, if we are only interested in reduction and not in equivalence transformation of a formula, we can even abstract Skolem terms from the ε-terms and just consider $z_4 = z^{\forall,\mathrm{w}}(y_1^{\exists})$ and $x_1 = x^{\forall,\mathrm{w}}$. Here $x^{\forall,\mathrm{w}}$ and $z^{\forall,\mathrm{w}}$ are *weak free universal variables*. These variables play the role of parameters and stand for an arbitrary object of which nothing is known. The small 'w' stands for "weak".

Note that with Skolemization we have no explosion at all and the same will be the case for our approach to ε-terms.

After a glimpse into the literature on a semantics for Hilbert's ε in the next two sections, we will further introduce into indefinite choice in Sect. 4.

3.3 Functional Semantics

Nearly all the semantics for Hilbert's ε in the literature is functional. Cf. [9] and the references there for an overview. There, a functional behavior is required for the ε, which contrasts the above suggested indefiniteness.

The omnipresence of the functionality requirement may have its historical justification in the fact that the dots "..." in the quotation preceding Ex. 2.2 actually contain the word "function", which could be understood in its mathematical sense to denote a (right-) unique relation. And, what kind of function could it be but a choice function, choosing an element from the set of objects that satisfy A? In [6], Hilbert even writes

> "Above that, the ε has the role of the choice function, i.e. in the case where $A(a)$ may hold for several things, $\varepsilon x.(A(x))$ is *an arbitrary one* of the things a for which A holds." (our translation)

Regarding the question of an intended unique behavior, these statements are ambiguous in the German original. Hilbert most probably wanted to have what today we call "*committed choice*", but, as this technical term did not exist at his time and he was not interested in semantics anyway, simply used the word "function". The strongest argument against Hilbert's intention of a functional behavior is that the implied formula

$$\forall x.\ (A \Leftrightarrow B) \quad \Rightarrow \quad \varepsilon x.\ A = \varepsilon x.\ B \tag{E2}$$

cannot be derived form Hilbert's axiomatization.

To be precise, the notion of a "choice function" must be generalized here because—in order to have a total function on the power set of the universe—a value must be supplied even on the empty set: f is defined to be a *generalized choice function* if $f : \mathrm{dom}(f) \to \bigcup(\mathrm{dom}(f))$ and $\forall x \in \mathrm{dom}(f).\ (x = \emptyset \ \vee \ f(x) \in x)$. Different possible choices for the value on the empty set are described in [9], but as the consequences of any special choice are quite queer, the only solution that is found to be sufficiently adequate in [9] is to consider validity in *any* model given by *each* generalized choice function on the power set of the universe. Note, however, that, in each single model, the behavior of the ε is still a *function* from the set of objects that satisfy A. This can be expressed by (E2) which, roughly speaking, together with (ε_1) and (ε_2) is shown in [9] to be complete for first-order ε predicate calculus w.r.t. this semantics.

In [5] the above treatment of [9] is called the *extensional* one because the value of $\varepsilon x.\ A$ in each semantical structure $\mathcal{A}$ is functionally dependent on the extension of the formula A, i.e. on $\{\ o \mid \mathrm{eval}(\mathcal{A} \uplus \{x \mapsto o\})(A)\ \}$, where 'eval' is the standard evaluation function that maps a structure (or interpretation) (possibly including a valuation of the free variables) to a function from terms (and formulas) to values. In order to get more freeness for the definition of a semantics of ε, in [5] the value of $\varepsilon x.\ A$ may depend on the syntax besides the semantics. It is given as a function depending on a structure and on the *term* $\varepsilon x.\ A$. In [5, p. 177] we read: "This definition contains no restriction whatsoever on the valuation of ε-terms." This is, however, not true because it imposes the restriction of a functional behavior, which denies the possibility of an indefinite behavior as we will see below.

3.4 Indefinite Semantics in the Literature

The only occurrences of an indefinite semantics for Hilbert's ε in the literature seem to be [3] and the references there. Consider $\varepsilon x.\,(x{=}x) = \varepsilon x.\,(x{=}x)$ from [3] or the even simpler

$$\varepsilon x.\ \mathsf{true} \ = \ \varepsilon x.\ \mathsf{true} \tag{R}$$

which may be valid or not, depending on the question whether the same object is taken on both sides of the equation or not. In natural language this like "Something is equal to something" whose truth is indefinite. If you do not think so, take $\varepsilon x.\ \mathsf{true} \ \neq \ \varepsilon x.\ \mathsf{true}$ instead, i.e. "Something is unequal to something". In [3], Kleene's strong three valued logic is taken as a mathematically elegant means to solve the problems with indefiniteness. From a practical point of view, however, we do not think this to be an adequate solution because it severely restricts the applicability of the logic: Logical arguments outside the academical world are never made explicit because the presence of logic is either not realized at all or taken to be trivial. In applications, a logic is not the object of investigation but a meta-logical tool. Thus, regarding applications in the western world we have to stick to our common meta-logic which is a subset of classical (modal) logic. A western court can accept that Lee Harvey Oswald killed John F. Kennedy as well as that he did not; but cannot accept a third possibility, a tertium, as required for Kleene's strong three valued logic, and esp. not the interpretation given in [3] that he *both* did and did not kill him.

4 Introduction to a New Indefinite Semantics for the ε

From the discussion above one could get the impression that an indefinite logical treatment of the ε is not easy to find. Indeed, there is the problem that the substitutivity $s = t \Rightarrow f(s) = f(t)$ (cf. (E2)) and the reflexivity axiom $t = t$ (cf. (R)) cannot be taken for granted. This means that it is not definitely okay to replace a subterm by an equal term and that even syntactically equal terms may not be definitely equal.

Therefore, it may be interesting to see that in programs we are quite used to an indefinite behavior of choosing including committed choice and that the violation of substitutivity and even reflexivity is no problem there: Consider the following ML function, returning the first element of a list that implements a set:

```
fun choose s = case s of Set (i :: _) => i | _ => raise Empty;
```

The behavior of the function `choose` is indefinite for a given set, but any time it is called for an implemented set, it chooses a special element and *commits to this choice*, i.e. when called again, it returns the same value. In this case we have `choose s` $=$ `choose s`, but `s` $=$ `t` does not imply `choose s` $=$ `choose t`. In an implementation where some parallel reordering of lists may take place, even `choose s` $=$ `choose s` may be wrong. The question of `choose s` $=$ `choose s` is indefinite until the choice steps have actually been performed. *This is exactly what we will do with the ε.* The steps that are performed in logic are proof steps. Thus, on the one hand, when we want to prove $\varepsilon x.\,\mathsf{true} = \varepsilon x.\,\mathsf{true}$ we can choose $\mathsf{0}$ for both occurrences of $\varepsilon x.\,\mathsf{true}$, get $\mathsf{0} = \mathsf{0}$, and the proof is successful. On the other hand, when we want to prove $\varepsilon x.\,\mathsf{true} \neq \varepsilon x.\,\mathsf{true}$ we can choose $\mathsf{0}$ for one occurrence and $\mathsf{1}$ for the other, get $\mathsf{0} \neq \mathsf{1}$, and the proof is successful, too. This procedure may seem wondrous again, but is very similar to something quite common with free existential variables, cf. Sect. 3.2: On the one hand, when we want to prove $x^{\exists} = y^{\exists}$ we can choose $\mathsf{0}$ to substitute for both $x^{\exists}$ and $y^{\exists}$, get $\mathsf{0} = \mathsf{0}$, and the proof is successful. On the other hand, when we want to prove $x^{\exists} \neq y^{\exists}$ we can choose $\mathsf{0}$ to substitute for $x^{\exists}$ and $\mathsf{1}$ to substitute for $y^{\exists}$, get $\mathsf{0} \neq \mathsf{1}$, and the proof is successful, too.

4.1 Replacing ε-terms with Strong Free Universal Variables

An important difference between the inequations $\varepsilon x.\,\mathsf{true} \neq \varepsilon x.\,\mathsf{true}$ and $x^{\exists} \neq y^{\exists}$ is that the latter does not violate the reflexivity axiom. And we are going to cure the violation of the former immediately by introducing a new kind of free variables, namely the *strong free universal variables*, which we write like $x^{\forall,\mathrm{S}}$.

Now, instead of $\varepsilon x.\ \mathsf{true} \neq \varepsilon x.\ \mathsf{true}$ we write $x^{\forall,\mathrm{s}} \neq y^{\forall,\mathrm{s}}$ and remember what these strong free universal variables stand for by storing this into a function C, called a *choice-condition*: $C(x^{\forall,\mathrm{s}}) := C(y^{\forall,\mathrm{s}}) := \mathsf{true}$.

At first, suppose that our ε-terms are not subordinate to any outside binder that binds variables inside the ε-term. In this case, we replace each ε-term $\varepsilon z.\ A$ with a new strong free universal variable $z^{\forall,\mathrm{s}}$ and extend the partial function C by $C(z^{\forall,\mathrm{s}}) := A\{z \mapsto z^{\forall,\mathrm{s}}\}$.

By this procedure we eliminate all ε-terms without loosing any syntactical information, the substitutivity and reflexivity axioms are immediately regained, and the problems discussed above disappear.

Another reason for replacing the ε-terms with strong free universal variables is that the latter can solve the question whether a committed choice is required: We can express—on the one hand—a committed choice by using a single strong free universal variable and—on the other hand—a choice without commitment by using several variables with the same choice-condition. Indeed, this solves our problems with committed choice of Ex. 3.2: Now, again using (ε_1), $\exists x.\ (x \neq x)$ reduces to $x^{\forall,\mathrm{s}} \neq x^{\forall,\mathrm{s}}$ with $C(x^{\forall,\mathrm{s}}) := (x^{\forall,\mathrm{s}} \neq x^{\forall,\mathrm{s}})$ and the proof attempt immediately fails due to the now regained reflexivity axiom.

We still have to explain what to do with subordinate ε-terms. If the ε-term $\varepsilon z.\ A$ contains occurrences of variables $v_0, \ldots, v_{l-1}$ that are not bound by quantifiers or other binders inside but outside the ε-term, then we have to replace it with $z^{\forall,\mathrm{s}}(v_0)\cdots(v_{l-1})$ for a new strong free universal variable $z^{\forall,\mathrm{s}}$ and to extend the choice-condition C by $C(z^{\forall,\mathrm{s}}) := \lambda v_0.\ \ldots \lambda v_{l-1}.\ (A\{z \mapsto z^{\forall,\mathrm{s}}(v_0)\cdots(v_{l-1})\})$.

Example 4.1 (Higher-Order Choice-Condition) *(continuing Ex. 3.3)*
In our framework the elimination of the $\forall$- and $\exists$-quantifiers turns $\forall x.\ \exists y.\ \forall z.\ \mathsf{P}(x,y,z)$ into $\mathsf{P}(x_1^{\forall,\mathrm{s}}, y_2^{\forall,\mathrm{s}}, z_4^{\forall,\mathrm{s}})$, and the subordinate ε-terms are replaced with the following higher-order choice-condition:

$C(z_4^{\forall,\mathrm{s}}) = \neg\mathsf{P}(x_1^{\forall,\mathrm{s}}, y_2^{\forall,\mathrm{s}}, z_4^{\forall,\mathrm{s}})$ $C(y_2^{\forall,\mathrm{s}}) = \mathsf{P}(x_1^{\forall,\mathrm{s}}, y_2^{\forall,\mathrm{s}}, z_3^{\forall,\mathrm{s}}(y_2^{\forall,\mathrm{s}}))$
$C(z_3^{\forall,\mathrm{s}}) = \lambda y.\ \neg\mathsf{P}(x_1^{\forall,\mathrm{s}}, y, z_3^{\forall,\mathrm{s}}(y))$ $C(x_1^{\forall,\mathrm{s}}) = \neg\mathsf{P}(x_1^{\forall,\mathrm{s}}, y_1^{\forall,\mathrm{s}}(x_1^{\forall,\mathrm{s}}), z_2^{\forall,\mathrm{s}}(x_1^{\forall,\mathrm{s}}))$
$C(z_2^{\forall,\mathrm{s}}) = \lambda x.\ \neg\mathsf{P}(x, y_1^{\forall,\mathrm{s}}(x), z_2^{\forall,\mathrm{s}}(x))$ $C(y_1^{\forall,\mathrm{s}}) = \lambda x.\ \mathsf{P}(x, y_1^{\forall,\mathrm{s}}(x), z_1^{\forall,\mathrm{s}}(x)(y_1^{\forall,\mathrm{s}}(x)))$
$C(z_1^{\forall,\mathrm{s}}) = \lambda x.\ \lambda y.\ \neg\mathsf{P}(x, y, z_1^{\forall,\mathrm{s}}(x)(y))$

4.2 Variable-Conditions

In some inference systems Skolemization is unsound, e.g. for higher-order systems like the one in [8] or the system in [10] for *descente infinie*. In these systems a δ-step (according to Smullyan's classification) substitutes a previously universally quantified variable (say z in $\forall z.\ \neg A$) not with a new Skolem term say $f(x_0^{\exists}, \ldots, x_{l-1}^{\exists})$ but with a new free universal variable $y^{\forall}$ (resulting in $\neg A\{z \mapsto y^{\forall}\}$) and the dependencies $\{\ (x_i^{\exists}, y^{\forall}) \mid i \prec l\ \}$ are stored in a so-called *variable-condition*, which disallows the instantiation of the $x_i^{\exists}$ with $y^{\forall}$ and takes over the part of the occur-check of unification in the Skolemizing inference systems.

While the benefit of the introduction of free existential variables in γ-rules is to delay the choice of a witnessing term, it is sometimes unsound to instantiate a free existential variable $x^{\exists}$ with a term containing a free universal variable $y^{\forall}$ that was introduced later than $x^{\exists}$:

Example 4.2 $\exists x.\ \forall y.\ (x = y)$
is not deductively valid. We can start a proof attempt via:
γ-step: $\forall y.\ (x^{\exists} = y)$
δ-step: $(x^{\exists} = y^{\forall})$
Now, if we were allowed to substitute the free existential variable $x^{\exists}$ with the free universal variable $y^{\forall}$, we would get the tautology $(y^{\forall} = y^{\forall})$, i.e. we would have proved an invalid formula. In order to prevent this, the δ-step has to record $(x^{\exists}, y^{\forall})$ in the variable-condition, which disallows the instantiation step.

In order to restrict the possible instantiations as little as possible, we should keep our variable-conditions as small as possible. While in [1] Wolfgang Bibel is quite generous in that he lets his variable-conditions become quite big, the variable-conditions of Michaël Kohlhase in [8] are too small for soundness, cf. [10]. Our variable-conditions of Sect. 5.1 do not have these problems.

Note that variable-conditions do not add unnecessary complexity to applications where Skolemization is no problem and variable-conditions are not needed: Firstly, if they are superfluous we can work with an empty variable-condition as if there would be no variable-condition at all. Secondly—and more importantly—we need the variable-conditions anyway to guarantee that our choice-conditions do not end up in meaningless (vicious) cycles, cf. Ex. 5.11 below.

The above δ-step on $\forall z.\ \neg A$ may be *liberalized* or not. If it is not liberalized, the introduced free universal variable is a weak one. If the δ-step is liberalized, however, it introduces a new strong free universal variable say $y^{\forall,\mathrm{s}}$ with choice-condition $C(y^{\forall,\mathrm{s}}) := A\{z \mapsto y^{\forall,\mathrm{s}}\}$. The advantage of the liberalized version is a smaller variable-condition which enables additional proofs on the same level of multiplicity, cf. [10].

4.3 Instantiating Strong Free Universal Variables

We now have to explain how to replace strong free universal variables with terms that satisfy their choice-condition. Just like the free existential variables, the strong free universal variables must be substituted globally! While a free existential variable $x^{\exists}$ can be replaced with nearly everything, the replacement of a strong free universal variable $y^{\forall,\mathrm{s}}$ requires some proof work, and a weak free universal variable $z^{\forall,\mathrm{w}}$ cannot be instantiated at all (unless for lemma or induction hypothesis application, when the formula becomes a tool instead of a task). More precisely, a free existential variable $x^{\exists}$ may be instantiated with any term that does not violate the current variable-condition. The instantiation of a strong free universal variable $y^{\forall,\mathrm{s}}$ additionally requires some proof work. In general, if a substitution σ substitutes—possibly among other free existential and strong free universal variables—the strong free universal variable $y^{\forall,\mathrm{s}}$ for which the choice-condition $C(y^{\forall,\mathrm{s}}) = \lambda v_0.\ \ldots \lambda v_{l-1}.\ B$ is given, then we must prove

$$\forall v_0.\ \ldots \forall v_{l-1}.\ \big(\ \exists y.\ (B\{y^{\forall,\mathrm{s}}\mapsto y\})\quad\Rightarrow\quad B\ \big)\sigma$$

in order to know that the global instantiation with σ is a sound proof step.

Example 4.3 Suppose that our domain is natural numbers and that $y^{\forall,\mathrm{s}}$ has the choice-condition $C(y^{\forall,\mathrm{s}}) = (z^{\forall,\mathrm{w}} = y^{\forall,\mathrm{s}}{+}1)$. Then, before we may replace $y^{\forall,\mathrm{s}}$ with $\mathsf{p}(z^{\forall,\mathrm{w}})$, we have to prove $\exists y.\ (z^{\forall,\mathrm{w}} = y{+}1) \Rightarrow (z^{\forall,\mathrm{w}} = \mathsf{p}(z^{\forall,\mathrm{w}}){+}1)$, which is true for the predecessor function p.

Example 4.4 (Canossa 1077) *(continuing Ex. 3.1)* The situation of Ex. 3.1 now reads

$$\mathsf{HolyGhost} = z_0^{\forall,\mathrm{s}} \quad\wedge\quad \mathsf{Joseph} = z_1^{\forall,\mathrm{s}} \tag{Ex. 4.4.1}$$

with $C(z_0^{\forall,\mathrm{s}}) = \mathsf{Father}(z_0^{\forall,\mathrm{s}}, \mathsf{Jesus})$ and $C(z_1^{\forall,\mathrm{s}}) = \mathsf{Father}(z_1^{\forall,\mathrm{s}}, \mathsf{Jesus})$.
On the one hand, this does not bring the Pope into trouble because we do not know whether $z_0^{\forall,\mathrm{s}} = z_1^{\forall,\mathrm{s}}$. On the other hand, from knowing (Ex. 2.1.2) we can prove (Ex. 4.4.1) by substituting $z_0^{\forall,\mathrm{s}}$ with $\mathsf{HolyGhost}$ because this solves $C(z_0^{\forall,\mathrm{s}})$, and $z_1^{\forall,\mathrm{s}}$ with Joseph because this solves $C(z_1^{\forall,\mathrm{s}})$.

5 Formal Discussion of the ε's New Indefinite Semantics

We use '$\uplus$' for the union of disjoint classes and 'id' for the identity function. For classes R, A, and B we define ${}_A{\upharpoonleft}R := \{\,(a,b)\in R \mid a\in A\,\}$; $R{\upharpoonright}_B := \{\,(a,b)\in R \mid b\in B\,\}$; $\langle A\rangle R := \{\,b \mid \exists a\in A.\ (a,b)\in R\,\}$; $R\langle B\rangle := \{\,a \mid \exists b\in B.\ (a,b)\in R\,\}$. A relation R [on A] is *wellfounded* if any non-empty class B [$\subseteq A$] has an R-minimal element, i.e. $\exists a\in B.\ \neg\exists a'\in B.\ a'\,R\,a$.

We define a *sequent* to be a (disjunctive) list of formulas. We assume the following four sets of symbols to be disjoint: $\mathrm{V}_\exists$, *free existential variables*, i.e. the free variables of [4]; $\mathrm{V}_\forall$, *free universal variables*, i.e. nullary parameters, our substitute for Skolem functions; $\mathrm{V}_{\mathrm{bound}}$, *bound variables*, i.e. variables for bound use only; Σ, *constants*, i.e. the function and predicate symbols from the signature. We split the free universal variables $\mathrm{V}_\forall$ into *weak free universal variables* $\mathrm{V}_{\forall,\mathrm{w}}$ and *strong free universal variables* $\mathrm{V}_{\forall,\mathrm{s}}$: $\mathrm{V}_\forall = \mathrm{V}_{\forall,\mathrm{w}} \uplus \mathrm{V}_{\forall,\mathrm{s}}$. We define the *free variables* by $\mathrm{V}_{\mathrm{free}} := \mathrm{V}_\exists \uplus \mathrm{V}_\forall$ and the *variables* by $\mathrm{V} := \mathrm{V}_{\mathrm{bound}} \uplus \mathrm{V}_{\mathrm{free}}$. For k being any of these kinds of variables and Γ being a term, formula, sequent, &c., we use '$\mathcal{V}_k(\Gamma)$' to denote the set of variables from V_k occurring in Γ. For a substitution σ we denote with '$\Gamma\sigma$' the result of replacing in Γ each occurrence of a variable $x \in \mathrm{dom}(\sigma)$ with $\sigma(x)$. In default situations, we tacitly assume that all occurrences of variables from $\mathrm{V}_{\mathrm{bound}}$ in terms and formulas on top level and in the ranges of substitutions are *bound occurrences* (i.e. that a variable $x \in \mathrm{V}_{\mathrm{bound}}$ occurs only in the scope of a binder on x) and that each substitution σ satisfies $\mathrm{dom}(\sigma) \subseteq \mathrm{V}_{\mathrm{free}}$, so that no bound occurrences of variables can be replaced and no additional variable occurrences can become bound (i.e. captured) when applying σ.

Validity is expected to be given with respect to some Σ-structure $\mathcal{A}$, assigning a non-empty universe (or "carrier" or "object domain"). For $\mathrm{X} \subseteq \mathrm{V}$ we denote the set of total $\mathcal{A}$-valuations of X (i.e. functions mapping variables to objects of the universe of $\mathcal{A}$) with $\mathrm{X} \to \mathcal{A}$ and the set of (possibly) partial $\mathcal{A}$-valuations of X with $\mathrm{X} \leadsto \mathcal{A}$. For $\tau \in \mathrm{X} \to \mathcal{A}$ we denote with '$\mathcal{A}{\uplus}\tau$'

the extension of $\mathcal{A}$ to the variables of X. More precisely, we assume the existence of some evaluation function 'eval' such that $\mathrm{eval}(\mathcal{A}\uplus\tau)$ maps any term whose free occurring symbols are from $\Sigma\uplus\mathrm{X}$ into the universe of $\mathcal{A}$ (respecting types) s.t. for all $x \in \mathrm{X}$: $\mathrm{eval}(\mathcal{A}\uplus\tau)(x) = \tau(x)$. Moreover, $\mathrm{eval}(\mathcal{A}\uplus\tau)$ maps any formula B whose free occurring symbols are from $\Sigma\uplus\mathrm{X}$ to TRUE or FALSE, such that B is valid in $\mathcal{A}\uplus\tau$ iff $\mathrm{eval}(\mathcal{A}\uplus\tau)(B) = \mathsf{TRUE}$. We assume that the *Substitution-Lemma* (cf. e.g. [4, p. 120]) holds in the sense that, for any substitution σ, Σ-structure $\mathcal{A}$, valuation $\tau \in \mathrm{V} \rightsquigarrow \mathcal{A}$, and term or formula B, if the variables that occur free in $B\sigma$ belong to $\mathrm{dom}(\tau)$: $\mathrm{eval}(\mathcal{A}\uplus\tau)(B\sigma) = \mathrm{eval}\big(\mathcal{A} \uplus (\sigma \uplus {}_{\mathrm{V}\setminus\mathrm{dom}(\sigma)}\upharpoonright\mathrm{id}) \circ \mathrm{eval}(\mathcal{A}\uplus\tau)\big)\big(B\big)$. Note that we assume the operator '$\circ$' to have higher priority than the operators '$\cup$' and '$\uplus$'. Finally, we assume that the value of the evaluation function on a term or formula B does not depend on the variables that do not occur free in B: For X being the set of the variables that occur free in B, we require: $\mathrm{eval}(\mathcal{A}\uplus\tau)(B) = \mathrm{eval}(\mathcal{A} \uplus {}_{\mathrm{X}}\upharpoonright\tau)(B)$. Further properties of validity or evaluation are definitely not needed. Note that we have left open what our formulas and what our Σ-structures are. All we need are the above basic requirements. We can assume any classical logic here, esp. forms of higher-order logic. Moreover, we can even include those modal logics where the Substitution-Lemma only holds for the substitution of rigid and rigidified (grounded, annotated) terms, although we do not take care of the then necessary restrictions in this paper.

5.1 Existential Substitutions and Valuations

Several binary relations on free variables will be introduced. The overall idea is that when (x, y) occurs in such a relation this means something like "x is older than y" or "the value of y depends on or is described in terms of x".

Definition 5.1 (Variable-Condition)
A *variable-condition* is a subset of $\mathrm{V}_{\mathrm{free}} \times \mathrm{V}_{\mathrm{free}}$.

Definition 5.2 (E_σ, U_σ)
For a substitution σ we define the *existential relation* to be

$$E_\sigma \ := \ \{\ (z^{\exists}, x) \mid z^{\exists} \in \mathcal{V}_{\exists}(\sigma(x)) \wedge x \in \mathrm{dom}(\sigma)\ \},$$

and the *universal relation* to be

$$U_\sigma \ := \ \{\ (y^{\forall}, x) \mid y^{\forall} \in \mathcal{V}_{\forall}(\sigma(x)) \wedge x \in \mathrm{dom}(\sigma)\ \}.$$

Definition 5.3 ([[Quasi-]Existential] R-Substitution)
Let R be a variable-condition, cf. Def. 5.1.
σ is an *R-substitution* if σ is a substitution and $R \cup E_\sigma \cup U_\sigma$ is wellfounded.
σ is *existential* if $\mathrm{dom}(\sigma) \subseteq \mathrm{V}_{\exists}$. σ is *quasi-existential* if $\mathrm{dom}(\sigma) \subseteq \mathrm{V}_{\exists} \uplus \mathrm{V}_{\forall,\mathrm{S}}$.

Note that, regarding syntax, $(x, y^{\forall}) \in R$ is intended to mean that an R-substitution σ may not replace x with a term in which $y^{\forall}$ occurs, i.e. $(y^{\forall}, x) \in U_\sigma$ must be disallowed, i.e. $R \cup U_\sigma$ must be wellfounded. Note that in practice w.l.o.g. R, E_σ, and U_σ can always be chosen to be finite, so that $R \cup E_\sigma \cup U_\sigma$ is wellfounded iff it is acyclic.

After application of an R-substitution σ, in case of $(x, y^{\forall}) \in R$, we have to update our variable-condition R in order to ensure that x is not replaced with $y^{\forall}$ via a future application of another R-substitution that replaces a free variable say $u^{\exists}$ occurring in $\sigma(x)$ with $y^{\forall}$. In this case, the transitive closure of the updated variable-condition has to contain $(u^{\exists}, y^{\forall})$. But we have $u^{\exists}\ E_{\sigma}\ x\ R\ y^{\forall}$. This means that $R \cup E_{\sigma}$ must be a subset of the updated variable-condition. Besides this, we have to add steps with U_{σ} again.

Definition 5.4 (σ-Update)
Let R be a variable-condition and σ be a substitution.
The *σ-update of R* is $R \cup E_{\sigma} \cup U_{\sigma}$.

Let $\mathcal{A}$ be some Σ-structure. We now define a semantical counterpart of our existential R-substitutions, which we will call "existential $(\mathcal{A}, R)$-valuation". Suppose that e maps each free existential variable not directly to an object of $\mathcal{A}$, but can additionally read the values of some free universal variables under an $\mathcal{A}$-valuation $\tau \in \mathrm{V}_{\forall} \to \mathcal{A}$, i.e. e gets some $\tau' \in \mathrm{V}_{\forall} \leadsto \mathcal{A}$ with $\tau' \subseteq \tau$ as a second argument; short: $e : \mathrm{V}_{\exists} \to (\mathrm{V}_{\forall} \leadsto \mathcal{A}) \leadsto \mathcal{A}$. Moreover, for each free existential variable $x^{\exists}$, we require the set of read free universal variables (i.e. $\mathrm{dom}(\tau')$) to be identical for all τ; i.e. there has to be some "semantical relation" $S_e \subseteq \mathrm{V}_{\forall} \times \mathrm{V}_{\exists}$ such that for all $x^{\exists} \in \mathrm{V}_{\exists}$: $e(x^{\exists}) : (S_e\langle\!\langle\{x^{\exists}\}\rangle\!\rangle \to \mathcal{A}) \to \mathcal{A}$. Note that, for each $e : \mathrm{V}_{\exists} \to (\mathrm{V}_{\forall} \leadsto \mathcal{A}) \leadsto \mathcal{A}$, at most one semantical relation exists, namely $S_e := \{\ (y^{\forall}, x^{\exists}) \mid y^{\forall} \in \mathrm{dom}(\bigcup(\mathrm{dom}(e(x^{\exists})))) \wedge x^{\exists} \in \mathrm{V}_{\exists}\ \}$.

Definition 5.5 (Semantical Relation (S_e))
Let $e : X \to (Y \leadsto \mathcal{A}) \leadsto \mathcal{A}$. The *semantical relation of* e is
$$S_e := \{\ (y, x) \mid y \in \mathrm{dom}(\bigcup(\mathrm{dom}(e(x)))) \wedge x \in X\ \}.$$
e is *semantical* if $e(x) : (S_e\langle\!\langle\{x\}\rangle\!\rangle \to \mathcal{A}) \to \mathcal{A}$ for all $x \in X$.

Definition 5.6 (Existential $(\mathcal{A}, R)$-Valuation) Let R be a variable-condition and $\mathcal{A}$ a Σ-structure. e is an *existential $(\mathcal{A}, R)$-valuation* if $e : \mathrm{V}_{\exists} \to (\mathrm{V}_{\forall} \leadsto \mathcal{A}) \leadsto \mathcal{A}$, e is semantical, and $R \cup S_e$ is wellfounded.

Definition 5.7 (ϵ) We define the function
$$\epsilon : \quad (X \to (Y \leadsto \mathcal{A}) \leadsto \mathcal{A}) \quad \leadsto \quad (Y \to \mathcal{A}) \to X \to \mathcal{A}$$
by ($e : X \to (Y \leadsto \mathcal{A}) \leadsto \mathcal{A}$ being semantical, $\tau \in Y \to \mathcal{A}$, $x \in X$)
$$\epsilon(e)(\tau)(x) := e(x)({}_{S_e\langle\!\langle\{x\}\rangle\!\rangle}\!\upharpoonright\!\tau).$$

We are now going to define R-validity of a set of sequents with free variables, in terms of validity of a formula.

Definition 5.8 (Validity) Let R be a variable-condition, $\mathcal{A}$ a Σ-structure, and G a set of sequents. G is *R-valid in $\mathcal{A}$* if there is an existential $(\mathcal{A}, R)$-valuation e such that G is $(e, \mathcal{A})$-valid. G is *$(e, \mathcal{A})$-valid* if G is $(\tau, e, \mathcal{A})$-valid for all $\tau \in \mathrm{V}_{\forall} \to \mathcal{A}$. G is *$(\tau, e, \mathcal{A})$-valid* if G is valid in $\mathcal{A} \uplus \epsilon(e)(\tau) \uplus \tau$. G is *valid in $\mathcal{A}$* if Γ is valid in $\mathcal{A}$ for all $\Gamma \in G$. A sequent Γ is *valid in $\mathcal{A}$* if there is some formula listed in Γ that is valid in $\mathcal{A}$.

Example 5.9 (Validity)
For $x^{\exists} \in \mathrm{V}_{\exists}$, $y^{\forall} \in \mathrm{V}_{\forall}$, the sequent $x^{\exists}{=}y^{\forall}$ is $\emptyset$-valid in any $\mathcal{A}$ because we can choose $S_e := \mathrm{V}_{\forall}{\times}\mathrm{V}_{\exists}$ and $e(x^{\exists})(\tau) := \tau(y^{\forall})$ for $\tau \in \mathrm{V}_{\forall} \to \mathcal{A}$, resulting in $\epsilon(e)(\tau)(x^{\exists}) = e(x^{\exists})({}_{S_e\langle\{x^{\exists}\}\rangle}{\upharpoonright}\tau) = e(x^{\exists})({}_{\mathrm{V}_{\forall}}{\upharpoonright}\tau) = \tau(y^{\forall})$. This means that $\emptyset$-validity of $x^{\exists}{=}y^{\forall}$ is the same as validity of $\forall y.\ \exists x.\ x{=}y$. Moreover, note that $\epsilon(e)(\tau)$ has access to the τ-value of $y^{\forall}$ just as a raising function f for x in the raised (i.e. dually Skolemized) version $f(y^{\forall}){=}y^{\forall}$ of $\forall y.\ \exists x.\ x{=}y$.

Contrary to this, for $R := \mathrm{V}_{\exists}{\times}\mathrm{V}_{\forall}$, the same formula $x^{\exists}{=}y^{\forall}$ is not R-valid in general because then the required irreflexivity of $S_e{\circ}R$ implies $S_e{=}\emptyset$, and $e(x^{\exists})({}_{S_e\langle\{x^{\exists}\}\rangle}{\upharpoonright}\tau) = e(x^{\exists})({}_{\emptyset}{\upharpoonright}\tau) = e(x^{\exists})(\emptyset)$ cannot depend on $\tau(y^{\forall})$ anymore. This means that $(\mathrm{V}_{\exists}{\times}\mathrm{V}_{\forall})$-validity of $x^{\exists}{=}y^{\forall}$ is the same as validity of $\exists x.\ \forall y.\ x{=}y$. Moreover, note that $\epsilon(e)(\tau)$ has no access to the τ-value of $y^{\forall}$ just as a raising function c for x in the raised version $c{=}y^{\forall}$ of $\exists x.\forall y.\ x{=}y$.

For a more general example let $G = \{\ A_{i,0} \ldots A_{i,n_i-1} \mid i \in I\ \}$, where for $i \in I$ and $j \prec n_i$ the $A_{i,j}$ are formulas with free existential variables from $\boldsymbol{x}$ and free universal variables from $\boldsymbol{y}$. Then $(\mathrm{V}_{\exists}{\times}\mathrm{V}_{\forall})$-validity of G means validity of

$$\exists \boldsymbol{x}.\ \forall \boldsymbol{y}.\ \forall i \in I.\ \exists j \prec n_i.\ A_{i,j}$$

whereas $\emptyset$-validity of G means validity of

$$\forall \boldsymbol{y}.\ \exists \boldsymbol{x}.\ \forall i \in I.\ \exists j \prec n_i.\ A_{i,j}$$

5.2 Choice-Conditions

Definition 5.10 (Choice-Condition, Extension)
C is an R-*choice-condition* if R is a wellfounded variable-condition, C is a partial function from $\mathrm{V}_{\forall\mathrm{s}}$ into the set of formulas, and $z\ R^{+}\ y^{\forall\mathrm{s}}$ for all $y^{\forall\mathrm{s}} \in \mathrm{dom}(C)$ and $z \in \mathcal{V}_{\mathrm{free}}(C(y^{\forall\mathrm{s}}))\setminus\{y^{\forall\mathrm{s}}\}$.
More generally, the values of C can be formula-valued λ-terms where,
for $y^{\forall\mathrm{s}} \in \mathrm{dom}(C)$ and $C(y^{\forall\mathrm{s}}) = \lambda v_0.\ \ldots \lambda v_{l-1}.\ B$, B is a formula whose free occurring variables from $\mathrm{V}_{\mathrm{bound}}$ are among $\{v_0, \ldots, v_{l-1}\} \subseteq \mathrm{V}_{\mathrm{bound}}$ and where, for $v_0 : \alpha_0, \ldots, v_{l-1} : \alpha_{l-1}$, we have $y^{\forall\mathrm{s}} : \alpha_0 \to \ldots \to \alpha_{l-1} \to \alpha_l$ for some type α_l, and any occurrence of $y^{\forall\mathrm{s}}$ in B must be of the form $y^{\forall\mathrm{s}}(v_0)\cdots(v_{l-1})$.
(C', R') is an *extension of* (C, R) if C is an R-choice-condition, C' is an R'-choice-condition, $C{\subseteq}C'$, and $R{\subseteq}R'$.

Example 5.11 (Choice-Condition) *(continuing Ex. 4.1)*
On the one hand, if R is a wellfounded variable-condition that satisfies

$$z_1^{\forall\mathrm{s}}\ R\ y_1^{\forall\mathrm{s}}\ R\ z_2^{\forall\mathrm{s}}\ R\ x_1^{\forall\mathrm{s}}\ R\ z_3^{\forall\mathrm{s}}\ R\ y_2^{\forall\mathrm{s}}\ R\ z_4^{\forall\mathrm{s}},$$

then the C of Ex. 4.1 is an R-choice-condition, indeed.
On the other hand, if some clever person would like to do the complete quantifier-elimination of Ex. 4.1 by the simpler $C'(z_4^{\forall\mathrm{s}}) = \neg\mathsf{P}(x_1^{\forall\mathrm{s}}, y_2^{\forall\mathrm{s}}, z_4^{\forall\mathrm{s}})$, $C'(y_2^{\forall\mathrm{s}}) = \mathsf{P}(x_1^{\forall\mathrm{s}}, y_2^{\forall\mathrm{s}}, z_4^{\forall\mathrm{s}})$, $C'(x_1^{\forall\mathrm{s}}) = \neg\mathsf{P}(x_1^{\forall\mathrm{s}}, y_2^{\forall\mathrm{s}}, z_4^{\forall\mathrm{s}})$, then he would—among other things—need $z_4^{\forall\mathrm{s}}\ R\ y_2^{\forall\mathrm{s}}\ R\ z_4^{\forall\mathrm{s}}$, which renders R non-wellfounded. Thus, this C' cannot be an R-choice-condition for any R.

Definition 5.12 (Compatibility)
Let C be an R-choice-condition, $\mathcal{A}$ a Σ-structure, and e an existential $(\mathcal{A},R)$-valuation. π is $(e,\mathcal{A})$-*compatible with* (C,R) if

1. $\pi : \mathrm{V}_{\forall,\mathrm{s}} \to (\mathrm{V}_{\forall,\mathrm{w}} \leadsto \mathcal{A}) \leadsto \mathcal{A}$ is semantical (cf. Def. 5.5) and $R \cup S_e \cup S_\pi$ is wellfounded.
2. For all $y^{\forall,\mathrm{s}} \in \mathrm{dom}(C)$ with $C(y^{\forall,\mathrm{s}}) = \lambda v_0.\ \ldots \lambda v_{l-1}.\ B$ and $\tau \in \mathrm{V}_{\forall,\mathrm{w}} \to \mathcal{A}$ and $\chi \in \{v_0,\ldots,v_{l-1}\} \to \mathcal{A}$:

   If, for some $\eta \in \{y^{\forall,\mathrm{s}}\} \to \mathcal{A}$,
   B is $\big({}_{\mathrm{V}_{\forall,\mathrm{s}}\setminus\{y^{\forall,\mathrm{s}}\}}\!\upharpoonright\!(\epsilon(\pi)(\tau)) \uplus \eta \uplus \tau \uplus \chi,\ \ e,\ \ \mathcal{A}\big)$-valid,
   then B is $(\epsilon(\pi)(\tau) \uplus \tau \uplus \chi, e, \mathcal{A})$-valid.

Item 1 of this definition is quite technical and needed for lemma application. Roughly speaking, it says that the flow of information between variables expressed in R, e, and π is acyclic.

Item 2 of $(e,\mathcal{A})$-compatibility of π with say $\big(\{(y^{\forall,\mathrm{s}},\ \lambda v_0.\ \ldots \lambda v_{l-1}.\ B)\},\ \ R\big)$ means that a different choice for the $\epsilon(\pi)(\tau)$-value of $y^{\forall,\mathrm{s}}$ cannot give rise to a previously not given validity of the formula B in $\mathcal{A} \uplus \epsilon\big(e\big)\big(\epsilon(\pi)(\tau) \uplus \tau\big) \uplus \epsilon(\pi)(\tau) \uplus \tau \uplus \chi$, or that $\epsilon(\pi)(\tau)(y^{\forall,\mathrm{s}})$ is chosen such that B becomes valid if such a choice is possible. This is closely related to Hilbert's ε-operator in the sense that $y^{\forall,\mathrm{s}}$ is given the value of $\lambda v_0.\ \ldots \lambda v_{l-1}.\ \varepsilon y.\ (B\{y^{\forall,\mathrm{s}}\mapsto y\})$ for $y \in \mathrm{V}_{\mathrm{bound}}\setminus\mathcal{V}(B)$.

Since the choice for $y^{\forall,\mathrm{s}}$ depends on the other free variables of $\lambda v_0.\ \ldots \lambda v_{l-1}.\ B$ (i.e. the free variables of $\lambda v_0.\ \ldots \lambda v_{l-1}.\ \varepsilon y.\ (B\{y^{\forall,\mathrm{s}}\mapsto y\})$), we required the inclusion of this dependency into the transitive closure of R in Def. 5.10.

Note that the empty function $\emptyset$ is an R-choice-condition for any wellfounded (which in the following will always be the case) variable-condition R. Moreover, any $\pi : \mathrm{V}_{\forall,\mathrm{s}} \to \{\emptyset\} \to \mathcal{A}$ is $(e,\mathcal{A})$-compatible with $(\emptyset, R)$ due to $S_\pi = \emptyset$. Indeed, a compatible π always exists:

Lemma 5.13 *Let C be an R-choice-condition, $\mathcal{A}$ a Σ-structure, e an existential $(\mathcal{A},R)$-valuation. Now, there is some π that is $(e,\mathcal{A})$-compatible with (C,R).*

Finally, we need means for expressing the requirement on a quasi-existential substitution that it replaces the strong free universal variables in a way that goes together well with the compatibility of Def. 5.12(2):

Definition 5.14 ($Q_{C,\sigma}$)
For a substitution σ and an R-choice-condition C, we require $Q_{C,\sigma}$ to be a function from $\mathrm{dom}(C) \cap \mathrm{dom}(\sigma)$ into the set of formulas such that for each $y^{\forall,\mathrm{s}} \in \mathrm{dom}(C) \cap \mathrm{dom}(\sigma)$ with $C(y^{\forall,\mathrm{s}}) = \lambda v_0.\ \ldots \lambda v_{l-1}.\ B$, we have $Q_{C,\sigma}(y^{\forall,\mathrm{s}}) = \forall v_0.\ \ldots \forall v_{l-1}.\ \big(\ \exists y.\ (B\{y^{\forall,\mathrm{s}}\mapsto y\})\ \ \Rightarrow\ \ B\ \big)\sigma$ for some $y \in \mathrm{V}_{\mathrm{bound}}\setminus\mathcal{V}(C(y^{\forall,\mathrm{s}}))$.

Definition 5.15 (Extended σ-Update) Let C be an R-choice-condition and σ a substitution. The *extended σ-update* (C',R') *of* (C,R) is given by:
$C' \ := \ \{\ (x, B\sigma) \mid (x,B) \in C \wedge x \notin \mathrm{dom}(\sigma)\ \}$, R' is the σ-update of R.

Lemma 5.16 (Extended σ-Update)
Let C be an R-choice-condition, σ an R-substitution, and (C',R') the extended σ-update of (C,R). Now: C' is an R'-choice-condition.

5.3 Strong Validity

Definition 5.17 (Strong Validity)
Let C be an R-choice-condition, $\mathcal{A}$ a Σ-structure, and G a set of sequents.
G is (C,R)*-strongly valid in* $\mathcal{A}$ if there is an existential $(\mathcal{A},R)$-valuation e such that G is (C,R)-strongly $(e,\mathcal{A})$-valid. G is (C,R)*-strongly* $(e,\mathcal{A})$*-valid* if G is strongly $(\pi,e,\mathcal{A})$-valid for some π that is $(e,\mathcal{A})$-compatible with (C,R). G is *strongly* $(\pi,e,\mathcal{A})$*-valid* if G is $(\epsilon(\pi)(\tau)\uplus\tau,e,\mathcal{A})$-valid for each $\tau\in\mathrm{V}_{\forall,\mathrm{w}}\to\mathcal{A}$. The rest is given by Def. 5.8.

Note that strong validity is called "strong" because it treats the strong free universal variables properly, whereas (weak) validity of Def. 5.8 does not. It is generally not the case, however, that strong validity is logically stronger than weak validity.

It should be pointed out that the "some π" in the above definition is something we can play around with, e.g. even "each π" is possible here.

Example 5.18 In our framework the formula (E2) of Sect. 3.3 looks like $\forall x.\,(A\Leftrightarrow B)\quad\Rightarrow\quad x_0^{\forall,\mathrm{s}}=x_1^{\forall,\mathrm{s}}$ for some new strong free universal variables $x_0^{\forall,\mathrm{s}}$ and $x_1^{\forall,\mathrm{s}}$ with $R=\mathcal{V}_{\mathrm{free}}(A)\times\{x_0^{\forall,\mathrm{s}}\}\cup\mathcal{V}_{\mathrm{free}}(B)\times\{x_1^{\forall,\mathrm{s}}\}$ and R-choice-condition C $=\{(x_0^{\forall,\mathrm{s}},A\{x\mapsto x_0^{\forall,\mathrm{s}}\}),\,(x_1^{\forall,\mathrm{s}},B\{x\mapsto x_1^{\forall,\mathrm{s}}\})\}$. On the one hand, for the definition via "some π", this formula becomes (C,R)-strongly valid because after choosing a value for $x_0^{\forall,\mathrm{s}}$ we can take the same value for $x_1^{\forall,\mathrm{s}}$ simply because $x_1^{\forall,\mathrm{s}}$ is new and can read all weak free universal variables. On the other hand, for the definition via "each π", this formula only becomes valid when there is a unique solution for the choice-condition on $x_0^{\forall,\mathrm{s}}$.

The reason why we prefer "some π" to "each π" here and in [10] is that the latter results in more valid formulas (e.g. (E2)) and makes theorem proving easier. Contrary to "each π" and to all semantics in the literature, the "some π" frees us from considering all possible choices: We just have to pick a single one and fix it in a proof step. As the major notion here and in [10] is not strong validity but reduction (cf. Def. 5.19), where the quantification of π must be universal no matter how we quantify in the notion of strong validity, changing the quantification of π in Def. 5.17 would only have very local consequences. Roughly speaking, only Theor. 5.20(a) for the case of $O\neq\emptyset$ becomes false for a different choice on the quantification of π in Def. 5.17.

5.4 Reduction

Definition 5.19 (Reduction)
Let C be an R-choice-condition, $\mathcal{A}$ a Σ-structure, and G_0, G_1 sets of sequents.
G_0 (C,R)*-reduces to* G_1 *in* $\mathcal{A}$ if for each existential $(\mathcal{A},R)$-valuation e and each π that is $(e,\mathcal{A})$-compatible with (C,R):

if G_1 is strongly $(\pi,e,\mathcal{A})$-valid, then G_0 is strongly $(\pi,e,\mathcal{A})$-valid.

Theorem 5.20 (Reduction) *Let C be an R-choice-condition; $\mathcal{A}$ a Σ-structure; G_0 and G_1 sets of sequents. For a quasi-existential R-substitution σ and the extended σ-update (C', R') of (C, R), and for O, N with $O \subseteq \mathrm{dom}(C) \cap \mathrm{dom}(\sigma) \subseteq O \uplus N$, $N \subseteq \mathrm{dom}(C) \setminus O$, $\mathrm{dom}(C) \cap \langle N \rangle R^+ \subseteq N$, $N \cap \mathcal{V}(G_0, G_1) = \emptyset$, and $\mathcal{V}_{\forall,\mathrm{s}}(\mathrm{ran}(\sigma)) \cap (\mathrm{dom}(\sigma) \cup N) = \emptyset$:*

(a) If $G_0\sigma \cup \langle O \rangle Q_{C,\sigma}$ is (C', R')-strongly valid in $\mathcal{A}$, then G_0 is (C, R)-strongly valid in $\mathcal{A}$.

(b) If G_0 (C, R)-reduces to G_1 in $\mathcal{A}$, then $G_0\sigma$ (C', R')-reduces to $G_1\sigma \cup \langle O \rangle Q_{C,\sigma}$ in $\mathcal{A}$.

Suppose that G_0 is a proposition we want to prove and G_1 represents the current state of our proof attempt with variable-condition R and R-choice-condition C. Then G_0 (C, R)-reduces to G_1. E.g., in [10, Sect. 14] G_0 is $z_0^{\exists}(x_0^{\forall,\mathrm{w}})(y_0^{\forall,\mathrm{w}}) \prec \mathsf{ack}(x_0^{\forall,\mathrm{w}}, y_0^{\forall,\mathrm{w}})$, which says that Ackermann's function has a lower bound that is to be determined during the proof, and the current state G_1 of the proof attempt reduces to a known lemma when we apply the substitution $\sigma := \{z_0^{\exists} \mapsto \lambda x.\ \lambda y.\ y\}$. Now, setting $O := N := \emptyset$, Theor. 5.20(b) says that the instantiated (and $\lambda\beta$-reduced) theorem $y_0^{\forall,\mathrm{w}} \prec \mathsf{ack}(x_0^{\forall,\mathrm{w}}, y_0^{\forall,\mathrm{w}})$ (C', R')-reduces to the instantiated proof state and thus is (C', R')-strongly valid.

Let us have a careful look at the assumptions of Theor. 5.20, which is by far the deepest result in this paper. When strong free universal variables are replaced besides free existential variables, we cannot set O and N to the empty set as above. We use the set O to form the set of formulas $\langle O \rangle Q_{C,\sigma}$ (cf. Def. 5.14) that—as explained around Ex. 4.3—captures the soundness of the substitution σ. The possibility that the set O is a proper subset of the set $\mathrm{dom}(C) \cap \mathrm{dom}(\sigma)$ (of strong free universal variables that are conditioned by C and replaced by σ) is not a sophistication but of practical importance as will be explained later. The set N that contains the other variables from $\mathrm{dom}(C) \cap \mathrm{dom}(\sigma)$ should be thought of as the set of variables that are not relevant for our local context. Thus, these variables must not occur in G_0 and G_1. Moreover, N must be closed according to $\mathrm{dom}(C) \cap \langle N \rangle R^+ \subseteq N$, which says that the C-conditioned variables that depend on the variables of N are in N again, i.e. also not relevant for our local context. Finally, the assumption of $\mathcal{V}_{\forall,\mathrm{s}}(\mathrm{ran}(\sigma)) \cap (\mathrm{dom}(\sigma) \cup N) = \emptyset$ simply says that the strong free universal variables that are globally replaced by σ are not at the same time introduced again by σ. Now we still have to explain why we need the set N.

Example 5.21 *(continuing Ex. 4.3)*
Suppose that $\exists y.\ (z^{\forall,\mathrm{w}} = y{+}1) \Rightarrow (z^{\forall,\mathrm{w}} = \mathsf{p}(z^{\forall,\mathrm{w}}){+}1)$ is one of our lemmas and that we want to use this lemma as justification for replacing $y^{\forall,\mathrm{s}}$ under R-choice-condition $C(y^{\forall,\mathrm{s}}) = (z^{\forall,\mathrm{w}} = y^{\forall,\mathrm{s}}{+}1)$ globally with $\mathsf{p}(z^{\forall,\mathrm{w}})$, which is possible because it is equal to $Q_{C,\sigma}(y^{\forall,\mathrm{s}})$ for $\sigma := \{y^{\forall,\mathrm{s}} \mapsto \mathsf{p}(z^{\forall,\mathrm{w}})\}$.

The problem we get here is that the lemma is applied globally. Esp. problematic is the possibility that $y^{\forall,\mathrm{s}}$ occurs in the proof of the lemma itself. Unless we are very careful, the lemma becomes a lemma of itself, which results in a cyclic lemma application relation of a globally applied lemma. Then our whole proof work is in vain because no validities can be inferred. Therefore, we either have to reintroduce the lemma as an open lemma and prove it again or had better take a closer look on which of our (possibly open) lemmas L really depend on the justifying lemma $Q_{C,\sigma}(y^{\forall,\mathrm{s}})$ after global application of σ.

Let us sketch how to do the latter now. Let M be the set of all strong free universal variables that occur in the lemma L, in its open goals (i.e. its state of proof attempt), or in the lemmas that have been applied in its proof attempt. Set $O := \mathrm{dom}(C) \cap \mathrm{dom}(\sigma) \cap R^*\langle M \rangle$ and $N := \mathrm{dom}(C) \cap \langle(\mathrm{dom}(C) \cap \mathrm{dom}(\sigma)) \setminus O\rangle R^*$. Now, setting $G_0 := \{L\}$ and G_1 to the open goals and applied lemmas of the proof attempt of L, as a consequence of Theor. 5.20(b)—provided that $\mathcal{V}_{\forall,\mathrm{s}}(\mathrm{ran}(\sigma)) \cap (\mathrm{dom}(\sigma) \cup N) = \emptyset$; the other conditions on O and N being automatically satisfied—we know that $L\sigma$ (which now depends on $G_1\sigma$ w.r.t. reduction) depends on the lemma $Q_{C,\sigma}(y^{\forall,\mathrm{s}})$ *only if* $y^{\forall,\mathrm{s}} \in O$. With these refined dependencies we get a smaller lemma application relation that breaks the cycles discussed in Ex. 5.21 for reasonable substitutions σ. Detailed and further information on this can be found in Theor. 10.6 of [10] and its proof.

6 Conclusion

We have proposed a semantics for Hilbert's ε which ...

- is similar to the referential use of the indefinite article in natural languages.
- makes obvious the requirements on a commitment of the choice, cf. Sect. 4.1.
- is based on an abstract formal approach that extends a semantics for closed formulas (satisfying only weak requirements, cf. Sect. 5) to a semantics with several kinds of free variables (universal, existential, ε-constrained), which is esp. designed for developing proofs in the style of a working mathematician.
- makes proof work most simple because we do not have to consider all proper choices (as in all other approaches) but only a single arbitrary one, which is fixed in a proof step just as choices are settled in program steps, cf. Sect. 4.

[1] Wolfgang Bibel (1987). *Automated Theorem Proving.* 2nd rev. ed., Vieweg, Braunschweig.

[2] Wolfgang Bibel, Peter H. Schmitt (eds.) (1998). *Automated Deduction — A Basis for Applications.* Kluwer.

[3] Andreas Blass, Yuri Gurevich (2000). *The Logic of Choice.* J. Symbolic Logic **65**, pp. 1264-1310.

[4] Melvin C. Fitting (1996). *First-Order Logic and Automated Theorem Proving.* 2nd extd. ed., Springer.

[5] Martin Giese, Wolfgang Ahrendt (1999). *Hilbert's ε-Terms in Automated Theorem Proving.* 8th TABLEAU 1999, LNAI 1617, pp. 171-185, Springer.

[6] David Hilbert (1928). *Die Grundlagen der Mathematik.* Abhandlungen aus dem mathematischen Seminar der Univ. Hamburg **6**, pp. 65-85.

[7] David Hilbert, Paul Bernays (1968/70). *Grundlagen der Mathematik.* 2nd ed., Springer.

[8] Michaël Kohlhase (1998). *Higher-Order Automated Theorem Proving.* In: [2], Vol. 1, pp. 431-462.

[9] A. C. Leisenring (1969). *Mathematical Logic and Hilbert's ε-Symbol.* Gordon and Breach, New York.

[10] Claus-Peter Wirth (2000). *Descente Infinie + Deduction.* Report 737/2000, FB Informatik, Univ. Dortmund. Extd. version, Feb. 1, 2002. `http://ags.uni-sb.de/~cp/p/tab99/new.html (Feb. 01, 2002).`