$\models \exists x.\, F_1$	By assumption
$\models [t/x]F_1$ for some t	By inversion
$\Gamma \stackrel{-}{\Longrightarrow} [t/x]A_1 \setminus [t/x]F_1$	By substitution for parameter a
$\Gamma \stackrel{-}{\Longrightarrow} [t/x]A_1$	By i.h.
$\Gamma \stackrel{-}{\Longrightarrow} \exists x.\, A_1$	By rule $\exists$R

Case:

$$\mathcal{R} = \dfrac{\begin{array}{c}\mathcal{R}_1 \\ \Gamma \stackrel{-}{\Longrightarrow} [a/x]A_1 \setminus [a/x]F_1\end{array}}{\Gamma \stackrel{-}{\Longrightarrow} \forall x.\, A_1 \setminus \forall x.\, F_1}\,\forall\mathrm{R}^a$$

$\models \forall x.\, F_1$	By assumption
$\models [b/x]F_1$ for a new parameter b	By inversion
$\models [a/x]F_1$	By substititution of a for b
$\Gamma \stackrel{-}{\Longrightarrow} [a/x]A_1$	By i.h.
$\Gamma \stackrel{-}{\Longrightarrow} \forall x.\, A_1$	By rule $\forall$R

□

The opposite direction is more difficult. The desired theorem:

If $\Gamma \stackrel{-}{\Longrightarrow} A$ *then* $\Gamma \stackrel{-}{\Longrightarrow} A \setminus F$ *for some* F *with* $\models F$

cannot be proved directly by induction, since the premisses of the two derivations are different in the $\exists$R and $\forall$L rules. However, one can be obtained from the other by substituting terms for parameters. Since this must be done simultaneously, we introduce a new notation.

$$\textit{Parameter Substitution} \quad \rho \quad ::= \quad \cdot \mid \rho, t/a$$

We assume all the parameters a substituted for by ρ are distinct to avoid ambiguity. We write $A[\rho]$, $F[\rho]$, and $\Gamma[\rho]$, for the result of applying the substitution ρ to a proposition, formula, or context, respectively.

Lemma 4.10 *If* $\Gamma \stackrel{-}{\Longrightarrow} A$ *where* $A = A'[\rho]$, $\Gamma = \Gamma'[\rho]$ *then* $\Gamma' \stackrel{-}{\Longrightarrow} A' \setminus F$ *for some* F *such that* $\models F[\rho]$.

Proof: The proof proceeds by induction on the structure of the given derivation $\mathcal{D}$. We show only two cases, the second of which required the generalization of the induction hypothesis.

Case:

$$\mathcal{D} = \dfrac{}{\Gamma_1, P \stackrel{-}{\Longrightarrow} P}\,\mathrm{init}$$

$\Gamma_1 = \Gamma_1'[\rho]$, $P = P'[\rho]$, and $P = P''[\rho]$	Assumption
$\Gamma_1', P' \stackrel{=}{\Longrightarrow} P'' \setminus P' \doteq P''$	By rule init
$\models P'[\rho] \doteq P''[\rho]$	By rule $\doteq$ I

Case:

$$\mathcal{D} = \cfrac{\begin{array}{c}\mathcal{D}_1 \\ \Gamma \stackrel{=}{\Longrightarrow} [t/x]A_1\end{array}}{\Gamma \stackrel{=}{\Longrightarrow} \exists x.\ A_1}\ \exists\mathrm{R}$$

$\exists x.\ A_1 = A'[\rho]$	Assumption
$A' = \exists x.\ A_1'$ for a new parameter a with	
$[a/x]A_1 = ([a/x]A_1')[\rho, a/a]$	By definition of substitution
$[t/x]A_1 = ([a/x]A_1')[\rho, t/a]$	By substitution for parameter a
$\Gamma = \Gamma'[\rho]$	Assumption
$\Gamma'[\rho] = \Gamma'[\rho, t/a]$	Since a is new
$\Gamma' \stackrel{=}{\Longrightarrow} [a/x]A_1' \setminus [a/x]F_1$, and	
$\models ([a/x]F_1)[\rho, t/a]$	By i.h.
$\Gamma' \stackrel{=}{\Longrightarrow} \exists x.\ A_1' \setminus \exists x.\ F_1$	By rule $\exists$R
$\models (\exists x.\ F_1)[\rho]$	By rule $\exists$R and definition of substitution

□

Theorem 4.11 (Completeness of Equality Residuation)
If $\Gamma \stackrel{=}{\Longrightarrow} A$ *then* $\Gamma \stackrel{=}{\Longrightarrow} A \setminus F$ *for some* F *and* $\models F$.

Proof: From Lemma 4.10 with $A' = A$, $\Gamma' = \Gamma$, and ρ the identity substitution on the parameters in Γ and A. □

Next we describe an algorithm for proving residuated formulas, that is, an algorithm for unification. We do this in two steps: first we solve the problem in the fragment without parameters and universal quantifiers and then we extend the solution to the general case.

There are numerous ways for describing unification algorithms in the literature. We describe the computation of the algorithm as the bottom-up search for the derivation of a judgment. We restrict the inference rules such that they are essentially deterministic, and the inference rules themselves can be seen as describing an algorithm. This algorithm is in fact quite close to the implementation of it in ML which is available together with these notes.

In order to describe the algorithm in this manner, we need to introduce *existential variables* (often called *meta-variables* or *logic variables*) which are place-holders for the terms to be determined by unification. We use X and Y to stand for existential variables. Existential variables are different from parameters which are interpreted *universally*: all instances of a derivation with a parameter are valid. Existential variables in a derivation require only one instance to arrive at a valid derivation. While parameters are always local to a subderivation, we consider existential variables to be global in a derivation.[4]

The second concept we need is that of a *substitution* for existential variables. We use a new notation, because this form of substitution is quite different from substitutions for bound variables x or parameters a.

$$\textit{Substitutions} \quad \theta \quad ::= \quad \cdot \mid \theta, X \mapsto t$$

We require that all variables X defined by a substitution are distinct. We write $\mathrm{dom}(\theta)$ for the variables defined by a substitution and $\mathrm{cod}(\theta)$ for all the variables occuring in the terms t. For a ground substitution $\mathrm{cod}(\theta)$ is empty. For the technical development it is convenient to assume that the domain and co-domain of a substitution share no variables. This rules out "circular" substitutions such as $X \mapsto f(X)$ and it also disallows identity substitutions $X \mapsto X$. The latter restriction can be dropped, but it does no harm and is closer to the implementation. As for contexts, we consider the order of the definitions in a substitution to be irrelevant.

We write $t[\theta]$, $A[\theta]$, and $\Gamma[\theta]$ for the application of a substitution to a term, proposition, or context. This is defined to be the identity on existential variables which are not explicitly defined in the substitution.

We also need an operation of composition, written as $\theta_1 \circ \theta_2$ with the property that $t[\theta_1 \circ \theta_2] = (t[\theta_1])[\theta_2]$ and similarly for propositions and contexts. Composition is defined by

$$\begin{aligned} (\cdot) \circ \theta_2 &= \theta_2 \\ (\theta_1, X \mapsto t) \circ \theta_2 &= (\theta_1 \circ \theta_2), X \mapsto t[\theta_2] \end{aligned}$$

In order for composition to be well-defined and have the desired properties we require that $\mathrm{dom}(\theta_1)$ and $\mathrm{dom}(\theta_2)$ are disjoint, but of course variables in the co-domain of θ_1 can be defined by θ_2.

Now we introduce the judgment which implicitly defines an algorithm for unification. We write

$$\models F \;/\; \theta \quad \theta \text{ is a most general unifier for } F.$$

The intent (to be proven later) is that if $\models F \;/\; \theta$ then $\models F[\theta]$, which means that θ is a *unifier* for F. Moreover, we show that whenever $\models F[\theta']$ then there exists a substitution θ'' such that $\theta' = \theta \circ \theta''$, which means that θ is a *most general unifier* for F (any unifier is an instance of θ).

Conjunction and Truth. Algorithmically, we impose a left-to-right order on the solution of F_1 and F_2, but this just fixes a don't care non-deterministic choice.

$$\frac{\models F_1 \;/\; \theta_1 \qquad \models F_2[\theta_1] \;/\; \theta_2}{\models F_1 \wedge F_2 \;/\; \theta_1 \circ \theta_2} \wedge I \qquad \frac{}{\models \top \;/\; \cdot} \top I$$

After all the rules have been shown, it will be easy to see that the side conditions on composition are satisfied and $\theta_1 \circ \theta_2$ is well-defined.

[4] [*example*]

Existential Quantification. Existential variables are introduced for existential quantifiers. They must be "new" (even though the judgment is not parametric). Because of the way existential variables are global to a derivation, this freshness requirement is a global requirement: in a complete derivation, the existential variables chosen for all existential quantifiers must be distinct. To be completely formal about this condition would require to thread a list of existential variables through a derivation. We will dispense with this complication here.

$$\frac{\models [X/x]F \;/\; (\theta, X \mapsto t) \quad X \text{ globally new}}{\models \exists x.\, F \;/\; \theta}\,\exists\mathrm{I}$$

Despite the strong requirement on X to be new, the derivation of the premise is not parametric in X. That is, we cannot substitute an arbitrary term t for X in a derivation of the permiss and obtain a valid derivation, since the vr, rv, vv$\neq$, and vv$=$ rules below require one or both sides of the equation to be an existential variable. Substituting for such a variables invalidates the application of these rules.

Predicate and Function Constants. An equation between the same function constant applied to arguments is decomposed into equations between the arguments. Unification fails if different function symbols are compared, but this is only indirectly reflected by an absence of an appropriate rule. Failure can also be explicitly incorporated in the algorithm (see Exercise **??**).

$$\frac{\models t_1 \doteq s_1 \wedge \cdots \wedge t_n \doteq s_n \;/\; \theta}{\models p(t_1, \ldots, t_n) \doteq p(s_1, \ldots, s_n) \;/\; \theta}\,\mathrm{pp} \qquad \frac{\models t_1 \doteq s_1 \wedge \cdots \wedge t_n \doteq s_n \;/\; \theta}{\models f(t_1, \ldots, t_n) \doteq f(s_1, \ldots, s_n) \;/\; \theta}\,\mathrm{rr}$$

These rules violate orthogonality by relying on conjunction in the premises for the sake of conciseness of the presentation. We could avoid this by introducing a separate judgment for the unification of lists of terms. When f or p have no arguments, the empty conjunction in the premise should be read as $\top$.

Existential Variables. There are four rules for variables. We write r for terms of the form $f(t_1, \ldots, t_n)$. Existential variables always range over terms (and not propositions), so we do not need rules for equations of the form $X \doteq P$ or $P \doteq X$.

$$\frac{X \text{ not in } r}{\models X \doteq r \;/\; (X \mapsto r)}\,\mathrm{vr} \qquad \frac{X \text{ not in } r}{\models r \doteq X \;/\; (X \mapsto r)}\,\mathrm{rv}$$

These two rules come with the proviso that the existential variable X does not occur in the term r. This is necessary to ensure that the substitution $X \mapsto r$ is indeed a unifier. Otherwise unification fails and we can recognize formulas such as $\exists x.\, x \doteq f(x)$ as false. This leaves equations of the form $X \doteq Y$ between two existential variables.

$$\frac{Y \neq X}{\models X \doteq Y \;/\; (X \mapsto Y)}\,\mathrm{vv}{\neq} \qquad \frac{}{\models X \doteq X \;/\; \cdot}\,\mathrm{vv}{=}$$

We now explore the soundness and completeness of these rules, and then analyze the rules as the basis of an algorithm. In the statement of the properties below we take some care so that for a judgment $\models F$, F never contains free existential variables which are seen only as part of the algorithm, not the definition of the logic. We write σ for ground substitutions.[5]

Theorem 4.12 (Soundness of Unification)
If $\models F \;/\; \theta$ then for any ground substitution σ defined on all existential variables in $F[\theta]$ we have $\models (F[\theta])[\sigma]$.

Proof: By induction on the structure of $\mathcal{U}$ of $\models F \;/\; \theta$. □

Lemma 4.13 (Completeness Lemma for Unification)
If $\models F[\sigma]$ for a ground substitution σ defined on all existential variables in F then $\models F \;/\; \theta$ for some θ and $\sigma = \theta \circ \sigma'$ for some σ'.

Proof: By induction on the structure of derivation of $\models F[\sigma]$. □

Theorem 4.14 (Completeness of Unification)
If $\models F$ (where F contains no existential variables) then $\models F \;/\; \cdot$.

Proof: From Lemma 4.13 using the empty substitution for σ. □

4.4 Exercises

Exercise 4.1 Give an alternative proof of the inversion properties (Theorem 4.1) which does not use induction, but instead relies on admissibility of cut in the sequent calculus (Theorem 3.11).

Exercise 4.2 Formulate one or several cut rules directly on inversion sequents as presented in Section 4.1 and prove that they are admissible. Does this simplify the development of the completeness result for inversion proofs? Show how admissibility might be used, or illustrate why it is not much help.

[5] *[warning: just reformulated the properties below and have not yet checked the proofs]*