Part I

Proof Theory

Material on proof theory is incomplete and experimental.

Chapter 1

The Sequent Calculus

1.1 Introduction

pt:seq:int:sec

Sequent calculi are derivation systems that were introduced in the 1930s by Gerhard Gentzen. They were intended less as a practical method for deriving things. Rather, Gentzen used them in order to be able to prove certain results about the possible structure of derivations and their property. These kinds of questions are exactly the questions with which proof theory is concerned.

Sequent calculi operate on *sequents*, as the name suggests, and not on formulas. A sequent is a pair of sequences or multisets of formulas. Gentzen's original systems (**LK** and **LJ**) used sequences; proof theorists now prefer to use multisets. A multiset is like a set, except that a multiset can contain an element more than once. Just as in sets, the order of elements does not matter, however. So the expressions $\{\varphi, \varphi, \psi\}$ and $\{\varphi, \psi, \varphi\}$ describe the same multiset, but that multiset is different from $\{\varphi, \psi\}$ and from $\{\varphi, \varphi, \psi, \psi\}$, since the latter contain different numbers of φ and ψ than the former.

Traditionally, multisets of formulas in proof theory are written using uppercase Greek letters. So when we write Γ, Δ, Π, Λ, or Θ, we intend these to be multisets of formulas of whatever logic we are working with. Traditionally also, instead of writing $\langle \Gamma, \Delta \rangle$ proof theorists separate the two components by a sequent symbol; we use $\Rightarrow$.

Definition 1.1 (Sequent). A *sequent* is an expression of the form

$$\Gamma \Rightarrow \Delta$$

where Γ and Δ are finite (possibly empty) multisets of formulas. Γ is called the *antecedent*, while Δ is the *succedent*.

Let γ be the conjunction of the formulas in Γ (taking the empty conjunction to be $\top$) and θ be the disjunction of the formulas in Δ (taking the empty disjunction to be $\bot$). Then the sequent $\Gamma \Rightarrow \Delta$ is true (in some structure $\mathfrak{M}$) iff $\gamma \to \theta$ is true. In classical logic, this amounts to the same as saying that either one of the formulas in Γ is false, or one of the formulas in Δ is true.

If Γ is a multiset of formulas, we write Γ, φ (or φ, Γ) for the result of adding φ to Γ. If Δ is also a multiset of formulas, then Γ, Δ is the multiset union of the two sequences—if Γ contains φ with *multiplicity* n (i.e., it contains n copies of φ) and Δ with multiplicity m, then Γ, Δ contains φ with multiplicity $n+m$. If the multiset on one side of a sequent is empty, we usually just leave a blank. For instance, $\Rightarrow \varphi$ is the sequent with an empty antecendent and the succedent contains φ with multiplicity 1.

A derivation in the sequent calculus is a tree of sequents, where the topmost (or initial) sequents are axioms in the system considered, and if a sequent stands below one or two other sequents, it must follow correctly by a rule of inference in the system. We will first consider two different sequent systems for classical logic, **G1c** and **G3c**. Their axioms and rules are given in Tables 1.1 and 1.2.

The systems are subtly different, and these differences mean that they have different proof theoretic properties. **G1c** has axioms $\varphi \Rightarrow \varphi$ with no other formulas allowed. **G3c** has axioms of the form $\chi, \Gamma \Rightarrow \Delta, \chi$ with arbitrary "side" formulas in Γ and Δ, but the formula χ appearing on both sides must be *atomic*. **G1c** has two versions each of the $\wedge$L and $\vee$R rules whereas **G3c** has just one each. And **G1c** has additional "structural rules" called "contraction" (CL, CR) and "weakening" (WL, WR).

The systems **G1c** and **G3c** happen to be sound and complete for classical logic. In other words, every sequent derivable in these systems is true in every structure, and every sequent that is true in every structure has a derivation. So the question *whether* a sequent has a derivation can also be answered by other methods of logic (e.g., those of model theory). And since both systems derive exactly the sequents true in every structure, they in particular prove the same sequents.

Proof theory, however, is not just concerned with *whether* something is provable, but *how*. Proof theorists consider questions such as "Can a certain rule always be avoided?," "Can this rule be added to the system without resulting in new sequents becoming provable?," or "Can derivations in one system be turned into derivations in another?" If the answer is "yes" to such a question, the proof of this fact is often proved by induction on the complexity of derivations, or, equivalently, by induction on the structure of derivations. This yields not just a proof of the fact (e.g., if **G1c** proves a sequent then **G3c** proves the same sequent), but a procedure (e.g., of turning derivations in **G1c** into derivations in **G3c**).

A central result in proof theory is the *cut-elimination theorem*. It shows that we can add the cut rule

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \varphi, \Pi \Rightarrow \Lambda}{\Gamma, \Pi \Rightarrow \Delta, \Lambda}\ \text{Cut}$$

to the sequent systems without changing which sequents are derivable. Moreover, the proof provides a method that transforms any derivation using the rules of **G1c** (or **G3c**) together with Cut into one that only uses the rules of **G1c** (or **G3c**). In other words, this procedure "eliminates" the Cut rule from derivations using it; hence the name.

Axioms

$$\varphi \Rightarrow \varphi \qquad\qquad \bot \Rightarrow$$

Structural Rules

$$\frac{\Gamma \Rightarrow \Delta}{\varphi, \Gamma \Rightarrow \Delta}\,\mathrm{WL} \qquad \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi}\,\mathrm{WR}$$

$$\frac{\varphi, \varphi, \Gamma \Rightarrow \Delta}{\varphi, \Gamma \Rightarrow \Delta}\,\mathrm{CL} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi, \varphi}{\Gamma \Rightarrow \Delta, \varphi}\,\mathrm{CR}$$

Logical Rules

$$\frac{\Gamma \Rightarrow \Delta, \varphi}{\neg\varphi, \Gamma \Rightarrow \Delta}\,\neg\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg\varphi}\,\neg\mathrm{R}$$

$$\frac{\varphi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi \qquad \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \wedge \psi}\,\wedge\mathrm{R}$$

$$\frac{\psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L}$$

$$\frac{\varphi, \Gamma \Rightarrow \Delta \qquad \psi, \Gamma \Rightarrow \Delta}{\varphi \vee \psi, \Gamma \Rightarrow \Delta}\,\vee\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\vee\mathrm{R}$$

$$\frac{\Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\vee\mathrm{R}$$

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \psi, \Gamma \Rightarrow \Delta}{\varphi \to \psi, \Gamma \Rightarrow \Delta}\,\to\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \to \psi}\,\to\mathrm{R}$$

$$\frac{\varphi(t), \Gamma \Rightarrow \Delta}{\forall x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\forall\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(a)}{\Gamma \Rightarrow \Delta, \forall x\, \varphi(x)}\,\forall\mathrm{R}$$

$$\frac{\varphi(a), \Gamma \Rightarrow \Delta}{\exists x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\exists\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(t)}{\Gamma \Rightarrow \Delta, \exists x\, \varphi(x)}\,\exists\mathrm{R}$$

Table 1.1: Rules of **G1c**. In ∀R and ∃L, the constant symbol a must not occur in the conclusion sequent. Γ and Δ are multisets of formulas.

pt:seq:int:
tab:G1c

Axioms

$$\chi, \Gamma \Rightarrow \Delta, \chi \qquad\qquad \bot, \Gamma \Rightarrow \Delta$$

Logical Rules

$$\frac{\Gamma \Rightarrow \Delta, \varphi}{\neg\varphi, \Gamma \Rightarrow \Delta}\,\neg\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg\varphi}\,\neg\mathrm{R}$$

$$\frac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi \qquad \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \wedge \psi}\,\wedge\mathrm{R}$$

$$\frac{\varphi, \Gamma \Rightarrow \Delta \qquad \psi, \Gamma \Rightarrow \Delta}{\varphi \vee \psi, \Gamma \Rightarrow \Delta}\,\vee\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi, \psi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\vee\mathrm{R}$$

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \psi, \Gamma \Rightarrow \Delta}{\varphi \to \psi, \Gamma \Rightarrow \Delta}\,\to\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \to \psi}\,\to\mathrm{R}$$

$$\frac{\varphi(t), \forall x\, \varphi(x) \Gamma \Rightarrow \Delta}{\forall x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\forall\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(a)}{\Gamma \Rightarrow \Delta, \forall x\, \varphi(x)}\,\forall\mathrm{R}$$

$$\frac{\varphi(a), \Gamma \Rightarrow \Delta}{\exists x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\exists\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \exists x\, \varphi(x), \varphi(t)}{\Gamma \Rightarrow \Delta, \exists x\, \varphi(x)}\,\exists\mathrm{R}$$

Table 1.2: Rules of **G3c**. In ∀R and ∃L, the constant symbol a must not occur in the conclusion sequent. Γ and Δ are multisets of formulas. In axioms, the formula χ must be atomic.

pt:seq:int:tab:G3c

1.2 Rules and Derivations

pt:seq:rp:sec

We begin by giving some definitions and fixing some terminology.

Consider, for example, the two rules in **G3c** that allow deriving sequents containing $\varphi \wedge \psi$ from sequents containing φ and ψ:

$$\frac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi \qquad \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \wedge \psi}\,\wedge\mathrm{R}$$

The sequents on top of the line are called the *premises* and the sequent below the line the *conclusion*. The formulas in Γ and Δ in each rule are called the *context* or the *side formulas*. The formula in the conclusion that isn't part of the context is called the *principal* formula, and the formulas in the premises not in the context the *active* (or *auxiliary*) formulas.

The quantifier rules are slightly more complicated. E.g., the rules in **G1c** for ∀ are:

$$\frac{\varphi(t), \Gamma \Rightarrow \Delta}{\forall x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\forall\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(a)}{\Gamma \Rightarrow \Delta, \forall x\, \varphi(x)}\,\forall\mathrm{R}$$

In the ∀L rule, the active formula is $\varphi(t)$ where t is a closed term, i.e., a term without variables. Such an inference is correct—that is, it matches the

schematic rule ∀L—if there is some formula φ, some variable x, and some closed term t such that the principal formula in the conclusion is $\forall x\, \varphi$ and the active formula in the premise is $\varphi[t/x]$. The ∀R rule works the same, except that instead of any term t the active formula in the premise must be $\varphi[a/x]$, where a is a constant symbol. This constant is called the *eigenvariable* of the ∀R inference. The inference is only correct if the lower sequent does not contain a, that is, if a does not occur in Γ, Δ, or φ. This is called the *eigenvariable condition*.

Sequent systems like **G1c** and **G3c** consists of a collection of such schematically specified rules together with some (also schematically specified) sequents that count as axioms. Derivations in such a system are finite trees of sequents that are inductively generated from axioms using the rules of inference. Each leaf is an axiom, and every other sequent follows from the sequents immediately above it by one of the inference rules of the system.

pt:seq:rp: defn:proof **Definition 1.2.** *Derivations* in a sequent system are finite trees of sequents inductively as follows:

pt:seq:rp: defn:prf:axiom 1. Every axiom is a derivation.

pt:seq:rp: defn:prf:one-prem 2. If π_1 is a derivation ending in a sequent S_1, and S is a sequent that follows from S_1 by a one-premise rule R, then

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\\ S_1\end{array}}{S}\, R$$

is a derivation.

pt:seq:rp: defn:prf:two-prem 3. If π_1 and π_2 are derivations ending in sequents S_1 and S_2, and S is a sequent that follows from S_1 andS_2 by a two-premise rule R, then

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\\ S_1\end{array} \qquad \begin{array}{c}\vdots\\ \pi_2\\ \vdots\\ S_2\end{array}}{S}\, R$$

is a derivation.

Definition 1.2 defines derivations as trees of sequents. We think of these trees as growing "upwards." The topmost (leaf) nodes of this tree are axioms and called *initial sequents*. The bottom-most sequent is called the *end-sequent*. If π is a derivation with end-sequent S we also write $\pi \vdash S$. If there is a derivation of a sequent S we write $\vdash S$, possibly indicating the system in which the proof exists on the left of the $\vdash$ symbol, e.g., $\mathbf{G1c} \vdash \varphi, \psi \Rightarrow \varphi \wedge \psi$. Every sequent in a derivation, with the exception of initial sequents, is the *conclusion*

of an inference by some rule R. The one or two sequents standing immediately above a sequent in a derivation (i.e., the parents of the node) are called the *premises* of the inference.

The derivations used in the inductive construction of a derivation are called its *sub-derivations*. The sub-derivations ending in the premises of an inference are called the *immediate sub-derivations* of the inference.

It is often necessary to measure the complexity of derivations as a natural number. We could simply count the number of sequents in a derivation, but often a more useful measure is the height of a derivation: it is the maximal number of sequents between the end-sequent and an initial sequent. We can also give an inductive definition.

Definition 1.3. The *height* $\mathrm{ht}(\pi)$ of a derivation π is inductively defined as follows.

1. If π consists of a single axiom, $\mathrm{ht}(\pi) = 0$.
2. If π ends in an inference with one premise with immediate sub-derivation π_1, then $\mathrm{ht}(\pi) = \mathrm{ht}(\pi_1) + 1$.
3. If π ends in an inference with two premises with immediate sub-derivations π_1 and π_2, then $\mathrm{ht}(\pi) = \max(\mathrm{ht}(\pi_1), \mathrm{ht}(\pi_2)) + 1$.

If a sequent S has a derivation of height $\leq n$, we write $\vdash_n S$.

Example 1.4. In **G3c**, any sequent of the form $\chi, \Gamma \Rightarrow \Delta, \chi$ is an axiom, where χ must be atomic. Suppose θ and α are atomic sentences. Then $\theta, \alpha \Rightarrow \theta$ and $\theta, \alpha \Rightarrow \alpha$ are instances of the axiom $\chi, \Gamma \Rightarrow \Delta, \chi$. E.g., if we take $\chi \equiv \alpha$, $\Gamma = \{\theta\}$ and $\Delta = \emptyset$ in $\chi, \Gamma \Rightarrow \Delta, \chi$ we get exactly the sequent $\theta, \alpha \Rightarrow \alpha$. (Remember that we are dealing with multi*sets* and so $\theta, \alpha \Rightarrow \alpha$ and $\alpha, \theta \Rightarrow \theta$ are the same sequent.)

Since $\theta, \alpha \Rightarrow \theta$ is an instance of an axiom of **G3c**, it (by itself) counts as a derivation in **G3c**, and so does $\theta, \alpha \Rightarrow \alpha$ by Definition 1.2(1). Let's call them π_1 and π_2. By (3), we can take these two derivations and generate a new one (π_3) using the two-premise rule $\wedge$R:

$$\dfrac{\theta, \alpha \Rightarrow \theta \qquad \theta, \alpha \Rightarrow \alpha}{\theta, \alpha \Rightarrow \theta \wedge \alpha} \wedge\mathrm{R}$$

From π_3 in turn we obtain the derivation π_4

$$\pi_3 \left\{ \dfrac{\dfrac{\theta, \alpha \Rightarrow \theta \qquad \theta, \alpha \Rightarrow \alpha}{\theta, \alpha \Rightarrow \theta \wedge \alpha} \wedge\mathrm{R}}{\theta \wedge \alpha \Rightarrow \theta \wedge \alpha} \wedge\mathrm{L} \right\} \pi_4$$

using the one-premise rule $\wedge$L (using Definition 1.2(2)). The immediate sub-derivation of the last inference in π_4 is π_3. The immediate sub-derivations of the $\wedge$R inference are π_1 and π_2. The sub-derivations of π_4 are all of π_1, π_2, π_3, and π_4 itself. We have $\mathrm{ht}(\pi_1) = \mathrm{ht}(\pi_2) = 0$, $\mathrm{ht}(\pi_3) = 1$ and $\mathrm{ht}(\pi_4) = 2$. We have shown that $\mathbf{G3c} \vdash_4 \theta \wedge \alpha \Rightarrow \theta \wedge \alpha$ for any atomic θ and α.

1.3 Examples of Proofs

pt:seq:ex: sec

To give derivations of sequents it is typically convenient to work your way upwards from the end-sequent by selecting a formula in it as the active formula and writing the premises of the corresponding rule on top of the sequent, then repeating the process until it leads to axioms. For instance, suppose we want to prove the sequent

$$(\chi \land \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)$$

Consider $\chi \to (\theta \to \alpha)$: it is of the form $\varphi \to \psi$ and occurs on the right of the sequent. Matching the entire sequent with the $\to$R rule of **G1c**,

$$\frac{\varphi, \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \to \psi} \to\text{R}$$

suggests we should take $\Gamma = \{\chi \to (\theta \to \alpha)\}$, $\Delta = \emptyset$, $\varphi \equiv \chi$ and $\psi \equiv \theta \to \alpha$. So the last inference of a proof could be

$$\frac{\chi, (\chi \land \theta) \to \alpha \Rightarrow \theta \to \alpha}{(\chi \land \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)} \to\text{R}$$

If we repeat this idea, focusing on $\theta \to \alpha$ as active formula, we generate

$$\dfrac{\dfrac{\theta, \chi, (\chi \land \theta) \to \alpha \Rightarrow \alpha}{\chi, (\chi \land \theta) \to \alpha \Rightarrow \theta \to \alpha} \to\text{R}}{(\chi \land \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)} \to\text{R}$$

Now we have to focus on $(\chi \land \theta) \to \alpha$ and use the $\to$L rule of **G1c**,

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \psi, \Gamma \Rightarrow \Delta}{\varphi \to \psi, \Gamma \Rightarrow \Delta} \to\text{L}$$

This leads us to

$$\dfrac{\dfrac{\dfrac{\theta, \chi \Rightarrow \alpha, \chi \land \theta \qquad \theta, \chi, \alpha \Rightarrow \alpha}{\theta, \chi, (\chi \land \theta) \to \alpha \Rightarrow \alpha} \to\text{L}}{\chi, (\chi \land \theta) \to \alpha \Rightarrow \theta \to \alpha} \to\text{R}}{(\chi \land \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)} \to\text{R}$$

Using the formula $\chi \land \theta$ in the left topmost sequent as principal and the $\land$R rule gives us

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\theta, \chi \Rightarrow \alpha, \chi \qquad \theta, \chi \Rightarrow \alpha, \theta}{\theta, \chi \Rightarrow \alpha, \chi \land \theta} \land\text{R} \qquad \theta, \chi, \alpha \Rightarrow \alpha}{\theta, \chi, (\chi \land \theta) \to \alpha \Rightarrow \alpha} \to\text{R}}{\chi, (\chi \land \theta) \to \alpha \Rightarrow \theta \to \alpha} \to\text{R}}{(\chi \land \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)} \to\text{R}$$

So far so good. Note that everything we did also works in **G3c**, since the rules we've made use of so far are the same in **G1c** and **G3c**. We've arrived at

a derivation fragment where all top sequents contain the same formula on the left as on the right. But they are not quite axioms yet.

In **G1c**, we have to give derivations of the topmost sequents from axioms of the form $\varphi \Rightarrow \varphi$. This is easily done using the weakening rules, e.g., the leftmost topmost sequent $\theta, \chi \Rightarrow \alpha, \chi$ can be derived as follows:

Overall, we have a derivation in **G1c** which looks like this:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\chi \Rightarrow \chi}{\theta, \chi \Rightarrow \chi}\,\mathrm{WL}}{\theta, \chi \Rightarrow \alpha, \chi}\,\mathrm{WR} \quad \dfrac{\dfrac{\theta \Rightarrow \theta}{\theta, \chi \Rightarrow \theta}\,\mathrm{WL}}{\theta, \chi \Rightarrow \alpha, \theta}\,\mathrm{WR}}{\theta, \chi \Rightarrow \alpha, \chi \wedge \theta}\,\wedge\mathrm{R} \quad \dfrac{\dfrac{\alpha \Rightarrow \alpha}{\chi, \alpha \Rightarrow \chi}\,\mathrm{WL}}{\theta, \chi, \alpha \Rightarrow \alpha}\,\mathrm{WL}}{\theta, \chi, (\chi \wedge \theta) \to \alpha \Rightarrow \alpha}\,{\to}\mathrm{L}}{\chi, (\chi \wedge \theta) \to \alpha \Rightarrow \theta \to \alpha}\,{\to}\mathrm{R}}{(\chi \wedge \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha)}\,{\to}\mathrm{R}$$

The structure of this proof (in particular, its height) is independent of what the formulas χ, θ, and α are. We can replace these symbols in the derivation above by any formulas and the result will still be a derivation in **G1c**. The proof in **G1c** is *schematic*.

In **G3c**, axioms are sequents of the form $\varphi, \Gamma \Rightarrow \Delta, \varphi$ where φ must be atomic. So although $\theta, \chi \Rightarrow \alpha, \chi$ looks like an axiom of **G3c** ($\Gamma = \{\theta\}$; $\Delta = \{\alpha\}$), it is only *really* an axiom if χ in fact is atomic. Our proof fragment is only a complete proof in **G3c** if χ, θ, and α are atomic formulas.

Although we haven't, strictly speaking, shown that

$$\mathbf{G1c} \vdash (\chi \wedge \theta) \to \alpha \Rightarrow \chi \to (\theta \to \alpha),$$

in practice this is all we need to do (and can do). That's because every sequent of the form $\varphi, \Gamma \Rightarrow \Delta, \varphi$ is provable in **G3c**, even if φ itself is not atomic. We can show this by induction on the depth $\mathrm{dp}(\varphi)$ of φ.

pt:seq:ex:
defn:depth

Definition 1.5. The *depth* $\mathrm{dp}(\varphi)$ of a formula is defined as follows:

1. If φ is atomic, $\mathrm{dp}(\varphi) = 0$.
2. If $\varphi \equiv \neg\psi$, $\varphi \equiv \forall x\, A$, or $\varphi \equiv \exists x\, A$, then $\mathrm{dp}(\varphi) = \mathrm{dp}(\psi) + 1$.
3. If $\varphi \equiv (\psi \wedge \chi)$, $(\psi \vee \chi)$, or $(\psi \to \chi)$, then $\mathrm{dp}(\varphi) = \max(\mathrm{dp}(\psi), \mathrm{dp}(\chi)) + 1$.

It is not hard to prove that $\mathrm{dp}(A(x)) = \mathrm{dp}(A(t))$ for any term t. The important fact for us is that in any rule, the depth of the principal formula is always greater than the depth of the active formulas. For instance, we have:

$$\begin{aligned} \mathrm{dp}(P(x)) = \mathrm{dp}(P(a)) &= 0, \\ \mathrm{dp}(\forall x\, P(x)) &= 1, \\ \mathrm{dp}((\forall x\, P(x) \to P(a))) &= \max(1, 0) + 1 = 2. \end{aligned}$$

We can use this to prove some results by induction on $\mathrm{dp}(\varphi)$, such as the following.

pt:seq:ex: prop:G1c-axioms

Proposition 1.6. *For any φ, there is a derivation in* **G1c** *of $\varphi \Rightarrow \varphi$ in which all axioms are atomic.*

Proof. By induction on $\mathrm{dp}(\varphi)$. If $\mathrm{dp}(\varphi) = 0$ then φ is atomic and $\varphi \Rightarrow \varphi$ already is a derivation in **G1c** of the required form. If $\mathrm{dp}(\varphi) > 0$, φ is constructed from formulas of lower depth, e.g, $\varphi \equiv \neg\psi$, $\varphi \equiv (\psi \wedge \chi)$, or $\varphi \equiv \forall x\, A(x)$. The inductive hypothesis is that for all θ with $\mathrm{dp}(\theta) < \mathrm{dp}(\varphi)$, $\mathbf{G1c} \vdash \theta \Rightarrow \theta$ from atomic axioms.

If $\varphi \equiv \neg\psi$, then $\mathrm{dp}(\psi) < \mathrm{dp}(\varphi)$ and by inductive hypothesis, there is a derivation π_1 in **G1c** from atomic axioms of $\psi \Rightarrow \psi$. We can extend this to a derivation of $\neg\psi \Rightarrow \neg\psi$ as follows:

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi \Rightarrow \psi}\,\neg\mathrm{L}}{\neg\psi, \psi \Rightarrow}\,\neg\mathrm{R}}{\neg\psi \Rightarrow \neg\psi}$$

If $\varphi \equiv (\psi \wedge \chi)$, then $\mathrm{dp}(\psi) < \mathrm{dp}(\varphi)$ and $\mathrm{dp}(\chi) < \mathrm{dp}(\varphi)$. By inductive hypothesis, there are derivations π_1 and π_2 in **G1c** from atomic axioms of $\psi \Rightarrow \psi$ and $\chi \Rightarrow \chi$. We can extend these to a derivation of $\psi \wedge \chi \Rightarrow \psi \wedge \chi$ as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi, \Rightarrow \psi}\,\mathrm{WL}}{\psi, \chi, \Rightarrow \psi}\,\wedge\mathrm{L}}{\psi \wedge \chi, \chi \Rightarrow \psi}\,\wedge\mathrm{L}}{\psi \wedge \chi, \psi \wedge \chi \Rightarrow \psi}\qquad \dfrac{\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\chi, \Rightarrow \chi}\,\mathrm{WL}}{\psi, \chi, \Rightarrow \chi}\,\wedge\mathrm{L}}{\psi \wedge \chi, \chi \Rightarrow \chi}\,\wedge\mathrm{L}}{\psi \wedge \chi, \psi \wedge \chi \Rightarrow \chi}\,\wedge\mathrm{R}}{\psi \wedge \chi \Rightarrow \psi \wedge \chi}\,\mathrm{CL}}{\psi \wedge \chi \Rightarrow \psi}$$

Now suppose $\varphi \equiv \forall x\, B(x)$. Take a constant symbol c that does not occur in $\forall x\, B(x)$. Then $\mathrm{dp}(B(c)) = \mathrm{dp}(\psi(x)) < \mathrm{dp}(\varphi)$ and by inductive hypothesis, there is a derivation π_1 in **G1c** from atomic axioms of $\psi(c) \Rightarrow \psi(c)$. We can extend this to a derivation of $\forall x\, B(x) \Rightarrow \forall x\, B(x)$ as follows:

$$\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi(c) \Rightarrow \psi(c)}\,\forall\mathrm{L}}{\forall x\, \psi(x) \Rightarrow \psi(c)}\,\forall\mathrm{R}}{\forall x\, \psi(x) \Rightarrow \forall x\, \psi(x)}$$

The remaining cases are left as exercises. □

Problem 1.1. Complete the proof of Proposition 1.6 by carrying out the cases where $\varphi \equiv (\psi \vee \chi)$, $\varphi \equiv (\psi \to \chi)$ and $\varphi \equiv \exists x\, B(x)$.

For our purposes, we could also have used the derivation

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi \Rightarrow \psi}}{\Rightarrow \psi, \neg\psi}\,\neg\mathrm{R}\qquad \dfrac{\Rightarrow \psi, \neg\psi}{\neg\psi \Rightarrow \neg\psi}\,\neg\mathrm{L}$$

in the case where $\varphi \equiv \neg\psi$. However, there are sequent systems in which this wouldn't have worked. The intuitionistic sequent calculus **G1i** is like **G1c**, except it restricts all sequents to at most one formula on the right side of a sequent. If $\Delta = \emptyset$, then the first derivation would also be a derivation in **G1i**, but the second would not be, since one sequent contains both ψ and $\neg\psi$ on the right.

Note that even in **G1c** we could not have used an alternate order of rules in the case where $\varphi \equiv \forall x\, \psi(x)$. The following is *not* a derivation

$$\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi(c) \Rightarrow \psi(c)}}{\psi(c) \Rightarrow \forall x\, \psi(x)}\,\forall\mathrm{R}}{\forall x\, \psi(x) \Rightarrow \forall x\, \psi(x)}\,\forall\mathrm{L}$$

since the eigenvariable condition on $\forall$R is violated.

Proposition 1.7. *Any sequent of the form $\varphi, \Gamma \Rightarrow \Delta, \varphi$ (with φ not necessarily atomic) is provable in* **G3c**. *pt:seq:ex: prop:G3c-axioms*

Proof. The proof is very similar to the one of Proposition 1.6, but we must be careful with the inductive hypothesis. The case where $n = \mathrm{dp}(\varphi) = 0$ is again easy. If $n = 0$, φ is atomic, and $\varphi, \Gamma \Rightarrow \Delta, \varphi$ is an axiom of **G3c**.

Now suppose $n > 0$. Again we distinguish cases. The inductive hypothesis is: for all D with $\mathrm{dp}(\theta) < n$, and all Π and Λ, **G3c** $\vdash \theta, \Pi \Rightarrow \Lambda, \theta$. Here is the case for $\varphi \equiv (\psi \wedge \chi)$. We use the inductive hypothesis twice: once for $\theta = \psi$, $\Pi = \chi, \Gamma$, and $\Lambda = \Delta$ to get a derivation π_1 of $\psi, \chi, \Gamma \Rightarrow \Delta, \psi$, and once for $\theta = \chi$, $\Pi = \psi, \Lambda$, and $\Lambda = \Delta$ to get a derivation π_2 of $\psi, \chi, \Gamma \Rightarrow \Delta, \chi$. Here is the derivation of $\psi \wedge \chi, \Gamma \Rightarrow \Delta, \psi \wedge \chi$:

$$\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi, \chi, \Gamma \Rightarrow \Delta, \psi}}{\psi \wedge \chi, \Gamma \Rightarrow \Delta, \psi}\,\wedge\mathrm{L} \qquad \dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\psi, \chi, \Gamma \Rightarrow \Delta, \chi}}{\psi \wedge \chi, \Gamma \Rightarrow \Delta, \chi}\,\wedge\mathrm{L}}{\psi \wedge \chi, \Gamma \Rightarrow \Delta, \psi \wedge \chi}\,\wedge\mathrm{R}$$

For $\varphi \equiv \forall x\, \psi(x)$, pick a constant symbol not in $\forall x\, \psi(x)$, Γ, or Δ. Apply the inductive hypothesis for $\theta = \psi(c)$, $\Pi = \forall x\, \psi(x), \Gamma$ and $\Lambda = \Delta$. This gives a derivation π_1 of $\psi(c), \forall x\, \psi(x), \Gamma \Rightarrow \Delta, \psi(c)$, from which we obtain:

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\psi(c), \forall x\, \psi(x), \Gamma \Rightarrow \Delta, \psi(c)}}{\dfrac{\forall x\, \psi(x), \Gamma \Rightarrow \Delta, \psi(c)}{\forall x\, \psi(x), \Gamma \Rightarrow \Delta, \forall x\, \psi(x)}\,\forall\mathrm{R}}\,\forall\mathrm{L}$$

The eigenvariable condition on $\forall$R is satisfied by the way we chose c. □

Problem 1.2. Continue the proof of Proposition 1.7 by carrying out the cases where $\varphi \equiv (\psi \to \chi)$ and $\varphi \equiv \exists x\, B(x)$.

1.4 Interpretation of Rules

pt:seq:irl:sec

Recall the interpretation of a sequent: A sequent $\Gamma \Rightarrow \Delta$ is true in a structure $\mathfrak{M}$ if either at least one $\varphi \in \Gamma$ is false or at least one $\varphi \in \Delta$ is true. We can read the logical rules of **G3c** as codifying the truth conditions for their principal formulas. Consider $\to$R. Its principal formula is $\varphi \to \psi$. It is true if either φ is false or ψ is true. Hence the sequent $\Rightarrow \varphi \to \psi$ is true iff $\varphi \Rightarrow \psi$ is true. If we add the context formulas Γ, Δ to the antecedents and succedents of these sequents, we get exactly the conclusion and premise of the $\to$R rule. The situation is similar for $\to$L. $\varphi \to \psi$ is false if φ is true and ψ is false. So $\varphi \to \psi \Rightarrow$ is true iff both $\Rightarrow \varphi$ and $\psi \Rightarrow$ are true. The conclusion of $\to$L is equivalent to the conjunction of its premises. This pattern ensures that the rules of **G3c** are sound (if all premises are true, the conclusion is true). They also ensure completeness.

In fact, the truth conditions for any truth-functional connective can be described in this way as a set of sequents. This follows from the fact that every truth-function in classical logic can be described by a formula in conjunctive normal form: a conjunction of disjunctions of atomic and negated atomic formulas. For instance, take $\varphi \oplus \psi$ to be true iff exactly one of φ and ψ is true, i.e., "exclusive or." Then $\varphi \oplus \psi$ is true iff $(\varphi \vee \psi) \wedge (\neg\varphi \vee \neg\psi)$ is, and it is false if $(\neg\varphi \vee \psi) \wedge (\varphi \vee \neg\psi)$ is true. Hence, we could introduce sequent calculus rules for $\oplus$ as

$$\dfrac{\varphi, \Gamma \Rightarrow \Delta, \psi \qquad \psi, \Gamma \Rightarrow \Delta, \varphi}{\varphi \oplus \psi, \Gamma \Rightarrow \Delta}\,\oplus\mathrm{L}$$

$$\dfrac{\Gamma \Rightarrow \Delta, \varphi, \psi \qquad \varphi, \psi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi \oplus \psi}\,\oplus\mathrm{R}$$

and these rules would also be sound and complete. Their conclusions are true iff all premises are true.

Problem 1.3. Suppose $\varphi \downarrow \psi$ is true iff neither φ nor ψ is true ("NOR"). What would its **G3c** rules be? (Find two sequents which are true at the same time iff $\varphi \downarrow \psi$ is true. That gives you the $\downarrow$R rule. Then find one sequent that is true iff $\varphi \downarrow \psi$ is false. That gives you the $\downarrow$L rule.)

In **G3c**, each connective and quantifier has exactly two rules: a left rule and a right rule. In **G1c**, however, there are two $\wedge$L and two $\vee$R rules. The reason is that Gentzen's original system **LK**, on which **G1c** is modeled, had a variant **LJ** for an important non-classical logic, viz., intuitionistic logic. To get a system that was sound and complete for intuitionistic logic, the succedents of any sequent in **LJ** is restricted to contain at most one formula. This makes it impossible to use the $\vee$R rule of **G3c**, as its premise contains both φ and ψ in the succedent. Gentzen thus used a pair of $\vee$R rules for **LJ**, and used the same for **LK**. The intuitionistic version **G1i** similarly is just **G1c** with all succedents restricted to at most one formula. (**G3c** also has an intuitionistic version, but it some of its rules differ from the corresponding rules in **G3c**. For instance, it, too, uses a pair of $\vee$R rules.)

That the structural rules of **G1c** are sound is easy to see. For instance, if $\Gamma \Rightarrow \Delta$ contains a false formula on the right or a true formula on the left, so does $\Gamma \Rightarrow \Delta, \varphi$: the same one, whether φ is true or false is irrelevant. So if the premise of WR is true, so is its conclusion. Note that here the converse is not true. The conclusion may be true because φ is true, and then the premise is not guaranteed to be true.

Why have structural rules at all? Contraction (CR, CL) is needed for instance to derive $\Rightarrow \varphi \vee \neg\varphi$. If you start with $\varphi \Rightarrow \varphi$ you first get $\Rightarrow \varphi, \neg\varphi$ by $\neg$R. Then you can use the two versions of $\vee$R in succession, once to get $\varphi \vee \neg\varphi$ from $\neg\varphi$, and once to get $\varphi \vee \neg\varphi$ from φ. What you get is $\Rightarrow \varphi \vee \neg\varphi, \varphi \vee \neg\varphi$. So to get $\Rightarrow \varphi \vee \neg\varphi$ in **G1c** you need to be able to "contract" two occurrences of the same formula in a derivation:

$$\dfrac{\dfrac{\dfrac{\dfrac{\varphi \Rightarrow \varphi}{\Rightarrow \varphi, \neg\varphi}\,\neg\mathrm{R}}{\Rightarrow \varphi, \varphi \vee \neg\varphi}\,\vee\mathrm{R}}{\Rightarrow \varphi \vee \neg\varphi, \varphi \vee \neg\varphi}\,\vee\mathrm{R}}{\Rightarrow \varphi \vee \neg\varphi}\,\mathrm{CR}$$

There are in fact valid sequents which you could not derive in **G1c** without weakening or contraction. For instance, a derivation of the sequent $\varphi \Rightarrow \psi \to \varphi$ in **G1c** has to end in $\to$R (since only $\psi \to \varphi$ can be principal), and that requires $\psi, \varphi \Rightarrow \varphi$ as a premise. This, in turn, requires WL to get it from the axiom $\varphi \Rightarrow \varphi$:

$$\dfrac{\dfrac{\varphi \Rightarrow \varphi}{\psi, \varphi \Rightarrow \varphi}\,\mathrm{WL}}{\varphi \Rightarrow \psi \to \varphi}\,\to\mathrm{R}$$

In **G3c** the contraction and weakening rules are not needed at all. Weakening is built into the axioms. Contraction can mostly be avoided by merging the two copies of the context Γ, Δ in rules with two premises into one. Both **G1c**

Axioms

$$\varphi \Rightarrow \varphi \qquad\qquad \bot \Rightarrow$$

Structural Rules

$$\frac{\Gamma \Rightarrow \Delta}{\varphi, \Gamma \Rightarrow \Delta}\,\mathrm{WL} \qquad \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi}\,\mathrm{WR}$$

$$\frac{\varphi, \varphi, \Gamma \Rightarrow \Delta}{\varphi, \Gamma \Rightarrow \Delta}\,\mathrm{CL} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi, \varphi}{\Gamma \Rightarrow \Delta, \varphi}\,\mathrm{CR}$$

Logical Rules

$$\frac{\Gamma \Rightarrow \Delta, \varphi}{\neg\varphi, \Gamma \Rightarrow \Delta}\,\neg\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg\varphi}\,\neg\mathrm{R}$$

$$\frac{\varphi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L} \qquad \frac{\Gamma_1 \Rightarrow \Delta_1, \varphi \qquad \Gamma_2 \Rightarrow \Delta_2, \psi}{\Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2, \varphi \wedge \psi}\,\wedge\mathrm{R}$$

$$\frac{\psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L}$$

$$\frac{\varphi, \Gamma_1 \Rightarrow \Delta_1 \qquad \psi, \Gamma_2 \Rightarrow \Delta_2}{\varphi \vee \psi, \Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2}\,\vee\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\vee\mathrm{R}$$

$$\frac{\Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\vee\mathrm{R}$$

$$\frac{\Gamma_1 \Rightarrow \Delta_1, \varphi \qquad \psi, \Gamma_2 \Rightarrow \Delta_2}{\varphi \to \psi, \Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2}\,\to\mathrm{L} \qquad \frac{\varphi, \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \to \psi}\,\to\mathrm{R}$$

$$\frac{\varphi(t), \Gamma \Rightarrow \Delta}{\forall x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\forall\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(a)}{\Gamma \Rightarrow \Delta, \forall x\, \varphi(x)}\,\forall\mathrm{R}$$

$$\frac{\varphi(a), \Gamma \Rightarrow \Delta}{\exists x\, \varphi(x), \Gamma \Rightarrow \Delta}\,\exists\mathrm{L} \qquad \frac{\Gamma \Rightarrow \Delta, \varphi(t)}{\Gamma \Rightarrow \Delta, \exists x\, \varphi(x)}\,\exists\mathrm{R}$$

Table 1.3: Rules of **G2c**. In ∀R and ∃L, the constant symbol a must not occur in the conclusion sequent. Γ and Δ are multisets of formulas.

pt:seq:irl:
tab:G2c

and **G3c** have such "context sharing" rules. In one-premise rules, contraction is only needed for ∨R, ∧L, ∃R, and ∀L. The versions of these rules in **G3c** make contraction entirely superfluous. There is a system **G2c** (see Table 1.3) in which the two-premise rules have independent contexts. In **G2c** contraction is even more important than in **G1c**.

1.5 Regular Proofs and Substitution

pt:seq:qua:sec

The rules for the quantifiers introduce some complications owing to the "eigenvariable conditions" that apply to the ∃L and ∀R rules. Most of these complications can be dealt with by ensuring that the constant symbol a in these inferences is never reused. Let's introduce the following definition.

Definition 1.8. A derivation π in **G1c** or **G3c** is *regular* if

1. no eigenvariable a is used as the eigenvariable of more than one ∃L or ∀R inference, and

2. if a is the eigenvariable of a ∃L or ∀R inference, it only occurs above that inference in π.

pt:seq:qua:lem:subst-regular

Lemma 1.9. *If π is regular, ends in $\Gamma \Rightarrow \Delta$, c is a constant symbol not used as an eigenvariable in π, and t is a term not containing any eigenvariables of π, then $\pi[t/c]$ resulting from π by replacing every occurrence of c by t is a regular derivation. The end-sequent of $\pi[t/c]$ is $\Gamma[t/c] \Rightarrow \Delta[t/c]$, where $\Gamma[t/c] = \{\varphi(t) : \varphi(c) \in \Gamma\}$.*

Proof. By induction on $\mathrm{ht}(\pi)$. If $\mathrm{ht}(\pi) = 0$, it consists simply of an axiom. Then $\pi[t/c]$ is also an axiom, since any formula $\varphi(c)$ in it turns into $\varphi(t)$.

If $\mathrm{ht}(\pi) > 0$ we distinguish cases according to the last inference in π. The cases for structural rules (in **G1c**) and logical rules for propositional connectives are easy and left as exercises. We consider the cases where π ends in a quantifier inference.

1. If the last inference is ∀R, π is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \pi'\\ \vdots\end{array}}{\Gamma(c) \Rightarrow \Delta'(c), \varphi(a,c)} \forall\mathrm{R}$$
$$\Gamma(c) \Rightarrow \Delta'(c), \forall x\, \varphi(x,c)$$

   By induction hypothesis, $\pi'[t/c]$ is a regular derivation of $\Gamma(t) \Rightarrow \Delta'(t), \varphi(a,t)$. Since a is an eigenvariable of π, a does not occur in t. Therefore, the last inference in $\pi[t/c]$,

$$\dfrac{\begin{array}{c}\vdots\\ \pi'[t/c]\\ \vdots\end{array}}{\Gamma(t) \Rightarrow \Delta'(t), \varphi(a,t)} \forall\mathrm{R}$$
$$\Gamma(t) \Rightarrow \Delta'(t), \forall x\, \varphi(x,t)$$

   is correct: the conclusion does not contain a.

2. If the last inference is ∀L in **G3c**, π is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \pi'\\ \vdots\end{array}}{\varphi(s(c),c),\forall x\,\varphi(x,c),\Gamma(c) \Rightarrow \Delta'(c)}\,\forall\mathrm{R}$$

By induction hypothesis, $\pi'[t/c]$ is a regular derivation of $\varphi(s(t),t),\forall x\,\varphi(x,t),\Gamma(t) \Rightarrow \Delta'(t)$. The last inference in $\pi[t/c]$,

$$\dfrac{\begin{array}{c}\vdots\\ \pi'[t/c]\\ \vdots\end{array}}{\varphi(s(t),t),\forall x\,\varphi(x,t),\Gamma(t) \Rightarrow \Delta'(t)}\,\forall\mathrm{R}$$

is correct. The case for ∀L in **G1c** is similar, but slightly less complicated.

3. The last inference is ∃R or ∃L: exercise.

In each case, we added a sequent to the end of a regular derivation (or two regular derivations in the case of two-premise rules). The new sequent differs from the end-sequent of π only in that c was replaced by the term t. Moreover, the eigenvariables of π and Since t was assumed not to contain an eigenvariable of π and $\pi[t/c]$ are the same, the new end-sequent does not contain an eigenvariable of $\pi[t/c]$. Consequently, $\pi[t/c]$ is regular. □

Proposition 1.10. *If π is a derivation in* **G1c** *or* **G3c**, *then there is a regular derivation π' with the same structure (in particular, the same height) which differs from π only by replacing eigenvariables.*

Proof. For the purpose of the proof only, call an eigenvariable inference of π *clean* if its eigenvariable is not the eigenvariable of another inference nor does it occur anywhere in π except the immediate sub-derivation of the inference, and *dirty* otherwise. Let n be the number of dirty eigenvariable inferences in π. We show, by induction on n, that π can be transformed into the required regular derivation π'. If $n = 0$, all eigenvariable inferences in π are clean, so and π is regular, and there is nothing to prove. Otherwise, pick a highest eigenvariable inference, say ∀R. The sub-derivation π_1 of π ending in it is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\Gamma \Rightarrow \Delta,\varphi(c)}\,\forall\mathrm{R}$$

Note that as c is an eigenvariable, it does not occur in Γ, Δ, or $\forall x\,\varphi(x)$. The proof π_2 is regular, since it contains no eigenvariable inferences at all. Let a

be a constant symbol not occuring in π. By Lemma 1.9, $\pi_2[a/c]$ is a regular proof of $\Gamma \Rightarrow \Delta, \varphi(a)$, to which we can apply $\forall$R. If we replace π_1 in π by the proof π_1',

$$\dfrac{\begin{array}{c}\vdots\\ \pi_2[a/c]\\ \vdots\end{array}}{\Gamma \Rightarrow \Delta, \varphi(a)}\,\forall\text{R}$$

we have replaced a dirty eigenvariable inference by a clean one, and the only difference is the replacement of the eigenvariable c by a. The resulting proof has $n-1$ dirty eigenvariable inferences, and the result follows from the inductive hypothesis. □

We will not appeal to the proposition just proved explicitly, but we will assume that all derivations considered are regular—if they are not, we can always replace eigenvariables to make them regular.

1.6 Admissible and Derivable Rules

pt:seq:adm:sec

Many results of proof theory concern—or use—results about whether everything that can be derived using a specific rule can also be derived without that rule. The cut-elimination theorem is the most well-known example. It shows that any sequent that can be derived using the rules of a sequent system with Cut can also be derived without Cut. Because this kind of property is so important, it has a name.

Definition 1.11. A rule R is *admissible* in a sequent system if any sequent that can be proved in that system together with the rule R can also be proved without R.

pt:seq:adm:ex:land-deriv

Example 1.12. The rule $\wedge$L takes different forms in **G1c** and **G3c**. In **G1c** there are two versions of the rule, one with the left conjunct φ of the principal formula $\varphi \wedge \psi$ in the premise, and one with the right conjunct ψ:

$$\dfrac{\varphi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\text{L}_1 \qquad \dfrac{\psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\text{L}_2$$

By contrast, **G3c** has only one rule:

$$\dfrac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\text{L}_3$$

We might now ask, "Is **G3c**'s rule $\wedge$L$_3$ admissible in **G1c**?" It is: we can directly simulate any inference using $\wedge$L$_3$ using only the rules of **G1c**. In addition to $\wedge$L$_1$ and $\wedge$L$_2$ we'll also need the structural rule CL:

$$\dfrac{\dfrac{\dfrac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \psi, \Gamma \Rightarrow \Delta}\,\wedge\text{L}_1}{\varphi \wedge \psi, \varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\text{L}_2}{\varphi \wedge \psi \Gamma \Rightarrow \Delta}\,\text{CL}$$

Simulating a inferences using one rule using other rules is one way to show that a rule is admissible. But it is not the only way, and some rules are admissible but can't be so simulated. Rules that can be are called *derivable*.

Definition 1.13. A rule R,

$$\frac{S_1 \qquad \dots \qquad S_n}{S}\, R$$

is called *derivable* in a system if there is a schematic derivation of the sequent S using the rules and axioms of the system together with additional initial sequents of the forms of $S_1, \dots, S_n$.

The derivation given in Example 1.12 shows that $\wedge\mathrm{L}_3$ is a derivable rule in **G1c**, as that derivation is a schematic derivation of the conclusion of $\wedge\mathrm{L}_3$ from the premises of $\wedge\mathrm{L}_3$ as initial sequents using only rules of **G1c**. (It didn't use any of the axioms of **G1c**.)

pt:seq:adm: prop:lor-G3-adm

Proposition 1.14. *The rules* $\wedge\mathrm{L}$ *and* $\vee\mathrm{R}$ *of* **G3c** *are derivable in* **G1c**.

Proof. The proof for $\wedge\mathrm{L}$ is given in Example 1.12. The proof for $\vee\mathrm{R}$ is left as an exercise. □

Problem 1.4. Show that the rule $\vee\mathrm{R}$ of **G3c** is derivable in **G1c**.

Problem 1.5. Suppose $\leftrightarrow$ is a defined logical operator: $\varphi \leftrightarrow \psi$ is short for $(\varphi \to \psi) \wedge (\psi \to \varphi)$. Show that the rules

1. $$\frac{\Gamma, \varphi, \psi \Rightarrow \Delta \qquad \Gamma \Rightarrow \Delta, \varphi, \psi}{\Gamma \Rightarrow \Delta, \varphi \leftrightarrow \psi} \leftrightarrow\mathrm{L}$$

2. $$\frac{\varphi, \Gamma \Rightarrow \Delta, \psi \qquad \psi, \Gamma \Rightarrow \Delta, \varphi}{\Gamma \Rightarrow \Delta, \varphi \leftrightarrow \psi} \leftrightarrow\mathrm{R}$$

are derivable in (a) **G1c** and (b) **G3c**.

pt:seq:adm: ex:weak-adm

Example 1.15. The rule WR of **G1c**,

$$\frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi}\,\mathrm{WR}$$

is admissible in **G3c**. The idea of the proof is simple: suppose there is a derivation π of the premise $\Gamma \Rightarrow \Delta$ in **G3c**. Add the formula φ to the succedent of every sequent in π. Axioms remain axioms; correct inferences are still correct inferences. The end-formula of the new derivation is $\Gamma \Rightarrow \Delta, \varphi$.

The rule WR is, however, not derivable in **G3c**. For suppose it were, i.e., suppose one could find a derivation of $\Gamma \Rightarrow \Delta, \varphi$ using only the axioms and rules of **G3c** plus possibly the additional initial sequent $\Gamma \Rightarrow \Delta$. If this derivation is, as would be required for WR to be derivable, purely schematic, we could turn it into a derivation of $\Rightarrow \varphi$ by deleting Γ and Δ. This would turn the

additional initial sequents $\Gamma \Rightarrow \Delta$ into the empty sequent $\quad \Rightarrow \quad$. Since the empty sequent contains no formulas at all, but all rules of **G3c** require at least one active formula in their premises, the empty sequent can't be a premise of any inference in this schematic derivation. Consequently, the derivation can only be schematic if it doesn't actually use the additional sequent $\Gamma \Rightarrow \Delta$ as an initial sequent. In other words, it would be a schematic derivation of $\Gamma \Rightarrow \Delta, \varphi$ using only the axioms and rules of **G3c** and we would have shown that any sequent of this form has a derivation in **G3c**. But this is impossible, since **G3c** is sound and so can only derive sequents that are true in every structure. Taking $\Gamma = \Delta = \emptyset$ and $\varphi \equiv \bot$, **G3c** would have to derive, e.g., the sequent $\Rightarrow \bot$, which it does not.

Let's give a nice inductive proof of the result that weakening is an admissible rule in **G3c**. In fact, a slightly stronger result holds, as the intuitive argument shows: by adding φ to the right side of every sequent of π, we do not change the structure of the derivation, in particular, we do not increase the height. Since this is important, it deserves its own definition.

Definition 1.16. A rule R,

$$\frac{S_1 \qquad \dots \qquad S_n}{S}\, R$$

is *height-preserving admissible* (or *hp-admissible*) in a system if, whenever $\vdash_{h_i} S_i$ for $i = 1, \dots, n$, then $\vdash_m S$ with $m = \max(h_1, \dots, h_n)$.

Note that it is enough for a rule to be height-preserving admissable that, whenever the premises S_i have derivations π_i with height $h_i = \mathrm{ht}(\pi_i)$, there is a derivation π of the conclusion of height $\mathrm{ht}(\pi)$ no greater than the maximum of the h_i. In particular, if there is only one premise S_1, it is enough that $\mathrm{ht}(\pi) \leq \mathrm{ht}(\pi_1)$. In the case of WR and WL we actually have $\mathrm{ht}(\pi) = \mathrm{ht}(\pi_1)$, but that is not necessary.

Proposition 1.17. *The rules* WR *and* WL *of* **G1c** *are height-preserving admissible in* **G3c**. *pt:seq:adm: prop:weak-G3c-adm*

Proof. We have to show that if **G3c** $\vdash \Gamma \Rightarrow \Delta$ with a derivation π of height $\mathrm{ht}(\pi) = n$ then there is a **G3c**-derivation π' of $\Gamma \Rightarrow \Delta, \varphi$ with $\mathrm{ht}(\pi') = n$. We prove this by induction on n. If $n = 0$, then π is just the axiom $\Gamma \Rightarrow \Delta$ (i.e., some atomic χ occurs in both Γ and Δ), and so is $\Gamma \Rightarrow \Delta, \varphi$. If $n = \mathrm{ht}(\pi) > 0$, the derivation π has a last inference. We distinguish cases according to which rule was used. We give just one example: suppose the last inference was $\wedge$R. Then $\Delta = \Delta', \psi \wedge \chi$ and the immediate sub-derivations π_1 and π_2 of the last inference have end-sequents $\Gamma \Rightarrow \Delta', \psi$ and $\Gamma \Rightarrow \Delta', \chi$, respectively. In other

words, π looks like this:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\end{array}}{\Gamma \Rightarrow \Delta', \psi} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_2\end{array}}{\Gamma \Rightarrow \Delta', \chi} \; \wedge\mathrm{R}$$

We also have that $n = \max(\mathrm{ht}(\pi_1), \mathrm{ht}(\pi_2)) + 1$. So $\mathrm{ht}(\pi_1) < n$ and $\mathrm{ht}(\pi_2) < n$, and the inductive hypothesis applies: there are derivations π_1' and π_2' of $\Gamma \Rightarrow \Delta', \psi, \varphi$ and $\Gamma \Rightarrow \Delta', \chi, \varphi$. Using the $\wedge$R rule, we obtain a derivation π' of $\Gamma \Rightarrow \Delta', \psi \wedge \chi, \varphi$:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1'\end{array}}{\Gamma \Rightarrow \Delta', \psi, \varphi} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_2'\end{array}}{\Gamma \Rightarrow \Delta', \chi, \varphi} \; \wedge\mathrm{R}$$

We have $\mathrm{ht}(\pi') = \max(\mathrm{ht}(\pi_1'), \mathrm{ht}(\pi_2')) + 1 = \max(\mathrm{ht}(\pi_1), \mathrm{ht}(\pi_2)) + 1 = \mathrm{ht}(\pi)$.□

Proposition 1.17 is extremely useful. We'll introduce a notation for repeated application of it: if π is a derivation of $\Gamma \Rightarrow \Delta$ then $\pi[\Pi \Rightarrow \Lambda]$ is the result of applying admissibility of weakening on the left with all the formulas in Π and admissibility of weakening on the right with all formulas in Λ. It is a derivation of $\Gamma, \Pi \Rightarrow \Delta, \Lambda$.

1.7 Invertibility of Rules

pt:seq:inv:sec

Definition 1.18. A rule R,

$$\dfrac{S_1 \qquad \dots \qquad S_n}{S} \; R$$

is *invertible* in a system if, whenever $\vdash S$ then $\vdash_{h_i} S_i$ for $i = 1, \dots, n$. It is *height-preserving invertible* (or *hp-invertible*) in a system if, whenever $\vdash_n S$ then $\vdash_n S_i$ for $i = 1, \dots, n$.

pt:seq:inv:lem:G3c-invert

Lemma 1.19. *The rules of* **G3c** *for* $\neg$, $\wedge$, $\vee$, *and* $\to$ *are height-preserving invertible, that is:*

1. $\neg$R*: If* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \neg\varphi$, *then* $\mathbf{G3c} \vdash_n \varphi, \Gamma \Rightarrow \Delta$.
2. $\neg$L*: If* $\mathbf{G3c} \vdash_n \Gamma, \neg\varphi \Rightarrow \Delta$, *then* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi$.
3. $\wedge$R*: If* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi \wedge \psi$, *then* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi$ *and* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \psi$.
4. $\wedge$L*: If* $\mathbf{G3c} \vdash_n \Gamma, \varphi \wedge \psi \Rightarrow \Delta$, *then* $\mathbf{G3c} \vdash_n \varphi, \psi, \Gamma \Rightarrow \Delta$.

5. $\vee$R: *If* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi \vee \psi$, *then* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi, \psi$.

6. $\vee$L: *If* $\mathbf{G3c} \vdash_n \varphi \vee \psi, \Gamma \Rightarrow \Delta$, *then* $\mathbf{G3c} \vdash_n \varphi, \Gamma \Rightarrow \Delta$ *and* $\mathbf{G3c} \vdash_n \psi, \Gamma \Rightarrow \Delta$.

7. $\to$R: *If* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi \to \psi$, *then* $\mathbf{G3c} \vdash_n \varphi, \Gamma \Rightarrow \Delta, \psi$.

8. $\to$L: *If* $\mathbf{G3c} \vdash_n \varphi \to \psi, \Gamma \Rightarrow \Delta$, *then* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi$ *and* $\mathbf{G3c} \vdash_n \psi, \Gamma \Rightarrow \Delta$.

Proof. We have to show that, for any rule R of **G3c**, if there is a derivation π of the conclusion S of R, then there are also derivations π_i of the premises S_i of R, with $\mathrm{ht}(\pi_i) \leq \mathrm{ht}(\pi)$. Consider the rule $\wedge$L: suppose there is a derivation π with end-sequent of the form $\varphi \wedge \psi, \Gamma \Rightarrow \Delta$. We have to show that there is a derivation π' of $\varphi, \psi, \Gamma \Rightarrow \Delta$ with $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$. We do this by induction on $n = \mathrm{ht}(\pi)$.

If $n = 0$, π consists of just an axiom: $\varphi \wedge \psi, \Gamma \Rightarrow \Delta$. In **G3c**, axioms are of the form $\chi, \Gamma \Rightarrow \Delta, \chi$ with χ atomic. In other words, χ can't be $\varphi \wedge \psi$. If $\varphi \wedge \psi, \Gamma \Rightarrow \Delta$ is an axiom, then some atomic χ must be an element of both Γ and Δ. But then $\varphi, \psi, \Gamma \Rightarrow \Delta$ is also an axiom, and we have found our derivation π' (which is also of height 0).

If $n > 0$ it ends in an inference. We have to prove the claim for n, but may assume that it is true for any derivation of height $< n$, even if the context is different. In other words, for any $k < n$ and any Π and Λ, if $\varphi \wedge \psi, \Pi \Rightarrow \Lambda$ has a derivation of height k, then there is a derivation of height $\leq k$ of $\varphi, \psi, \Pi \Rightarrow \Lambda$.

We have two cases: either $\varphi \wedge \psi$ on the left is principal, or it is not. If it is principal, the last inference is $\wedge$L and the immediate sub-derivation π_1 ends in $\varphi, \psi, \Gamma \Rightarrow \Delta$. The result holds in this case as $\mathrm{ht}(\pi_1) < \mathrm{ht}(\pi)$.

If $\varphi \wedge \psi$ is not principal, some other formula is and $\varphi \wedge \psi$ belongs to the context of the last inference. Then the inductive hypothesis applies to the immediate sub-derivations of the last rule R, which yields a derivation π'_1 or two derivations π'_1 and π'_2 of height no greater than the height(s) of the immediate sub-derivations of π. The end-sequents of π'_1 and π'_2 contain φ, ψ in the context on the left instead of $\varphi \wedge \psi$. Apply the rule R to obtain the proof π'. For instance, suppose π ends in $\to$L:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2 \\ \varphi \wedge \psi, \Gamma' \Rightarrow \Delta, \chi & \varphi \wedge \psi, \Gamma', \theta \Rightarrow \Delta\end{array}}{\varphi \wedge \psi, \Gamma', \chi \to \theta \Rightarrow \Delta} \to\mathrm{L}$$

Here $\chi \to \theta$ on the left is principal, and the context on the left is $\varphi \wedge \psi, \Gamma'$ (i.e., $\Gamma = \Gamma', \chi \to \theta$). The inductive hypothesis yields derivations π'_1 and π'_2 of $\varphi \wedge \psi, \Gamma' \Rightarrow \Delta, \chi$ and $\varphi \wedge \psi, \Gamma', \theta \Rightarrow \Delta$, respectively. These two sequents can

again serve as premises for the $\to$L rule to get π':

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2' \\ \varphi, \psi, \Gamma' \Rightarrow \Delta, \chi & \varphi, \psi, \Gamma', \theta \Rightarrow \Delta\end{array}}{\varphi, \psi, \Gamma', \chi \to \theta \Rightarrow \Delta} \to\text{L}$$

We have

$$\begin{aligned}\mathrm{ht}(\pi) &= \max(\mathrm{ht}(\pi_1), \mathrm{ht}(\pi_2)) + 1\\ \mathrm{ht}(\pi') &= \max(\mathrm{ht}(\pi_1'), \mathrm{ht}(\pi_2')) + 1\end{aligned}$$

by definition, and

$$\begin{aligned}\mathrm{ht}(\pi_1') &\leq \mathrm{ht}(\pi_1)\\ \mathrm{ht}(\pi_2') &\leq \mathrm{ht}(\pi_2)\end{aligned}$$

by inductive hypothesis. So,

$$\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi).$$

□

Problem 1.6. Give the proof of Lemma 1.19 for $\to$R.

pt:seq:inv: lem:invert-quant

Lemma 1.20. *The rules* $\forall$R *and* $\exists$L *are height-preserving invertible in the following sense:*

1. *If* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \forall x\, \varphi(x)$ *then* $\mathbf{G3c} \vdash_n \Gamma \Rightarrow \Delta, \varphi(c)$*, and*
2. *if* $\mathbf{G3c} \vdash_n \exists x\, \varphi(x), \Gamma \Rightarrow \Delta$ *then* $\mathbf{G3c} \vdash_n \varphi(c), \Gamma \Rightarrow \Delta$*,*

for any c *not occuring in* Γ*,* Δ*, and* $\forall x\, \varphi(x)$ *or* $\exists x\, \varphi(x)$*, respectively.*

Proof. We show (1) and leave (2) as an exercise. The argument is similar to the proof of Lemma 1.19. In the inductive step, there are again two cases. The case where $\forall x\, \varphi(x)$ is principal again is quick: the premise of the last $\forall$R inference is of the required form $\Gamma \Rightarrow \Delta, \varphi(a)$ and the immediate sub-derivation π_1 ending in it has $\mathrm{ht}(\pi_1) < \mathrm{ht}(\pi)$. Moreover, the eigenvariable condition gives us that a does not occur in Γ, Δ, or $\forall x\, \varphi(x)$. Since we may assume that π_1 is regular, $\pi_1[c/a]$ is a derivation of the same height of $\Gamma \Rightarrow \Delta, \varphi(c)$ by Lemma 1.9.

In the second case, $\forall x\, \varphi(x)$ is not principal in the last inference; some other formula is. We again consider a case by way of example. Suppose the last inference is $\wedge$R. Some formula in Δ is of the form $\psi \wedge \chi$ and is principal. The derivation π ends in

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2 \\ \Gamma \Rightarrow \Delta', \forall x\, \varphi(x), \psi & \Gamma \Rightarrow \Delta', \forall x\, \varphi(x), \chi\end{array}}{\Gamma \Rightarrow \Delta', \forall x\, \varphi(x), \psi \wedge \chi} \wedge\text{R}$$

where $\Delta = \Delta', \psi \wedge \chi$. Let c be a constant symbol that does not appear in Δ; then clearly it also does not occur in Δ', ψ or in Δ', χ. So the inductive hypothesis applies to π_1 and π_2; we have derivations π_1' and π_2' with $\mathrm{ht}(\pi_1') \leq \mathrm{ht}(\pi_1)$ and $\mathrm{ht}(\pi_2') \leq \mathrm{ht}(\pi_2)$. The required derivation π' of $\Gamma \Rightarrow \Delta, \varphi(c)$ is:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2' \\ \Gamma \Rightarrow \Delta', \varphi(c), \psi & \Gamma \Rightarrow \Delta', \varphi(c), \chi\end{array}}{\Gamma \Rightarrow \Delta', \varphi(c), \psi \wedge \chi}\,\wedge\mathrm{R}$$

with $\mathrm{ht}(\pi') = \max(\mathrm{ht}(\pi_1'), \mathrm{ht}(\pi_2')) + 1 \leq \mathrm{ht}(\pi)$. □

Lemma 1.19 plays an important role in the proof of the cut-elimination theorem. But it can also be used to show the following:

Proposition 1.21. *The contraction rules* CL *and* CR *of* **G1c** *are height-preserving admissible in* **G3c**. *pt:seq:inv: prop:G3c-cont-adm*

Proof. We give the argument for CR and CL at the same time. Suppose π is a derivation of $\Gamma \Rightarrow \Delta, \varphi, \varphi$ or of $\varphi, \varphi, \Gamma \Rightarrow \Delta$ in **G3c**. We have to show that there is also a **G3c**-derivation π' of $\Gamma \Rightarrow \Delta, \varphi$ or $\varphi, \Gamma \Rightarrow \Delta$, respectively, with $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$. We do this by induction on $n = \mathrm{ht}(\pi)$.

If $n = 0$, π is just an axiom. But if $\Gamma \Rightarrow \Delta, \varphi, \varphi$ is an axiom of **G3c**, so is $\Gamma \Rightarrow \Delta, \varphi$: either some atomic χ is an element of both Γ and Δ, or φ is atomic and an element of Γ. The same reasoning applies if the end-sequent is $\varphi, \varphi, \Gamma \Rightarrow \Delta$.

If $n > 0$, π ends in some rule R. If φ is not principal in this last inference, we can obtain a derivation π' as in the proof of Lemma 1.19 by applying the inductive hypothesis to the immediate sub-derivations of π and applying the same rule R to the resulting derivation(s). The inductive hypothesis is that if there is a derivation of height $k < n$ of either $\Pi \Rightarrow \Lambda, \theta, \theta$ or $\theta, \theta, \Pi \Rightarrow \Lambda$, then there is a derivation of height $\leq k$ of $\Pi \Rightarrow \Lambda, \theta$ or $\theta, \Pi \Rightarrow \Lambda$, respectively. (Note that the inductive hypothesis applies if the contexts or contracted formula are different from Γ, Δ, or φ!)

For instance, suppose R is $\wedge$R and π ends in

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2 \\ \Gamma \Rightarrow \Delta', \psi, \varphi, \varphi & \Gamma \Rightarrow \Delta', \chi, \varphi, \varphi\end{array}}{\Gamma \Rightarrow \Delta', \psi \wedge \chi, \varphi, \varphi}\,\wedge\mathrm{R}$$

where $\Delta = \Delta', \psi \wedge \chi$. The inductive hypothesis yields derivations π_1' and π_2' of $\Gamma \Rightarrow \Delta', \psi, \varphi$ and $\Gamma \Rightarrow \Delta', \chi, \varphi$, respectively, with $\mathrm{ht}(\pi_1') \leq \mathrm{ht}(\pi_1)$ and

$\mathrm{ht}(\pi_2') \leq \mathrm{ht}(\pi_2)$. The derivation π' is:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2' \\ \Gamma \Rightarrow \Delta', \psi, \varphi & \Gamma \Rightarrow \Delta', \chi, \varphi\end{array}}{\Gamma \Rightarrow \Delta', \psi \wedge \chi, \varphi}\,\wedge\mathrm{R}$$

If φ is principal in the last inference, we apply the inversion lemma (Lemma 1.19) to the immediate sub-derivations of π, then the inductive hypothesis, and then rule R again. For instance, suppose π proves $\Gamma \Rightarrow \Delta, \varphi, \varphi$, $\varphi \equiv (\psi \wedge \chi)$ and principal in the last inference. Then π looks as follows:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2 \\ \Gamma \Rightarrow \Delta, \psi \wedge \chi, \psi & \Gamma \Rightarrow \Delta, \psi \wedge \chi, \chi\end{array}}{\Gamma \Rightarrow \Delta, \psi \wedge \chi, \psi \wedge \chi}\,\wedge\mathrm{R}$$

Apply Lemma 1.19 to π_1 to get a derivation π_1' of $\Gamma \Rightarrow \Delta, \psi, \psi$. (It also yields a derivation of $\Gamma \Rightarrow \Delta, \chi, \psi$, but we don't need that.) Similarly, by applying Lemma 1.19 to π_2 we obtain a derivation π_2' of $\Gamma \Rightarrow \Delta, \chi, \chi$. As the inversion of $\wedge$R is height-preserving, $\mathrm{ht}(\pi_1') \leq \mathrm{ht}(\pi_1) < \mathrm{ht}(\pi)$ and $\mathrm{ht}(\pi_2') \leq \mathrm{ht}(\pi_2) < \mathrm{ht}(\pi)$. Thus, the inductive hypothesis applies to π_1' and π_2': there are derivations π_1'' and π_2'' of $\Gamma \Rightarrow \Delta, \psi$ and $\Gamma \Rightarrow \Delta, \chi$, respectively, with $\mathrm{ht}(\pi_1'') \leq \mathrm{ht}(\pi_1') \leq \mathrm{ht}(\pi_1)$ and $\mathrm{ht}(\pi_2'') \leq \mathrm{ht}(\pi_2') \leq \mathrm{ht}(\pi_2)$. The derivation π' then is:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1'' & \vdots\,\pi_2'' \\ \Gamma \Rightarrow \Delta, \psi & \Gamma \Rightarrow \Delta, \chi\end{array}}{\Gamma \Rightarrow \Delta, \psi \wedge \chi}\,\wedge\mathrm{R}$$

with $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$.

We must take special care for the quantifier rules in the cases where the formula is principal.

Suppose $\varphi \equiv \forall x\, \psi(x)$ is principal on the right. Then π ends in

$$\dfrac{\begin{array}{c}\vdots\,\pi_1 \\ \Gamma \Rightarrow \Delta, \forall x\, \psi(x), \psi(a)\end{array}}{\Gamma \Rightarrow \Delta, \forall x\, \psi(x), \forall x\, \psi(x)}\,\forall\mathrm{R}$$

By Lemma 1.20, there is a derivation π_1' with height $\leq \mathrm{ht}(\pi_1)$ of $\Gamma \Rightarrow \Delta, \psi(c), \psi(a)$ for any c not occuring in $\Gamma \Rightarrow \Delta, \forall x\, \psi(x), \psi(a)$. We may assume that π_1' is regular; so by Lemma 1.9 the derivation $\pi_1'[a/c]$ is a derivation of $\Gamma \Rightarrow \Delta, \psi(a), \psi(a)$ also of height $\leq \mathrm{ht}(\pi_1)$. Now the inductive hypothesis applies

and gives us a derivation π_1'' of $\Gamma \Rightarrow \Delta, \psi(a)$. Apply the $\forall$R rule to get π':

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1''\\ \vdots\end{array}}{\Gamma \Rightarrow \Delta, \psi(a)}\,\forall\mathrm{R}$$
$$\Gamma \Rightarrow \Delta, \forall x\, \psi(x)$$

with $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$.

Now suppose $\varphi \equiv \exists x\, \psi(x)$ and principal on the right. Then π looks as follows:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\Gamma \Rightarrow \Delta, \exists x\, \psi(x), \exists x\, \psi(x), \psi(t)}\,\exists\mathrm{R}$$
$$\Gamma \Rightarrow \Delta, \exists x\, \psi(x), \exists x\, \psi(x)$$

The inductive hypothesis applies to π_1 and there is a derivations π_1' of $\Gamma \Rightarrow \Delta, \exists x\, \psi(x), \psi(t)$. (Note that no inversion lemma is needed here!) The derivation π' then is:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1'\\ \vdots\end{array}}{\Gamma \Rightarrow \Delta, \exists x\, \psi(x), \psi(t)}\,\exists\mathrm{R}$$
$$\Gamma \Rightarrow \Delta, \exists x\, \psi(x)$$

with $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$. □

Problem 1.7. Sketch the proof of Proposition 1.21 for CL. Describe the cases where $\varphi \equiv (\psi \wedge \chi)$ and $\varphi \equiv (\psi \to \chi)$.

Problem 1.8. Sketch the proof of Proposition 1.21 for CL for the remaining quantifier cases, i.e., where $\varphi \equiv \forall x\, \psi(x)$ and $\varphi \equiv \exists x\, \psi(x)$.

1.8 Translating between G1c and G3c

pt:seq:tr:sec

We mentioned that both **G1c** and **G3c** are sound and complete for classical logic, i.e., they each prove all and only the sequents true in every structure. In particular, they prove the same sequents. We are now in a position to prove this fact without the detour through the soundness and completeness theorems, but using purely proof-theoretically. Such a proof provides *transformations* of derivations in **G1c** into derivations in **G3c** and vice versa.

Proposition 1.22. *If* $\mathbf{G1c} \vdash S$ *then* $\mathbf{G3c} \vdash S$.

Proof. We show that if π is a derivation in **G1c**, then there is a derivation π' of S in **G3c**. We do this by induction on $n = \mathrm{ht}(\pi)$.

If $n = 0$, then π is an axiom. The axiom $\bot \Rightarrow$ is also an axiom in **G3c** (take the context Γ, Δ to be empty). In **G3c**, axioms $\varphi \Rightarrow \varphi$ with arbitrary,

not just atomic φ are allowed. But we proved in Proposition 1.7 that all such sequents are derivable in **G3c**.

If $n > 0$, π ends in a rule R. The inductive hypothesis guarantees that there are derivations of the premises of that rule in **G3c**, i.e., that the immediate sub-derivations have translations in **G3c**. If the rule R is also a rule in **G3c**, we apply it to the translated derivations of the premises. For those rules that are no also rules of **G3c**, we have already shown that they are at least admissible. So, if rule R was $\wedge$L or $\vee$R, apply Proposition 1.14. If it was a weakening rule, apply Proposition 1.17. If it was a contraction rule, apply Proposition 1.21. □

Proposition 1.23. *If* $\mathbf{G3c} \vdash S$ *then* $\mathbf{G1c} \vdash S$.

Proof. We again proceed by induction on the height of a derivation π of S. Every axiom of **G3c** can be derived in **G1c** from an axiom plus several applications of the WL and WR rules:

$$\dfrac{\dfrac{\chi \Rightarrow \chi}{\vdots}\,\mathrm{WL}}{\chi, \Gamma \Rightarrow \chi}\,\mathrm{WR} \qquad \dfrac{\dfrac{\bot \Rightarrow}{\vdots}\,\mathrm{WL}}{\bot, \Gamma \Rightarrow}\,\mathrm{WR}$$

$$\chi, \Gamma \Rightarrow \Delta, \chi \qquad \bot, \Gamma \Rightarrow \Delta$$

If π ends in a rule R, the inductive hypothesis yields derivations in **G1c** of the premises, and we can apply the same rule R to those derivations. The exceptions are the rules $\wedge$L and $\vee$R, which differ between **G1c** and **G3c**. The **G3c** versions are derivable in **G1c** as follows:

$$\dfrac{\dfrac{\dfrac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi, \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L}}{\varphi \wedge \psi, \varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\wedge\mathrm{L}}{\varphi \wedge \psi, \Gamma \Rightarrow \Delta}\,\mathrm{CL} \qquad \dfrac{\dfrac{\dfrac{\Gamma \Rightarrow \Delta, \varphi, \psi}{\Gamma \Rightarrow \Delta, \varphi \vee \psi, \psi}\,\vee\mathrm{R}}{\Gamma \Rightarrow \Delta, , \varphi \vee \psi, \varphi \vee \psi}\,\vee\mathrm{R}}{\Gamma \Rightarrow \Delta, \varphi \vee \psi}\,\mathrm{CR}$$

□

Chapter 2

Cut Elimination

2.1 Introduction

pt:cut:int:sec

The Cut rule is the rule

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \varphi, \Pi \Rightarrow \Lambda}{\Gamma, \Pi \Rightarrow \Delta, \Lambda}\ \text{Cut}$$

(The formula φ is called the *cut formula*.)

Cut was included in Gentzen's original sequent calculus **LK**. Gentzen's *Hauptsatz* ("main theorem") states that every sequent derivable in **LK** is also derivable without using the Cut rule, i.e., that Cut can be "eliminated" from **LK**-proofs. Our systems **G1c** and **G3c** do not contain Cut, but the question can be raised for them as well: can Cut be eliminated from derivations that use Cut plus the other rules of **G1c** (or those of **G3c**)? In our established terminology it can be asked equivalently as: Is Cut an admissible rule of **G1c** (or **G3c**).

The cut elimination theorem is perhaps the most important and celebrated result of proof theory. It first of all establishes that **G1c** and **G3c** don't need a rule which you might, at first blush, think is clearly required. For when you try to formalize actual proofs, you'll notice that you need a way to deal with the use of lemmas. Mathematicians like to prove results, and then appeal to these results in proofs of other results. So you might first formalize a proof of result L (the lemma) as the derivation of a sequent $\Rightarrow L$. Then you formalize the proof of the second result R, the proof of which uses L, as the derivation of a sequent $L \Rightarrow R$. How do you know that there is a proof of R, i.e., a derivation of just $\Rightarrow R$? You need a way to combine the two derivations, and the Cut rule would allow you to do that.

We mentioned that **G1c** and **G3c** are complete, and so they prove any sequent you might expect to be able to prove—all valid sequents. This can be proved directly. But at the time Gentzen was writing, only axiomatic derivation systems were known to be complete. To show **LK** complete, Gentzen provided a translation of derivations in the axiomatic system into **LK**-derivations. This

translation used Cut essentially. Cut elimination is not just an important result for classical logic: many other logics have sequent calculi. For some logics, a direct proof of completeness for their sequent calculi is hard or impossible, so translations from other systems (using Cut) are the only way to show completeness. Then proving cut elimination proof-theoretically is necessary to establish that Cut is superfluous and that the system without Cut is complete.

Cut elimination is also important because it can be applied to obtain other, independently interesting results. Two cases we'll consider are the Craig interpolation lemma and Herbrand' theorem.

The ideas used to prove cut elimination typically involve showing how a derivation that uses Cut can be transformed into one without. These transformations are usually "local," i.e., we take a part of a derivation (usually a sub-derivation ending in a Cut inference) and replace it with another sub-derivation of the same sequent but in which the Cut inference has been removed or replaced by others. Such arguments provide step-by-step algorithms for transforming derivations with Cut into ones without. This provides more information when cut elimination procedures are used in applications. For instance, there are model theoretic proofs of the interpolation lemma the form "if $\varphi \to \psi$" is valid then an interpolant exists (never mind what and interpolant is for the moment). A proof-theoretic argument using cut elimination provides an algorithm that takes a derivation of $\varphi \to \psi$ and returns an interpolant. In contrast to the model theoretic argument, the proof theoretic argument is *constructive*.

There are two popular ways to prove the cut elimination theorem. One (going back to Gentzen) shows that you can remove top-most cuts in a proof, and then just removes top-most cuts until none are left. The other (originally due to Schütte and Tait) focuses on the most complex cuts in a derivation, and successively removes cuts that are maximal in the sense that no other cut has a cut formula of higher depth. They each have their advantages and drawbacks. For some systems, one approach works while another is hard to implement or even impossible. So, we describe both approaches. We'll prove cut elimination for **G3c**. In fact, we'll show not that Cut can be eliminated, but *context-sharing cut* $\mathrm{Cut}_{\mathrm{CS}}$:

$$\frac{\Gamma \Rightarrow \Delta, \varphi \qquad \varphi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}$$

If a sequent calculus allows weakening and contraction (either as actual rules like in **G1c** or as admissible rules like in **G3c**), Cut can simulate $\mathrm{Cut}_{\mathrm{CS}}$, and vice versa. We choose $\mathrm{Cut}_{\mathrm{CS}}$ over Cutbecause the first cut elimination strategy is easier to describe for $\mathrm{Cut}_{\mathrm{CS}}$, and the second strategy only works for $\mathrm{Cut}_{\mathrm{CS}}$.

2.2 Removing Topmost Cuts

pt:cut:top:sec

Definition 2.1. Consider a derivation π ending in a $\mathrm{Cut}_{\mathrm{CS}}$ rule,

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \Gamma \Rightarrow \Delta, \varphi & \varphi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

and containing no other $\mathrm{Cut}_{\mathrm{CS}}$ inferences. The *cut height* of π is $\mathrm{ch}(\pi) = \mathrm{ht}(\pi_1) + \mathrm{ht}(\pi_2)$. The *cut rank* of π is $\mathrm{cr}(\pi) = \mathrm{dp}(\varphi)$.

Lemma 2.2. *The* $\mathrm{Cut}_{\mathrm{CS}}$ *rule is admissible in* **G3c**. *In other words, if* π *is a derivation ending in a single* $\mathrm{Cut}_{\mathrm{CS}}$ *and otherwise using only rules of* **G3c**, *then there is a derivation* π' *of the same end-sequent in* **G3c**. *pt:cut:top: lem:cut-adm-G3c*

Proof. The proof is by double induction on cut height and rank. The derivation π ends in:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \Gamma \Rightarrow \Delta, \varphi & \varphi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

Induction basis: $\mathrm{ch}(\pi) = 0$ and $\mathrm{cr}(\pi) = 0$. Since $\mathrm{ch}(\pi) = \mathrm{ht}(\pi_1) + \mathrm{ht}(\pi_2) = 0$, both $\Gamma \Rightarrow \Delta, \varphi$ and $\varphi, \Gamma \Rightarrow \Delta$ are axioms. There are two cases: (a) either φ is not principal in one of them, or (b) φ is principal in both of them. In both cases we'll show below that $\Gamma \Rightarrow \Delta$ must be an axiom (and we don't need the inductive hypothesis to do that).

We'll now distinguish a number of possible cases according to whether π_1 and π_2 are axioms or end in an inference, and whether φ is principal in that inference or not. Cases A and B cover the induction basis. In the inductive step we assume $\mathrm{ch}(\pi) > 0$ or $\mathrm{cr}(\pi) > 0$ (or both). We can use the inductive hypothesis, namely that the last $\mathrm{Cut}_{\mathrm{CS}}$ can be removed from any derivation which ends in a single $\mathrm{Cut}_{\mathrm{CS}}$ which either has cut rank $= \mathrm{cr}(\pi)$ and a cut height $< \mathrm{ch}(\pi)$, or has cut rank $< \mathrm{cr}(\pi)$. The case where $\mathrm{ch}(\pi) = 0$ (both premises are axioms) are covered by cases A and B. The cases where $\mathrm{ch}(\pi) > 0$ are covered by cases C–D.

Case A: Some ψ occurs in both Γ and Δ. Then $\Gamma \Rightarrow \Delta$ is an axiom (if ψ is atomic), or has a derivation π' in **G3c** without $\mathrm{Cut}_{\mathrm{CS}}$ by Proposition 1.7. This covers case (a) of the induction basis, since if φ is not principal in $\Gamma \Rightarrow \Delta, \varphi$ or not principal in $\varphi, \Gamma \Rightarrow \Delta$ and these are axioms, then some other atomic ψ must occur in both Γ and Δ. It also covers the inductive step if one of the premises is an axiom in which φ is not principal (even if the other premise is the conclusion of a rule).

Case B: Both premises are axioms, and φ is principal and hence atomic. Consider the left premise, $\Gamma \Rightarrow \Delta, \varphi$. Since φ is principal, it must occur in Γ (otherwise $\Gamma \Rightarrow \Delta, \varphi$ wouldn't be an axiom). Similarly, φ must occur in Δ, or else $\varphi, \Gamma \Rightarrow \Delta$ wouldn't be an axiom. Thus φ is atomic and occurs in both Γ and Δ, i.e., $\Gamma \Rightarrow \Delta$ is an axiom. This covers case (b) of the induction basis.

Alternative cases A and B

Case A: One of the premises of the cut is an axiom of the form $\chi, \Gamma' \Rightarrow \Delta', \chi$ with χ atomic. If χ is not the cut formula φ, then χ must occur in both Γ and Δ. In that case, $\Gamma \Rightarrow \Delta$ is an axiom and we can take it as π'. If $\chi \equiv \varphi$ is the cut formula then either (if the left premise is an axiom) $\varphi \in \Gamma$ and $\Gamma = \Gamma', \varphi$, or (if the right premise is an axiom) $\varphi \in \Delta$ and $\Delta = \Delta', \varphi$. In the first case, π_2 is a derivation of $\varphi, \varphi, \Gamma' \Rightarrow \Delta$ and by admissibility of CL (Proposition 1.21) there is a derivation of $\Gamma \Rightarrow \Delta$. In the second case, π_2 is a derivation of $\Gamma \Rightarrow \Delta', \varphi, \varphi$ and by admissibility of CR we get a derivation of $\Gamma \Rightarrow \Delta$.

Case B: One of the premises is an axiom of the form $\bot, \Gamma' \Rightarrow \Delta'$. If the left premise is such an axiom, then $\bot \in \Gamma$ and $\Gamma \Rightarrow \Delta$ is also an axiom. If the right premise is such an axiom, then either $\bot \in \Gamma$ (in which case $\Gamma \Rightarrow \Delta$ again is an axiom), or $\varphi \equiv \bot$. In the latter case, π_1 is a derivation of $\Gamma \Rightarrow \Delta, \bot$. We can turn it into a derivation of $\Gamma \Rightarrow \Delta$ by deleting one copy of $\bot$ from the consequent of every sequent in π_1. (Exercise: verify this by induction on $\mathrm{ht}(\pi_1)$.)

Now suppose that neither case A nor case B applies. In the remaining cases, $\mathrm{ch}(\pi) > 0$, i.e., at least one of the premises is the conclusion of an inference.

Cases C and D: Permuting cuts. The possibilities considered in cases C and D all have a common pattern. We begin with a derivation ending in $\mathrm{Cut_{CS}}$ where one premise is obtained by a rule R in which the cut formula φ is not principal (some other formula θ is). In this derivation, $\mathrm{Cut_{CS}}$ follows the rule R, and schematically looks as follows:

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma \Rightarrow \Delta', \theta, \varphi\end{array} \qquad \left.\dfrac{\begin{array}{c}\vdots\,\pi_3\\ \varphi, \Gamma, \Pi \Rightarrow \Delta', \Lambda\end{array}}{\varphi, \Gamma \Rightarrow \Delta', \theta}\, R\right\}\pi_2}{\Gamma \Rightarrow \Delta', \theta}\,\mathrm{Cut_{CS}}$$

This only covers the case where R is a right rule with one premise, A is not the principal formula of R, and the formula θ which *is* principal occurs in the succedent of the right premise of $\mathrm{Cut_{CS}}$. If θ occurs in the antecedent (i.e., R is a left rule) or in the left premise of $\mathrm{Cut_{CS}}$, it would look different, of course. The formulas in Π and Λ represent the active formulas of rule R.

We turn this order around and consider a derivation where R follows Cut_{CS}:

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1'\end{array}}{\Gamma, \Pi \Rightarrow \Delta', \theta, \Lambda, \varphi} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_3'\end{array}}{\varphi, \Gamma, \Pi \Rightarrow \Delta', \theta, \Lambda}}{\dfrac{\Gamma, \Pi \Rightarrow \Delta', \theta, \Lambda}{\Gamma, \Rightarrow \Delta', \theta, \theta}\, R}\, \text{Cut}_{\text{CS}}$$

We use height-preserving weakening to obtain π_1' and π_3' from π_1 and π_3, respectively.

Because we switch the order of Cut_{CS} and R, this is referred to as "permuting a cut upwards across R." Doing so decreases the height of the derivation ending in the right premise of Cut_{CS}, and so decreases the cut height. That is, the sub-derivations ending in the new Cut_{CS} can be assumed to have height no larger than π_1 and π_3, so the cut height is reduced (we've removed one inference on the right). Now the induction hypothesis applies, which tells us that the Cut_{CS} inference can be removed. Finally, we remove the extra copy of θ using the admissibility of CR.

Case C: At least one of π_1 and π_2 ends in a one-premise rule R in which φ is not principal. There are a total of 18 cases, depending on whether φ is not principal in the left or the right premise of Cut_{CS}, and depending on which of the nine one-premise rules ($\neg$L, $\neg$R, $\vee$R, $\wedge$L, $\to$R, and the four quantifier rules) the rule R is. We only consider two examples: where φ is not principal in the last inference of π_2, which ends in $\to$R or in $\exists$L. The other cases are analogous and left as exercises.

Suppose π ends in:

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\end{array}}{\Gamma \Rightarrow \Delta', \psi \to \chi, \varphi} \qquad \left.\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_3\end{array}}{\varphi, \psi, \Gamma \Rightarrow \Delta', \chi}}{\varphi, \Gamma \Rightarrow \Delta', \psi \to \chi}\, \to\text{R}\right\}\pi_2}{\Gamma \Rightarrow \Delta', \psi \to \chi}\, \text{Cut}_{\text{CS}}$$

where $\Delta = \Delta', \psi \to \chi$. We apply height-preserving weakening (Proposition 1.17) to π_1 to obtain a derivation π_1' of $\psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi, \varphi$ (we weaken with ψ on the left and χ on the right). Similarly, we apply Proposition 1.17 to π_3 to get a derivation π_3' of $\varphi, \psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi$ (by weakening with $\psi \to \chi$ on the right). The induction hypothesis applies to

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1'\end{array}}{\psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi, \varphi} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_3'\end{array}}{\varphi, \psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi}}{\psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi}\, \text{Cut}_{\text{CS}}$$

with cut formula φ: This derivation has the same cut rank and a lower cut height than π. This gives us a derivation π_4 in **G3c** without $\mathrm{Cut_{CS}}$ of $\psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi$. Apply $\to$R to get a derivation of $\Gamma \Rightarrow \Delta', \psi \to \chi, \psi \to \chi$:

$$\dfrac{\begin{array}{c}\vdots\,\pi_4\\ \psi, \Gamma \Rightarrow \Delta', \psi \to \chi, \chi\end{array}}{\Gamma \Rightarrow \Delta', \psi \to \chi, \psi \to \chi}\,{\to}\mathrm{R}$$

To get a derivation π' of $\Gamma \Rightarrow \Delta$, i.e., $\Gamma \Rightarrow \Delta', \psi \to \chi$, apply Proposition 1.21 (the admissibility of contraction).

Now suppose π ends in:

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \exists x\,\psi(x), \Gamma' \Rightarrow \Delta, \varphi\end{array} \qquad \left.\dfrac{\begin{array}{c}\vdots\,\pi_3\\ \varphi, \psi(c), \Gamma' \Rightarrow \Delta\end{array}}{\varphi, \exists x\,\psi(x), \Gamma' \Rightarrow \Delta}\,\exists\mathrm{L}\right\}\pi_2}{\Gamma' \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

We again apply the inductive hypothesis to

$$\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi_1[\psi(c) \Rightarrow]\\ \exists x\,\psi(x), \psi(c), \Gamma' \Rightarrow \Delta, \varphi\end{array} \qquad \begin{array}{c}\vdots\,\pi_3[\exists x\,\psi(x) \Rightarrow]\\ \varphi, \exists x\,\psi(x), \psi(c), \Gamma' \Rightarrow \Delta\end{array}}{\exists x\,\psi(x), \psi(c), \Gamma' \Rightarrow \Delta}\,\mathrm{Cut_{CS}}}{\exists x\,\psi(x), \exists x\,\psi(x), \Gamma' \Rightarrow \Delta}\,\exists\mathrm{L}$$

Note that the eigenvariable condition is satisfied, since the eigenvariable condition on $\exists$L in π_2 guarantees that c does not occur in $\exists x\,\psi(x)$, Γ', or Δ. Finally, apply Proposition 1.21.

Case D: At least one of π_1 and π_2 ends in a two-premise rule R in which φ is not principal. There are 6 cases. We consider the case where φ is not principal in the last inference of π_2, which ends in $\wedge$R; the other cases are analogous.

Assume π ends in:

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma \Rightarrow \Delta', \psi \wedge \chi, \varphi\end{array} \qquad \left.\dfrac{\begin{array}{c}\vdots\,\pi_3\\ \varphi, \Gamma \Rightarrow \Delta', \psi\end{array} \qquad \begin{array}{c}\vdots\,\pi_4\\ \varphi, \Gamma \Rightarrow \Delta', \chi\end{array}}{\varphi, \Gamma \Rightarrow \Delta', \psi \wedge \chi}\,\wedge\mathrm{R}\right\}\pi_2}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

The inductive hypothesis applies to the derivations

$$\dfrac{\begin{array}{c}\vdots\,\pi_1'\\ \Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi, \varphi\end{array} \qquad \begin{array}{c}\vdots\,\pi_3'\\ \varphi, \Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi\end{array}}{\Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi}\,\mathrm{Cut_{CS}}$$

and

$$\dfrac{\begin{array}{c}\vdots\,\pi_1''\\ \Gamma \Rightarrow \Delta, \psi \wedge \chi, \chi, \varphi\end{array} \qquad \begin{array}{c}\vdots\,\pi_4'\\ \varphi, \Gamma \Rightarrow \Delta', \psi \wedge \chi, \chi\end{array}}{\Gamma \Rightarrow \Delta', \psi \wedge \chi, \chi}\,\mathrm{Cut_{CS}}$$

Here, the **G3c** derivations π_1' and π_1'' are obtained by height-preserving weakening (Proposition 1.17) from π_1 (once weakening on the right with ψ, and once with χ). The derivations π_3' and π_4' are obtained from π_3 and π_4 also by height-preserving weakening from π_3 and π_4, respectively (by weakening with $\psi \wedge \chi$ on the right).

The inductive hypothesis gives us derivations π_5 and π_6 in **G3c** of $\Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi$ and $\Gamma \Rightarrow \Delta', \psi \wedge \chi, \chi$. We combine these into a derivation of $\Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi \wedge \chi$ using the rule $\wedge$R:

$$\dfrac{\begin{array}{c}\vdots\,\pi_5\\ \Gamma \Rightarrow \Delta', \psi \wedge \chi, \psi\end{array} \qquad \begin{array}{c}\vdots\,\pi_6\\ \Gamma \Rightarrow \Delta', \psi \wedge \chi, \chi\end{array}}{\Gamma \Rightarrow \Delta, \psi \wedge \chi, \psi \wedge \chi}\,\wedge\mathrm{R}$$

To get a derivation of $\Gamma \Rightarrow \Delta$, i.e., $\Gamma \Rightarrow \Delta', \psi \wedge \chi$, apply Proposition 1.21 (the admissibility of contraction).

Case E: Reducing cuts. The final case concerns the situation where both π_1 and π_2 end in inferences in which φ is principal. There are six cases, one for each possible main operator of φ.

In each case, we turn a $\mathrm{Cut_{CS}}$ with cut formula φ into one or more cuts on sub-formulas of φ, namely on the active formulas of the inferences of which φ is the principal formula. These cases are therefore often referred to as "reduction" or "conversion" of the cut inference.

1. If $\varphi \equiv \neg\psi$ then π ends in

$$\pi_1\left\{\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi_1'\\ \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta, \neg\psi}\,\neg\mathrm{R} \qquad \dfrac{\begin{array}{c}\vdots\,\pi_2'\\ \Gamma \Rightarrow \Delta, \psi\end{array}}{\neg\psi, \Gamma \Rightarrow \Delta}\,\neg\mathrm{L}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}\right\}\pi_2$$

By inductive hypothesis, $\mathrm{Cut}_{\mathrm{CS}}$ can be eliminated from

$$\dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_1' \\ \Gamma \Rightarrow \Delta, \psi & \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

2. If $\varphi \equiv (\psi \vee \chi)$ then π ends in

$$\pi_1\left\{\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi_1' \\ \Gamma \Rightarrow \Delta, \psi, \chi\end{array}}{\Gamma \Rightarrow \Delta, \psi \vee \chi}\,\vee\mathrm{R} \qquad \dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_2'' \\ \psi, \Gamma \Rightarrow \Delta & \chi, \Gamma \Rightarrow \Delta\end{array}}{\psi \vee \chi, \Gamma \Rightarrow \Delta}\,\vee\mathrm{L}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}\right\}\pi_2$$

Consider the derivation ending in $\mathrm{Cut}_{\mathrm{CS}}$,

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2''' \\ \Gamma \Rightarrow \Delta, \psi, \chi & \chi, \Gamma \Rightarrow \Delta, \psi\end{array}}{\Gamma \Rightarrow \Delta, \psi}\,\mathrm{Cut}_{\mathrm{CS}}$$

where π_2''' is obtained from π_2'' by height-preserving weakening. It has cut rank $< \mathrm{cr}(\pi)$ as $\mathrm{dp}(\chi) < \mathrm{dp}(\psi \vee \chi)$, so the inductive hypothesis yields a derivation in **G3c** of π_1'' of $\Gamma \Rightarrow \Delta, \psi$. The derivation ending in $\mathrm{Cut}_{\mathrm{CS}}$,

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1'' & \vdots\,\pi_2' \\ \Gamma \Rightarrow \Delta, \psi & \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

has cut rank $= \mathrm{dp}(\psi) < \mathrm{dp}(\psi \vee \chi)$, so the inductive hypothesis applies again, and yields a derivation π' of $\Gamma \Rightarrow \Delta$.

3. $\varphi \equiv (\psi \wedge \chi)$. Exercise.

4. If $\varphi \equiv (\psi \to \chi)$ then π ends in

$$\pi_1\left\{\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi_1' \\ \psi, \Gamma \Rightarrow \Delta, \chi\end{array}}{\Gamma \Rightarrow \Delta, \psi \to \chi}\,\to\mathrm{R} \qquad \dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_2'' \\ \Gamma \Rightarrow \Delta, \psi & \chi, \Gamma \Rightarrow \Delta\end{array}}{\psi \to \chi, \Gamma \Rightarrow \Delta}\,\to\mathrm{L}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}\right\}\pi_2$$

The derivation ending in $\mathrm{Cut}_{\mathrm{CS}}$,

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2'' \\ \psi, \Gamma \Rightarrow \Delta, \chi & \chi, \Gamma \Rightarrow \Delta\end{array}}{\psi, \Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

has lower cut rank than π. The inductive hypthesis yields a derivation π_1'' of $\psi, \Gamma \Rightarrow \Delta$. Now consider the derivation ending in $\mathrm{Cut}_{\mathrm{CS}}$,

$$\dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_1'' \\ \Gamma \Rightarrow \Delta, \psi & \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

Its cut rank is also less than the cut rank of π, so the induction hypothesis yields a **G3c**-derivation π' of $\Gamma \Rightarrow \Delta$.

5. If $\varphi \equiv \forall x\, \psi(x)$ then π ends in

$$\dfrac{\pi_1\left\{\dfrac{\begin{array}{c}\vdots\,\pi_1' \\ \Gamma \Rightarrow \Delta, \psi(c)\end{array}}{\Gamma \Rightarrow \Delta, \forall x\, \psi(x)}\,\forall\mathrm{R} \quad \dfrac{\begin{array}{c}\vdots\,\pi_2' \\ \psi(t), \forall x\, \psi(x), \Gamma \Rightarrow \Delta\end{array}}{\forall x\, \psi(x), \Gamma \Rightarrow \Delta}\,\forall\mathrm{L}\right\}\pi_2}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

By inductive hypothesis, $\mathrm{Cut}_{\mathrm{CS}}$ can be eliminated from

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1'' & \vdots\,\pi_2' \\ \psi(t), \Gamma \Rightarrow \Delta, \forall x\, \psi(x) & \psi(t), \forall x\, \psi(x), \Gamma \Rightarrow \Delta\end{array}}{\psi(t), \Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

where π_1'' is obtained from π_1 (not π_1'!) by height-preserving weakening. (The derivation has the same cut rank but lower cut height.) This gives us a derivation π_2'' of $\psi(t), \Gamma \Rightarrow \Delta$.

The derivation π_1' has end-sequent $\Gamma \Rightarrow \Delta, \psi(c)$, where c is the eigenvariable of a $\forall$R inference. So c does not occur in $\psi(x)$, Γ, or Δ. We assume the derivation π is regular, so c is the eigenvariable of only the following $\forall$R inference and only occurs above it, i.e., only occurs in π_1' and is not an eigenvariable of an inference in π_1'. Since c does not occur in π_2, it cannot occur in t. By Lemma 1.9, the derivation $\pi_1'[t/c]$ is a derivation of $\Gamma \Rightarrow \Delta, \psi(t)$. We now apply the inductive hypothesis to the derivation

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1'[t/c] & \vdots\,\pi_2'' \\ \Gamma \Rightarrow \Delta, \psi(t) & \psi(t), \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

(The induction hypothesis applies since the cut formula is $\psi(t)$, so the derivation has lower cut rank than π.)

6. We leave the case $\varphi \equiv \exists x\, \psi(x)$ as an exercise. □

Problem 2.1. Verify the sub-case of case D in the proof of Lemma 2.2 where φ is not principal in the last inference of π_1, which ends in $\to$L. (Note: part of this exercise is to clearly state what you have to prove, i.e., what form π has in this case.)

Problem 2.2. Verify the sub-case of case D in the proof of Lemma 2.2 where φ is not principal in the last inference of π_1, which ends in $\forall$L.

Problem 2.3. Verify the sub-case of case E in the proof of Lemma 2.2 where φ is principal in the last inferences of both π_1 and π_2 and $\varphi \equiv (\psi \wedge \chi)$.

Problem 2.4. Verify the sub-case of case E in the proof of Lemma 2.2 where φ is principal in the last inferences of both π_1 and π_2 and $\varphi \equiv \exists x\, \psi(x)$.

Note that in case E the cut height of the derivation ending in $\mathrm{Cut_{CS}}$ to which we apply the inductive hypothesis can be larger than the cut height of the original proof. E.g., in the case where $\varphi \equiv (\psi \to \chi)$ we turned π into a derivation involving two $\mathrm{Cut_{CS}}$ inferences,

$$\dfrac{\dfrac{\Gamma \Rightarrow \Delta, \psi \quad (\pi_2')}{\psi, \Gamma \Rightarrow \Delta} \mathrm{Cut_{CS}} \qquad \left.\dfrac{\psi, \Gamma \Rightarrow \Delta, \chi \quad (\pi_1') \qquad \chi, \Gamma \Rightarrow \Delta \quad (\pi_2'')}{\psi, \Gamma \Rightarrow \Delta} \mathrm{Cut_{CS}}\right\} \pi_1''}{\psi, \Gamma \Rightarrow \Delta} \mathrm{Cut_{CS}}$$

The top $\mathrm{Cut_{CS}}$ between π_1' and π_2'' has both lower cut rank and lower cut height than π. But the derivation π_1''—the result of applying the inductive hypothesis—no longer has lower cut height than π. Since the cut formula of the lower $\mathrm{Cut_{CS}}$ is ψ, its cut *rank*, however, is lower than the cut rank of π.

Lemma 2.2 establishes that we can eliminate a single topmost $\mathrm{Cut_{CS}}$ inference from a derivation. We can show that we can eliminate any number of $\mathrm{Cut_{CS}}$ inferences from a derivation by successively applying it to the topmost sub-derivations ending in $\mathrm{Cut_{CS}}$.

pt:cut:top: thm:cut-elim-G3c-topmost

Theorem 2.3. *Any derivation π in* **G3c** $+ \mathrm{Cut_{CS}}$ *can be transformed into a derivation in* **G3c**.

Proof. By induction on the number n of $\mathrm{Cut_{CS}}$ inferences in π. If $n = 0$, there is nothing to do: π is already a derivation in **G3c**. Otherwise, consider the sub-derivation π' ending in a top-most $\mathrm{Cut_{CS}}$ inference. It contains a single $\mathrm{Cut_{CS}}$ rule as the last inference. Apply Lemma 2.2 to transform it into a $\mathrm{Cut_{CS}}$-free proof π'', and replace the sub-derivation π' by π''. The result is a derivation containing $n - 1$ $\mathrm{Cut_{CS}}$ inferences, and the induction hypothesis applies. □

2.3 Removing Largest Cuts

pt:cut:inv:sec

The proof of Lemma 1.19 shown that if **G3c** derives the conclusion of a logical rule (other than ∃R and ∀L), it also derives each of the premises of the rule without increasing the height of the proof. We can do better and show that this holds also for $\mathbf{G3c} + \mathrm{Cut}_{\mathrm{CS}}$.

As before, we take the *cut rank* of a $\mathrm{Cut}_{\mathrm{CS}}$ inference to be the depth of its cut formula.

pt:cut:inv:lem:inv-G3c-cut

Lemma 2.4. *If* $\mathbf{G3c} + \mathrm{Cut}_{\mathrm{CS}}$ *derives the conclusion of a logical rule* R *other than* ∃R *or* ∀L *using a derivation* π *then it also derives each premise with derivation* π' *(or derivations* π', π'' *depending on whether the rule has one or two premises). Moreover, (a)* $\mathrm{ht}(\pi') \leq \mathrm{ht}(\pi)$ *(and* $\mathrm{ht}(\pi'') \leq \mathrm{ht}(\pi)$*) and (b) the number of* $\mathrm{Cut}_{\mathrm{CS}}$ *inferences as well as their ranks is the same in* π' *(and* π''*) as in* π*.*

Proof. By inspection of the proofs of Lemmata 1.19 and 1.20. That proof was by induction on $\mathrm{ht}(\pi)$. In the base case, we replaced an axiom by another axiom. So conditions (a) and (b) are trivially satisfied. If the last rule of π was R, its immediate sub-derivations are the derivations π_i we want, and (a) and (b) are both satisfied. In all other cases, we applied the inductive hypothesis to the immediate sub-derivations of π, i.e., those leading to the premise(s) of the last inference R, and then applied that same rule R to the results. The conditions (a) and (b) hold for the immediate sub-derivations of the new proof, and so (a) and (b) hold for the entire proof. It remains to verify the case where the last inference is $\mathrm{Cut}_{\mathrm{CS}}$. We do this again for one example only: suppose π is a derivation of the conclusion $\psi \wedge \chi, \Gamma \Rightarrow \Delta$ of ∧L, and ends in $\mathrm{Cut}_{\mathrm{CS}}$:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \psi \wedge \chi, \Gamma \Rightarrow \Delta, \varphi & \varphi, \psi \wedge \chi, \Gamma \Rightarrow \Delta\end{array}}{\psi \wedge \chi, \Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

The induction hypothesis applies to π_1 and π_2 (which are also derivations of instances of the conclusion of ∧L) and yields derivations π_1' and π_2' of $\psi, \chi, \Gamma \Rightarrow \Delta, \varphi$ and $\varphi, \psi, \chi, \Gamma \Rightarrow \Delta$, respectively. We combine these using $\mathrm{Cut}_{\mathrm{CS}}$ to obtain π':

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2'\\ \psi, \chi, \Gamma \Rightarrow \Delta, \varphi & \varphi, \psi, \chi, \Gamma \Rightarrow \Delta\end{array}}{\psi, \chi, \Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

Clearly, (a) and (b) are satisfied. □

Note that this would not work if we used Cut instead of $\mathrm{Cut}_{\mathrm{CS}}$, as then the conclusion of the last derivation would have two copies of ψ, χ, Γ in the antecedent and two copies of Δ in the succedent. We would have to appeal

to the admissibility of contraction, but the proof of Proposition 1.21 uses the invertibility lemma. We'd thus be arguing in a circle.

We can use Lemma 2.4 to give an alternative proof of cut elimination. In this proof we do not successively remove topmost occurrences of $\mathrm{Cut}_{\mathrm{CS}}$ inferences, but occurrences of $\mathrm{Cut}_{\mathrm{CS}}$ inferences of maximal rank. That is, we start with a derivation π in $\mathbf{G3c} + \mathrm{Cut}_{\mathrm{CS}}$ and remove a cut anywhere in it (possibly replacing it with cuts of smaller rank)—and then keep doing that until no cuts remain. The only condition is that above the cut we deal with, no cuts of the same or higher rank occur.

pt:cut:inv: lem:max-cut-red-G3c

Lemma 2.5. *If π is a derivation in $\mathbf{G3c}+\mathrm{Cut}_{\mathrm{CS}}$ ending in a $\mathrm{Cut}_{\mathrm{CS}}$ inference of rank n and otherwise using only rules of $\mathbf{G3c}$ and $\mathrm{Cut}_{\mathrm{CS}}$ inferences of rank $< n$, then there is a proof π' of the same end-sequent in $\mathbf{G3c} + \mathrm{Cut}_{\mathrm{CS}}$ where every $\mathrm{Cut}_{\mathrm{CS}}$ inference has rank $< n$.*

Proof. The derivation π ends in:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \Gamma \Rightarrow \Delta, \varphi & \varphi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut}_{\mathrm{CS}}$$

We distinguish cases according to the form of φ.

Suppose φ is atomic. Observe that by the condition that all $\mathrm{Cut}_{\mathrm{CS}}$ inferences in π_1 and π_2 must have rank $< \mathrm{dp}(\varphi) = 0$, there can be no $\mathrm{Cut}_{\mathrm{CS}}$ inferences in π_1 or π_2. We show that $\Gamma \Rightarrow \Delta$ has a cut-free proof by induction on $\mathrm{ht}(\pi_2)$. If $\mathrm{ht}(\pi_2) = 0$, then $\varphi, \Gamma \Rightarrow \Delta$ is an axiom. If φ is not principal, some atomic χ is an element of both Γ and Δ, and $\Gamma \Rightarrow \Delta$ itself is an axiom. If φ is principal, then $\Delta = \Delta, \varphi$. In this case, π_1 ends in $\Gamma \Rightarrow \Delta', \varphi, \varphi$. We obtain a cut-free proof of $\Gamma \Rightarrow \Delta', \varphi$ by applying Proposition 1.21 (admissibility of contraction). If $\mathrm{ht}(\pi_2) > 0$ it ends in a rule R. Take a premise of that rule: it has the form $\varphi, \Gamma', \Pi \Rightarrow \Delta', \Lambda$, where Π and Λ are the active formulas of R. Proposition 1.17 applied to π_1 and π_2 yields derivations of $\Gamma, \Gamma', \Pi \Rightarrow \Delta, \Delta', \Lambda, \varphi$ and $\varphi, \Gamma, \Gamma', \Pi \Rightarrow \Delta, \Delta', \Lambda$, to which the induction hypothesis applies. This gives us a derivation of $\Gamma, \Gamma', \Pi \Rightarrow \Delta, \Delta', \Lambda$, for each premise of R. Applying R yields $\Gamma, \Gamma \Rightarrow \Delta, \Delta$, and Proposition 1.21 provides the cut-free proof of $\Gamma \Rightarrow \Delta$.

If φ is not atomic, we distinguish cases according to the main operator of φ. In each case, we apply Lemma 2.4 to obtain derivations of the premises of the left- and right-rules with φ as principal formula. We then proceed as in case E of the proof of Lemma 2.2.

1. $\varphi \equiv \neg\psi$. The derivations π_1 and π_2 end in $\Gamma \Rightarrow \Delta, \neg\psi$ and $\neg\psi, \Gamma \Rightarrow \Delta$. By Lemma 2.4, there are derivations π_1' and π_2' of $\psi, \Gamma \Rightarrow \Delta$ and $\Gamma \Rightarrow$

Δ, ψ. The derivation π' is:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_1' \\ \Gamma \Rightarrow \Delta, \psi & \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

2. $\varphi \equiv (\psi \vee \chi)$. The derivations π_1 and π_2 end in $\Gamma \Rightarrow \Delta, \psi \vee \chi$ and $\psi \vee \chi, \Gamma \Rightarrow \Delta$. By Lemma 2.4 (inversion for for $\vee$R applied to π_1), there is a derivation π_1' of $\Gamma \Rightarrow \Delta, \psi, \chi$. By Lemma 2.4 (inversion for $\vee$L applied to π_2), there is a derivation π_2' of $\psi, \Gamma \Rightarrow \Delta$ and a derivation π_2'' of $\chi, \Gamma \Rightarrow \Delta$. By applying Proposition 1.17 to π_2'' we also get a derivation π_2''' of $\chi, \Gamma \Rightarrow \Delta, \psi$. The derivation π' is:

$$\dfrac{\dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2''' \\ \Gamma \Rightarrow \Delta, \psi, \chi & \chi, \Gamma \Rightarrow \Delta, \psi\end{array}}{\Gamma \Rightarrow \Delta, \psi}\,\mathrm{Cut_{CS}} \qquad \begin{array}{c}\vdots\,\pi_2' \\ \psi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

3. $\varphi \equiv (\psi \wedge \chi)$. Exercise.

4. $\varphi \equiv (\psi \to \chi)$. The derivations π_1 and π_2 end in $\Gamma \Rightarrow \Delta, \psi \to \chi$ and $\psi \to \chi, \Gamma \Rightarrow \Delta$. By Lemma 2.4 (inversion for $\to$R applied to π_1), there is a derivation π_1' of $\psi, \Gamma \Rightarrow \Delta, \chi$. By Lemma 2.4 (inversion for $\to$L applied to π_2), there is a derivation π_2' of $\Gamma \Rightarrow \Delta, \psi$ and a derivation π_2'' of $\chi, \Gamma \Rightarrow \Delta$. By applying Proposition 1.17 to π_2'' we also get a derivation π_2''' of $\chi, \psi, \Gamma \Rightarrow \Delta$. The derivation π' is:

$$\dfrac{\begin{array}{cc}\vdots\,\pi_2' & \dfrac{\begin{array}{cc}\vdots\,\pi_1' & \vdots\,\pi_2''' \\ \psi, \Gamma \Rightarrow \Delta, \chi & \chi, \psi, \Gamma \Rightarrow \Delta\end{array}}{\psi, \Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}} \\ \Gamma \Rightarrow \Delta, \psi & \end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

5. $\varphi \equiv \forall x\, \psi(x)$. The derivations π_1 and π_2 end in $\Gamma \Rightarrow \Delta, \forall x\, \psi(x)$ and $\forall x\, \psi(x), \Gamma \Rightarrow \Delta$. By Lemma 2.4, there is a derivation π_1' of $\Gamma \Rightarrow \Delta, \psi(c)$.

   Now consider the derivation π_2. It ends in $\forall x\, \psi(x), \Gamma \Rightarrow \Delta$. Moving upward from the end-sequent of π_2, mark every occurrence of $\forall x\, \psi(x)$ on the left of a sequent that "leads to" the occurrence of $\forall x\, \psi(x)$ in the end-sequent of π_2, that is: (a) Mark $\forall x\, \psi(x)$ in the end-sequent of π_2. (b) If $\forall x\, \psi(x)$ occurs in the context of the conclusion of a rule and is marked, mark a corresponding occurrence in the premise(s) of the rule. (c) If $\forall x\, \psi(x)$ is principal in the conclusion of a $\forall$L rule and is marked,

mark the corresponding active occurrences of $\forall x\, \psi(x)$ and $\psi(t)$ in the premise.

Now consider all these marked occurrences of $\forall x\, \psi(x)$. None of them are active in a rule other than $\forall$L (only active formulas of $\forall$L get marked). Delete all marked occurrences of $\forall x\, \psi(x)$ in π_2. If this is a correct derivation, the marked occurrences of $\forall x\, \psi(x)$ were never active; they all came directly from axioms. The derivation ends in $\Gamma \Rightarrow \Delta$ and we can replace π by it. We have a removed a cut of maximal rank.

Otherwise, what we have is not a correct proof because we have deleted formulas $\forall x\, \psi(x)$ which were active in some $\forall$L inferences. We have turned these into "inferences" of the form

$$\dfrac{\begin{array}{c}\vdots\\ \pi_t\\ \vdots\\ \psi(t), \Pi \Rightarrow \Lambda\end{array}}{\Pi \Rightarrow \Lambda}\,\forall\mathrm{L}$$

Replace each topmost such "inference" by

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1'[t/c][\Pi \Rightarrow \Lambda]\\ \vdots\\ \Gamma, \Pi \Rightarrow \Delta, \Lambda, \psi(t)\end{array} \qquad \begin{array}{c}\vdots\\ \pi_t[\Gamma \Rightarrow \Delta]\\ \vdots\\ \psi(t), \Gamma, \Pi \Rightarrow \Delta, \Lambda\end{array}}{\Gamma, \Pi \Rightarrow \Delta, \Lambda}\,\mathrm{Cut_{CS}}$$

and add Γ to the left and Δ to the right of every sequent below it. Repeat until all $\forall$L "inferences" with a marked $\psi(t)$ in the premise have been replaced by a $\mathrm{Cut_{CS}}$ inference on some $\psi(t)$. The end-sequent of the resulting derivation (and it now is a correct derivation) is of the form $\Gamma, \dots, \Gamma \Rightarrow \Delta, \dots, \Delta$. Apply admissibility of contraction to turn it into a derivation of $\Gamma \Rightarrow \Delta$ and replace π with it. We have replaced one cut on $\forall x\, \psi(x)$ with (possibly many) cuts on $\psi(t)$, but these are all of lower rank. Any cuts already occurring in π_1 or π_2 remain present, but they are also all of lower rank. □

Problem 2.5. Verify the case where $\varphi \equiv (\psi \wedge \chi)$ in the proof of Lemma 2.4. That is, show that if there is a derivation π ending in

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\\ \Gamma \Rightarrow \Delta, \psi \wedge \chi\end{array} \qquad \begin{array}{c}\vdots\\ \pi_2\\ \vdots\\ \psi \wedge \chi, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta}\,\mathrm{Cut_{CS}}$$

with π_1 and π_2 containing only cuts of rank $< \mathrm{dp}(\psi \wedge \chi)$ then there is a derivation π' of $\Gamma \Rightarrow \Delta$ containing only cuts of rank $< \mathrm{dp}(\psi \wedge \chi)$.

Lemma 2.2 establishes that we can replace a Cut_{CS} inference of highest rank in a derivation by Cut_{CS} inferences of lower rank. We can again use this lemma to show that we can eliminate any number of Cut_{CS} inferences from a derivation by successively applying it to the topmost sub-derivations ending in maximal Cut_{CS}.

Theorem 2.6. *Any derivation π in* **G3c** $+ \text{Cut}_{\text{CS}}$ *can be transformed into a proof in* **G3c**. pt:cut:inv: thm:cut-elim-G3c-largest

Proof. First we prove that all cuts of maximal rank in π can be removed. Let d be the maximal rank of cuts occuring in π, and let n be the number of cuts of rank d. The proof is by induction on n. If $n = 0$ there is nothing to prove. If $n > 0$, pick a topmost cut of rank d, and let π_1 be the sub-derivation ending in it. By Lemma 2.5, there is a derivation π_1' of the same end-sequent with all cuts of rank $< d$. Replace π_1 by π_1' in π. This results in a derivation containing $n - 1$ cuts of rank d, and the induction hypothesis applies.

Now the result follows by induction on d. If $d = 0$, all cuts are atomic, and by the previous result can all be removed. If $d > 0$, the previous result yields a proof in which all cuts of rank d have been replaced by cuts of rank $< d$. Since d was the maximum rank of cuts in π, we have proof in which the maximal rank of cuts is $< d$, and the induction hypothesis applies. □

2.4 The Midsequent Theorem and Herbrand's Theorem

An important consequence of the cut-elimination theorem is the midsequent theorem. It states that if there is a derivation π of $\Gamma \Rightarrow \Delta$ in which every formula is prenex, then there is a derivation π' containing a sequent $S \equiv \Pi \Rightarrow \Lambda$ (called the *midsequent*) such that all inferences above S in π' are propositional and all inferences below S are quantifier inferences. (A formula φ is *prenex* (or in *prenex form*) if it is of the form pt:cut:mid: sec

$$\mathsf{Q}_1 x_1 \dots \mathsf{Q}_n x_n \, \psi(x_1, \dots, x_n)$$

where each Q_i is either $\forall$ or $\exists$.) An important consequence of the midsequent theorem is Herbrand's theorem. Its general case is a little hard to describe, but very generally it takes the form: for every derivable sentence φ of first-order logic, there is a propositional tautology φ' from which φ can be derived. Here is the version for sentences of the form $\exists x \, \psi(x)$ where ψ is quantifier-free:

Theorem 2.7 (Herbrand's Theorem). *Suppose that $\psi(x)$ is quantifier-free and that $\vdash \exists x \, \psi(x)$. Then there are terms $t_1, \dots, t_n$ such that $\vdash \psi(t_1) \vee \dots \vee \psi(t_n)$.*

The formula $\psi(t_1) \vee \dots \vee \psi(t_n)$ is called a *Herbrand disjunction* of $\exists x \, \psi(x)$. Since ψ is quantifier-free, this disjunction is also quantifier-free. So if it is derivable, it is derivable by a cut-free proof and consequently with a derivation that only uses propositional inferences. Hence, it is a tautology.

The difficult part of Herbrand's Theorem is showing that a Herbrand disjunction exists for every derivable formula $\exists x\, \psi(x)$. The converse, i.e., that if $\exists x\, \psi(x)$ has a Herbrand disjunction, then it is provable, is easy. In fact, the Herbrand disjunction entails $\exists x\, \psi(x)$. E.g., here is a derivation in **G3c** for the case where $n = 2$:

$$\dfrac{\dfrac{\dfrac{\psi(t_1) \Rightarrow \exists x\, \psi(x), \psi(t_1), \psi(t_2) \qquad \psi(t_2) \Rightarrow \exists x\, \psi(x), \psi(t_1), \psi(t_2)}{\psi(t_1) \vee \psi(t_2) \Rightarrow \exists x\, \psi(x), \psi(t_1), \psi(t_2)} \vee\mathrm{L}}{\psi(t_1) \vee \psi(t_2) \Rightarrow \exists x\, \psi(x), \psi(t_2)} \exists\mathrm{R}}{\psi(t_1) \vee \psi(t_2) \Rightarrow \exists x\, \psi(x)} \exists\mathrm{R}$$

To extract the Herbrand disjunction from a derivation π of $\Rightarrow \exists x\, \psi(x)$, consider the case where all $\exists$R inferences in a cut-free derivation π are at the end:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\\ \Rightarrow \exists x\, \psi(x), \psi(t_1), \ldots, \psi(t_n)\end{array}}{\Rightarrow \exists x\, \psi(x), \psi(t_2), \ldots, \psi(t_n)} \exists\mathrm{R}$$
$$\vdots$$
$$(n-1) \times \exists\mathrm{R}$$
$$\vdots$$
$$\Rightarrow \exists x\, \psi(x)$$

If these are all the $\exists$R inferences in π with $\exists x\, \psi(x)$ active, then there are no other quantifier inferences in π_1'. That's because $\psi(x)$ contains no other quantifiers, and the end-sequent doesn't contain any formulas on the left. In other words, $\Rightarrow \exists x\, \psi(x), \psi(t_1), \ldots, \psi(t_n)$ is the mid-sequent of π. Moreover, since there are no quantifier inferences above the mid-sequent, all sequents above it contain $\exists x\, \psi(x)$ and these are neither active nor principal. So they can be deleted from the derivation π_1, and we get a derivation of $\Rightarrow \psi(t_1), \ldots, \psi(t_n)$, and using $n - 1$ applications of $\vee$R, a derivation of the Herbrand disjunction $\Rightarrow \psi(t_1) \vee \cdots \vee \psi(t_n)$.

The trouble is that we can't assume that even a cut-free derivation of $\Rightarrow \exists x\, \psi(x)$ has all its $\exists$R inferences in one block at the end. There might be inferences in between these $\exists$R inferences in which some $\psi(t_i)$ is principal. Hence, the need for the midsequent theorem.

pt:cut:mid:thm:midsequent

Theorem 2.8 (Midsequent Theorem). *If there is a* **G3c***-derivation π of $\Gamma \Rightarrow \Delta$ where all formulas in Γ and Δ are prenex, then there is a derivation π' containing a sequent $\Pi \Rightarrow \Lambda$ such that all inferences below it are quantifier inferences and all inferences above it are propositional inferences.*

Proof. Define the *order* of a quantifier inference in π as the number of propositional inferences below it, and the order $o(\pi)$ of π as the sum of the orders of all its quantifier inferences. If $o(\pi) = 0$, then no quantifier inference has a propositional inference below it, and we are done: Since all quantifier inferences

have one premise, there is a single topmost quantifier inference. Its premise is the midsequent.

If $o(\pi) > 0$, pick a bottom-most quantifier inference with order > 0. Since it has order > 0, there must be a propositional inference below it. Since it is a bottom-most one, its conclusion can't be the premise of a quantifier inference: that second, lower quantifier inference would also have a propositional inference below it otherwise. We distinguish (many!) cases according to what quantifier inference we have, and what the propositional inference below it is. Because the end-sequent consists of prenex formulas only, the principal formula of the quantifier inference can't be active in the following propositional inference. Otherwise the principal formula of that inference would have a quantifier in the scope of a propositional connective. By the sub-formula property, this would have to be a sub-formula of a formula in the end-sequent. So, the principal formula of the quantifier inference is an element of the context of the lower propositional inference.

Here are two examples:

First, suppose we have a $\forall$R which ends in the right premise of $\to$L. Then the relevant sub-derivation of π ends in the sub-derivation π_1.

$$\dfrac{\begin{array}{cc}\vdots\,\pi_2 & \dfrac{\vdots\,\pi_3}{\chi, \Gamma' \Rightarrow \Delta', \varphi(c)}\,\forall\text{R} \\ \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x), \psi & \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x)\end{array}}{\psi \to \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x)}\,\to\text{L}$$

We can assume that π is regular, and so the eigenvariable c does not occur outside π_3; in particular, it does not occur in $\psi \to \chi$ or in π_2. Now consider the proof π_1':

$$\dfrac{\dfrac{\begin{array}{cc}\vdots\,\pi_2' & \vdots\,\pi_3 \\ \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x), \varphi(c), \psi & \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x), \varphi(c)\end{array}}{\psi \to \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x), \varphi(c)}\,\to\text{L}}{\psi \to \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x), \forall x\,\varphi(x)}\,\forall\text{R}$$

where π_2' is obtained from π_2 by weakening with $\varphi(c)$ on the right. (This does not add any inferences, in particular, no quantifier inferences which might affect the order. It also does not violate any eigenvariable conditions in π_2, since no constants in $\varphi(c)$ are used as eigenvariables in π_2.) The eigenvariable condition on $\forall$R is still satisfied, since c does not occur in $\psi \to \chi$.

We invoke the admissibility of contraction to obtain a derivation π_1'' of $\psi \to \chi, \Gamma' \Rightarrow \Delta', \forall x\,\varphi(x)$. The proof of the admissibility of CR for $\forall x\,\varphi(x)$ and the proof of invertibility of $\forall$R used in that proof did not change the order of any inference: it, at most, deleted inferences. Thus, there are no more quantifier inferences in π_1'' than in π_1, and each quantifier inference in π_1''

has the same number of proposional inferences below it as its corresponding inference in π_1—with the exception of the last $\forall$R: it had a $\to$L below it in π, which we have moved above it in π_1''. Replace π_1 in π by π_1''. The resulting derivation has a lower order, since at least the order of the $\forall$R inference has been reduced by (at least) 1. Now apply the inductive hypothesis.

The cases where the quantifier inference is $\exists$R is even easier, since contraction is not needed and no eigenvariable condition needs to be worried about. E.g., suppose $\exists$R is followed by $\wedge$R (this time as the left premise). The relevant sub-derivation is π_1:

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \varphi(t), \psi}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \psi}\,\exists\mathrm{R} \qquad \begin{array}{c}\vdots\\ \pi_3\\ \vdots\end{array}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \chi}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \psi \wedge \chi}\,\wedge\mathrm{R}$$

Replace π_1 in π by π_1':

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \varphi(t), \psi} \qquad \begin{array}{c}\vdots\\ \pi_3'\\ \vdots\end{array}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \varphi(t), \chi}\,\wedge\mathrm{R}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \varphi(t), \psi \wedge \chi}}{\Gamma' \Rightarrow \Delta', \exists x\, \varphi(x), \psi \wedge \chi}\,\exists\mathrm{R}$$

where π_3' is obtained from π_3 by weakening with $\varphi(t)$ on the right. This weakening involves merely adding $\varphi(t)$ to the right side of every sequent in π_3, so does not change the order of any quantifier inference. The order of quantifier inferences is the same in π' as in π, except the order of the $\exists$R inference permuted with $\wedge$R has decreased by 1. The inductive hypothesis applies. □

Problem 2.6. Describe the cases in the proof of Theorem 2.8 where $\exists$L ends in the premise of $\to$R, and where $\forall$L ends in the left premise of $\vee$L.

Proof of Herbrand's Theorem. Suppose π ends in $\Rightarrow \exists x\, \varphi(x)$. By Theorem 2.8, we can transform π into a derivation π' with a midsequent S, which must be of the form

$$\Rightarrow \exists x\, \varphi(x), \varphi(t_1), \ldots, \varphi(t_n).$$

Since all inferences above S are propositional, $\exists x\, \varphi(x)$ cannot be principal in an inference above S. Since it is not atomic, it cannot be principal in an axiom. Since $\exists x\, \varphi(x)$ cannot be a proper subformula of $\varphi(t_1)$ (remember that $\varphi(x)$ is quantifier-free), it can also not be active above S. Deleting $\exists x\, \varphi(x)$ from the sub-derivation ending in S therefore yields a correct derivation of $\Rightarrow \varphi(t_1), \ldots, \varphi(t_n)$, from which a derivation of $\Rightarrow \varphi(t_1) \vee \cdots \vee \varphi(t_n)$ can be obtained by $n - 1$ $\vee$R inferences. □

2.5 Maehara's Lemma and Craig's Interpolation Theorem

pt:cut:itp:sec

The interpolation theorem, due to William Craig, states that if $\varphi \vDash \psi$ then there is a sentence χ (the *interpolant*) such that $\varphi \vDash \chi$ and $\chi \vDash \psi$. Importantly, the interpolant χ can be chosen so that the constant symbols and predicate symbols occurring in χ all occur in both φ and ψ. (We assume that no function symbols are involved.) Without this restriction, φ and ψ could trivially serve as interpolants.

The interpolation theorem holds for classical first-order logic and has been called the "last important property of first-order logic" to be discovered. It has important applications in model theory, and automated reasoning. Its original motivation and application was in the philosophy of science, since it can be used to study the primitives by which a theory can or cannot be axiomatized. It also implies Beth's definability theorem, which states that if a theory can implicitly define a concept, it can define it explicitly.

Like many results in mathematical logic, the interpolation theorem can be proved both using the methods of model theory and those of proof theory. The advantage of proof theoretic methods is that not only do they show the existence of an interpolant, but provide an algorithm that computes the interpolant from a derivation that $\varphi \vDash \psi$, e.g., a derivation in **G3c** of $\varphi \Rightarrow \psi$.

Let $\mathcal{L}(\varphi)$ be the set of constant symbols and predicate symbols in φ and $\mathcal{L}(\Gamma) = \bigcup\{\mathcal{L}(\varphi) : \varphi \in \Gamma\}$. Throughout this section, we assume we are dealing with formulas without function symbols.

pt:cut:itp:thm:interpolation

Theorem 2.9 (Craig's Interpolation Theorem). *Let the language of φ be $\mathcal{L}(\varphi)$ and that of ψ be $\mathcal{L}(\psi)$. If* **G3c** $\vdash \varphi \Rightarrow \psi$ *then there is a formula χ in the language $\mathcal{L}(\chi) \subseteq \mathcal{L}(\varphi) \cap \mathcal{L}(\psi)$ such that* **G3c** $\vdash \varphi \Rightarrow \chi$ *and* **G3c** $\vdash \chi \Rightarrow \psi$.

We prove a generalization of the interpolation theorem to make use of the structure of derivations in **G3c**. To this end, imagine the antecendent and succedents of sequents to be separated into two parts each. We write $\Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2$ to indicate that Γ_1 and Δ_1 belong to the first part and Γ_2 and Δ_2 to the second part. The interpolation theorem then follows immediately from the following lemma.

pt:cut:itp:lem:maehara

Lemma 2.10 (Maehara's Lemma). *If* **G3c** $\vdash \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2$ *then there is a formula χ in the language $\mathcal{L}(\chi) \subseteq \mathcal{L}(\Gamma_1, \Delta_1) \cap \mathcal{L}(\Gamma_2, \Delta_2)$ such that* **G3c** $\vdash \Gamma_1 \Rightarrow \Delta_1, \chi$ *and* **G3c** $\vdash \chi, \Gamma_2 \Rightarrow \Delta_2$.

Proof. Assume that $\pi \vdash \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2$. We prove the claim by induction on $n = \mathrm{ht}(\pi)$.

If $n = 0$, then $\Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2$ is an axiom, and some atomic formula φ occurs in both Γ_1, Γ_2 and Δ_1, Δ_2. There are four cases depending on whether φ belongs to the first part or the second on the left, and whether it belongs to the first or second part on the right.

1. $\varphi \in \Gamma_1$ and $\varphi \in \Delta_1$. We want a χ such that both $\Gamma_1 \Rightarrow \Delta_1, \chi$ and $\chi, \Gamma_2 \Rightarrow \Delta_2$ are derivable. The first is already an axiom regardless of how we pick χ since φ occurs on both the left and right. Our only option to guarantee that $\chi, \Gamma_2 \Rightarrow \Delta_2$ is derivable is to let $\chi \equiv \bot$.

2. $\varphi \in \Gamma_1$ and $\varphi \in \Delta_2$. Then $\Gamma_1 \Rightarrow \Delta_1, \varphi$ and $\varphi, \Gamma_2 \Rightarrow \Delta_2$ are both axioms, so we can pick $\chi \equiv \varphi$.

3. $\varphi \in \Gamma_2$ and $\varphi \in \Delta_1$. Then $\varphi, \Gamma_1 \Rightarrow \Delta_1$ and $\Gamma_2 \Rightarrow \Delta_2, \varphi$ would be axioms, but they have φ on the wrong sides. So φ isn't the interpolant. However, using $\neg$R and $\neg$L we can derive $\Gamma_1 \Rightarrow \Delta_1, \neg\varphi$ and $\neg\varphi, \Gamma_2 \Rightarrow \Delta_2$.

4. $\varphi \in \Gamma_2$ and $\varphi \in \Delta_2$. Exercise.

The cases where $\Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2$ is an axiom because $\bot \in \Gamma_1$ or $\bot \in \Gamma_2$ are left as exercises.

If $n > 0$, π ends in a logical inference. We have to distinguish cases cases according to which rule was used, and to which part the principal formula belongs. In each case we can use the induction hypothesis that the lemma holds for all derivations π_1 with $\mathrm{ht}(\pi_1) < n$, and however the end-sequent of π_1 may be partitioned into two parts. Let's consider some specific cases.

1. π ends in $\neg$R with principal formula $\neg\varphi$. We have two cases: depending on whether $\neg\varphi$ belongs to the first or second part, π ends in either

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \varphi, \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2\end{array}}{\Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1, \neg\varphi \,;\, \Delta_2}\,\neg\mathrm{R}$$

   or

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma_1 \,;\, \varphi, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2\end{array}}{\Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2, \neg\varphi}\,\neg\mathrm{R}$$

   Note that if the principal formula $\neg\varphi$ belongs to the first part we assigned the active formula φ also to the first part the premise. This is a choice we get to make. The inductive hypothesis would also apply if we assigned the active formula to the other part, but then we would not be able to find an interpolant (exercise: try it).

   Consider the first situation first, where $\neg\varphi$ belongs to the first part. By inductive hypothesis, there is an interpolant of the end-sequent of π_1, i.e., a formula χ_1 such that both

$$\varphi, \Gamma_1 \Rightarrow \Delta_1, \chi_1 \qquad \text{and} \qquad \chi_1, \Gamma_2 \Rightarrow \Delta_2$$

are provable. What we need is a formula χ such that both of the following are derivable:

$$\Gamma_1 \Rightarrow \Delta_1, \neg\varphi, \chi \qquad \text{and} \qquad \chi, \Gamma_2 \Rightarrow \Delta_2$$

Clearly, we can just take $\chi \equiv \chi_1$: the inductive hypothesis already told us that the sequent on the right is derivable, and the sequent on the left follows from the sequent we got from the inductive hypothesis by $\neg$R. The inductive hypothesis also tells us that $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\varphi, \Gamma_1, \Delta_1) \cap \mathcal{L}(\Gamma_2, \Delta_2)$. Since $\mathcal{L}(\neg\varphi) = \mathcal{L}(\varphi)$ and $\mathcal{L}(\chi) = \mathcal{L}(\chi_1)$ we have that $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \neg\varphi) \cap \mathcal{L}(\Gamma_2, \Delta_2)$.

In the second scenario the situation is slightly different in that $\neg\varphi$ belongs to the second part. We assign the active formula in the premise also to the second part and apply the inductive hypothesis to

$$\Gamma_1 \Rightarrow \Delta_1, \chi_1 \qquad \text{and} \qquad \chi_1, \varphi, \Gamma_2 \Rightarrow \Delta_2$$

are provable. If we again apply the $\neg$R rule to the second of these, we get derivations of

$$\Gamma_1 \Rightarrow \Delta_1, \chi_1 \qquad \text{and} \qquad \chi_1, \Gamma_2 \Rightarrow \Delta_2, \neg\varphi$$

These are again exactly the sequents we want to be derivable if χ_1 is an interpolant; we again let $\chi \equiv \chi_1$. $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1) \cap \mathcal{L}(\Gamma_2, \Delta_2, \neg\varphi)$ as before.

2. π ends in $\rightarrow$R with principal formula $\varphi \rightarrow \psi$. We again have two cases, depending on whether $\varphi \rightarrow \psi$ belongs to the first or second part. The derivation π ends in either

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1 \\ \vdots \\ \varphi, \Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1, \psi \,;\, \Delta_2\end{array}}{\Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1, \varphi \rightarrow \psi \,;\, \Delta_2} \rightarrow\text{R}$$

or

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1 \\ \vdots \\ \Gamma_1 \,;\, \varphi, \Gamma_2 \;\Rightarrow\; \Delta_1 \,;\, \Delta_2, \psi\end{array}}{\Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1 \,;\, \Delta_2, \varphi \rightarrow \psi} \rightarrow\text{R}$$

Again we have assigned the active formulas to the same part in the premise as the principal formula in the conclusion. In the first secenario, we apply the inductive hypothesis and get derivations of

$$\varphi, \Gamma_1 \Rightarrow \Delta_1, \psi, \chi_1 \qquad \text{and} \qquad \chi_1, \Gamma_2 \Rightarrow \Delta_2$$

The sequent on the left is a suitable premise for →R. So we have the following derivable sequents:

$$\Gamma_1 \Rightarrow \Delta_1, \varphi \to \psi, \chi_1 \qquad \text{and} \qquad \chi_1, \Gamma_2 \Rightarrow \Delta_2$$

This means that χ_1 is also an interpolant for the end-sequent of π and we can take $\chi \equiv \chi_1$ once more.

The other scenario is handled the same way.

3. π ends in →L with principal formula $\varphi \to \psi$. If $\varphi \to \psi$ belongs to the first part, the derivation π ends in

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2 \\ \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1, \varphi \,;\, \Delta_2 & \psi, \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2\end{array}}{\varphi \to \psi, \Gamma_1 \,;\, \Gamma_2 \Rightarrow \Delta_1 \,;\, \Delta_2}$$

The inductive hypothesis now gives us four derivable sequents, two for the interpolant χ_1 of the left premise and two for the interpolant χ_2 of the second premise:

$$\begin{array}{rlrl} \Gamma_1 \Rightarrow \Delta_1, \varphi, \chi_1 & (1) & \chi_1, \Gamma_2 \Rightarrow \Delta_2 & (2) \\ \psi, \Gamma_1 \Rightarrow \Delta_1, \chi_2 & (3) & \chi_2, \Gamma_2 \Rightarrow \Delta_2 & (4) \end{array}$$

Sequents (1) and (3) can serve as premises of →L, if we apply weakening and add χ_2 on the right to (1) and χ_1 on the right to (3):

$$\dfrac{\begin{array}{cc}\vdots & \vdots \\ \Gamma_1 \Rightarrow \Delta_1, \varphi, \chi_1, \chi_2 & \psi, \Gamma_1 \Rightarrow \Delta_1, \chi_1, \chi_2\end{array}}{\varphi \to \psi, \Gamma_1 \Rightarrow \Delta_1, \chi_1, \chi_2}\,{\to}\mathrm{L}$$

By applying ∨R we get the sequent

$$\varphi \to \psi, \Gamma_1 \Rightarrow \Delta_1, \chi_1 \vee \chi_2$$

This suggests that $\chi_1 \vee \chi_2$ might be our interpolant in this case. And indeed, we can derive the other required sequent from the remaining sequents (2) and (4):

$$\dfrac{\begin{array}{cc}\vdots & \vdots \\ \chi_1, \Gamma_2 \Rightarrow \Delta_2 & \chi_2, \Gamma_2 \Rightarrow \Delta_2\end{array}}{\chi_1 \vee \chi_2, \Gamma_2 \Rightarrow \Delta_2}\,{\vee}\mathrm{L}$$

It remains to verify that $\chi_1 \vee \chi_2$ is in the correct language. The inductive hypothesis gives us that

$$\begin{array}{ll} \mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \varphi) \cap \mathcal{L}(\Gamma_2, \Delta_2) & \text{and} \\ \mathcal{L}(\chi_2) \subseteq \mathcal{L}(\psi, \Gamma_1, \Delta_1) \cap \mathcal{L}(\Gamma_2, \Delta_2) & \end{array}$$

Since

$$\begin{aligned}\mathcal{L}(\varphi) &\subseteq \mathcal{L}(\varphi \to \psi) \qquad \text{and}\\ \mathcal{L}(\psi) &\subseteq \mathcal{L}(\varphi \to \psi)\end{aligned}$$

we have both

$$\begin{aligned}\mathcal{L}(\chi_1) &\subseteq \mathcal{L}(\Gamma_1, \Delta_1, \varphi \to \psi) \cap \mathcal{L}(\Gamma_2, \Delta_2) \qquad \text{and}\\ \mathcal{L}(\chi_2) &\subseteq \mathcal{L}(\Gamma_1, \Delta_1, \varphi \to \psi) \cap \mathcal{L}(\Gamma_2, \Delta_2)\end{aligned}$$

and consequently, $\mathcal{L}(\chi_1 \vee \chi_2) =$

$$\mathcal{L}(\chi_1) \cup \mathcal{L}(\chi_2) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \varphi \to \psi) \cap \mathcal{L}(\Gamma_2, \Delta_2),$$

which is as required.

In the case where $\varphi \to \psi$ belongs to the second part, the situation is different, but handled in a similar way. The inductive hypothesis again gives us four derivable sequents:

$$\begin{array}{llll}\Gamma_1 \Rightarrow \Delta_1, \chi_1 & (1) & \chi_1, \Gamma_2 \Rightarrow \Delta_2, \varphi & (2)\\ \Gamma_1 \Rightarrow \Delta_1, \chi_2 & (3) & \chi_2, \psi, \Gamma_2 \Rightarrow \Delta_2 & (4)\end{array}$$

Now sequents (2) and (4)—plus some weakened formulas—serve as premises of $\to$L:

$$\dfrac{\begin{array}{cc}\vdots & \vdots\\ \chi_1, \chi_2, \Gamma_2 \Rightarrow \Delta_2, \varphi & \psi, \chi_1, \chi_2, \Gamma_2 \Rightarrow \Delta_2\end{array}}{\varphi \to \psi, \chi_1, \chi_2, \Gamma_2 \Rightarrow \Delta_2} \to\text{L}$$

We again want to turn this into a sequent that has a single formula built from χ_1 and χ_2, but this time we need it in the antecedent (since we're now dealing with the second part). Applying $\wedge$L does the trick; we should take $\chi \equiv (\chi_1 \wedge \chi_2)$ as the interpolant. Fortuitously, we can use the remaining sequents (1) and (3) to derive the corresponding sequent for the other part:

$$\dfrac{\begin{array}{cc}\vdots & \vdots\\ \Gamma_1 \Rightarrow \Delta_1, \chi_1 & \Gamma_1 \Rightarrow \Delta_1, \chi_2\end{array}}{\Gamma_1 \Rightarrow \Delta_1, \chi_1 \wedge \chi_2} \wedge\text{R}$$

As before, we have $\mathcal{L}(\chi_1 \wedge \chi_2) \subseteq \mathcal{L}(\Gamma_1, \Delta_1) \cap \mathcal{L}(\Gamma_2, \Delta_2, \varphi \to \psi)$.

4. π ends in $\forall$R with principal formula $\forall x\, \varphi(x)$:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1, \varphi(c) \,;\, \Delta_2}\,\forall\mathrm{R}$$

or

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1 \,;\, \Delta_2, \varphi(c)}\,\forall\mathrm{R}$$

In the first scenario, the induction hypothesis gives us derivations of $\Gamma_1 \Rightarrow \Delta_1, \varphi(c), \chi_1$ and $\chi_1, \Gamma_2 \Rightarrow \Delta_2$. We can apply $\forall$R to the first to obtain $\Gamma_1 \Rightarrow \Delta_1, \forall x\, \varphi(x), \chi_1$. (The eigenvariable condition is satisfied, since it was already satisfied in the $\forall$R rule at the end of π.) Thus χ_1 would be an interpolant if the condition that $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \forall x\, \varphi(x)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$ is satisfied. By the inductive hypothesis, we have $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \varphi(c)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$. So we have to worry about c—can c occur in χ_1? It can't. If c occurs in χ_1, we would have both $c \in \mathcal{L}(\Gamma_1, \Delta_1, \varphi(c))$ (which is true) and also $c \in \mathcal{L}(\Gamma_2, \Delta_2)$. But then c would occur in the conclusion of the $\forall$R inference in π, violating the eigenvariable condition.

The other scenario is similar.

5. π ends in $\exists$R with principal formula $\exists x\, \varphi(x)$. We again have two scenarios. We consider the scenario where $\exists x\, \varphi(x)$ belongs to the first part, and leave the other as an exercise.

In the first scenario, π ends in

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\Gamma_1 \,;\, \Gamma_2 \;\Rightarrow\; \Delta_1, \exists x\, \varphi(x), \varphi(t) \,;\, \Delta_2}\,\exists\mathrm{R}$$

The induction hypothesis gives us derivations of $\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \varphi(t), \chi_1$ and $\chi_1, \Gamma_2 \Rightarrow \Delta_2$. We can apply $\exists$R to the first to obtain $\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \chi_1$. Thus χ_1 would again be an interpolant if the condition that $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$ is satisfied. Here the inductive hypothesis gives us $\mathcal{L}(\chi_1) \subseteq \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x), \varphi(t)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$. Now remember that we assumed that there are no function symbols. That means the term t can only be a constant symbol, i.e., $t \equiv b$. If $b \notin \mathcal{L}(\chi_1)$ or $b \in \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$ there is nothing to worry about: we can take χ_1 as the interpolant. But what if $b \in \mathcal{L}(\chi_1)$ and $b \notin \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x)) \cap \mathcal{L}(\Gamma_2, \Delta_2)$? First of all, by replacing all occurrences

of b in χ_1 by a variable y we can write $\chi_1 \equiv \chi_1'[b/y]$ where χ_1' does not contain b. Furthermore, either $b \notin \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x))$ or $b \notin \mathcal{L}(\Gamma_2, \Delta_2)$. We can take $\chi \equiv \forall y\, \chi_1'(y)$ or $\chi \equiv \exists y\, \chi_1'(y)$, respectively. In the first case, we have:

$$\dfrac{\dfrac{\dfrac{\vdots}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \varphi(b), \chi_1'(b)}\,\exists\mathrm{R}}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \chi_1'(b)}\,\forall\mathrm{R}}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \forall y\, \chi_1'(y)}$$

and

$$\dfrac{\dfrac{\vdots}{\chi_1'(b), \forall y\, \chi_1'(y), \Gamma_2 \Rightarrow \Delta_2}\,\forall\mathrm{L}}{\forall y\, \chi_1'(y), \Gamma_2 \Rightarrow \Delta_2}$$

The derivations of the top sequents are from the inductive hypothesis (in the second one, we apply admissibility of weakening with $\forall y\, \chi_1'(y)$ in the antecedent). The eigenvariable condition on $\forall\mathrm{R}$ in the first derivation is satisfied because $b \notin \mathcal{L}(\Gamma_1, \Delta_1, \exists x\, \varphi(x))$, and χ_1' was defined so it does not contain b.

In the second case, we have (from the inductive hypothesis plus admissibility of weakening by $\exists y\, \chi_1'(y)$):

$$\dfrac{\dfrac{\dfrac{\vdots}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \varphi(b), \exists y\, \chi_1'(y), \chi_1'(b)}\,\exists\mathrm{R}}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \exists y\, \chi_1'(y), \chi_1'(b)}\,\exists\mathrm{R}}{\Gamma_1 \Rightarrow \Delta_1, \exists x\, \varphi(x), \exists y\, \chi_1'(y)}$$

and

$$\dfrac{\dfrac{\vdots}{\chi_1'(b), \Gamma_2 \Rightarrow \Delta_2}\,\exists\mathrm{L}}{\exists y\, \chi_1'(y), \Gamma_2 \Rightarrow \Delta_2}$$

The eigenvariable condition on $\exists\mathrm{L}$ in the second derivation is satisfied because $b \notin \mathcal{L}(\Gamma_2, \Delta_2)$, and χ_1' was defined so it does not contain b.

The other cases are similary; we leave them as exercises. □

Problem 2.7. Suppose we had a connective $\uparrow$ with rules

$$\dfrac{\Gamma \Rightarrow \Delta, \varphi \qquad \Gamma \Rightarrow \Delta, \psi}{\varphi \uparrow \psi, \Gamma \Rightarrow \Delta}\,\uparrow\mathrm{L} \qquad \dfrac{\varphi, \psi, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi \uparrow \psi}\,\uparrow\mathrm{R}$$

Show that Maehara's Lemma still holds by carrying out the cases in the proof of Lemma 2.10 where π ends in one of these rules.

Having established Maehara's Lemma, we can use it to prove Craig's Interpolation Theorem.

Proof of Theorem 2.9. Suppose $\varphi \Rightarrow \psi$ is derivable. Apply Lemma 2.10 to $\varphi \,;\, \Rightarrow \,;\, \psi$ to obtain the interpolant χ. We have $\vdash \varphi \Rightarrow \chi$ and $\vdash \chi \Rightarrow \psi$. □

Suppose Γ is a theory (set of sentences) in a language $\mathcal{L}$ containing an n-place predicate R. We say that Γ *implicitly defines* R, iff

$$\Gamma, \Gamma' \vDash \forall x_1 \ldots \forall x_n \, (R(x_1, \ldots, x_n) \leftrightarrow R'(x_1, \ldots, x_n)),$$

where Γ' is Γ with all occurrences of R replaced by $R' \notin \mathcal{L}$. Γ *explicitly defines* R iff there is some $\varphi(x_1, \ldots, x_n)$ not containing R such that

$$\Gamma \vDash \forall x_1 \ldots \forall x_n \, (R(x_1, \ldots, x_n) \leftrightarrow \varphi(x_1, \ldots, x_n)).$$

Clearly, if Γ explicitly defines R it implicitly defines R. The converse is not obvious, but nevertheless holds:

Lemma 2.11 (Beth Definability Theorem). *If Γ implicitly defines R then it also explicitly defines it.*

Proof. Suppose Γ implicitly defines R. Let R', Γ' be as above and $\mathcal{L}' = \mathcal{L}(\Gamma') = \mathcal{L} \setminus \{R\} \cup \{R'\}$. By hypothesis, we have

$$\Gamma, \Gamma' \Rightarrow \forall x_1 \ldots \forall x_n \, (R(x_1, \ldots, x_n) \leftrightarrow R'(x_1, \ldots, x_n))$$

Let $c_1, \ldots, c_n$ be constant symbols not occurring in Γ. By the invertibility lemma we have

$$\Gamma, \Gamma', R(c_1, \ldots, c_n) \Rightarrow R'(c_1, \ldots, c_n)$$

Apply Maehara's Lemma to

$$\Gamma, R(c_1, \ldots, c_n) \,;\, \Gamma' \Rightarrow \,;\, R'(c_1, \ldots, c_n)$$

to find an $\varphi \equiv \varphi(x_1, \ldots, x_n)$ such that

$$\begin{aligned} \Gamma, R(c_1, \ldots, c_n) &\Rightarrow \varphi(c_1, \ldots, c_n) \\ \varphi(c_1, \ldots, c_n), \Gamma' &\Rightarrow R'(c_1, \ldots, c_n) \end{aligned}$$

have derivations. Since $\mathcal{L}(\varphi) = \mathcal{L}_1 \cap \mathcal{L}_2 = \mathcal{L}(\Gamma) \setminus \{R\}$, φ does not contain R. Replace R' everywhere by R in the derivation of the second sequent to get a derivation of

$$\varphi(c_1, \ldots, c_n), \Gamma \Rightarrow R(c_1, \ldots, c_n)$$

Apply $\to$R and $\forall$R to get a derivation of

$$\Gamma \Rightarrow \forall x_1 \ldots \forall x_n \, (R(x_1, \ldots, x_n) \leftrightarrow \varphi(x_1, \ldots, x_n)).$$

This shows that $\varphi(x_1, \ldots, x_n)$ explicitly defines ρ in Γ. □

A third result which is equivalent to interpolation (and, like it, can be proved using Maehara's Lemma) is the following theorem: If two consistent theories Γ_1, Γ_2 contain a complete theory Γ in the language $\mathcal{L}(\Gamma) \subseteq \mathcal{L}(\Gamma_1) \cap \mathcal{L}(\Gamma_2)$, then $\Gamma_1 \cup \Gamma_2$ is consistent.

Recall that a complete theory is one such that for every sentence φ in its language either $\Gamma \vdash \varphi$ or $\Gamma \vdash \neg\varphi$. Since Γ_1 and Γ_2 are consistent extensions of Γ, Γ must be consistent. It follows that if φ is a sentence in the language $\mathcal{L}(\Gamma_1) \cap \mathcal{L}(\Gamma_2)$, we can't have $\Gamma_1 \models \varphi$ and $\Gamma_2 \models \neg\varphi$. To see this, suppose there were such an φ. If $\Gamma_1 \models \varphi$ then $\Gamma \models \varphi$. Otherwise, by completeness of Γ, $\Gamma \models \neg\varphi$, and since $\Gamma \subseteq \Gamma_1$, $\Gamma_1 \models \neg\varphi$, making Γ_1 inconsistent. So $\Gamma \models \varphi$ and since $\Gamma \subseteq \Gamma_2$, $\Gamma_2 \models \varphi$. So $\Gamma \not\models \neg\varphi$, or Γ_2 would be inconsistent.

The non-existence of a sentence φ in the joint language $\mathcal{L}(\Gamma_1) \cap \mathcal{L}(\Gamma_2)$ with $\Gamma_1 \models \varphi$ and $\Gamma_2 \models \neg\varphi$ is all that's needed for the proof, which we give here proof-theoretically using Maehara's Lemma.

Theorem 2.12 (Robinson's Joint Consistency Theorem). *Suppose Γ_1, Γ_2 are consistent theories and for no φ in the language $\mathcal{L}(\Gamma_1) \cap \mathcal{L}(\Gamma_2)$ do we have $\Gamma_1 \models \varphi$ and $\Gamma_2 \models \neg\varphi$. Then $\Gamma_1 \cup \Gamma_2$ is consistent.*

Proof. Suppose for the contrary that $\Gamma_1 \cup \Gamma_2$ is inconsistent. Then the sequent $\Gamma_1, \Gamma_2 \Rightarrow$ has a derivation. Apply Maehara's Lemma to $\Gamma_1 \,;\, \Gamma_2 \Rightarrow \;;$ to get a sentence φ in the language $\mathcal{L}(\Gamma_1) \cap \mathcal{L}(\Gamma_2)$ such that $\vdash \Gamma_1 \Rightarrow \varphi$ and $\vdash \varphi, \Gamma_2 \Rightarrow$. From the latter we get $\vdash \Gamma_2 \Rightarrow \neg\varphi$. But we assumed that no such sentence exists. □

Above we have tacitly assumed that the theories Γ, Γ_1, Γ_2 are finite. By compactness, the results also hold for infinite theories.

Chapter 3

Natural Deduction

3.1 Introduction

pt:nat:int:sec

Natural deduction is a family of derivation systems in which each logical connective or quantifier has a pair of inferences, one that "introduces" the operator (an introduction rule) and one that "eliminates" the operator (an elimination rule). For instance, $\rightarrow$Intro has formulas of the form $\varphi \rightarrow \psi$ as conclusion, and $\rightarrow$Elim has formulas of the form $\varphi \rightarrow \psi$ as premises.

The first two versions of natural deduction were introduced in the 1930s by Gerhard Gentzen and Stanisław Jaśkowski. Gentzen's is the system proof theorists discuss mostly, and Jaśkowski gave rise to the systems most widely used in teaching. In both systems, one can start with *assumptions*—formulas that don't need to be derived, but that can be used in deriving other formulas. An entire derivation may depend on such assumptions, and since those are unproven, the derivation does not establish its conclusion (its *end-formula*), but only that its conclusion follows from the assumptions.

Some inference rules are such that their conclusions no longer depend on some assumptions: they "discharge" assumptions. For instance, the $\rightarrow$Intro rule takes a derivation of the conclusion ψ from an assumption φ, and has $\varphi \rightarrow \psi$ as its conclusion. The derivation ending in $\varphi \rightarrow \psi$ no longer depends on φ; the assumption φ is discharged. In Gentzen's systems, derivations are trees of formulas, the assumptions are the topmost formulas, and rules like $\rightarrow$Intro carry labels that indicate which assumptions are discharged. In Jaśkowski's systems, parts of the derivation that depend on an assumption are set off in some way, e.g, indented behind a vertical line or in a box, with the assumption φ discharged and the consequent ψ of the conditional the first and last line in the box, and the conclusion $\varphi \rightarrow \psi$ of $\rightarrow$Intro outside the box.

The systems are "natural" because the inference rules mirror how we (or mathematicians, at least) reason naturally. To establish a hypothetical result (a conditional), we assume the antecedent and use it to prove the consequent. That's $\rightarrow$Intro. When we know a conditional $\varphi \rightarrow \psi$ and its antecedent φ, we conclude ψ. That's $\rightarrow$Elim. We use proof-by-cases to show that, assuming a

disjunction $\varphi \vee \psi$, some conclusion holds. That's $\vee$Elim. To prove a general claim $\forall x\, \varphi(x)$ we prove that $\varphi(c)$ holds, for an arbitrary c. That's $\forall$Intro. And so on.

For instance here's a simple derivation of $\varphi \to (\varphi \vee \psi)$ in Gentzen-style natural deduction:

$$\dfrac{\dfrac{[\varphi]^x}{\varphi \vee \psi}\,\vee\text{Intro}}{\varphi \to (\varphi \vee \psi)}\,{}^x\,{\to}\text{Intro}$$

It starts with the assumption φ, uses $\vee$Intro to infer $\varphi \vee \psi$, and then $\to$Intro to infer $\varphi \to (\varphi \vee \psi)$. The $\to$Intro inference discharges the assumption φ; that's indicated by the label x.

Proof theorists prefer Gentzen's original systems working with trees of formulas, although one variant is also important in proof theory. A derivation of φ from assumptions Γ shows that φ follows from Γ, i.e., $\Gamma \vDash \varphi$. The inference rules of natural deduction can naturally be read as telling us something about such entailments, e.g., $\to$Intro says that if $\Gamma + \varphi \vDash \psi$ then $\Gamma \vDash \varphi \to \psi$. It is only natural to formulate the rules directly so that they operate on both the assumptions and the conclusion that follow from them in the premises and conclusion of the inference rule, i.e., they operate on sequents $\Gamma \Rightarrow \varphi$. The simple proof above would look like this in such a system:

$$\dfrac{\dfrac{x : \varphi \;\Rightarrow\; \varphi}{x : \varphi \;\Rightarrow\; \varphi \vee \psi}\,\vee\text{Intro}}{\Rightarrow\; \varphi \to (\varphi \vee \psi)}\,{}^x\,{\to}\text{Intro}$$

As you see, the rules are the same: they work on the formula to the right of $\Rightarrow$, the formulas on the left just list the assumptions they depend on at each step.

The complete set of introduction/elimination rule pairs is not complete for classical logic by itself. We need to add two additional rules that have to do with $\bot$. The first is the "intuitionistic absurdity rule" $\bot_I$, or "explosion," which says that from $\bot$ we can infer anything. The second is the classical principle of indirect proof $\bot_C$ (or the "classical absurdity rule") which says that if we can derive $\bot$ from $\neg\varphi$, we have derived φ. If we leave out $\bot_C$, we get a system appropriate for intuitionsitic logic. If we leave out both, we have what's known as "minimal logic."

We will discuss Gentzen-style natural deduction for classical and intuitionsitic logic. These are the systems **N1c** and **N1i**. The systems **N2c** and **N2i** are the corresponding systems working on sequents $\Gamma \Rightarrow \varphi$ rather than single formulas φ.

3.2 Rules and Derivations

We begin by giving some definitions and fixing some terminology. pt:nat:rp:sec

Consider the two rules in **N1c** and **N1i** that mention $\varphi \to \psi$.

$$x\ \dfrac{\begin{array}{c}[\varphi]^x \\ \vdots \\ \psi\end{array}}{\varphi \to \psi}\ {\to}\text{Intro} \qquad \dfrac{\varphi \to \psi \qquad \varphi}{\psi}\ {\to}\text{Elim}$$

The →Elim rule is simple enough. The formulas above the line are the *premises*. The formula on the left is the *principal* formula $\varphi \to \psi$. All elimination rules have one such premise, always on the far left. It is called the *major* premise of the inference. If there are others, such as an this case, they are called *minor* premises. The premise(s) of introduction rules are also called major premise(s). The formula below the line is the *conclusion*.

The →Intro rule only has one premise, and here the conclusion is the principal formula $\varphi \to \psi$. All introduction rules are like this: the conclusion is the principal formula; its main operator is the one being "introduced."

The notation $[\varphi]^x$ means the following. In a derivation, the →Intro inference is applied to a formula ψ as premise. The sub-derivation ending in ψ has a number of initial formulas. These are called *assumptions*. A derivation of ψ from one or more assumptions of the form φ is a derivation that φ entails ψ. When we apply the →Intro rule, we conclude $\varphi \to \psi$, and the conclusion no longer depends on the assumption φ. To indicate this, the assumptions of the form φ in the derivation including the →Intro inference are *discharged* or *closed*. The rule states which assumptions can be discharged. In the case of →Intro with conclusion $\varphi \to \psi$, that's assumptions that have the form φ, i.e., agree with the antecedent of the conditional. It is sometimes important to indicate which assumptions are discharged at which rule, and this is accomplished with the discharge label x. Importantly, rules that discharge assumptions don't *have to* discharge *all* assumptions of the requisite form, or indeed any at all. For instance, the following is correct, even though the first →Intro does not discharge any assumptions:

$$y\ \dfrac{x\ \dfrac{[\varphi]^y}{\psi \to \varphi}\ {\to}\text{Intro}}{\varphi \to (\psi \to \varphi)}\ {\to}\text{Intro}$$

When a rule does not discharge any assumptions, we just leave out the discharge label (as we did here).

In the derivation

$$y\ \dfrac{x\ \dfrac{\dfrac{[\varphi]^x \qquad [\varphi]^y}{\varphi \wedge \varphi}\ {\wedge}\text{Intro}}{\varphi}\ {\wedge}\text{Elim}}{\varphi \to \varphi}\ {\to}\text{Intro}}{\varphi \to (\varphi \to \varphi)}\ {\to}\text{Intro}$$

the first →Intro inference discharges the left assumption φ^x, the second the right assumption φ^y. That is, all assumptions labelled x are discharged by the

$\to$Intro with label x, and all assumptions with label y by the one labelled y. For this to work, it is important that all assumptions with the same label are also the same formula.

The quantifier rules are slightly more complicated. E.g., the rules in **N1i** for $\exists$ are:

$$\frac{\varphi(t)}{\exists x\, \varphi(x)}\,\exists\text{Intro} \qquad \frac{\exists x\, \varphi(x) \qquad \begin{array}{c}[\varphi(c)]^x \\ \vdots \\ \psi\end{array}}{\psi}\,x\,\exists\text{Elim}$$

In the $\exists$Intro rule, the premise is $\varphi(t)$ where t is a closed term, i.e., a term without variables. Such an inference is correct—that is, it matches the schematic rule $\exists$Intro—if there is some formula φ, some variable x, and some closed term t such that the conclusion is $\exists x\, \varphi$ and premise is $\varphi[t/x]$. The $\exists$Elim allows us to discharge assumptions of the form $\varphi(c)$, i.e., for every inference there is some constant symbol c and assumptions of the form $\varphi[c/x]$ can be discharged. The discharged assumptions of one inference all have to use the same c. This constant is called the *eigenvariable* of the $\exists$Elim inference. The inference is only correct if c does not occur in an open assumption other than the discharged $\varphi(c)$, does not occur in ψ, and does not occur in φ. This is called the *eigenvariable condition*.

Derivations in **N1c** and **N1i** are finite trees of sequents that are inductively generated from axioms using the rules of inference. Each leaf is an assumption, possibly labelled with a discharge label, and every other formula follows from the formulas immediately above it by one of the inference rules of the system. If an occurence of a formula as an assumption in the tree above an inference has not been discharged by an inference above it, it is called *open* (at the inference).

Definition 3.1. *Derivations* in **N1c** and **N1i** are finite trees of formulas inductively as follows: pt:nat:rp: defn:proof

1. Every formula labelled by a discharge label such as x or y, by itself is a derivation. In a derivation consisting of a single formula, the formula counts as an open assumption. pt:nat:rp: defn:prf:assumption

2. If R is a rule with one, two, or three premises φ_1, φ_2, φ_3, conclusion φ, and δ_1, δ_2, δ_3 are derivations ending in φ_1, φ_2, φ_3, respectively, the pt:nat:rp: defn:prf:one-prem

$$\frac{\begin{array}{c}\vdots\,\delta_1 \\ \varphi_1\end{array}}{\varphi}\,R \qquad \frac{\begin{array}{c}\vdots\,\delta_1 \\ \varphi_1\end{array} \quad \begin{array}{c}\vdots\,\delta_2 \\ \varphi_2\end{array}}{\varphi}\,R \qquad \frac{\begin{array}{c}\vdots\,\delta_1 \\ \varphi_1\end{array} \quad \begin{array}{c}\vdots\,\delta_2 \\ \varphi_2\end{array} \quad \begin{array}{c}\vdots\,\delta_3 \\ \varphi_2\end{array}}{\varphi}\,R$$

   is a derivation, respectively, provided the assumption labels in the derivations are compatible (no two different formulas have the same discharge label) and the inferences are correct applications of R.

3. Topmost formula occurrences in the new derivation count as *open assumptions* of the derivation if they are open assumptions of δ_1, δ_2, or δ_3, except those occurrences discharged by the rule. If the rule is labelled x (or $x\,y$), then for $\to$Intro and $\neg$Intro, that is all open assumptions in δ_1 labelled x; for $\exists$Elim, it is all open assumptions labelled x in δ_2; for $\vee$Elim, it is all open assumptions labelled x in δ_2 and all open assumptions labelled y in δ_3.

The derivations used in the inductive construction of a derivation are called its *sub-derivations*. The sub-derivations ending in the premises of an inference are called the *immediate sub-derivations* of the inference. The rules of **N2c** and **N2i** are given in Table 3.1

Definition 3.2. The *height* $\mathrm{ht}(\delta)$ of a derivation δ is inductively defined as follows.

1. If δ consists of an assumption, $\mathrm{ht}(\delta) = 0$.
2. If δ ends in an inference with one premise with immediate sub-derivation δ_1, then $\mathrm{ht}(\delta) = \mathrm{ht}(\delta_1) + 1$.
3. If δ ends in an inference with two premises with immediate sub-derivations δ_1 and δ_2, then $\mathrm{ht}(\delta) = \max(\mathrm{ht}(\delta_1), \mathrm{ht}(\delta_2)) + 1$.
4. If δ ends in $\vee$Elim with immediate sub-derivations δ_1, δ_2, and δ_3, then $\mathrm{ht}(\delta) = \max(\mathrm{ht}(\delta_1), \mathrm{ht}(\delta_2), \mathrm{ht}(\delta_3)) + 1$.

If a sequent S has a derivation of height $\leq n$, we write $\vdash_n S$.

Example 3.3. In **N1i**, any formula counts as a derivation from itself as an open assumption. So

$$\varphi \wedge \psi^x$$

is a derivation δ_1. Its height is 0. From δ_1 we can obtain two derivations δ_2 and δ_3 by applying one of the two version of $\wedge$Elim

$$\frac{\varphi \wedge \psi^x}{\psi}\,\wedge\mathrm{Elim}_2 \qquad \frac{\varphi \wedge \psi^x}{\varphi}\,\wedge\mathrm{Elim}_1$$

The immediate sub-derivations of the last inference in δ_2 and in δ_3 is the same: $\varphi \wedge \psi$. Both derivations have the occurrence of $\varphi \wedge \psi$ as an open assumption. We have $\mathrm{ht}(\delta_1) = \mathrm{ht}(\delta_2) = 1$.

The rule $\wedge$Intro requires two premises of the form ψ and φ to justify a conclusion of the form $\psi \wedge \varphi$. Thus, the following is a derivation δ_4:

$$\delta_2\Big\{\frac{\dfrac{\varphi \wedge \psi^x}{\psi}\,\wedge\mathrm{Elim}_2 \qquad \dfrac{\varphi \wedge \psi^x}{\varphi}\,\wedge\mathrm{Elim}_1}{\psi \wedge \varphi}\,\wedge\mathrm{Intro}\Big\}\delta_3$$

Its height $\mathrm{ht}(\delta_4) = \max(\mathrm{ht}(\delta_2), \mathrm{ht}(\delta_3)) + 1 = \max(1, 1) + 1 = 2$. It has two occurences of $\varphi \wedge \psi$ as assumptions, both are open.

$$\dfrac{\begin{array}{c}[\varphi]^x\\ \vdots\\ \psi\end{array}}{\varphi \to \psi}\ {\scriptstyle x}\ \to\text{Intro} \qquad \dfrac{\varphi \to \psi \qquad \varphi}{\psi}\ \to\text{Elim}$$

$$\dfrac{\varphi \qquad \psi}{\varphi \wedge \psi}\ \wedge\text{Intro} \qquad \dfrac{\varphi \wedge \psi}{\varphi}\ \wedge\text{Elim}_1 \qquad \dfrac{\varphi \wedge \psi}{\psi}\ \wedge\text{Elim}_2$$

$$\dfrac{\varphi}{\varphi \vee \psi}\ \vee\text{Intro}_1 \qquad \dfrac{\psi}{\varphi \vee \psi}\ \vee\text{Intro}_2 \qquad \dfrac{\varphi \vee \psi \qquad \begin{array}{c}[\varphi]^x\\ \vdots\\ \chi\end{array} \qquad \begin{array}{c}[\psi]^y\\ \vdots\\ \chi\end{array}}{\chi}\ {\scriptstyle x\,y}\ \vee\text{Elim}$$

$$\dfrac{\begin{array}{c}[\varphi]^x\\ \vdots\\ \bot\end{array}}{\neg\varphi}\ {\scriptstyle x}\ \neg\text{Intro} \qquad \dfrac{\neg\varphi \qquad \varphi}{\bot}\ \neg\text{Elim}$$

$$\dfrac{\varphi(a)}{\forall x\, \varphi(x)}\ \forall\text{Intro} \qquad \dfrac{\forall x\, \varphi(x)}{\varphi(t)}\ \forall\text{Elim}$$

$$\dfrac{\varphi(t)}{\exists x\, \varphi(x)}\ \exists\text{Intro} \qquad \dfrac{\exists x\, \varphi(x) \qquad \begin{array}{c}[\varphi(a)]^x\\ \vdots\\ \chi\end{array}}{\chi}\ {\scriptstyle x}\ \exists\text{Elim}$$

$$\dfrac{\bot}{\varphi}\ \bot_I \qquad \dfrac{\begin{array}{c}[\neg\varphi]^x\\ \vdots\\ \bot\end{array}}{\varphi}\ {\scriptstyle x}\ \bot_C$$

Table 3.1: Rules of **N1i** and **N1c**. In ∀Intro the constant symbol a must not occur in the conclusion or an open assumption. In ∃Elim, a must not occur in the premises or an open assumption. **N1i** does not include $\bot_C$.

pt:nat:rp:
tab:N1

Finally we can apply $\to$Intro to obtain δ_5:

$$\dfrac{\dfrac{\delta_2\left\{\dfrac{[\varphi\wedge\psi]^x}{\psi}\wedge\text{Elim}_2 \qquad \dfrac{[\varphi\wedge\psi]^x}{\varphi}\wedge\text{Elim}_1\right\}\delta_3}{\psi\wedge\varphi}\wedge\text{Intro}}{(\varphi\wedge\psi)\to(\psi\wedge\varphi)}\,{\scriptstyle x}\,\to\text{Intro}$$

Its height is $\text{ht}(\delta_4)+1=3$. The $\to$Intro inference discharges all occurrences of assumptions of the form $\varphi\wedge\psi$ labelled x, i.e., both previously open assumptions of the immediate subproof. Since those are the only assumptions of δ_5, it has no open assumptions, and so is a derivation *of* its end-formula $(\varphi\wedge\psi)\to(\psi\wedge\varphi)$.

3.3 Natural Deduction with Sequents: N2c and N2i

pt:nat:seq: sec

A derivation δ in **G1c** or **G1i** shows that its end-formula φ follows from its open assumptions. If Γ is the set of formulas ψ, where ψ^x occurs as an open assumptions in δ, we can write this as $\delta\vdash\Gamma\Rightarrow\varphi$. This suggests that we can formulate a natural deduction system also for sequents. In such a system, each sequent carries the information of which open assumptions the formula on the right depends on, i.e., which are open in the sub-derivation ending in a sequent. This is natural deduction with sequents, or natural deduction "in sequent style," and we call the resulting systems **N2c** and **N2i**.

To make the correspondence between **N1** and **N2** closer, let's take the sets Γ on the left to consist of labelled formulas $x:\psi$. Since rules in **N1** only operate on the formulas in the premises, i.e., the end-formulas of the immediate sub-derivations, rules of **N2** only operate on the succedent of the sequent. We can thus refer to the labelled formulas on the left of sequents simply as the context.

Instead of assumptions, **N2** proofs have as topmost sequents axioms of the form

$$x:\varphi\Rightarrow\varphi.$$

The rules correspond exactly to the rules of **N1** and are given in Table 3.2.

Because we must allow (as we do in **N1c** and **N1i**) that the rules $\to$Intro, $\neg$Intro, $\vee$Elim, and $\exists$Elim may, but need not actually discharge assumptions, we allow correct applications of the corresponding rules in **N2c** and **N2i** also when the corresponding context formulas φ, ψ, and $\varphi(a)$ are not present in the context of the relevant premise.

Proposition 3.4. *If δ is a derivation in* **N1c** *or* **N1i** *of φ, then there is a derivation δ' of $\Gamma\Rightarrow\varphi$ in* **N2c** *(***N2i***), where Γ is the set of all $x:\psi$ such that ψ^x is an open assumption of δ.*

Proof. By induction on $n=\text{ht}(\delta)$. If $n=0$, δ is a single open assumption φ^x. Then $\Gamma=\{x:\varphi\}$ and $\Gamma\Rightarrow\varphi$ is an axiom of **N2c** or **N2i**.

If $n>0$ we distinguish cases according to the last inference in δ. We discuss just a one case; the rest are left as exercises.

$$\frac{x : \varphi, \Gamma \Rightarrow \psi}{\Gamma \Rightarrow \varphi \to \psi} \to\text{Intro} \qquad \frac{\Gamma_1 \Rightarrow \varphi \to \psi \qquad \Gamma_2 \Rightarrow \varphi}{\Gamma_1, \Gamma_2 \Rightarrow \psi} \to\text{Elim}$$

$$\frac{\Gamma \Rightarrow \varphi \wedge \psi}{\Gamma \Rightarrow \varphi} \wedge\text{Elim}_1$$

$$\frac{\Gamma_1 \Rightarrow \varphi \qquad \Gamma_2 \Rightarrow \psi}{\Gamma_1, \Gamma_2 \Rightarrow \varphi \wedge \psi} \wedge\text{Intro} \qquad \frac{\Gamma \Rightarrow \varphi \wedge \psi}{\Gamma \Rightarrow \psi} \wedge\text{Elim}_2$$

$$\frac{\Gamma \Rightarrow \varphi}{\Gamma \Rightarrow \varphi \vee \psi} \vee\text{Intro}_1 \qquad \frac{\Gamma \Rightarrow \psi}{\Gamma \Rightarrow \varphi \vee \psi} \vee\text{Intro}_2$$

$$\frac{\Gamma_1 \Rightarrow \varphi \vee \psi \qquad x : \varphi, \Gamma_2 \Rightarrow \chi \qquad y : \psi, \Gamma_3 \Rightarrow \chi}{\Gamma_1, \Gamma_2, \Gamma_3 \Rightarrow \chi} \vee\text{Elim}$$

$$\frac{x : \varphi, \Gamma \Rightarrow \bot}{\Gamma \Rightarrow \neg\varphi} \neg\text{Intro} \qquad \frac{\Gamma_1 \Rightarrow \neg\varphi \qquad \Gamma_2 \Rightarrow \varphi}{\Gamma_1, \Gamma_2 \Rightarrow \bot} \neg\text{Elim}$$

$$\frac{\Gamma \Rightarrow \varphi(a)}{\Gamma \Rightarrow \forall x\, \varphi(x)} \forall\text{Intro} \qquad \frac{\Gamma \Rightarrow \forall x\, \varphi(x)}{\Gamma \Rightarrow \varphi(t)} \forall\text{Elim}$$

$$\frac{\Gamma \Rightarrow \varphi(t)}{\Gamma \Rightarrow \exists x\, \varphi(x)} \exists\text{Intro} \qquad \frac{\Gamma_1 \Rightarrow \exists x\, \varphi(x) \qquad x : \varphi(a), \Gamma_2 \Rightarrow \chi}{\Gamma_1, \Gamma_2 \Rightarrow \chi} \exists\text{Elim}$$

$$\frac{\Gamma \Rightarrow \bot}{\Gamma \Rightarrow \varphi} \bot_I \qquad \frac{x : \neg\varphi, \Gamma \Rightarrow \bot}{\Gamma \Rightarrow \varphi} \bot_C$$

Table 3.2: Rules of **N2c** and **N2i**. In ∀Intro and ∃Elim, the constant symbol a must not occur in the conclusion sequent. Γ_i are sets of labelled formulas $x : \varphi$. The rules ∀Intro and ∃Elim must satisfy the eigenvariable condition: a must not occur in the conclusion sequent. **N2i** does not include $\bot_C$.

pt:nat:seq:tab:N2

Suppose the last inference is →Intro. Then δ ends in

$$x\,\frac{\begin{array}{c}[\psi]^x \\ \vdots\ \delta_1 \\ \chi\end{array}}{\psi \to \chi} \to\text{Intro}$$

By induction hypothesis there is a derivation δ_1' in **N2c** (**N2i**) of $\Gamma' \Rightarrow \chi$, where Γ' are the labelled assumptions open in δ_1. If ψ^x is an open assumption of δ_1, then it is discharged at the →Intro inference, and the open assumptions of δ are $\Gamma = \Gamma' \setminus \{x : \psi\}$. In other words, δ_1' is a derivation of $x : \psi, \Gamma \Rightarrow \chi$,

and we can take δ' to be

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1'\\ \vdots\\ x : B, \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ {\to}\text{Intro}$$

If ψ^x is not an open assumption of δ_1, then $\to$Intro does not discharge an assumption, and the open assumptions of δ and δ_1 are the same. Since the context formula $x : \psi$ is not required to be present in the premise of $\to$Intro in **N2c** and **N2i**, we can take δ' to be

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1'\\ \vdots\\ \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ {\to}\text{Intro}$$

The cases where δ ends in another rule are similar. □

Proposition 3.5. *If δ is a derivation in* **N2c** *or* **N2i** *of $\Gamma \Rightarrow \varphi$, then there is a derivation δ' in* **N1c** *(***N1i***) with end-formula φ and open assumptions ψ^x for every $x : \psi \in \Gamma$.*

Proof. We again proceed by induction on the height n of δ. If $n = 0$, δ is an axiom, $x : \varphi \Rightarrow \varphi$. The derivation δ' is the assumption φ^x.

If $n > 0$, we distinguish cases according to the last inference in δ. For instance, if that inference is $\to$Intro, δ ends in

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ x : \psi, \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ {\to}\text{Intro} \quad \text{or} \quad \dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ {\to}\text{Intro}$$

By inductive hypothesis, there is a derivation in **N1c** (**N1i**) of χ. In the first case, ψ^x is an open assumption of δ'; in the second case it is not. We can take the derivation δ' to be

$$\dfrac{\begin{array}{c}[\psi]^x\\ \vdots\\ \delta_1\\ \vdots\\ \chi\end{array}}{\psi \to \chi}\ {\to}\text{Intro}\ {}^x \quad \text{or} \quad \dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ \chi\end{array}}{\psi \to \chi}\ {\to}\text{Intro}$$

respectively. □

3.4 Regular Derivations and Substitution

The rules for the quantifiers introduce some complications owing to the "eigenvariable conditions" that apply to the ∃Elim and ∀Intro rules. Most of these complications can be dealt with by ensuring that the constant symbol c in these inferences is never reused. Let's introduce the following definition. pt:nat:qua:sec

Definition 3.6. A natural deduction derivation δ is *regular* if

1. no eigenvariable a is used as the eigenvariable of more than one ∃Elim or ∀Intro inference, and

2. if a is the eigenvariable of an ∃Elim or ∀Intro inference, it only occurs in the immediate sub-derivation leading to the minor premise, or the premise, respectively.

Lemma 3.7. *If δ is regular, ends in χ, c is a constant symbol not used as an eigenvariable in δ, and t is a term not containing any eigenvariables of δ, then $\delta[t/c]$ resulting from δ by replacing every occurrence of c by t is a regular derivation. The end-formula of $\delta[t/c]$ is $\chi[t/c]$. If φ^x is an open assumption of δ then $\varphi[t/c]^x$ is an open assumption of $\delta[t/c]$.* pt:nat:qua:lem:subst-regular

Proof. By induction on $\mathrm{ht}(\delta)$. If $\mathrm{ht}(\delta) = 0$, it consists simply of an assumption φ^x. Then $\delta[t/c]$ is $\varphi[t/c]^x$.

If $\mathrm{ht}(\delta) > 0$ we distinguish cases according to the last inference in δ. The cases for rules for propositional connectives are easy and left as exercises. We consider the cases where δ ends in a quantifier inference.

1. If the last inference is ∀Intro, δ is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \delta'\\ \vdots\\ \varphi(a,c)\end{array}}{\forall x\, \varphi(x,c)}\,\forall\text{Intro}$$

   By induction hypothesis, $\delta'[t/c]$ is a regular derivation of $\varphi(a,t)$. Since a is an eigenvariable of δ, a does not occur in t. Therefore, the last inference in $\delta[t/c]$,

$$\dfrac{\begin{array}{c}\vdots\\ \delta'[t/c]\\ \vdots\\ \varphi(a,t)\end{array}}{\forall x\, \varphi(x,t)}\,\forall\text{Intro}$$

   is correct: since t does not contain a, neither the conclusion nor any open assumption contains a.

2. If the last inference is ∀Elim, δ is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \delta'\\ \vdots\end{array}}{\dfrac{\forall x\,\varphi(x,c)}{\varphi(s(c),c)}\,\forall\text{Elim}}$$

By induction hypothesis, $\delta'[t/c]$ is a regular derivation of $\forall x\,\varphi(x,t)$. (The term $s(c)$ in the conclusion may, but need not, contain c.) The last inference in $\delta[t/c]$,

$$\dfrac{\begin{array}{c}\vdots\\ \delta'[t/c]\\ \vdots\end{array}}{\dfrac{\forall x\,\varphi(x,t)}{\varphi(s(t),t)}\,\forall\text{Elim}}$$

is correct, and $\varphi(s(t),t) \equiv \varphi(s(c),c)[t/c]$. The case for ∃Intro is similar.

3. If the last inference is ∃Elim, δ is of the form

$$\dfrac{\exists x\,\varphi(x,c)\quad\begin{array}{c}[\varphi(a,c)]^x\\ \vdots\\ \chi\end{array}}{\chi}\,\exists\text{Elim}\;x$$

By induction hypothesis, $\delta'_1[t/c]$ and $\delta'_2[t/c]$ are regular derivations of $\exists x\,\varphi(x,c)$ and of χ from the open assumption $\varphi(a,t)$. Since a is an eigenvariable of δ, a does not occur in t. Therefore, the last inference in $\delta[t/c]$,

$$\dfrac{\begin{array}{c}\vdots\\ \delta'[t/c]\\ \vdots\end{array}}{\dfrac{\varphi(a,t)}{\forall x\,\varphi(x,t)}\,\forall\text{Intro}}$$

is correct: $\varphi(x,t)[a/x] \equiv \varphi(a,c)[t/c]$, neither the conclusion χ nor any open assumption contains a. □

pt:nat:qua: prop:regular **Proposition 3.8.** *If δ is a derivation in* **N1c** *or* **N1i**, *then there is a regular derivation δ' with the same structure (in particular, the same height) which differs from δ only by replacing eigenvariables.*

Proof. For the purpose of the proof only, call an eigenvariable inference of δ *clean* if its eigenvariable is not the eigenvariable of another inference nor does it occur anywhere in δ except the immediate sub-derivation of the inference, and *dirty* otherwise. Let n be the number of dirty eigenvariable inferences in δ. We

show, by induction on n, that δ can be transformed into the required regular derivation δ'. If $n = 0$, all eigenvariable inferences in δ are clean, and so δ is regular, and there is nothing to prove. Otherwise, pick a highest eigenvariable inference, say ∀Intro. The sub-derivation δ_1 of δ ending in it is of the form

$$\dfrac{\begin{array}{c}\vdots\\ \delta_2\\ \vdots\\ \varphi(c)\end{array}}{\forall x\, \varphi(x)}\,\forall\text{Intro}$$

Note that as c is an eigenvariable, it does not occur in $\forall x\, \varphi(x)$ or an open assumption. The proof δ_2 is regular, since it contains no eigenvariable inferences at all. Let a be a constant symbol not occurring in δ. By Lemma 3.7, $\delta_2[a/c]$ is a regular proof of $\varphi(a)$. By the eigenvariable condition, no open assumption of δ_2 contains c, hence no open assumption of $\delta_2[a/c]$ contains a. Thus we can apply ∀Intro. If we replace δ_1 in δ by the proof δ_1',

$$\dfrac{\begin{array}{c}\vdots\\ \delta_2[a/c]\\ \vdots\\ \varphi(a)\end{array}}{\forall x\, \varphi(x)}\,\forall\text{Intro}$$

we have replaced a dirty eigenvariable inference by a clean one, and the only difference is the replacement of the eigenvariable c by a. The resulting proof has fewer than n dirty eigenvariable inferences, and the result follows from the inductive hypothesis. □

Problem 3.1. Describe the case in the proof of Proposition 3.8 where the topmost eigenvariable inference is ∃Elim.

We will not appeal to the proposition just proved explicitly, but we will assume that all derivations considered are regular—if they are not, we can always replace eigenvariables to make them regular.

Lemma 3.7 and Proposition 3.8 also hold for **N2c** and **N2i**. We leave the proofs as exercises.

3.5 Grafting Derivations

pt:nat:gra:sec

Derivations in **N1c** and **N1i** are derivations of formulas φ from assumptions ψ. Suppose ψ itself has a derivation. How do we combine the derivation of ψ with the derivation of φ from ψ to get a derivation of φ that doesn't use the assumption ψ?

One option is to remove the assumption ψ from the derivation of φ by applying →Intro. Then we can combine the resulting derivation of $\psi \to \varphi$ with

the derivation of ψ to obtain

$$\dfrac{\dfrac{\begin{array}{c}[\psi]^x\\ \vdots\\ \varphi\end{array}}{\psi \to \chi}\, x \to\text{Intro} \qquad \begin{array}{c}\vdots\\ \psi\end{array}}{\varphi} \to\text{Elim}$$

This is straightforward, but somewhat ugly: why introduce $\to$ only to eliminate it right away? It should be possible to avoid this kind of "detour" through $\psi \to \varphi$. (That they can always be avoided is in fact the content of the normalization theorem.)

In **N1c** and **N1i** we can actually avoid it in this case quite easily: we can simply replace the assumptions ψ^x by the entire derivation of ψ. This "grafting" of derivations onto the assumptions of others is quite useful and so we record it as a definition.

Definition 3.9. If δ_1 is a **N1c**-derivation (or **N1i**-derivation) of φ from open assumption ψ^x, and δ_2 is a **N1c**-derivation (**N1i**-derivation) of ψ, then $\delta_1[\delta_2/\psi^x]$ is the result of replacing every undischarged occurrence of ψ^x in δ_1 by δ_2.

Provided that δ_1 and δ_2 don't have different assumption formulas labelled with the same label, and no open assumption of δ_2 contains an eigenvariable of δ_1, the resulting $\delta_1[\delta_2/\psi^x]$ is a correct derivation, and its open assumptions are those of δ_1 together with those of δ_2. When grafting derivations, these conditions will typically be satisfied. The first is satisfied whenever δ_1 and δ_2 are subproofs of a larger derivation. If the second is not satisfied, we can rename the eigenvariables of δ_1 so that it is.

A corresponding definition applies to **N2c** and **N2i**.

Definition 3.10. If δ_1 is a **N2c**-derivation (or **N2i**-derivation) of $x : \psi, \Gamma_1 \Rightarrow \varphi$, and δ_2 is a **N1c**-derivation (**N1i**-derivation) of $\Gamma_2 \Rightarrow \psi$, then $\delta_1[\delta_2/x : \psi]$ is the result of replacing every axiom $x : \psi \Rightarrow \psi$ in δ_1 by δ_2, and adding Γ_2 to the context of every sequent below such axioms.

Proposition 3.11. *If $\delta_1 \vdash x : \psi, \Gamma_1 \Rightarrow \varphi$ and $\delta_2 \vdash \Gamma_2 \Rightarrow \psi$, then $\delta_1[\delta_2/x : \psi] \vdash \Gamma_1, \Gamma_2 \Rightarrow \varphi$, provided no eigenvariable of δ_1 occurs in Γ_2.*

Proof. By induction on $\mathrm{ht}(\delta)$. □

3.6 Translating from G2i to N2i

pt:nat:tgi:sec

The antecedents of sequents in **G2i** are multisets Γ of formulas, while those of sequents in **N2i** are sets Γ' of labelled formulas with the property that every label occurs only once. Let's say a set Γ' of labelled formulas *corresponds to* the multiset Γ if, for every formula φ, the number of elements of Γ' of the form

$x : \varphi$ is less than or equal to the multiplicity of φ (i.e., the number of copies of φ) in Γ.

Proposition 3.12. *If* $\mathbf{G2i} + \mathrm{Cut} \vdash \Gamma \Rightarrow \Delta$ *then* $\mathbf{N2i} \vdash \Gamma' \Rightarrow \chi$, *where* $\chi \equiv \bot$ *if* Δ *is empty and* $\Delta = \{\varphi\}$ *otherwise, and* Γ' *corresponds to* Γ. *pt:nat:tgi: prop:G2i-to-N2i*

Proof. We show that if π is a derivation of $\Gamma \Rightarrow \Delta$ in **G2ci**, then there is a Γ' corresponding to Γ and a derivation δ of $\Gamma' \Rightarrow \chi$ in **N2i**. We do this by induction on $n = \mathrm{ht}(\pi)$.

If $n = 0$, then π is an axiom. If it is the axiom $\bot \Rightarrow$, δ is the axiom $x : \bot \Rightarrow \bot$, i.e., $\Gamma' = \{x : \bot\}$ and $\chi \equiv \bot$. If π is an axiom of the form $\varphi \Rightarrow \varphi$, then δ is $x : \varphi \Rightarrow \varphi$.

If $n > 0$, π ends in a rule R. The inductive hypothesis guarantees that there are derivations of sequents corresponding to the premises of that rule in **N2c**. We can assume that all discharge labels in those proofs are distinct, and that a label x only occurs above an inference that discharges it, if it is discharged at all (by relabelling if necessary). We then show how these derivations can be combined to provide proofs of a suitable $\Gamma' \Rightarrow \chi$. We distinguish cases according to R.

1. If R is WR, π ends in

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma \Rightarrow\end{array}}{\Gamma \Rightarrow \varphi}\ \mathrm{WR}$$

   The induction hypothesis applied to π_1 yields a derivation of $\Gamma' \Rightarrow \bot$. We apply $\bot_I$ to get a derivation of $\Gamma' \Rightarrow \varphi$.

2. If R is WL, π ends in

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma \Rightarrow \Delta\end{array}}{\varphi, \Gamma \Rightarrow \Delta}\ \mathrm{WR}$$

   The induction hypothesis applied to π_1 yields a derivation of $\Gamma' \Rightarrow \chi$. Which we take as δ. Note that if Γ' corresponds to Γ, it also corresponds to φ, Γ as the number of $x : \varphi$ in Γ' is no greater than the number of copies of φ in Γ, and so also no greater than the number of copies of φ in $\Gamma \cup \{\varphi\}$.

3. If R is CL, π ends in

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \varphi, \varphi, \Gamma \Rightarrow \Delta\end{array}}{\varphi, \Gamma \Rightarrow \Delta}\ \mathrm{CL}$$

The induction hypothesis applied to π_1 yields a derivation of $\Pi, \Gamma' \Rightarrow \chi$, where $\Pi \subseteq \{x : \varphi, y : \varphi\}$. If $\Pi = \{x : \varphi, y : \varphi\}$, replace all labelled formulas $y : \varphi$ in π_1 by $x : \varphi$. Since x occurs in the context of the end-sequent of π_1 we can assume that no inference in π_1 discharges a formula of the form $x : \psi$, i.e., no $x : \psi$ on the left side of a sequent in π_1 is active. Consequently, the result is still a derivation in **N2i**.

4. If R is $\to$R, π ends in

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \varphi, \Gamma \Rightarrow \psi\end{array}}{\Gamma \Rightarrow \varphi \to \psi}\,{\to}\mathrm{R}$$

   By the inductive hypothesis, there is a derivation δ_1 of $x : \varphi, \Gamma' \Rightarrow \psi$ or $\Gamma' \Rightarrow \psi$ where Γ' corresponds to Γ. We can take δ_1 to be δ plus an application of $\to$Intro.

5. If R is $\to$L, π ends in

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \Gamma_1 \Rightarrow \varphi & \psi, \Gamma_2 \Rightarrow \Delta\end{array}}{\varphi \to \psi, \Gamma_1, \Gamma_2 \Rightarrow \Delta}\,{\to}\mathrm{L}$$

   The inductive hypothesis yields derivations δ_1 of $\Gamma_1' \Rightarrow \varphi$ and δ_2 of $x : \psi, \Gamma_2' \Rightarrow \chi$ or of of $\Gamma_2' \Rightarrow \chi$. If δ_2 is of the latter form, δ_2 can serve as δ. Otherwise, we can relabel all formulas and discharge labels in π_1 to ensure that no label occurs in both δ_1 and δ_2. Then $\Gamma_1' \cup \Gamma_2'$ corresponds to $\Gamma_1 \cup \Gamma_2$. Since $x : \psi$ occurs in the end-sequent of δ_2, $x : \psi$ cannot be active in any inference with discharge label x in δ_2. Replace every axiom $x : \psi \Rightarrow \psi$ in δ_2 with the derivation δ_3

$$\dfrac{\begin{array}{cc} & \vdots\,\delta_1\\ x : \varphi \to \psi \Rightarrow \varphi \to \psi & \Gamma_1' \Rightarrow \varphi\end{array}}{x : \varphi \to \psi \Rightarrow \psi}\,{\to}\mathrm{Elim}$$

   and every occurrence of $x : \psi$ in δ_2 by $x : \varphi \to \psi$. The result $\delta_2[\delta_3/x : \psi]$ is a derivation of $x : \varphi \to \psi, \Gamma_1', \Gamma_2' \Rightarrow \chi$.

6. If R is $\wedge$R, π ends in

$$\dfrac{\begin{array}{cc}\vdots\,\pi_1 & \vdots\,\pi_2\\ \Gamma_1 \Rightarrow \varphi & \Gamma_2 \Rightarrow \psi\end{array}}{\Gamma_1, \Gamma_2 \Rightarrow \varphi \wedge \psi}\,{\wedge}\mathrm{R}$$

The inductive hypothesis yields proof δ_1 of $\Gamma_1' \Rightarrow \varphi$ and δ_2 of $\Gamma_2' \Rightarrow \varphi$. By relabelling all formulas and discharge labels in π_2 we can ensure that no label occurs in both δ_1 and δ_2. Then $\Gamma_1' \cup \Gamma_2'$ corresponds to $\Gamma_1 \cup \Gamma_2$. We obtain a derivation δ by applying $\wedge$Intro to the end-sequents of δ_1 and δ_2.

7. If R is $\wedge$L, π ends in, e.g.,

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\begin{array}{c}\varphi, \Gamma \Rightarrow \Delta\\ \hline \varphi \wedge \psi, \Gamma \Rightarrow \Delta\end{array}} \wedge\mathrm{L}$$

By the inductive hypothesis, there is a derivation δ_1 of $x : \varphi, \Gamma' \Rightarrow \chi$ or $\Gamma' \Rightarrow \chi$ where Γ' corresponds to Γ. In the latter case, we can take δ_1 to be δ. In the former case, observe that since $x : \varphi$ occurs in the end-sequent, $x : \varphi$ is not active in any inference with discharge label x in δ_1. Replace every axiom $x : \varphi \Rightarrow \varphi$ with the derivation δ_2

$$\dfrac{x : \varphi \wedge \psi \Rightarrow \varphi \wedge \psi}{x : \varphi \wedge \psi \Rightarrow \varphi} \wedge\mathrm{Elim}$$

and every occurrence of $x : \varphi$ by $x : \varphi \wedge \psi$, i.e., consider $\delta_1[\delta_2/x : \varphi]$. This is a derivation of $x : \varphi \wedge \psi, \Gamma' \Rightarrow \chi$.

8. If R is Cut, π ends in

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array} \qquad \begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\begin{array}{c}\Gamma_1 \Rightarrow \varphi \qquad \varphi, \Gamma_2 \Rightarrow \Delta\\ \hline \Gamma_1, \Gamma_2 \Rightarrow \Delta\end{array}} \mathrm{Cut}$$

The inductive hypothesis yields proof δ_1 of $\Gamma_1' \Rightarrow \varphi$ and δ_2 of $x : \varphi, \Gamma_2' \Rightarrow \chi$ or of $\Gamma_2' \Rightarrow \chi$. In the latter case, take δ to be δ_2. Otherwise, replace every axiom $x : \varphi \Rightarrow \varphi$ in δ_2 by δ_1 to obtain δ as the derivation $\delta_2[\delta_1/x : \varphi]$ of $\Gamma_1', \Gamma_2' \Rightarrow \chi$. □

3.7 Translating from N2 to G2

pt:nat:tni: sec
pt:nat:tni: prop:N2-to-G2

Proposition 3.13. *If* **N2c** (**N2i**) $\vdash \Gamma \Rightarrow \varphi$ *then* **G2c** (**N2i**) + Cut $\vdash \Gamma' \Rightarrow \Delta$ *where* Γ' *is the multiset of formulas resulting from* Γ *by removing labels, and* $\Delta = \{\varphi\}$ *if* φ *is not* $\bot$*, and* $\Delta = \emptyset$ *if it is.*

Proof. The proof is again by induction on the height of the derivation δ of $\Gamma \Rightarrow \varphi$. If δ is an axiom $x : \varphi \Rightarrow \varphi$, π is the corresponding axiom $\varphi \Rightarrow \varphi$ or $\bot \Rightarrow$ if $\varphi \equiv \bot$.

If δ ends in an inference rule, we distinguish cases:

1. If R is $\to$Intro with principal formula $\psi \to \chi$ and $x : \psi$ is present in the context of the premise but not the conclusion, then δ ends in

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ x : \psi, \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ \to\text{Intro}$$

   The inductive hypothesis yields a derivation π_1 of $\psi, \Gamma' \Rightarrow \chi$ (or of $\psi, \Gamma' \Rightarrow$ if $\chi \equiv \bot$), from which we obtain π by applying $\to$R (after applying WR in case $\chi \equiv \bot$).

2. If R is $\to$Intro, but $x : \psi$ is not present in the context of the premise, then δ ends in

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ \Gamma \Rightarrow \chi\end{array}}{\Gamma \Rightarrow \psi \to \chi}\ \to\text{Intro}$$

   The inductive hypothesis yields a derivation π_1 of $\Gamma' \Rightarrow \chi$, from which we obtain π by applying WL with ψ, and then $\to$R.

3. If R is $\to$Elim, δ ends in

$$\dfrac{\begin{array}{c}\vdots\\ \delta_1\\ \vdots\\ \Gamma_1 \Rightarrow \psi \to \varphi\end{array} \qquad \begin{array}{c}\vdots\\ \delta_2\\ \vdots\\ \Gamma_2 \Rightarrow \psi\end{array}}{\Gamma_1, \Gamma_2 \Rightarrow \varphi}\ \to\text{Elim}$$

   where $\Gamma = \Gamma_1 \cup \Gamma_2$. The inductive hypothesis yields derivations π_1 of $\Gamma_1' \Rightarrow \psi \to \varphi$ and π_2 of $\Gamma_2' \Rightarrow \psi$. We obtain the derivation π as

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\\ \Gamma_1' \Rightarrow \psi \to \varphi\end{array} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\\ \Gamma_2' \Rightarrow \psi\end{array} \qquad \varphi \Rightarrow \varphi}{\psi \to \varphi, \Gamma_2' \Rightarrow \varphi}\ \to\text{L}}{\Gamma_1', \Gamma_2' \Rightarrow \varphi}\ \text{Cut}$$

   The multiset $\Gamma_1' \cup \Gamma_2'$ may not equal $\Gamma' = (\Gamma_1 \cup \Gamma_2)'$: e.g., if both Γ_1 and Γ_2 contain $x : \chi$ then $\Gamma_1' \cup \Gamma_2'$ contains two copies of χ but $(\Gamma_1 \cup \Gamma_2)'$ contains only one copy. We can correct this by adding suitable CL inferences.

   In the event that $\varphi \equiv \bot$, the inductive hypothesis applied to δ_2 yields a derivation of $\Gamma_2' \Rightarrow$, from which we can obtain the derivation π by adding suitable WL inferences and a WR with φ.

4. If R is $\bot_C$, δ ends in

$$\dfrac{\begin{array}{c}\vdots\,\delta_1\\ x : \neg\varphi, \Gamma \Rightarrow \bot\end{array}}{\Gamma \Rightarrow \varphi}\,\bot_C$$

By inductive hypothesis, there is a derivation π_1 of $\neg\varphi, \Gamma' \Rightarrow$. We take the derivation π to be

$$\dfrac{\dfrac{\varphi \Rightarrow \varphi}{\Rightarrow \varphi, \neg\varphi}\,\neg\mathrm{R} \qquad \begin{array}{c}\vdots\,\pi_1\\ x : \neg\varphi, \Gamma' \Rightarrow\end{array}}{\Gamma' \Rightarrow \varphi}\,\mathrm{Cut}$$

Note that only the translation of $\bot_C$ results in a derivation π with more than one formula on the right. That is, the translation of a **N2i** derivation is a **G2i**-derivation. □

Chapter 4

Normalization

4.1 Introduction

pt:nor:int: sec

If you have derived a conditional $\varphi \to \psi$, you might be in a position to use it in a derivation of ψ by using $\to$Elim. This would of course also require a derivation δ_1 of φ. Your derivation δ_1 of $\varphi \to \psi$, more likely than not, ends in $\to$Intro, which turns a derivation δ_2 of ψ from the open assumption φ into one of $\varphi \to \psi$. If it does, your final derivation has the form

$$\dfrac{\dfrac{\begin{array}{c}[\varphi]^x\\ \vdots\,\delta_2\\ \psi\end{array}}{\varphi \to \psi}\,x \to\text{Intro} \qquad \begin{array}{c}\vdots\,\delta_2\\ \varphi\end{array}}{\psi} \to\text{Elim}$$

Although your strategy is efficient in that you re-used something you already had (a derivation of $\varphi \to \psi$), your derivation isn't. It seems like a detour to introduce $\varphi \to \psi$ only to immediately eliminate it. There should be a more direct way to derive ψ.

In fact, there always is. "Detours" like the above—an introduction rule immediately followed by an elimination rule—can always be removed from natural deduction derivations. This is called *normalization* and the result is a *normal* derivation (or a derivation *in normal form*).

The proof of this result is, like the cut-elimination theorem in the sequent calculus, somewhat involved and requires us to consider lots of cases. It is harder to prove for classical **N1c** than it is for intuitionistic **N1i**. Just like cut-elimination shows that you can't derive more using the Cut rule in the sequent calculus, normalization shows that you can't prove more things using detours than you can while avoiding them. There are other similarities: normal derivation can be much longer than the shortest derivation with detours of the same result, just like derivations in the sequent calculus can be much shorter with cuts than the shortest cut-free derivation. The sub-formula property,

that you can always derive a result using only sub-formulas of the result in your derivation, follows from normalization for natural deduction just as it follows from cut-elimination for the sequent calculus.

4.2 Segments and Cuts

pt:nor:seg:sec

An introduction rule followed by an elimination rule is a detour in a derivation, and to normalize a derivation is to eliminate such detours. However, the $\vee$Elim and $\exists$Elim make even the definition of "detour we want to eliminate" a bit complicated. The problem is this. Because of these rules, it's not guaranteed that in a detour the elimination rule *immediately* follows the introduction rule. There could be an $\vee$Elim or $\exists$Elim in between, e.g.,

$$\dfrac{\dfrac{\dfrac{\exists x\,\chi(x)}{\varphi \to \psi}\,y \qquad \dfrac{\overset{[\varphi]^x}{\vdots}\,\psi}{\varphi \to \psi}\,x\ \to\text{Intro}}{\psi}\ \exists\text{Elim} \qquad \vdots\,\varphi}{\psi}\ \to\text{Elim}$$

In fact, any number of $\vee$Elim or $\exists$Elim can separate the $\to$Intro from the $\to$Elim, e.g.,

$$\dfrac{\dfrac{\dfrac{\exists x\,\chi(x)}{\varphi \to \psi}\,y \qquad \dfrac{\dfrac{\theta \vee \alpha \qquad \dfrac{\overset{[\varphi]^x}{\vdots}\,\psi}{\varphi \to \psi}\,x\ \to\text{Intro} \qquad \dfrac{\overset{[\varphi]^x}{\vdots}\,\psi}{\varphi \to \psi}\,x\ \to\text{Intro}}{\varphi \to \psi}\ \vee\text{Elim}}{\varphi \to \psi}\ \exists\text{Elim} \qquad \vdots\,\varphi}{\psi}\ \to\text{Elim}$$

This example shows that the same $\to$Elim might be involved in multiple detours, even, as it follows more than one $\to$Intro. To account for all these possibilities, we define detours using certain kinds of sequences of formula occurrences in an **N1i**-proof called *segments*. In our example, it's a sequence of occurrences of $\varphi \to \psi$, where a minor premise of $\vee$Elim or $\exists$Elim follows its conclusion. In the example, there are two such segments, each containing three occurrences of $\varphi \to \psi$. Here is the official definition.

Definition 4.1. A *segment* in a **N1i** derivation δ is a sequence of occurrences $\varphi_1, \ldots, \varphi_n$ of the same formula φ in δ where

1. φ_1 is not the conclusion of $\vee$Elim or $\exists$Elim;

2. φ_n is not a minor premise of $\vee$Elim or $\exists$Elim;

3. For $i = 1, \ldots, n-1$, φ_i is a minor premise of $\vee$Elim or $\exists$Elim and A_{i+1} is the conclusion of that inference.

The *length* of that segment is n.

Not every segment is a target for elimination (a detour). The ones that are are called *cuts*.

Definition 4.2. A *cut* is a segment in which φ_n is the major premise of an elimination inference, and either $n = 1$ and $\varphi_1 = \varphi_n$ is the conclusion of an introduction inference or of $\bot_I$, or $n > 1$.

Note that a cut segment need not begin with the conclusion of an introduction rule. All that's required is that it ends with the major premise of an elimination rule. However, a segment of length 1 does have to be the conclusions of introduction rules to count as a cut.

Definition 4.3. The *rank* of a cut is $\mathrm{dp}(\varphi)$. The *cut rank* $\mathrm{cr}(\delta)$ of a derivation δ is the maximal rank of cuts in δ. A cut is *maximal* in δ if its rank is $\mathrm{cr}(\delta)$, i.e., it has maximal rank. The *cut length* of δ is the sum of the lengths of all its maximal cuts.

4.3 Permutation Conversions

pt:nor:per: sec

Cuts in **N1i** are segments that end in the major premise of an elimination rule. One case in the normalization proof is to reduce the length of maximal cuts, and with it the cut length of the derivation. A cut of length > 1 can end in many ways depending on whether the last formula occurrence of φ in the cut is the conclusion of ($\vee$Elim or $\exists$Elim) and depending on the elimination inference it is the major premise of.

For instance, if the inferences involved are $\vee$Elim and $\wedge$Elim, $\varphi \equiv \psi \vee \chi$ and there is a sub-derivation δ_1 that ends in:

$$\dfrac{\begin{array}{ccc} & [\theta]^x & [\alpha]^y \\ \vdots\,\delta_2 & \vdots\,\delta_3 & \vdots\,\delta_4 \\ \theta \vee \alpha & \psi \wedge \chi & \psi \wedge \chi \end{array}}{\dfrac{\psi \wedge \chi}{\psi}\,\wedge\mathrm{Elim}_1}\,{}^{x\,y}\vee\mathrm{Elim}$$

The last formula occurrence φ_n in the cut is the conclusion of $\vee$Elim, the lower $\psi \wedge \chi$. The second-to-last formula occurrence φ_{n-1} is one of the minor premises of $\vee$Elim.

To reduce the length of this cut, we can switch the order ("permute") of the $\vee$Elim and $\wedge$Elim inferences. That is, we replace the subproof δ_1 in δ by the following:

$$
\dfrac{\begin{array}{c}\vdots\,\delta_2\\ \theta \vee \alpha\end{array} \qquad \dfrac{\begin{array}{c}[\theta]^x\\ \vdots\,\delta_3\\ \psi \wedge \chi\end{array}}{\psi}\,\wedge\mathrm{Elim}_1 \qquad \dfrac{\begin{array}{c}[\alpha]^y\\ \vdots\,\delta_4\\ \psi \wedge \chi\end{array}}{\psi}\,\wedge\mathrm{Elim}_1}{\psi}\,\vee\mathrm{Elim}\;{\scriptstyle x\,y}
$$

The cuts that ended in the premise of $\wedge$Elim below $\vee$Elim now end in the premise of one of the $\wedge$Elim inferences above the minor premises of $\vee$Elim, i.e., the length of each has decreased by 1.

Moreover, any cut that lies entirely within δ_2, δ_3, or δ_4, or entirely outside δ_1 in δ is unchanged: it still involves the same formulas, goes through the same inferences, so in particular has the same rank and length as the corresponding cut in δ. So the cut rank of the new derivation has not increased, and the cut length of the new derivation is less than that of the old derivation.

Problem 4.1. After a permutation conversion, the length of at least one cut is reduced by 1. If we permute $\wedge$Elim across an $\vee$Elim, there are actually two cut segments the length of which decrease, so in this case the cut length of the derivation decreases by at least 2. Can the cut length of a derivation decrease by more than 2 by a single such permutation? Can the cut *rank* decrease?

The process of moving an elimination rule above an $\vee$Elim or $\exists$Elim rule is called a *permutation conversion*. The patterns for some of the rule pairs is given in Table 4.1. (The remaining are left as exercises.)

$$
\dfrac{\dfrac{\theta \vee \alpha \qquad \begin{array}{c}[\theta]^x\\ \vdots\\ \psi \wedge \chi\end{array} \qquad \begin{array}{c}[\alpha]^y\\ \vdots\\ \psi \wedge \chi\end{array}}{\psi \wedge \chi}\,\vee\mathrm{Elim}}{\psi}\,\wedge\mathrm{Elim}_1 \;\longrightarrow
$$

$$
\dfrac{\theta \vee \alpha \qquad \dfrac{\begin{array}{c}[\theta]^x\\ \vdots\\ \psi \wedge \chi\end{array}}{\psi}\,\wedge\mathrm{Elim}_1 \qquad \dfrac{\begin{array}{c}[\alpha]^y\\ \vdots\\ \psi \wedge \chi\end{array}}{\psi}\,\wedge\mathrm{Elim}_1}{\psi}\,\vee\mathrm{Elim}\;{\scriptstyle x\,y}
$$

$$
\dfrac{\dfrac{\theta \vee \alpha \qquad \begin{array}{c}[\theta]^x\\ \vdots\\ \psi \to \chi\end{array} \qquad \begin{array}{c}[\alpha]^y\\ \vdots\\ \psi \to \chi\end{array}}{\psi \to \chi}\,\vee\mathrm{Elim} \qquad \psi}{\chi}\,\to\mathrm{Elim}\;\longrightarrow
$$

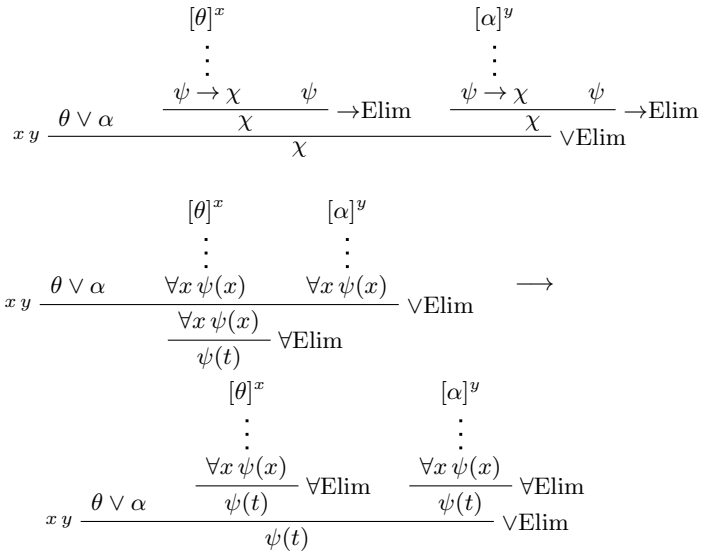

Table 4.1: Some permutation conversions for **N1i**.

pt:nor:per:
tab:permutation-conversions

Problem 4.2. Give the permutation conversions for ∨Elim followed by ∨Elim or ¬Elim.

Problem 4.3. Give the permutation conversions for ∃Elim followed by ∃Elim. Explain why no eigenvariable condition is violated in the latter.

We should convince ourselves that when applying a permutation conversion, we don't perhaps violate an eigenvariable condition. Consider

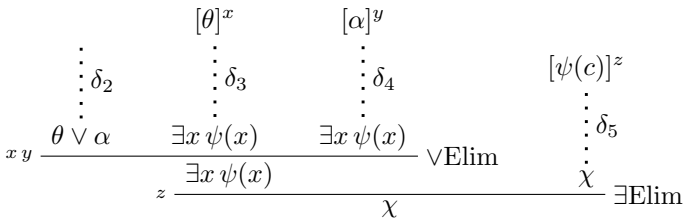

If we replace this by

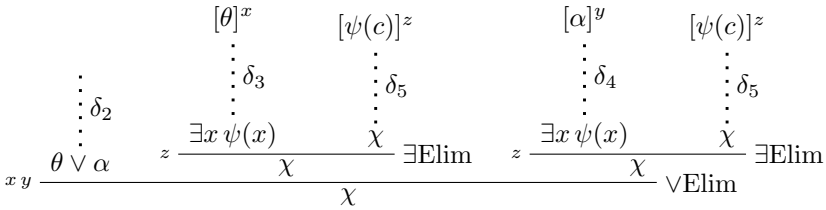

we might worry that the eigenvariable c might occur in, say, θ. After all, the assumption θ is discharged above the original ∃Elim and so in the original

derivation it does not occur in an assumption that's open above the major premise of the ∃Elim inference, and so there the eigenvariable condition is satisfied. But in the new derivation, θ is not discharged until after the ∃Elim inferences, and so the eigenvariable condition would be violated. Recall, however, that we assume our derivations are regular: every eigenvariable is the eigenvariable of a single inference and only occurs above the minor premise of ∃Elim in the derivation. In particular, c cannot occur in δ_3 or δ_4.

This example shows something else: Note that the new derivation has two copies of δ_5. This means that if δ_5 contains a maximal cut, we have *not* reduced the cut length!Thus we must take care that we apply a permutation conversion only when we know that there is no maximal cut in δ_5. We will ensure this by always selecting a *topmost* maximal cut, that is a maximal cut such that the sub-derivation (in this case, δ_1) contains no maximal cut that ends above the last inference in the sub-derivation. This rules out the possibility that there are other maximal cuts in δ_5 (or indeed in δ_2, δ_3, or δ_4).

4.4 Reduction Conversions

pt:nor:red:sec

If a cut in a derivation in **N1i** has length 1, it is a segment consisting of a single formula occurrence, and that formula is both the conclusion of an introduction inference and the major premise of an elimination inference. In the normalization proof, such cuts are "reduced" by replacing the sub-derivation ending in such an introduction/elimination pair by a new derivation in which the pair of inferences has been eliminated. Such a replacement may result in new cuts. However, we can convince ourselves that the new cuts have lower rank than the cut being reduced. Thus, the new derivation either has lower cut length (because a maximal cut of length 1 has been removed) or the cut rank of the derivation has been reduced (if the maximal cut in question was the only maximal cut).

For instance, if the inferences involved are →Intro and →Elim, $\varphi \equiv \psi \to \chi$ and there is a sub-derivation δ_1 that ends in:

$$\dfrac{\dfrac{\begin{array}{c}[\psi]^x\\ \vdots\,\delta_2\\ \chi\end{array}}{\psi \to \chi}\,{\scriptstyle x}\,{\to}\text{Intro} \qquad \begin{array}{c}\vdots\,\delta_3\\ \psi\end{array}}{\chi}\,{\to}\text{Elim}$$

To remove this cut we replace the subproof δ_1 in δ by the following:

$$\begin{array}{l}\vdots\\ \delta_3\\ \vdots\\ \psi\\ \vdots\\ \delta_2[\delta_3/\psi^x]\\ \vdots\\ \chi\end{array}$$

By $\delta_2[\delta_3/\psi^x]$ we mean the derivation that results from δ_2 by replacing every assumption ψ that's labelled x with the entire derivation δ_3. This derivation no longer has open assumptions of the form ψ with label x. It has no other assumptions labelled x, since we assume, as usual, that no two formulas occur as assumptions with the same label. Hence, the new derivation has the same open assumptions as δ_1, every inference in δ_2 that discharges assumptions discharges the same assumptions in $\delta_2[\delta_3/\psi^x]$, and every inference in δ_3 that discharges assumptions results in possibly multiple copies, which discharge the corresponding assumptions in $\delta_2[\delta_3/\psi^x]$.

Any cut that lies entirely within δ_2, δ_3, or δ_4, or entirely outside δ_1 in δ is unchanged: it still involves the same formulas, goes through the same inferences, so in particular has the same rank and length as the corresponding cut in δ. We have removed a maximal cut of length 1 and so either have reduced the cut length, or reduced the cut rank (if the cut length of δ was 1 and this was the only maximal cut).

However, two things may happen:

1. By replacing all assumptions ψ^x by δ_3, we have multiplied all the cuts in δ_3. If some of those were maximal cuts, the cut length of the new derivation may be greater than the cut length of the old derivation. We have to guarantee that this does not happen. We do so by applying reduction conversions only to *rightmost* maximal cuts, that is, to cuts which are such that no maximal cut lies on a branch through δ to the right of the cut being reduced. If there were a cut in δ_3 it would be to the right of the cut being reduced here, and so we wouldn't be working on a rightmost cut. By always selecting rightmost cuts, we guarantee that either we decrease the cut length of the derivation, or even reduce the cut rank.

2. If an assumption ψ^x in δ_2 is the major premise of an elimination inference, and the last inference of δ_3 is $\vee$Elim, $\exists$Elim, or an introduction inference, we have generated a new cut by grafting δ_3 onto δ_2 at such an assumption. What was just an assumption in δ_2 is now a cut of length 1 (the conclusion of an introduction inference and major premise of an elimination inference) or the last formula occurrence of a cut segment (major premise of an elimination inference and conclusion of $\vee$Elim or $\exists$Elim). If B^x was a minor premise of $\vee$Elim or $\exists$Elim, it already was the first

formula occurrence of a segment, and possibly of a cut segment. In the new derivation that segment may now be longer. However, observe that the formula making up this new cut or existing cut segment is ψ, and $\mathrm{dp}(\psi) < \mathrm{dp}(\psi \to \chi)$. Hence, although we have potentially added cuts or increased the length of cuts, none of these are maximal cuts. Thus we have either decreased the cut rank, or, if there are cuts of the same rank remaining, have decreased the cut length.

The process of reducing cut of length 1 is called a *reduction conversion* (or *detour conversion*). The patterns for some of the rule pairs is given in Table 4.2. (The remaining are left as exercises.)

Problem 4.4. Give the reduction conversion for ¬Intro followed by ¬Elim.

There is one case still to be considered. The definition of cut includes the case where the length of the cut is 1 and φ is the conclusion of $\bot_I$ as well as the major premise of an elimination inference. This case is dealt with in a similar way as when φ is the conclusion of an introduction rule. For instance, replace

$$\dfrac{\dfrac{\begin{array}{c}\vdots\,\delta_2\\ \bot\end{array}}{\psi \to \chi} \quad \begin{array}{c}\vdots\,\delta_3\\ \psi\end{array}}{\chi} \to\text{Elim} \quad \text{by} \quad \begin{array}{c}\vdots\,\delta_2\\ \bot\\ \chi\end{array}$$

We must be a little careful in the cases where $\varphi \equiv (\psi \vee \chi)$ or $\varphi \equiv \exists x\, \psi(x)$. E.g., if we replaced

$$\dfrac{\dfrac{\begin{array}{c}\vdots\,\delta_2\\ \bot\end{array}}{\psi \vee \chi}\,\bot_I \quad \begin{array}{c}[\psi]^x\\ \vdots\,\delta_3\\ \theta\end{array} \quad \begin{array}{c}[\chi]^x\\ \vdots\,\delta_4\\ \theta\end{array}}{\theta} \vee\text{Elim} \quad \text{by} \quad \dfrac{\begin{array}{c}\vdots\,\delta_2\\ \bot\end{array}}{\theta}\,\bot_I$$

we might be introducing a new cut with cut formula θ, and there is no guarantee that $\mathrm{dp}(\theta) < \mathrm{dp}(\psi \vee \chi)$! So we must proceed the same way here as we did for an ∨Intro/∨Elim reduction and replace with

$$\begin{array}{c}\vdots\,\delta_2\\ \dfrac{\bot}{\psi}\,\bot_I\\ \vdots\,\delta_3\\ \theta\end{array} \quad \text{or} \quad \begin{array}{c}\vdots\,\delta_2\\ \dfrac{\bot}{\chi}\,\bot_I\\ \vdots\,\delta_4\\ \theta\end{array}$$

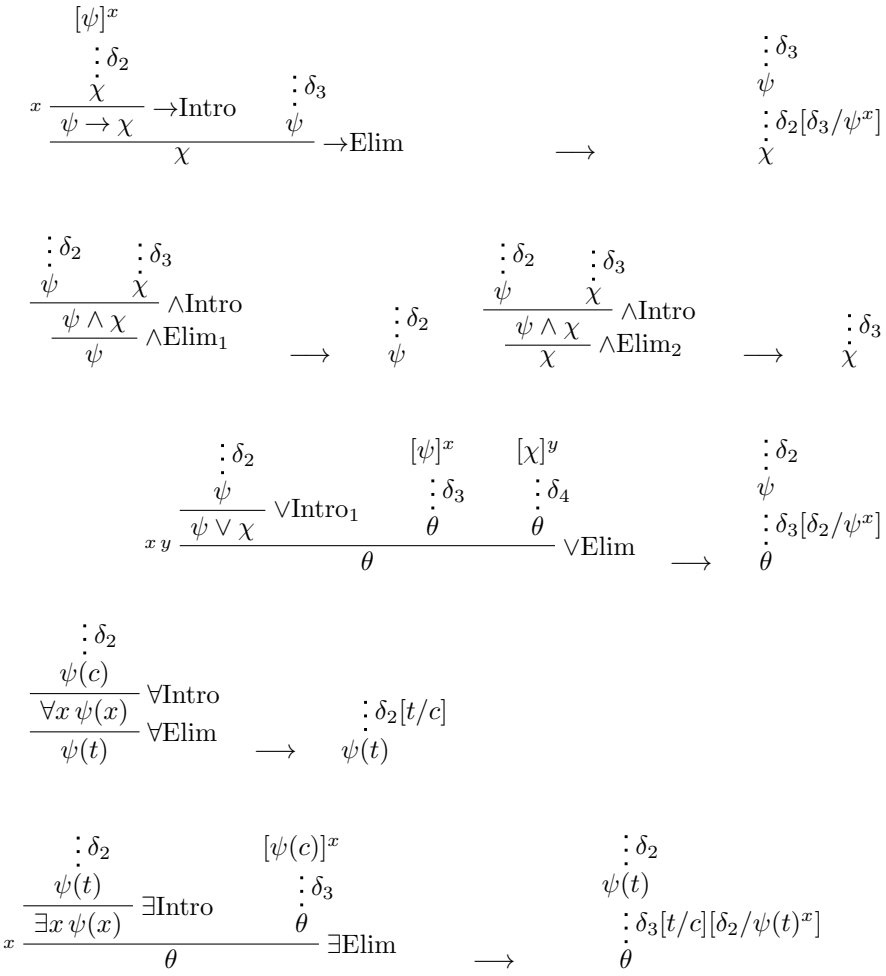

Table 4.2: Some reduction conversions for **N1i**.

pt:nor:red:
tab:reduction-conversions

4.5 The Normalization Theorem

A derivation is in *normal* form if it has no cut segments. Clearly this is the case iff we can apply neither a permutation conversion nor a reduction conversion to it. To *normalize* a derivation means to transform it into a normal proof by applying permutation and reduction conversions until no cuts remain. That this is always possible is the *normalization theorem*. pt:nor:thm:sec

Theorem 4.4. *Any* **N1i**-*derivation* δ *normalizes, i.e., can be transformed into a derivation* δ' *without cuts.* pt:nor:thm:thm:normalization

Proof. By induction on the cut rank and length of δ. The induction hypothesis is that any derivation which either has lower cut rank, or the same cut rank but lower cut length, normalizes.

If the cut length is 0, there are no cuts, δ is already normal. So assume that cut rank and cut length of δ are > 0. Pick a topmost, rightmost maximal cut segment. If the cut segment has length > 1, apply a permutation conversion. If the cut segment has length 1, apply the corresponding reduction conversion. This results in a new derivation which either has the same cut rank but a lower cut length, or a lower cut rank, and the inductive hypothesis applies. □

The strategy used in the above proof (always deal with a topmost rightmost maximal cut) suggests a specific order in which conversions must be applied to a derivation in order to result in a normal derivation. Surprisingly enough, the order actually does not matter. Applying permutation and reduction conversions in any order eventually results in a normal proof.

The definition of segment and cut can easily be formulated for **N2i**, and the permutation and reduction conversions for **N1i** directly translate to **N2i**. Consequently, the normalization theorem also holds for **N2i**.

Theorem 4.5. *Any* **N2i**-*derivation* δ *normalizes, i.e., can be transformed into a derivation* δ' *without cuts.* pt:nor:thm:thm:normalization-N2i

4.6 Translating Between Normal N2i-Derivations and G2i

Derivations in **N2i** can be translated to derivations in **G2i** and vice versa. First of all, the result of translating cut-free **2i** derivations results in normal **N2i**-derivations. pt:nor:tr:sec

To state the result we need a bit of terminology to let us relate antecendents of sequents in **N2i** and in **G2i** more easily: a set Γ' of labelled formulas *corresponds to* the multiset Γ of unlabelled formulas if, for every formula φ, the number of elements of Γ' of the form $x : \varphi$ is less than or equal to the multiplicity of φ in Γ.

Corollary 4.6. *If* $\mathbf{G2i} \vdash \Gamma \Rightarrow \Delta$ *then there is a normal* **N2c**-*derivation* δ *of* $\Gamma' \Rightarrow \chi$*, where the set* Γ' *of labelled formulas corresponds to the multiset* Γ*,*

i.e., for every formula φ, the number of elements of Γ' of the form $x : \varphi$ is less than or equal to the multiplicity of φ (i.e., the number of copies of φ) in Γ.

Proof. By inspection of the proof of Proposition 3.12: we add introduction rules to the end of derivations, and elimination rules only at the top of derivations, so never introduce an introduction rule above an elimination rule. □

The translation of the Cut rule given in Proposition 3.12, however, can introduce cuts in the **N2i**-derivation δ. If a **G2i** derivation ends in Cut with immediate sub-derivations π_1 and π_2 of $\Gamma_1 \Rightarrow \varphi$ and $\varphi, \Gamma_2 \Rightarrow \Delta_2$, respectively, we graft the translation δ_1 of π_1 onto the axioms $x : \varphi \Rightarrow \varphi$ occurring in the translation δ_2 of π_2. If δ_1 ends in an introduction rule with principal formula φ and the axiom $x : \varphi \Rightarrow \varphi$ in δ_2 is the major premise of an elimination rule, this would result in a cut in δ.

It is also possible to translate derivations in **N2i** to derivations in **G2i**. In the proof of Proposition 3.13, the translation of ∧Elim, →Elim, ¬Elim, ∀Elim results in Cut inferences in the resulting **G2i** derivation. If the **N2i**-derivation we start from, this can, however, be avoided.

Let's call a list $S_1, \ldots, S_n$ of occurrences of sequents in a **N2i**-derivation δ a *branch* if S_1 is an axiom, and whenever S_i is a major premise of an inference, S_{i+1} is its conclusion. In other words, a branch traces sequents downwards in δ from axioms, but stops at minor premises of elimination rules. A *main branch* of δ is a branch where S_n is the end-sequent. Every derivation δ must have at least one main branch, since every rule has at least one major premise (only ∧Intro has two).

pt:nor:tr: prop:normal-N2i-to-G2i

Proposition 4.7. *If δ is a normal* **N2c** *or* **N2i**-*derivation of $\Gamma \Rightarrow \varphi$ then there is a (cut-free)* **G2c**-*derivation (***G2i**-*derivation) π of $\Gamma' \Rightarrow \Delta$, where Γ' is the multiset of formulas resulting from Γ by removing labels, and $\Delta = \{\varphi\}$ if φ is not $\bot$, and $\Delta = \emptyset$ if it is.*

Proof. This time, induction is on the number of inferences in the derivation δ of $\Gamma \Rightarrow \varphi$. If there is no inference in δ, it is an axiom $x : \varphi \Rightarrow \varphi$, and π is the corresponding axiom $\varphi \Rightarrow \varphi$ or $\bot \Rightarrow$ if $\varphi \equiv \bot$.

If δ ends in an inference rule, we distinguish cases:

1. If R is an introduction rule or $\bot_I$, the translations proceeds exactly as in the proof of Proposition 3.13, and does not require Cut.

   If the last rule R of δ is an elimination rule, observe the following:

   a) A main branch of δ passes through no introduction rules, as otherwise there would be an introduction rule or $\bot_I$ followed by an elimination rule along the main branch, and δ would not be normal.

   b) Since there are no introduction rules along main branches of δ, there is only one main branch. (Only ∧Intro has two major premises, so multiple main branches would have to go through the premises of a ∧Intro rule.)

c) If the main branch includes the major premise of ∨Elim or ∃Elim, there is only one such rule and it is the last one, i.e., R. (Otherwise the main branch would contain an ∨Elim or ∃Elim rule where the conclusion of the rule is major premise of an elimination rule, and the proof would not be normal as we could apply a permutation conversion.)

d) No elimination rules other than ∨Elim and ∃Elim discharge assumptions, i.e., change the left context of sequents. ∨Elim and ∃Elim only discharge assumptions above a minor premise, i.e., change the left context only of their minor premises. In other words, if a sequent S_i on the main branch of δ has context Γ_1, the conclusion S_{i+1} of the inference of which S_i is the major premise has context $\Gamma_1' \supseteq \Gamma_1$.

e) Consequently, if $x : \chi \Rightarrow \chi$ is the topmost sequent S_1 in the main branch, $x : \chi \in \Gamma$.

We now distiguish cases according to the form of χ. Since every S_i in the main branch is major premise of an elimination rule, this also applies to $x : \chi \Rightarrow \chi$.

2. $\chi \equiv \theta \wedge \alpha$. Then δ has the form

$$\dfrac{\dfrac{x : \theta \wedge \alpha \;\Rightarrow\; \theta \wedge \alpha}{x : \theta \wedge \alpha \;\Rightarrow\; \theta}\,\wedge\text{Elim}_1}{\begin{array}{c}\vdots\\ x : \theta \wedge \alpha, \Gamma_1 \;\Rightarrow\; \varphi\end{array}}$$

The main branch includes $x : \theta \wedge \alpha \Rightarrow \theta$ contains no major premise of →Intro, ¬Intro nor a minor premise of ∨Elim or ∃Elim, the context formulas in $x : \theta \wedge \alpha$ remain unchanged along the main branch. This explains why the context of the end-sequent must contain $x : \theta \wedge \alpha$. Moreover, we can turn the derivation δ into a derivation δ_1 by replacing the sub-derivation ending in ∧Elim by $x : \theta \Rightarrow \theta$ and replacing $x : \theta \wedge \alpha$ in the context of every sequent along the main branch by $x : \theta$, i.e., turn

$$\dfrac{\dfrac{x : \theta \wedge \alpha \;\Rightarrow\; \theta \wedge \alpha}{x : \theta \wedge \alpha \;\Rightarrow\; \theta}\,\wedge\text{Elim}_1}{\begin{array}{c}\vdots\\ x : \theta \wedge \alpha, \Gamma_1 \;\Rightarrow\; \varphi\end{array}} \qquad \text{into} \qquad \dfrac{\begin{array}{c}x : \theta \;\Rightarrow\; \theta\\ \vdots\\ x : \theta, \Gamma_1 \;\Rightarrow\; \varphi\end{array}}{}$$

The induction hypothesis applies to δ_1, since the number of inferences in δ is the number of inferences in δ' plus 1.

By induction hypothesis, there is a **G2i**-derivation π_1 of $x : \theta, \Gamma_1 \Rightarrow \varphi$. We obtain the derivation π by applying $\wedge$L:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\dfrac{\theta, \Gamma_1' \Rightarrow \varphi}{\theta \wedge \alpha, \Gamma_1' \Rightarrow \varphi}\,\wedge\mathrm{L}}$$

If $\varphi \equiv \bot$ we insert a WR, e.g,

$$\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\dfrac{\theta, \Gamma_1' \Rightarrow}{\theta \wedge \alpha, \Gamma_1' \Rightarrow}\,\wedge\mathrm{L}}}{\theta \wedge \alpha, \Gamma_1' \Rightarrow \varphi}\,\mathrm{WR}$$

3. If $\chi \equiv \theta \vee \alpha$, it must be the major premise of $\vee$Elim, and this inference must be the last inference R on the main branch. In other words, δ has the form:

$$\dfrac{x : \theta \vee \alpha \Rightarrow \theta \vee \alpha \qquad \begin{array}{c}\vdots\\ \delta_1\\ \vdots\end{array} \quad y : \theta, \Gamma_1 \Rightarrow \varphi \qquad \begin{array}{c}\vdots\\ \delta_2\\ \vdots\end{array} \quad z : \alpha, \Gamma_2 \Rightarrow \varphi}{x : \theta \vee \alpha, \Gamma_1, \Gamma_2 \Rightarrow \varphi}\,\vee\mathrm{Elim}$$

The induction hypothesis applies to the derivations δ_1 and δ_2; we obtain **G2i**-derivations π_1 and π_2 of $\theta, \Gamma_1' \Rightarrow \varphi$ and $\alpha, \Gamma_2' \Rightarrow \varphi$. We apply $\vee$L to get π:

$$\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array} \quad \theta, \Gamma_1' \Rightarrow \varphi \qquad \begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array} \quad \alpha, \Gamma_2' \Rightarrow \varphi}{\theta \vee \alpha, \Gamma_1', \Gamma_2' \Rightarrow \varphi}\,\vee\mathrm{L}$$

Should $\vee$Elim not discharge $y : \theta$ or $z : \alpha$ (i.e., these aren't present in the context of the minor premises), or should φ happen to be $\bot$, we have to add some WL and WR. For instance,

$$\dfrac{\dfrac{\dfrac{\begin{array}{c}\vdots\\ \pi_1\\ \vdots\end{array}}{\dfrac{\Gamma_1' \Rightarrow}{\theta, \Gamma_1' \Rightarrow}\,\mathrm{WL}} \qquad \dfrac{\begin{array}{c}\vdots\\ \pi_2\\ \vdots\end{array}}{\dfrac{\Gamma_2' \Rightarrow}{\alpha, \Gamma_2' \Rightarrow}\,\mathrm{WL}}}{\theta \vee \alpha, \Gamma_1', \Gamma_2' \Rightarrow}\,\vee\mathrm{L}}{\theta \vee \alpha, \Gamma_1', \Gamma_2' \Rightarrow \varphi}\,\mathrm{WR}$$

4. $\chi \equiv \theta \to \alpha$. Then δ has the form

$$\dfrac{\dfrac{\begin{array}{c}\vdots\,\delta_1\end{array}}{x:\theta\to\alpha \;\Rightarrow\; \theta\to\alpha \qquad \Gamma_1 \;\Rightarrow\; \theta}\;{\to}\text{Elim}}{\begin{array}{c}x:\theta\to\alpha,\Gamma_1 \;\Rightarrow\; \alpha\\ \vdots\,\delta_2\\ x:\theta\to\alpha,\Gamma_1,\Gamma_2 \;\Rightarrow\; \varphi\end{array}}$$

Here note that since the main branch, which includes $x:\theta\to\alpha,\Gamma_1\Rightarrow\alpha$, contains no major premise of $\to$Intro, $\neg$Intro nor a minor premise of $\vee$Elim or $\exists$Elim, the context formulas in $x:\theta\to\alpha,\Gamma_1$ remain unchanged along the main branch. This explains why the context of the end-sequent must contain $x:\theta{\to}\alpha,\Gamma_1$. Moreover, we can turn the derivation fragment δ_2 into a correct derivation δ_2' by replacing the sub-derivation ending in $\to$Elim by $x:\alpha\Rightarrow\alpha$ and replacing $x:\theta\to\alpha,\Gamma_1$ in the context of every sequent along the main branch by $x:\alpha$, i.e., turn

$$\begin{array}{c}x:\theta\to\alpha,\Gamma_1 \;\Rightarrow\; \alpha\\ \vdots\,\delta_2\\ x:\theta\to\alpha,\Gamma_1,\Gamma_2 \;\Rightarrow\; \varphi\end{array} \quad\text{into}\quad \begin{array}{c}x:\alpha \;\Rightarrow\; \alpha\\ \vdots\,\delta_2'\\ x:\alpha,\Gamma_2 \;\Rightarrow\; \varphi\end{array}$$

The induction hypothesis applies to δ_1 and δ_2', since the number of inferences in δ is the sum of the numbers of inferences in δ_1 and δ_2, plus 1 to account for the $\to$Elim inference at the start of the main branch. (If that inference is also the last inference R in δ then δ_1 contains 0 inferences.)

By induction hypothesis, there are **G2i**-derivations π_1 of $\Gamma_1\Rightarrow\theta$ and π_2 of $\alpha,\Gamma_1'\Rightarrow\varphi$. We obtain the derivation π by applying $\to$L:

$$\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma_1' \;\Rightarrow\; \theta\end{array} \qquad \begin{array}{c}\vdots\,\pi_2\\ \alpha,\Gamma_2' \;\Rightarrow\; \varphi\end{array}}{\theta\to\alpha,\Gamma_1',\Gamma_2' \;\Rightarrow\; \varphi}\;{\to}\text{L}$$

If θ and/or $\varphi\equiv\bot$ we insert one or two WR, e.g,

$$\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi_1\\ \Gamma_1' \;\Rightarrow\end{array}}{\Gamma_1' \;\Rightarrow\; \theta}\;\text{WR} \qquad \dfrac{\begin{array}{c}\vdots\,\pi_2\\ \alpha,\Gamma_2' \;\Rightarrow\end{array}}{\alpha,\Gamma_2' \;\Rightarrow\; \varphi}\;\text{WR}}{\theta\to\alpha,\Gamma_1',\Gamma_2' \;\Rightarrow\; \varphi}\;{\to}\text{L}$$

The remaining cases are treated similarly and are left as exercises. □

Corollary 4.8. *Every formula occuring in a normal* **N1i** *or* **N2i**-*derivation is a sub-formula of the end-formula or end-sequent.*

Proof. By Proposition 4.7, every normal **N2i**-derivation can be translated into a **G2i**-derivation. An inspection of the proof shows that every formula occurring in the **N2i**-derivation also occurs in the **G2i**-derivation. **G2i**, lacking the Cut rule, has the sub-formula property. □

Chapter 5

Propositions as Types

This is a *very experimental* draft of a chapter on the Curry–Howard correspondence. It needs more explanation and motivation, and there are probably errors and omissions. The proof of normalization should be reviewed and expanded. There are no examples for the product type. Permutation and simplification conversions are not covered. It will make a lot more sense once there is also material on the (typed) lambda calculus which is basically presupposed here. Use with extreme caution.

5.1 Introduction

int:pty:int:sec

The lambda calculus is a simple model of computation. It operates on terms built from variables. If N and M are terms, then (NM) is also a term ("N applied to M"). If N is a term, then $\lambda x.\, N$ is a term. If N contains the variable x, then the corresponding occurrences of x in $\lambda x.\, N$ are bound by the λx operator, and so are no longer free. A term of the form $(\lambda x.\, NM)$ is said to β-contract to $N[M/x]$ (the result of replacing every free occurrence of x in N by M). A term N β-converts to N' if N' results from N by β-contracting a subterm, and we write $N \to N'$. A computation in the λ-calculus is a sequence $N \to N_1 \to N_2 \to \ldots$. If this sequence ends in a term N_k where no subterm can be β-contracted, it is called *normal* or a normal form of N. Think of computation in the λ-calculus as such sequences. The computation halts if it results in a normal form, and the normal form is the result of the computation. This simple model of computations turns out to be Turing complete: every partial recursive function can be computated by a term in the λ-calculus.

Haskell Curry and William Howard independently discovered a close similarity between natural deduction and the λ-calculus. Intuitively, a term $\lambda x.\, N$ is a function: x is the argument; N tells us how to compute the value given x. A term $(\lambda x.\, NM)$ then is applying the function $\lambda x.\, N$ to an argument M, and β-contraction tells us how to evaluate it: replace x in N by M (and keep evaluating the resulting term). Now think of these as functions with a fixed domain

and range. If the domain of $\lambda x.\, N$ is A and and its range B, then we should think of the variable x as only taking arguments from A and N as producing values from B. Consequently, $(\lambda x.\, NM)$ only should make sense if $\lambda x.\, N$ is a function $A \to B$, and $M \in A$. If we want to add this information to the λ-calculus we obtain the *typed* λ-calculus. In the typed λ-calculus not every expression $(\lambda x.\, NM)$ is legal, but only those where the "types" of $\lambda x.\, N$ and M fit together, i.e., $\lambda x.\, N$ has type $A \to B$ and M has type A. The type of $\lambda x.\, N$ is $A \to B$ only if the type of x is A and the type of N is B. In general, not every expression $(M_1 M_2)$ is legal. Only those terms $(M_1 M_2)$ are legal where M_1 has type $A \to B$ and M_2 has type A; if so, then $(M_1 M_2)$ has type B.

These typing rules now clearly correspond to derivations in natural deduction, where the types are represented by formulas, and the derivations themselves to λ-terms. For instance, a derivation of $\varphi \to \psi$ corresponds to a term $\lambda x^{\varphi}.\, N$, which has the function type $\varphi \to \psi$. The inference rules of natural deduction correspond to typing rules in the typed lambda calculus, e.g.,

$$\dfrac{\begin{array}{c}[\varphi]^x \\ \vdots \\ \psi\end{array}}{\varphi \to \psi}\, {\scriptstyle x}\ {\to}\text{Intro} \qquad \text{corresponds to} \qquad \dfrac{x : \varphi \;\Rightarrow\; N : \psi}{\Rightarrow\; \lambda x^{\varphi}.\, N : \varphi \to \psi}$$

where the rule on the right means that if x is of type φ and N is of type ψ, then $\lambda x^{\varphi}.\, N$ is of type $\varphi \to \psi$.

The $\to$Elim rule corresponds to the typing rule for composition terms, i.e.,

$$\dfrac{\varphi \to \psi \qquad \varphi}{\psi}\ {\to}\text{Elim} \quad \text{corresponds to} \quad \dfrac{\Rightarrow\; M_1 : \varphi \to \psi \qquad \Rightarrow\; M_2 : \varphi}{\Rightarrow\; (M_1 M_2) : \psi}$$

If a $\to$Intro rule is followed immediately by a $\to$Elim rule, the derivation can be simplified:

$$\dfrac{\dfrac{\begin{array}{c}[\varphi]^x \\ \vdots \\ \psi\end{array}}{\varphi \to \psi}\, {\scriptstyle x}\ {\to}\text{Intro} \qquad \begin{array}{c}\vdots \\ \varphi\end{array}}{\psi}\ {\to}\text{Elim} \qquad \to \qquad \begin{array}{c}\vdots \\ \varphi \\ \vdots \\ \psi\end{array}$$

which corresponds to β-contraction of lambda terms

$$(\lambda x^{\varphi}.\, NM) \quad \to \quad N[M/x].$$

Similar correspondences hold between the rules for $\wedge$ and "product" types, and between the rules for $\vee$ and "sum" types.

This correspondence between terms in the simply typed lambda calculus and natural deduction derivations is called the "Curry–Howard", "proofs-as-programs," or "propositions as types" correspondence. In addition to formulas

(propositions) corresponding to types, and proofs to terms, we can summarize the correspondences as follows:

logic	program
proposition/formula	type
proof	term
assumption	variable
discharged assumption	bound variable
undischarged assumption	free variable
$\varphi \to \psi$	function type
$\varphi \wedge \psi$	product type
$\varphi \vee \psi$	sum type
$\bot$	empty type

The Curry–Howard correspondence is one of the cornerstones of automated proof assistants and type checkers for programs, since checking a proof witnessing a proposition (as we did above) amounts to checking if a program (term) has the declared type.

5.2 Proof Terms

int:pty:ter:sec

We label assumptions in **N1i** with variables (e.g., φ^x) as discharge labels. In **N2i**, every sequent $\Gamma \Rightarrow \varphi$ has a list of variable–formula pairs in the context Γ on the left. Thus, in **N2i**, every sequent in a derivation carries not only the information what the derivation so far has derived (i.e., φ), but also which assumptions are open at each point (i.e., Γ). What if we also included in each sequent the information *how* φ was derived. To this end, we can define a *term-labelled* system **tN2i**, in which sequents have the form $\Gamma \Rightarrow N : \varphi$. Γ and φ are as before: a set of pairs of the form $x : \psi$ and the formula the sub-derivation ending in the sequent has shown to follow from the assumptions in Γ. N will be a certain term that encodes the derivation ending in the sequent.

The terms we assign to each sequent are built from variables x, y, ..., using function symbols which encode the inferences carried out so far. We need a separate function symbol for each inference rule. The function symbol corresponding to introduction rules are called "constructors" and those corresponding to elimination rules "destructors." Each function symbol takes as many arguments as the rule has premises. For instance, the function symbols for $\wedge$Intro and $\wedge\text{Elim}_1$ might be $c^\wedge$ (2-place) and $d_1^\wedge$ (1-place). In **tN2i** these rules then become

$$\frac{\Gamma_1 \Rightarrow N : \varphi \qquad \Gamma_2 \Rightarrow M : \psi}{\Gamma_1, \Gamma_2 \Rightarrow c^\wedge(N, m) : \varphi \wedge \psi}\,\wedge\text{Intro} \qquad \frac{\Gamma \Rightarrow N : \varphi \wedge \psi}{\Gamma \Rightarrow d_1^\wedge(N) : \varphi}\,\wedge\text{Elim}_1$$

If an inference rule discharges an assumption, i.e., the rule has a premise that indicates a labelled formula $x : \varphi$ in a premise that is not in the context of the conclusion, we have to indicate this in the proof term as well. The function symbols associated with such rules *bind* the corresponding variables.

For instance, the $\to$Intro rule goes from the premise $x : \varphi, \Gamma \Rightarrow \psi$ to $\Gamma \Rightarrow \varphi \to \psi$. The constructor $c^{\to}$ takes one argument, but also must indicate that the variable x becomes bound. This is the proof term's way of expressing that the assumption with label x has been discharged. For instance, if the proof term for the premise is N, we could write the proof term for the conclusion as $c^{\to}([x]N)$ and the rule as

$$\frac{x : \varphi, \Gamma \Rightarrow N : \psi}{\Gamma \Rightarrow c^{\to}([x]N) : \varphi \to \psi} \to\text{Intro}$$

In order to be able to fully reconstruct a derivation, we need some additional information. If the premise formulas of a rule don't fully determine what the conclusion formula is, we have to add that information to the term. This is the case in $\to$Intro, so we will write $c^{\to}([x : \varphi]N)$. It is also the case in $\vee$Intro and $\bot_I$, so we use function symbols that carry that information, e.g.,

$$\frac{\Gamma \Rightarrow N : \varphi}{\Gamma \Rightarrow c_1^{\vee\psi} : \varphi \vee \psi} \vee\text{Intro}_1 \qquad \frac{\Gamma \Rightarrow N : \bot}{\Gamma \Rightarrow d_{\varphi}^{\bot}(N) : \varphi} \bot_I$$

With these ideas it is possible to construct a sequent-style natural deduction system in which every sequent is of the form $\Gamma \Rightarrow N : \varphi$, where N is a *proof term*.

It turns out that there is a close connection between such proof terms and terms in the so-called *typed lambda calculus*. To respect this connection, we'll use symbols used in the literature on the lambda calculus rather than the generic constructor and destructor symbols above. E.g., instead of writing $c^{\to}([x : \varphi]N)$ we write $\lambda x^{\varphi}.\, N$, instead of $d^{\to}(N, M)$ we write (NM), $\langle N, M\rangle$ instead of $c^{\wedge}(N, M)$, and $\pi_i(N)$ instead of $d_i^{\wedge}(N)$.

Definition 5.1. Proof terms are inductively generated by the following rules:

1. Every variable x is a proof term (Axioms).
2. If x is a variable, φ is a formula, and N is a proof term, then $\lambda x^{\varphi}.\, N$ is also a proof term ($\to$Intro).
3. If N and M are proof terms, then (NM) is also a proof term ($\to$Elim).
4. If N and M are proof terms, then $\langle N, M\rangle$ is a proof term ($\wedge$Intro).
5. If N is a proof term, then $\pi_i(N)$ is also a proof term, where i is 1 or 2 ($\wedge\text{Elim}_i$).
6. If N is a proof term, and φ is a formula, then $\iota_i^{\varphi}(M)$ is a proof term, where i is 1 or 2 ($\vee\text{Intro}_i$).
7. If M, N_1, N_2 are proof terms, and x_1, x_2 are variables, then $\delta\, M\, x_1.N_1\, x_2.N_2$ is a proof term ($\vee$Elim).
8. If N is a proof term and φ is a formula, then $\varepsilon^{\varphi}(N)$ is proof term ($\bot_I$).

Axioms:

$$x : \varphi \Rightarrow x : \varphi$$

Inference Rules:

$$\frac{x : \varphi, \Gamma \Rightarrow N : \psi}{\Gamma \Rightarrow \lambda x^{\varphi}.\, N : \varphi \to \psi} \to\text{Intro}$$

$$\frac{\Gamma_1 \Rightarrow N : \varphi \to \psi \qquad \Gamma_2 \Rightarrow M : \varphi}{\Gamma_1, \Gamma_2 \Rightarrow (NM) : \psi} \to\text{Elim}$$

$$\frac{\Gamma_1 \Rightarrow N_1 : \varphi \qquad \Gamma_2 \Rightarrow N_2 : \psi}{\Gamma_1, \Gamma_2 \Rightarrow \langle N_1, N_2 \rangle : \varphi \wedge \psi} \wedge\text{Intro}$$

$$\frac{\Gamma \Rightarrow N : \varphi \wedge \psi}{\Gamma \Rightarrow \pi_1(N) : \varphi} \wedge\text{Elim}_1 \qquad \frac{\Gamma \Rightarrow N : \varphi \wedge \psi}{\Gamma \Rightarrow \pi_2(N) : \psi} \wedge\text{Elim}_2$$

$$\frac{\Gamma \Rightarrow N : \varphi}{\Gamma \Rightarrow \iota_1^{\psi}(N) : \varphi \vee \psi} \vee\text{Intro}_1 \qquad \frac{\Gamma \Rightarrow N : \psi}{\Gamma \Rightarrow \iota_2^{\varphi}(N) : \varphi \vee \psi} \vee\text{Intro}_2$$

$$\frac{\Gamma_1 \Rightarrow M : \varphi \vee \psi \qquad x : \varphi, \Gamma_2 \Rightarrow N_1 : \chi \qquad y : \psi, \Gamma_3 \Rightarrow N_2 : \chi}{\Gamma_1, \Gamma_2, \Gamma_3 \Rightarrow \delta\, M\, x.N_1\, y.N_2 : \chi} \vee\text{Elim}$$

$$\frac{\Gamma \Rightarrow N : \bot}{\Gamma \Rightarrow \varepsilon^{\varphi}(N) : \varphi} \bot_I$$

Table 5.1: Rules of the propositional term-labelled intuitionistic natural deduction system **tN2ip**.

int:pty:ter:tab:tN2ip

Not every proof term is the proof term of a derivation. For instance, a term $\langle N, M \rangle$ encodes a proof ending in a $\wedge$Intro inference, so its end-formula is a conjunction $\varphi \wedge \psi$. This can't be the major premise of a $\to$Elim rule, so $(\langle N, M \rangle P)$ can't be the proof term of a correct proof. Only those proof terms that follow the rules of **tN2ip** are *correct* proof terms. These rules are given in Table 5.1.

5.3 Converting Derivations to Proof Terms

int:pty:pt:sec

The process of converting natural deduction derivations to proof terms, i.e., translating derivations in **N2i** to derivations in **tN2i**, is extremely straightforward. It simply amounts to writing a proof term to the left of each formula on the right of a sequent in the derivation, resulting in sequents of the form $\Gamma \Rightarrow M : \varphi$. We'll then say that M *witnesses* φ in context Γ. Which proof term to write is entirely determined by the rules of **tN2i**.

Proposition 5.2. *Every derivation δ of $\Gamma \Rightarrow \varphi$ in* **N2i** *corresponds to a derivation δ' of $\Gamma \Rightarrow N : \varphi$ in* **tN2i**. *The proof term N is uniquely determined by δ.*

Proof. We proceed by induction on the height of δ. If δ is the axiom $x : \varphi \Rightarrow \varphi$, then δ' is the axiom $x : \varphi \Rightarrow x : \varphi$. If δ ends in an inference rule, the inductive hypothesis provides proof terms for each premise, and **tN2i**-derivations of sequents including these proof terms corresponding to each premise. We apply the corresponding rule of **tN2i**; the corresponding proof term is determined by the **tN2i** inference rule. □

Example 5.3. Consider the very simple derivation in **N1i** of $(\varphi \wedge \psi) \to (\psi \wedge \varphi)$:

$$\dfrac{\dfrac{\dfrac{[\varphi \wedge \psi]^x}{\psi}\wedge\text{Elim}_2 \qquad \dfrac{[\varphi \wedge \psi]^x}{\varphi}\wedge\text{Elim}_1}{\psi \wedge \varphi}\wedge\text{Intro}}{(\varphi \wedge \psi) \to (\psi \wedge \varphi)}\,{\scriptstyle x}\,\to\text{Intro}$$

In **N2i** it looks like this:

$$\dfrac{\dfrac{\dfrac{x : \varphi \wedge \psi \Rightarrow \varphi \wedge \psi}{x : \varphi \wedge \psi \Rightarrow \psi}\wedge\text{Elim}_2 \qquad \dfrac{x : \varphi \wedge \psi \Rightarrow \varphi \wedge \psi}{x : \varphi \wedge \psi \Rightarrow \varphi}\wedge\text{Elim}_1}{x : \varphi \wedge \psi \Rightarrow \psi \wedge \varphi}\wedge\text{Intro}}{\Rightarrow (\varphi \wedge \psi) \to (\psi \wedge \varphi)}\to\text{Intro}$$

We assign proof terms from the top down. The proof terms for the axioms are just the same as the variable in the context, x. The proof terms associated with $\wedge\text{Elim}_i$ are then determined as $\pi_2(x)$ and $\pi_1(x)$. When we apply $\wedge$Intro, we combine these two proof terms: $\langle \pi_2(x), \pi_1(x) \rangle$. Finally, the $\to$Intro rule applies a λ-abstraction, binding x and recording that x labelled the assumption $\varphi \wedge \psi$: $\lambda x^{\varphi \wedge \psi}. \langle \pi_2(x), \pi_1(x) \rangle$. The derivation in **tN2i** then is:

$$\dfrac{\dfrac{\dfrac{x : \varphi \wedge \psi \Rightarrow x : \varphi \wedge \psi}{x : \varphi \wedge \psi \Rightarrow \pi_2(x) : \psi}\wedge\text{Elim}_2 \qquad \dfrac{x : \varphi \wedge \psi \Rightarrow x : \varphi \wedge \psi}{x : \varphi \wedge \psi \Rightarrow \pi_1(x) : \varphi}\wedge\text{Elim}_1}{x : \varphi \wedge \psi \Rightarrow \langle \pi_2(x), \pi_1(x) \rangle : \psi \wedge \varphi}\wedge\text{Intro}}{\Rightarrow \lambda x^{\varphi \wedge \psi}. \langle \pi_2(x), \pi_1(x) \rangle : (\varphi \wedge \psi) \to (\psi \wedge \varphi)}\to\text{Intro}$$

5.4 Recovering Derivations from Proof Terms

int:pty:tp: sec

Now let us consider the other direction: translating a proof term N back to a derivation in **N2i**. In general, it is only possible to do this relative to a context that contains pairs $x : \varphi$ for every variable x that occurs free in the given proof term. But let's start with an example where there are no free variables, namely the term

$$\lambda x^{\varphi \wedge \psi}. \langle \pi_2(x), \pi_1(x) \rangle$$

It is of the form $\lambda x^{\varphi\wedge\psi}.\, N_1$. Since the only rule of **tN2i** that produces a term of this form is $\to$Intro, we know that the derivation encoded by N has to end with $\to$Intro, i.e., it ends in

$$\dfrac{x : \varphi \wedge \psi \;\Rightarrow\; \langle \pi_2(x), \pi_1(x)\rangle : \chi}{\Rightarrow\; \lambda x^{\varphi\wedge\psi}.\, \langle \pi_2(x), \pi_1(x)\rangle : (\varphi \wedge \psi) \to \chi}\,{\to}\text{Intro}$$

We don't yet know what χ is, but we do know that the antecedent is $\varphi \wedge \psi$, and that the context of the premise has to include $x : \varphi \wedge \psi$.

Now the proof term associated with the premise is already determined: $\langle \pi_2(x), \pi_1(x)\rangle$. The derivation *it* encodes must end with $\wedge$Intro. We still don't know what χ is, but because it is the conclusion of $\wedge$Intro, we know it must be a conjunction, i.e., $\chi \equiv (\chi_1 \wedge \chi_2)$. We continue the reconstruction:

$$\dfrac{\dfrac{x : \varphi \wedge \psi \;\Rightarrow\; \pi_2(x) : \chi_1 \qquad x : \varphi \wedge \psi \;\Rightarrow\; \pi_1(x) : \chi_2}{x : \varphi \wedge \psi \;\Rightarrow\; \langle \pi_2(x), \pi_1(x)\rangle : \chi_1 \wedge \chi_2}\,{\wedge}\text{Intro}}{\Rightarrow\; \lambda x^{\varphi\wedge\psi}.\, \langle \pi_2(x), \pi_1(x)\rangle : (\varphi \wedge \psi) \to (\chi_1 \wedge \chi_2)}\,{\to}\text{Intro}$$

The proof term $\pi_2(x)$ means that top sequent on the left was obtained by $\wedge\text{Elim}_2$, i.e., the premise is of the form $\theta \wedge \chi_1$. On the right, we can determine that the inference leading to χ_2 must be $\wedge\text{Elim}_1$, and that the premise is of the form $\chi_2 \wedge \alpha$.

$$\dfrac{\dfrac{\dfrac{x : \varphi \wedge \psi \;\Rightarrow\; x : \theta \wedge \chi_1}{x : \varphi \wedge \psi \;\Rightarrow\; \pi_2(x) : \chi_1}\,{\wedge}\text{Elim}_2 \qquad \dfrac{x : \varphi \wedge \psi \;\Rightarrow\; x : \chi_2 \wedge \alpha}{x : \varphi \wedge \psi \;\Rightarrow\; \pi_1(x) : \chi_2}\,{\wedge}\text{Elim}_1}{x : \varphi \wedge \psi \;\Rightarrow\; \langle \pi_2(x), \pi_1(x)\rangle : \chi_1 \wedge \chi_2}\,{\wedge}\text{Intro}}{\Rightarrow\; \lambda x^{\varphi\wedge\psi}.\, \langle \pi_2(x), \pi_1(x)\rangle : (\varphi \wedge \psi) \to (\chi_1 \wedge \chi_2)}\,{\to}\text{Intro}$$

Since the proof terms of the topmost sequents are now variables, we know that they must be axioms. This determines χ_1, χ_2, θ, and α: $\theta \equiv \varphi$ and $\chi_1 \equiv \psi$ or else the left sequent is not an axiom, and $\chi_2 \equiv \varphi$ and $\alpha \equiv \psi$, or else the right sequent isn't. We have recovered the proof fully:

$$\dfrac{\dfrac{\dfrac{x : \varphi \wedge \psi \;\Rightarrow\; x : \varphi \wedge \psi}{x : \varphi \wedge \psi \;\Rightarrow\; \pi_2(x) : \psi}\,{\wedge}\text{Elim}_2 \qquad \dfrac{x : \varphi \wedge \psi \;\Rightarrow\; x : \varphi \wedge \psi}{x : \varphi \wedge \psi \;\Rightarrow\; \pi_1(x) : \varphi}\,{\wedge}\text{Elim}_1}{x : \varphi \wedge \psi \;\Rightarrow\; \langle \pi_2(x), \pi_1(x)\rangle : \psi \wedge \varphi}\,{\wedge}\text{Intro}}{\Rightarrow\; \lambda x^{\varphi\wedge\psi}.\, \langle \pi_2(x), \pi_1(x)\rangle : (\varphi \wedge \psi) \to (\psi \wedge \varphi)}\,{\to}\text{Intro}$$

5.5 Types

int:pty:typ:sec

A proof term such as (MN) by itself does not uniquely determine the proof it is a proof term of. That's because the free variables, i.e., the proof terms of axioms, do not specify the formulas in the corresponding axioms. However, if we specify these in a context Γ, the proof *is* uniquely determined.

Definition 5.4. A *context* for a proof term N is a set Γ such that for every free variable of N, Γ contains exactly one pair $x : \varphi$.

Note that the definition does not require that Γ contains $x : \varphi$ *only* for free variables of N. This, e.g., $\Gamma = \{x : \varphi, y : \psi, z : \chi\}$ is a context for $\langle x, y\rangle$.

int:pty:typ: defn:type **Definition 5.5.** Suppose Γ is a context for N. The *type* of N (relative to Γ) is inductively defined as follows:

1. The type of x is φ iff $x : \varphi \in \Gamma$.
2. If N has type ψ relative to $x : \varphi, \Gamma$, then the type of $\lambda x^{\varphi}. N$ is $\varphi \to \psi$.
3. If N has type $\varphi \to \psi$ and M has type φ then (NM) has type ψ
4. If N has type φ and M has type ψ, then $\langle N, M\rangle$ has type $\varphi \wedge \psi$.
5. If N has type $\varphi_1 \wedge \varphi_2$, then $\pi_i(N)$ has type φ_i.
6. If N has type φ, then $\iota_i^{\psi}(N)$ has type $\varphi \vee \psi$ if $i = 1$ and type $\psi \vee \varphi$ if $i = 2$.
7. If M has type $\varphi \vee \psi$, N_1 has type χ relative to $x_1 : \varphi, \Gamma$, and N_2 has type χ relative to $x_2 : \psi, \Gamma$, then $\delta\, M\, x_1.N_1\, x_2.N_2$ has type χ.
8. If N has type $\bot$, then $\varepsilon^{\varphi}(N)$ has type φ.

Proposition 5.6. *Every proof term N has exactly one type relative to a context Γ for N, if it has a type at all (i.e., is a correct proof term).*

Proof. By induction on N. □

We write $\Gamma \Rightarrow N : \varphi$ if the type of N relative to Γ is φ. The clauses in Definition 5.5 should seem familiar: they are precisely the rules of **tN2i** with one slight difference: the context Γ is the same throughout. We can codify the definition as a calculus **tN3i** with the rules given in Table 5.2.

Proof terms are terms in a version of the λ-calculus, the *typed λ-calculus with products, sums, and empty type*. The classical λ-calculus only has term built from variables, λ-*abstracts* $\lambda x. N$ and application (MN), and it lacks the annotation by φ of the typed λ-abstract $\lambda x^{\varphi}. N$. The λ-calculus is a model of computation which is Turing complete (i.e., every partial computable function can be represented in it by a term). The version of the λ-calculus given by proof terms and typing rules is a restricted model of computation, but the basic ideas are the same. It is restricted by the type assignment: only *well-typed* terms, i.e., proof terms with types (relative to a context Γ) are legal terms. For instance, $\lambda x. (xx)$ is a legal term in the classical untyped λ-calculus, but $\lambda x^{\varphi}. (xx)$ is not well-typed for any φ. That's because the typing rules require that in order for (xx) to have a type, the first x would have to have type $\varphi \to \psi$ and the second x type φ, and that's impossible since φ is a proper sub-formula of $\varphi \to \psi$.

We can think of types intuitively as the kind of thing named by a term of that type. So a term of type $\varphi \wedge \psi$ picks out a pair with first component of type φ and second component of type ψ, i.e., an element of the product $\varphi \times \psi$.

Axioms:

$$\Gamma \Rightarrow x : \varphi \text{ if } x : \varphi \in \Gamma$$

Inference Rules:

$$\frac{x : \varphi, \Gamma \Rightarrow N : \psi}{\Gamma \Rightarrow \lambda x^{\varphi}.\, N : \varphi \to \psi} \to\text{Intro}$$

$$\frac{\Gamma \Rightarrow N : \varphi \to \psi \qquad \Gamma \Rightarrow M : \varphi}{\Gamma \Rightarrow (NM) : \psi} \to\text{Elim}$$

$$\frac{\Gamma \Rightarrow N_1 : \varphi \qquad \Gamma \Rightarrow N_2 : \psi}{\Gamma \Rightarrow \langle N_1, N_2 \rangle : \varphi \wedge \psi} \wedge\text{Intro}$$

$$\frac{\Gamma \Rightarrow N : \varphi \wedge \psi}{\Gamma \Rightarrow \pi_1(N) : \varphi} \wedge\text{Elim}_1 \qquad \frac{\Gamma \Rightarrow N : \varphi \wedge \psi}{\Gamma \Rightarrow \pi_2(N) : \psi} \wedge\text{Elim}_2$$

$$\frac{\Gamma \Rightarrow N : \varphi}{\Gamma \Rightarrow \iota_1^{\psi}(N) : \varphi \vee \psi} \vee\text{Intro}_1 \qquad \frac{\Gamma \Rightarrow N : \varphi}{\Gamma \Rightarrow \iota_2^{\psi}(N) : \psi \vee \varphi} \vee\text{Intro}_2$$

$$\frac{\Gamma \Rightarrow M : \varphi \vee \psi \qquad x : \varphi, \Gamma \Rightarrow N_1 : \chi \qquad y : \psi, \Gamma \Rightarrow N_2 : \chi}{\Gamma \Rightarrow \delta\, M\, x.N_1\, y.N_2 : \chi} \vee\text{Elim}$$

$$\frac{\Gamma \Rightarrow N : \bot}{\Gamma \Rightarrow \varepsilon^{\varphi}(N) : \varphi} \bot_I$$

Table 5.2: Rules of the propositional term-labelled intuitionistic natural deduction system **tN3ip** serving as a type assignment system.

int:pty:typ:
tab:tN3ip

A term of type $\varphi \to \psi$ is a function from things of type φ to things of type ψ. A term of type $\varphi \vee \psi$ is either something of type φ or something of type ψ, plus the information which one it is. This is like the disjoint union of the two sets φ and ψ defined as $\{\langle i, x\rangle : i = 1 \text{ or } 2, x \in \varphi \text{ if } i = 1, x \in \psi \text{ if } i = 2\}$. The type $\bot$ is the empty set, i.e., if a term has that type it can't name anything. Since natural deduction (including **tN3i**) is consistent, there is no derivation of $\bot$ unless there are open assumptions, and so there is no closed proof term (term without free variables) that has type $\bot$.

The operators of the typed λ-calculus produce terms that intuitively name the right sort of thing, provided the terms it is applied to name certain things. E.g., "if $N_1 : \varphi$ and $N_2 : \psi$, then $\langle N_1, N_2\rangle : \varphi \wedge \psi$" says that if N_1 names somethign of type φ and N_2 names something of type ψ, then $\langle N_1, N_2\rangle$ names something of type $\varphi \wedge \psi$.

5.6 Reduction

int:pty:red: sec

In natural deduction derivations, an introduction rule that is followed by an elimination rule is a detour that can be "reduced". For instance, an ∧Intro followed by ∧Elim "cancel out," since the conclusion of ∧Elim is always one of the conjuncts of the conjunction that ∧Intro produces, and the premises of ∧Intro are both conjuncts. Thus if we need a derivation of either conjunct, we can simply use that derivation without introducing the conjunction followed by eliminating it. We have

$$\dfrac{\dfrac{\begin{array}{cc}\vdots\,\delta_1 & \vdots\,\delta_2\\ \varphi_1 & \varphi_2\end{array}}{\varphi_1 \wedge \varphi_2}\,\wedge\text{Intro}}{\varphi_i}\,\wedge\text{Elim}_i \quad\to\quad \begin{array}{c}\vdots\,\delta_i\\ \varphi_i\end{array}$$

We can consider a similar relation of reduction on the associated proof terms suggested by this simplification. If N_1 is the proof term for δ_1 and N_2 the one for δ_2, then we can say the proof term for the proof on the left reduces (in one step) to the proof term for the proof on the right:

$$\pi_i(\langle N_1, N_2, \rangle) \to N_i$$

In the typed lambda calculus, this is the *beta reduction* rule for the product type.

In natural deduction, a pair of inferences such as a ∧Intro folloed by ∧Elim, i.e., a pair that is subject to reduction, is called a *cut*. In the typed lambda calculus the term on the left of → is called a *redex*, and the term to the right is called the *reductum*. Unlike untyped lambda calculus, where only $(\lambda x.\, N)Q$ is considered to be redex, in the typed lambda calculus the syntax is extended to terms involving $\langle N, M\rangle$, $\pi_i(N)$, $\iota_i(\varphi)N$, $\delta\, N\, x_1.M_1\, x_2.M_2$, and $\varepsilon^{\varphi}(N)$, with corresponding redexes.

Similarly we have reduction for disjunction:

$$\dfrac{\dfrac{\begin{array}{ccc}\vdots\,\delta & [\varphi_1]^{x_1} & [\varphi_2]^{x_2}\\ \varphi_i & \vdots\,\delta_1 & \vdots\,\delta_2\end{array}}{\varphi_1 \vee \varphi_2}\,\vee\text{Intro} \qquad \chi \qquad \chi}{\chi}\,{\scriptstyle x_1, x_2}\,\vee\text{Elim} \quad\to\quad \begin{array}{c}\vdots\,\delta\\ \varphi_i\\ \vdots\,\delta_i\\ \chi\end{array}$$

This corresponds to a reduction on proof terms:

$$\delta\, \iota_i^{\varphi_i}(M)\, x_1.N_1\, x_2.N_2 \to N_i[M/x_i]$$

where M is the proof term of δ, and N_i is the proof term of δ_i. This is the beta reduction rule of for sum types. Here, $M[N_i/x]$ means replacing all free occurrences of the variable x in M by N_i.

The reduction of the function type corresponds to removal of a detour of →Intro followed by a →Elim.

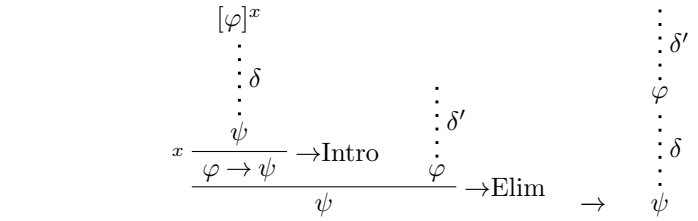

For proof terms, this amounts to ordinary beta reduction:

$$(\lambda x^{\varphi}.\, NM) \to N[M/x]$$

In addition to reduction conversions on derivations, we can consider permutation conversions. These apply to (sub)derivations that end in a ∨Elim followed by another elimination rule, e.g.:

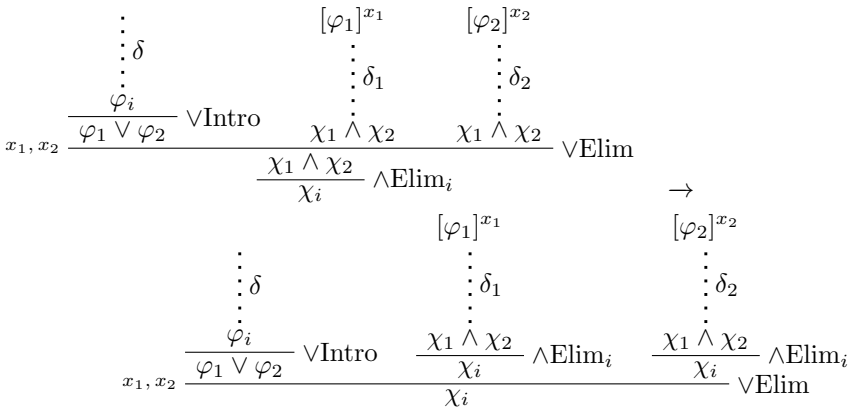

If the proof terms of δ, δ_1, and δ_2 are M, N_1, and N_2, respectively, the corresponding reduction rule for proof terms, i.e., λ terms, is:

$$\pi_i(\delta\, M\, x_1.N_1\, x_2.N_2) \to \delta\, M\, x_1.\pi_i(N_1)\, x_2.\pi_i(N_2)$$

Of course we can eliminate detours not just at the ends of derivations, but also within derivations, by applying the reduction conversion to a sub-derivation ending in a detour. Similarly, β-reduction applies to subterms of λ-terms. If N is a subterm of M, $N \to N'$, we can replace the subterm N in M by N' to obtain M'. In that case we also write $M \to M'$. If there is a sequence $M = M_1 \to M_2 \to M_2 \to \ldots \to M_k = M'$ (including where $k = 1$, i.e., $M = M'$) we write $M \twoheadrightarrow M'$.

A derivation where no conversion can be applied to any sub-derivation is called *normal*. Similarly, a λ-term where no conversion can be applied (to any sub-term) is called *normal*. A normal λ-term is also called a *normal form*.

5.7 Type Preservation

int:pty:tpr:sec

The reduction and permutation conversions considered can also be defined directly for the systems **tN2i** and **tN3i**. The system **tN3i** defines the type of a λ-term N relative to the context Γ as the formula φ such that $\Gamma \Rightarrow N : \varphi$ has a derivation in **tN3i**. The fact that the β-reduction rules for λ-terms exactly mirror the reduction and permutation conversions of derivations has an important consequence: the type of a λ-term relative to a context Γ remains the same after β-conversion. This property is called *type preservation*.

Proposition 5.7. *If N has type φ relative to Γ and $N \to N'$, then N' has type φ relative to Γ.*

Proof. By induction on N and the height of derivations δ of $\Gamma \Rightarrow N : \varphi$ we can easily establish the following: If M is a sub-term of N, then δ contains a sub-derivation δ_1 ending in $\Gamma' \Rightarrow M : \psi$. If M is a redex, then a conversion applies to δ_1. By applying the reduction to the δ_1 we obtain a derivation δ_1' of $\Gamma' \Rightarrow M' : \psi$. By inspection of the conversions of **tN3i**, M' is the reductum of M. Replacing δ_1 in δ by δ_1' results in a derivation δ' of $\Gamma \Rightarrow N' : \varphi$, where N' results from N by replacing the subterm M by M'.

For instance, consider the reduction conversion for $\wedge$Intro followed by $\wedge$Elim:

$$\dfrac{\dfrac{\vdots\ \Gamma \Rightarrow N_1 : \psi_1 \qquad \vdots\ \Gamma \Rightarrow N_2 : \psi_2}{\Gamma \Rightarrow \langle N_1, N_2\rangle : \psi_1 \wedge \psi_2}\wedge\text{Intro}}{\Gamma \Rightarrow \pi_i(\langle N_1, N_2\rangle) : \psi_i}\wedge\text{Elim}_i \quad \to \quad \vdots\ \Gamma \Rightarrow N_i : \psi_i$$

We see that the proof term of the derivation on the right, i.e., δ_1' is N_i. It is exactly the result of applying β-conversion to the proof term $\pi_i(\langle N_1, N_2\rangle)$ of the derivation on the left (δ_1).

Or consider a $\to$Intro followed by an $\to$Elim inference, and the reduction

conversion applied to it:

$$\dfrac{\dfrac{\begin{array}{c}\delta_2\\ \vdots\\ x:\psi_1,\Gamma \Rightarrow N:\psi_2\end{array}}{\Gamma \Rightarrow \lambda x^{\psi_1}.N:\psi_1\to\psi_2}\to\text{Intro} \qquad \begin{array}{c}\delta_3\\ \vdots\\ \Gamma \Rightarrow M:\psi_1\end{array}}{\Gamma \Rightarrow (\lambda x^{\psi_1}.NM):\psi_2}\to\text{Elim} \quad \to$$

$$\begin{array}{c}\delta_3\\ \vdots\\ \Gamma \Rightarrow M:\psi_1\\ \vdots\\ \delta_2[\delta_3/x:\psi_1]\\ \vdots\\ \Gamma \Rightarrow N[M/x]:\psi_2\end{array}$$

Again the proof term of the derivation on the right is exactly the result of β-reducing the proof term of the derivation on the left. Of course one would have to carefuly verify (by induction on the height of δ_2) that if we replace all axioms of the form $\Gamma' \Rightarrow x:\psi_1$ in δ_2 by the derivation δ_3 of $\Gamma \Rightarrow M:\psi_1$ actually results in a derivation $\delta_2[\delta_3/x:\psi_1]$ of $\Gamma \Rightarrow N[M/x]:\psi_2$. □

The tight correspondence between derivations and λ-terms has another consequence. If N is a λ-term of type φ relative to a context Γ and contains a redex M (i.e., it is not a normal form), then the corresponding **tN3i** proof δ of $\Gamma \Rightarrow N:\varphi$ contains a sub-derivation (ending in $\Gamma' \Rightarrow M:\psi$) to which a conversion can be applied. If $N \to N'$ by replacing M by its reductum, the corresponding result of applying the reduction to δ is a derivation of $\Gamma \Rightarrow N':\varphi$. Hence every λ-term not in normal form β-reduces to another. This property is called *progress*.

Progress together with type preservation ensure that β-reduction of λ-terms never "gets stuck": if a term is well-typed and isn't in normal form, it reduces to another well-typed term (and that term is of the same type).

5.8 Normalization

int:pty:nor:sec

In this section we prove that, via some reduction order, any deduction can be reduced to a normal deduction, which is called the *normalization property*. We will make use of the propositions-as-types correspondence: we show that every proof term can be reduced to a normal form; normalization for natural deduction derivations then follows.

Firstly we define some functions that measure the complexity of terms. The *length* $\mathrm{len}(\varphi)$ of a formulas is defined by

$$\begin{aligned}\mathrm{len}(p) &= 0\\ \mathrm{len}(\varphi \wedge \psi) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1\\ \mathrm{len}(\varphi \vee \psi) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1\\ \mathrm{len}(\varphi \to \psi) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1.\end{aligned}$$

The complexity of a redex M is measured by its *cut rank* $\mathrm{cr}(M)$:

$$\begin{aligned}\mathrm{cr}((\lambda x^{\varphi}.\, N^{\psi})Q) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1\\ \mathrm{cr}(\mathrm{p}_i(\langle M^{\varphi}, N^{\psi}\rangle)) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1\\ \mathrm{cr}(\mathrm{case}(\mathrm{in}_i^{\varphi_i}(M^{\varphi_i}), x_1^{\varphi_1}.N_1^{\chi}, x_2^{\varphi_2}.N_2^{\chi})) &= \mathrm{len}(\varphi) + \mathrm{len}(\psi) + 1\end{aligned}$$

The complexity of a proof term is measured by the most complex redex in it, and 0 if it is normal:

$$\mathrm{mr}(M) = \max\{\mathrm{cr}(N) | N \text{ is a sub term of } M \text{ and is redex}\}$$

int:pty:nor: lem:subst **Lemma 5.8.** *If $M[N^{\varphi}/x^{\varphi}]$ is a redex and $M \not\equiv x$, then one of the following cases holds:*

1. *M is itself a redex, or*
2. *M is of the form $\mathrm{p}_i(x)$, and N is of the form $\langle P_1, P_2\rangle$*
3. *M is of the form $\mathrm{case}(i, x_1.P_1, x_2.P_2)$, and N is of the form $\mathrm{in}_i(Q)$*
4. *M is of the form xQ, and N is of the form $\lambda x.\, P$*

In the first case, $\mathrm{cr}(M[N/x]) = \mathrm{cr}(M)$; in the other cases, $\mathrm{cr}(M[N/x]) = \mathrm{len}(\varphi))$.

Proof. Proof by induction on M.

1. If M is a single variable y and $y \not\equiv x$, then $y[N/x]$ is y, hence not a redex.
2. If M is of the form $\langle N_1, N_2\rangle$, or $\lambda x.\, N$, or $\mathrm{in}_i^{\varphi}(N)$, then $M[N^{\varphi}/x^{\varphi}]$ is also of that form, and so is not a redex.
3. If M is of the form $\mathrm{p}_i(P)$, we consider two cases.
   a) If P is of the form $\langle P_1, P_2\rangle$, then $M \equiv \mathrm{p}_i(\langle P_1, P_2\rangle)$ is a redex, and clearly
   $$M[N/x] \equiv \mathrm{p}_i(\langle P_1[N/x], P_2[N/x]\rangle)$$
   is also a redex. The cut ranks are equal.

b) If P is a single variable, it must be x to make the substitution a redex, and N must be of the form $\langle P_1, P_2\rangle$. Now consider

$$M[N/x] \equiv \mathrm{p}_i(x)[\langle P_1, P_2\rangle/x],$$

which is $\mathrm{p}_i(\langle P_1, P_2\rangle)$. Its cut rank is equal to $\mathrm{cr}(x)$, which is $\mathrm{len}(\varphi)$.

The cases of $\mathrm{case}(N, x_1.N_1, x_2.N_2)$ and PQ are similar. □

Lemma 5.9. *If M contracts to M', and $\mathrm{cr}(M) > \mathrm{cr}(N)$ for all proper redex sub-terms N of M, then $\mathrm{cr}(M) > \mathrm{mr}(M')$.*

Proof. Proof by cases.

1. If M is of the form $\mathrm{p}_i(\langle M_1, M_2\rangle)$, then M' is M_i; since any sub-term of M_i is also proper sub-term of M, the claim holds.

2. If M is of the form $(\lambda x^\varphi.\, N)Q^\varphi$, then M' is $N[Q^\varphi/x^\varphi]$. Consider a redex in M'. Either there is corresponding redex in N with equal cut rank, which is less than $\mathrm{cr}(M)$ by assumption, or the cut rank equals $\mathrm{len}(\varphi)$, which by definition is less than $\mathrm{cr}((\lambda x^\varphi.\, N)Q)$.

3. If M is of the form

$$\mathrm{case}(\mathrm{in}_i(N^{\varphi_i}), x_1^{\varphi_1}.N_1^\chi, x_2^{\varphi_2}.N_2^\chi),$$

then $M' \equiv N_i[N/x_i^{\varphi_i}]$. Consider a redex in M'. Either there is corresponding redex in N_i with equal cut rank, which is less than $\mathrm{cr}(M)$ by assumption; or the cut rank equals $\mathrm{len}(\varphi_i)$, which by definition is less than $\mathrm{cr}(\mathrm{case}(\mathrm{in}_i(N^{\varphi_i}), x_1^{\varphi_1}.N_1^\chi, x_2^{\varphi_2}.N_2^\chi))$. □

Theorem 5.10. *All proof terms reduce to normal form; all derivations reduce to normal derivations.*

Proof. The second follows from the first. We prove the first by complete induction on $m = \mathrm{mr}(M)$, where M is a proof term.

1. If $m = 0$, M is already normal.

2. Otherwise, we proceed by induction on n, the number of redexes in M with cut rank equal to m.

   a) If $n = 1$, select any redex N such that $m = \mathrm{cr}(N) > \mathrm{cr}(P)$ for any proper sub-term P which is also a redex of course. Such a redex must exist, since any term only has finitely many subterms.

   Let N' denote the reductum of N. Now by the lemma $\mathrm{mr}(N') < \mathrm{mr}(N)$, thus we can see that n, the number of redexes with $\mathrm{cr}(=)m$ is decreased. So m is decreased (by 1 or more), and we can apply the inductive hypothesis for m.

b) For the induction step, assume $n > 1$. the process is similar, except that n is only decreased to a positive number and thus m does not change. We simply apply the induction hypothesis for n. □

The fact that typed λ-terms normalize means that every such term *has a normal form*. If we think of β-reducing λ-terms to their normal form as executing a computation and the normal form as the value of the computation, this means that every computation in the typed λ-calculus eventually halts an produces a value. Moreover, type preservation ensures that the value has the right type. E.g., if N has type $\varphi \to \psi$ (i.e., it is a function that is supposed to work for arguments of type φ and produces values of type ψ) and we apply it to a term M of type φ (the input), (NM) reduces to a term N' (the output) and that output is guaranteed to be of type ψ.

In fact, terms in the typed λ-calculus don't just reduce to normal forms for one way of applying β-reductions to subterms, but *any* way of applying β-reductions to subtems eventually terminates in a normal form. This property is called *strong normalization*. It means that not only is there a way to guarantee that a computation eventually halts and produces a value, but *any* way to execute the computation eventually produces a value.

How do we know that different ways of executing a computation produce the *same* value? So far we just know (from type preservation) that all values it produces are of the same type. The typed λ-calculus has another important property: it is *confluent* or *Church-Rosser*. If $N \twoheadrightarrow N_1$ and $N \twoheadrightarrow N_2$ then there is some term N' such that $N_1 \twoheadrightarrow N'$ and $N_2 \twoheadrightarrow N'$. Recall that $N \twoheadrightarrow N'$ means: N' results from zero or more $\to$-reduction steps. If N_1 is in normal form then the only term N' such that $N_1 \twoheadrightarrow N'$ is N_1 itself (via zero $\to$-reduction steps). It follows that if N_1 and N_2 are both in normal form, then $N_1 = N' = N_2$. In other words, since $\twoheadrightarrow$ is strongly normalizing, every way of reducing a term eventually results in a normal form. Since $\twoheadrightarrow$ is Church-Rosser, any two normal forms are equal. Hence, computation in the typed λ-calculus always terminates and the value (normal form) is the same regardless of how the computation proceeds.

Proving that a relation is confluent is typically difficult. However, we can easily establish the following weaker condition:

Proposition 5.11 (Weak Church-Rosser Property). *If $N \to N_1$ and $N \to N_2$ then there is a term N' such that $N_1 \twoheadrightarrow N'$ and $N_2 \twoheadrightarrow N'$.*

Proof. The two terms N_1 and N_2 result from N by replacing the redex M_1 or M_2 by its reductum, respectively. The redexes M_1 and M_2 are subterms of N, and so cannot overlap. This means that either one is a subterm of the other, or they occur in different non-overlapping places in N. Consider the latter case: N is of the form $N(M_1, M_2)$. Replacing M_1 by its reductum M_1' results in N_1 which is $N(M_1', M_2)$, and replacing M_2 by its reductum M_2' results in N_2 which is $N(M_1, M_2')$. Both reduce to $N(M_1', M_2')$. For the other case, suppose M_2 is a sub-term of M_1 (the other case is symmetric) and distinguish cases according

to what kind of redex M_1 is. E.g., if $M_1 \equiv \pi_1(\langle O_1, O_2\rangle)$ it reduces to O_1. If M_2 is a subterm of O_1, then converting M_2 results in O_1'. So converting M_1 first and then M_2 results in O_1'. But converting M_2 first results in $\pi_1(\langle O_1', O_2\rangle)$ and this reduces also to O_2. □

Lemma 5.12 (Newman's Lemma). *If $\twoheadrightarrow$ is strongly normalizing and has the weak Church-Rosser property, it is confluent.*

Proof. Since $\twoheadrightarrow$ is strongly normalizing, any reduction sequence results in a normal form. Normal forms are unique iff $\twoheadrightarrow$ is confluent. So suppose some term N has two distinct normal forms N' and N'', the result of reduction sequences

$$N = N_1' \to N_2' \to \ldots \to N_k' = N' \text{ and}$$
$$N = N_1'' \to N_2'' \to \ldots \to N_l'' = N''$$

(The reduction sequence must have length > 0 since otherwise N would already be in normal form.)

If $N_1' = N_1''$ then it has two distinct normal forms (N' and N''). If $N_1' \neq N_1''$ then by weak Church-Rosser there is some N''' such that $N_1' \twoheadrightarrow N'''$ and $N_1'' \twoheadrightarrow N'''$. Since $N' \neq N''$, either $N''' \neq N'$ or $N''' \neq N''$. Consequently, either N_1' or N_1'' has two distinct normal forms. We've shown that if N has two distinct normal forms, there is an N_1 such that $N \to N_1$ and N_1 has two distinct normal forms. Obviously we can keep going and produce $N \to N_1 \to N_2 \to \ldots$, but this is impossible because $\to$ is strongly normalizing. □

Photo Credits

Bibliography