A Tutorial on the Curry-Howard Correspondence

Darryl McAdams

April 9, 2013

Abstract

Typical introductions to the Curry-Howard Correspondence employ explanations that take the form "a proof of $A \wedge B$ is just a proof of A together with a proof of B, so $A \wedge B$ is just a pair type". This tutorial instead provides a from-scratch explanation, starting with the Natural Deduction proof system for Intuitionistic Propositional Logic with connectives $\top\ \bot\ \wedge\ \vee\ \Rightarrow$. From that, we can easily provide a *meta*-proof system that lets us prove things *about Natural Deduction proofs*, which will give us precisely the typing rules of the simply typed λ calculus. We'll also look at how the Natural Deduction rules for proof simplification/normalization give rise directly to the computation rules of the λ calculus.

1 Natural Deduction for Intuitionistic Propositional Logic (IPL)

Natural Deduction is a format for doing proofs using tree-like proof objects. A proof of a proposition P is a tree with P as the root node, where each node is a proposition with a horizontal line drawn between it and its child nodes. Child nodes are placed against the line of the parent node. In this tutorial the root is at the bottom, with child nodes above parent nodes, and with rule names on the right. For a given horizontal line, the propositions of its child nodes are said to be the "premises" of the inference, and the proposition of its parent node is said to be the "conclusion" of the inference. For example:

$$\frac{P \quad Q \quad R}{S}\,\textsf{rule}^{1}$$

[1]Martin-Löf [ITT] introduced the use of judgements into proof theory, where proofs don't involve bare propositions but rather judgments about propositions, e.g. that a

The inference rule $\textsf{rule}$ here is given three premises, P, Q, and R, and produces the conclusion S. We can read this as "from P, Q, and R, we can use $\textsf{rule}$ to conclude S". Inference rules can have zero or more premises, and when they have none they are called axioms. These proofs can be read top down, in which case you think of them as a discovery process, finding out what you can build from the top-most propositions, or bottom up, in which case you think of them as searches for a proof of the root proposition.

Natural Deduction for Intuitionistic Propositional Logic (IPL), like for most logics, also has a second, distinct use of horizontal lines that goes adjacent to propositional nodes, called a "hypothesis". This kind of horizontal line is labeled with an arbitrary name that's not an inference rule name. We'll represent this using a different font like so:

$$\frac{}{P}\,p$$

This is used to mean "assume p is a proof of P". We'll also use a subscripted letter on rule names to relate the use of the rule to a hypothetical in the proof(s) above it. A hypothesis is said to be "active" if there is no rule below it that's superscripted with the name of the hypothesis, and "inactive" otherwise (we say the rule "discharges" the hypothesis).

With this core, let's now define the inference rules for IPL, which constitute a collection of valid node shapes that can appear inside a proof for IPL. Let's start with conjunction ($\wedge$), since it exemplifies a duality of rules quite nicely.

Introduction Rules	**Elimination Rules**
$\dfrac{A \quad B}{A \wedge B}\,\wedge\textsf{I}$	$\dfrac{A \wedge B}{A}\,\wedge\textsf{E}^{\textsf{L}} \qquad \dfrac{A \wedge B}{B}\,\wedge\textsf{E}^{\textsf{R}}$

What we see here is that there are two kinds of rules for $\wedge$: one that introduces a new occurrence of the symbol, and one that eliminates an

proposition is true. This schema would be written using judgments as

$$\frac{P\ \text{true} \quad Q\ \text{true} \quad R\ \text{true}}{S\ \text{true}}\,\textsf{rule}$$

Many logicians don't employ judgments like this, but it's implicit in proof systems that we're always making judgments about propositions. This is useful to keep in mind, though not entirely necessary, because later we'll be using explicit judgements of the form $\Gamma \vdash \mathcal{P} : P$.

occurrence of the symbol, hence the names of the two classes of rules. Many logical symbols employ both introduction and elimination rules, but some use only one class of rules. You'll also notice that we can have multiple elimination rules. As we'll see when we get to disjunction ($\vee$), you can also have multiple introduction rules.

Let's move on now to the simplest propositions: truth ($\top$) and falsity ($\bot$):

Introduction Rules	**Elimination Rules**
$\dfrac{}{\top}\ \top\mathsf{I}$	$\top$ has no elimination rules
$\bot$ has no introduction rules	$\dfrac{\bot}{C}\ \bot\mathsf{E}$

Proving $\top$ should be trivial, since it's true, and so it's no surprise that $\top\mathsf{I}$ has no premises, because we don't need to rely on any other proofs to get $\top$. We also have no elimination rules for $\top$, because proving $\top$ doesn't let us get any new information. Similarly, because $\bot$ is supposed to be unprovable, we don't want to be able to get $\bot$ into a proof unless something else gives it to us to use via an eliminator, so we have no introduction rules. But that means we shouldn't be able to prove $\bot$ from scratch, so we can eliminate it if we do ever have it in the proof. Doing so lets us prove any proposition C we like, because we know that to get $\bot$ in the first place required doing something that was impossible, so we haven't magically gained any new insights. If the impossible is true, anything is true.

Moving on to implication ($\Rightarrow$), we have some fairly simple rules:

Introduction Rules	**Elimination Rules**
$\dfrac{\dfrac{}{A}\ a \\ \vdots \\ B}{A \Rightarrow B}\ \Rightarrow\mathsf{I}_a$	$\dfrac{A \Rightarrow B \quad A}{B}\ \Rightarrow\mathsf{E}$

The introduction rule is somewhat mysterious. What it means is, if we can prove B from the hypothesis that A holds, then we can conclude $A \Rightarrow B$. Prior to adding the last inference, we just have a proof with an active hypothesis:

$$\begin{array}{c} \overline{}^{\,a} \\ A \\ \vdots \\ B \end{array}$$

Here, the hypothesis really is an assumption we've made, and from that concluded B. But after using the $\Rightarrow\mathsf{I}_a$ inference rule, we've discharged that hypothesis, so that $A \Rightarrow B$ has been proven to hold *even if A doesn't actually hold*. The dots in this should be taken to mean that the proof that leads to B might use the hypothesis any number of times, and might have other hypotheses (that is to say, the hypothesis that A is true is sufficient to prove B, but it may not be necessary, and may not even appear in an actual proof above B when we use this inference rule).

Our last set of inference rules are for disjunction ($\vee$):

Introduction Rules	**Elimination Rules**
$\dfrac{A}{A \vee B}\,\vee\mathsf{I}^{\mathsf{L}} \qquad \dfrac{B}{A \vee B}\,\vee\mathsf{I}^{\mathsf{R}}$	$\dfrac{A \vee B \qquad \begin{array}{c} \overline{}^{\,a} \\ A \\ \vdots \\ C \end{array} \qquad \begin{array}{c} \overline{}^{\,b} \\ B \\ \vdots \\ C \end{array}}{C}\,\vee\mathsf{E}_{a,b}$

The introduction rules are simple enough: if we can prove either disjunct, we can prove the whole disjunction. The elimination rule is slightly less obvious; though we have a proof of $A \vee B$, we don't know whether it's A that's true or B that's true, so what we have to provide is two hypothetical proofs: one that proves C holds if A were true, and one that proves C holds if B were true. If we can show that C follows from each of them independently, then we can conclude C follows from the disjunction of them, because the disjunction tells us that at least one of them is true. Note that the rule name for elimination is subscripted, again keeping track of the assumptions we've discharged.

2 From Natural Deduction to the λ Calculus

In the previous section, we considered inference rules as static parts of proofs. That is to say, these inference rules are just things that show up in proof trees, such as the following proof of $(A \wedge B) \Rightarrow (B \wedge A)$:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{A \wedge B}\,p}{B}\,\wedge\mathsf{ER} \qquad \dfrac{\dfrac{}{A \wedge B}\,p}{A}\,\wedge\mathsf{E^L}}{B \wedge A}\,\wedge\mathsf{I}}{(A \wedge B) \Rightarrow (B \wedge A)}\,\Rightarrow\mathsf{I}_p$$

These inference rules don't describe the *process* of building a proof, they only describe the local appearance of the proof trees: if every node in the tree, leaf or branch node, can be formed using one of the inference rules from the last section, then the proposition at the bottom of the tree must follow from the propositions at the leaves (and there might be none of those, if the leaves consist entirely of discharged hypotheses!).

What we'd like to do is develop a *meta*-system for describing the proof objects themselves and how they get constructed, not the local properties of the proof objects. To that end, let's consider some visualizations of how we, the provers, go about using these inference rules. Consider conjunction introduction. We start out having constructed two proof trees, $\mathcal{A}$ and $\mathcal{B}$, say, each with some hypotheses at the leaves (which we'll represent with any set Γ which contains at least the hypotheses in $\mathcal{A}$ and $\mathcal{B}$, but possibly more):

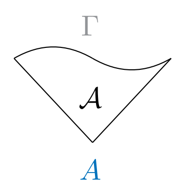

Maybe $\mathcal{A}$ and $\mathcal{B}$ are big complicated proofs, maybe not. Maybe Γ includes lots of hypotheses, maybe it includes only the hypotheses in $\mathcal{A}$ and $\mathcal{B}$. Either way, we can of course form a new proof tree that has these as subtrees:

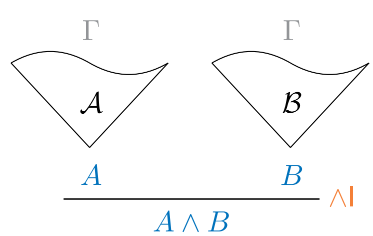

The resulting proof tree has the hypotheses of both input proof trees. What we want to do, now, is give Natural Deduction style inference rules that describe *these manipulations of proof trees* — we want a meta-proof system. To do this, we'll use just the inferential core of Natural Deduction — there are no meta-hypotheses. We'll also utilize a new ternary symbol, $\cdot \vdash \cdot : \cdot$, which will have hypotheses on the left, proof trees in the middle, and propositions on the right. So, in place of proofs with assumptions like

$$\Gamma$$

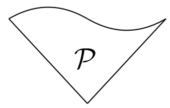

$$P$$

we'll instead use

$$\Gamma \ \vdash \ \mathcal{P} : P$$

which should be read as "assuming any of the hypotheses in Γ, the tree $\mathcal{P}$ is a proof that P holds". The Greek letters above the triangular parts of these trees are to be understood as collections of hypothesis, and the proofs denoted by the name inside the triangular part contains only hypotheses above the triangular part (possibly none of them). Consequently, the boring hypothetical proof

$$\dfrac{\dfrac{}{A}\,a}{A \vee A}\,\vee\mathsf{I}^{\mathsf{L}}$$

would be represented in this notation as

$$\Gamma \quad \vdash \quad \dfrac{\dfrac{}{A}\,a}{A \vee A}\,\vee\mathsf{I}^{\mathsf{L}} \quad : \quad A \vee A$$

where Γ is any collection of hypothesis that includes a hypothesis of A named a, such as $\{\dfrac{}{A}\,a, \dfrac{}{B}\,b\}$. The fact that a Γ that is "above" a proof $\mathcal{P}$ can contain more propositions than just those used in $\mathcal{P}$ is important, because it allows us to express the idea that some assumptions may not matter to a proof. Variations on the nature of Γ and how we're allowed to manipulate it give rise to different kinds of logics.

When we discharge a hypothesis p with an inference rule, the resulting tree is considered to have one less assumption (the one corresponding to p), and the assumption can't be used anywhere that's not above the discharging inference. This leaves open the possibility that two discharging inferences might use the same name for hypotheses. For convenience, assume this never happens — hypothesis names are really supposed to be means to connect discharging inferences with hypotheses, and we could just as easily use arrows to point from a discharging inference to its hypotheses instead of using names, but arrows are unconventional.

We will also mix Greek letters with actual hypotheses to emphasize some hypotheses as relevant. When we do this, the distinguished hypotheses are assumed to *not* appear in the Greek letter beside it.

Because manipulating trees themselves is a bit bulky, we'll also make use of a convenient shorthand. The translation into the shorthand, which we will represent with $\lfloor - \rfloor$, is defined recursively as follows.

$$\left\lfloor \frac{}{P}\, p \right\rfloor = p$$

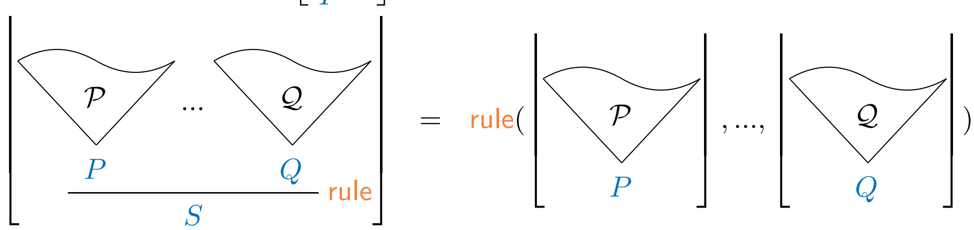

Notice that our translation lacks the propositions along the way. This can be recovered fully from the context that the proof is given in.

We'll call the hypotheticals "variables", the operators for introduction rules "constructors" and the operators for elimination rules "eliminators" (sometimes called "destructors" in other literature). Operators that have subscripts are called "binders" because they discharge hypotheses (i.e. bind variables).

As an example, let's translate the proof of $(P \wedge Q) \Rightarrow (Q \wedge P)$ into this new linear notation.

$$
\left\lfloor \dfrac{\dfrac{\dfrac{\dfrac{}{A \wedge B}\,p}{B}\,\wedge\mathsf{E}^{\mathsf{R}} \qquad \dfrac{\dfrac{}{A \wedge B}\,p}{A}\,\wedge\mathsf{E}^{\mathsf{L}}}{B \wedge A}\,\wedge\mathsf{I}}{(A \wedge B) \Rightarrow (B \wedge A)}\,\Rightarrow\mathsf{I}_p \right\rfloor
= \Rightarrow\mathsf{I}_p\left(\left\lfloor \dfrac{\dfrac{\dfrac{}{A \wedge B}\,p}{B}\,\wedge\mathsf{E}^{\mathsf{R}} \qquad \dfrac{\dfrac{}{A \wedge B}\,p}{A}\,\wedge\mathsf{E}^{\mathsf{L}}}{B \wedge A}\,\wedge\mathsf{I} \right\rfloor\right)
$$

$$
= \Rightarrow\mathsf{I}_p\left(\wedge\mathsf{I}\left(\left\lfloor \dfrac{\dfrac{}{A \wedge B}\,p}{B}\,\wedge\mathsf{E}^{\mathsf{R}} \right\rfloor, \left\lfloor \dfrac{\dfrac{}{A \wedge B}\,p}{A}\,\wedge\mathsf{E}^{\mathsf{L}} \right\rfloor\right)\right)
$$

$$
= \Rightarrow\mathsf{I}_p\left(\wedge\mathsf{I}\left(\wedge\mathsf{E}^{\mathsf{R}}\left(\left\lfloor \dfrac{}{A \wedge B}\,p \right\rfloor\right), \wedge\mathsf{E}^{\mathsf{L}}\left(\left\lfloor \dfrac{}{A \wedge B}\,p \right\rfloor\right)\right)\right)
$$

$$
= \Rightarrow\mathsf{I}_p(\wedge\mathsf{I}(\wedge\mathsf{E}^{\mathsf{R}}(p), \wedge\mathsf{E}^{\mathsf{L}}(p)))
$$

This now gives us a sense of the objects that we're working with. It also gives us a sense of the tree rules we have to write. The tree rules that correspond to the original inference rules are fairly straight forward, and we'll prefix their names with M to indicate that they're *meta*-rules, rules about how proofs rules are use to build proofs.

Introduction Rules | **Elimination Rules**

$$\frac{}{\Gamma \vdash \top\mathsf{I} : \top}\ \mathsf{M}\top\mathsf{I}$$

—

—

$$\frac{\Gamma \vdash \mathcal{L} : \bot}{\Gamma \vdash \bot\mathsf{E}(\mathcal{L}) : C}\ \mathsf{M}\bot\mathsf{E}$$

$$\frac{\Gamma,\ a : A \vdash \mathcal{B} : B}{\Gamma \vdash \Rightarrow\mathsf{I}_a(\mathcal{B}) : A \Rightarrow B}\ \mathsf{M}{\Rightarrow}\mathsf{I}$$

$$\frac{\Gamma \vdash \mathcal{F} : A \Rightarrow B \qquad \Gamma \vdash \mathcal{A} : A}{\Gamma \vdash \Rightarrow\mathsf{E}(\mathcal{F}, \mathcal{A}) : B}\ \mathsf{M}{\Rightarrow}\mathsf{E}$$

$$\frac{\Gamma \vdash \mathcal{A} : A \qquad \Gamma \vdash \mathcal{B} : B}{\Gamma \vdash \wedge\mathsf{I}(\mathcal{A}, \mathcal{B}) : A \wedge B}\ \mathsf{M}\wedge\mathsf{I}$$

$$\frac{\Gamma \vdash \mathcal{P} : A \wedge B}{\Gamma \vdash \wedge\mathsf{E}^{\mathsf{L}}(\mathcal{P}) : A}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{L}} \qquad \frac{\Gamma \vdash \mathcal{P} : A \wedge B}{\Gamma \vdash \wedge\mathsf{E}^{\mathsf{R}}(\mathcal{P}) : B}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{R}}$$

$$\frac{\Gamma \vdash \mathcal{A} : A}{\Gamma \vdash \vee\mathsf{I}^{\mathsf{L}}(\mathcal{A}) : A \vee B}\ \mathsf{M}\vee\mathsf{I}^{\mathsf{L}}$$

$$\frac{\Gamma \vdash \mathcal{B} : B}{\Gamma \vdash \vee\mathsf{I}^{\mathsf{R}}(\mathcal{B}) : A \vee B}\ \mathsf{M}\vee\mathsf{I}^{\mathsf{R}}$$

$$\frac{\Gamma \vdash \mathcal{D} : A \vee B \qquad \Gamma,\ a : A \vdash \mathcal{C} : C \qquad \Gamma,\ b : B \vdash \mathcal{C}' : C}{\Gamma \vdash \vee\mathsf{E}_{a,b}(\mathcal{D}, \mathcal{C}, \mathcal{C}') : C}\ \mathsf{M}\vee\mathsf{E}$$

We also have to add one axiom for hypothesis nodes (where the order on the left-hand side of $p : P$ relative to the context of Γ is irrelevant):

$$\frac{}{\Gamma,\ p : P \vdash p : P}\ \mathsf{hyp}$$

This now lets us give a meta-proof used to show the validity of the old Natural Deduction proof tree for $(A \wedge B) \Rightarrow (B \wedge A)$:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{p : A \wedge B \vdash p : A \wedge B}\ \mathsf{hyp}}{p : A \wedge B \vdash \wedge\mathsf{E}^{\mathsf{R}}(p) : B}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{R}} \qquad \dfrac{\dfrac{}{p : A \wedge B \vdash p : A \wedge B}\ \mathsf{hyp}}{p : A \wedge B \vdash \wedge\mathsf{E}^{\mathsf{L}}(p) : A}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{L}}}{p : A \wedge B \vdash \wedge\mathsf{I}(\wedge\mathsf{E}^{\mathsf{R}}(p), \wedge\mathsf{E}^{\mathsf{L}}(p)) : B \wedge A}\ \mathsf{M}\wedge\mathsf{I}}{\vdash \Rightarrow\mathsf{I}_p(\wedge\mathsf{I}(\wedge\mathsf{E}^{\mathsf{R}}(p), \wedge\mathsf{E}^{\mathsf{L}}(p))) : (A \wedge B) \Rightarrow (B \wedge A)}\ \mathsf{M}{\Rightarrow}\mathsf{I}$$

Unfortunately, this notation is a bit unwieldy for normal use. If this were a functional programming language like Haskell, for instance, we'd have a semi-readable representation because we'd probably stick to using ASCII to represent our trees and write out these names as something like Lambda and AndIntro, etc. Let's give the inference rules again, but using the more familiar notation of the λ calculus:

Introduction Rules

$$\frac{}{\Gamma \vdash \langle\rangle : \top}\ \mathsf{M}\top\mathsf{I}$$

—

$$\frac{\Gamma,\ a : A \vdash \mathcal{B} : B}{\Gamma \vdash \lambda a.\, \mathcal{B} : A \Rightarrow B}\ \mathsf{M}{\Rightarrow}\mathsf{I}$$

$$\frac{\Gamma \vdash \mathcal{A} : A \qquad \Gamma \vdash \mathcal{B} : B}{\Gamma \vdash \langle \mathcal{A}, \mathcal{B} \rangle : A \wedge B}\ \mathsf{M}\wedge\mathsf{I}$$

$$\frac{\Gamma \vdash \mathcal{A} : A}{\Gamma \vdash \mathsf{left}\ \mathcal{A} : A \vee B}\ \mathsf{M}\vee\mathsf{I}^{\mathsf{L}}$$

$$\frac{\Gamma \vdash \mathcal{B} : B}{\Gamma \vdash \mathsf{right}\ \mathcal{B} : A \vee B}\ \mathsf{M}\vee\mathsf{I}^{\mathsf{R}}$$

Elimination Rules

—

$$\frac{\Gamma \vdash \mathcal{L} : \bot}{\Gamma \vdash \bot\mathsf{elim}\ \mathcal{L} : C}\ \mathsf{M}\bot\mathsf{E}$$

$$\frac{\Gamma \vdash \mathcal{F} : A \Rightarrow B \qquad \Gamma \vdash \mathcal{A} : A}{\Gamma \vdash \mathcal{F}\mathcal{A} : B}\ \mathsf{M}{\Rightarrow}\mathsf{E}$$

$$\frac{\Gamma \vdash \mathcal{P} : A \wedge B}{\Gamma \vdash \mathsf{fst}\ \mathcal{P} : A}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{L}} \qquad \frac{\Gamma \vdash \mathcal{P} : A \wedge B}{\Gamma \vdash \mathsf{snd}\ \mathcal{P} : B}\ \mathsf{M}\wedge\mathsf{E}^{\mathsf{R}}$$

$$\frac{\Gamma \vdash \mathcal{D} : A \vee B \qquad \Gamma,\ a : A \vdash \mathcal{C} : C \qquad \Gamma,\ b : B \vdash \mathcal{C}' : C}{\Gamma \vdash \mathsf{case}\ \mathcal{D}\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathcal{C}\ ;\ \mathsf{right}\ b \mapsto \mathcal{C}'\ \} : C}\ \mathsf{M}\vee\mathsf{E}$$

I should stress that the case expression here does not have genuine pattern matching. Instead, it is a single, mixfix operator that *binds* the variables a and b, in the same way that a λ in $\lambda x.\, y$ binds x. The left and right in the case expression are not constructors, they're just parts of the mixfix name that happen to be identical to the names of the constructors for the disjunction. This is done intentionally to make the case expression comprehensible immediately upon inspection, but it's important to keep in mind that the expression doesn't have pattern matching any more than $\vee\mathsf{E}_{a,b}(\mathcal{D}, \mathcal{C}, \mathcal{C}')$ does.

Our meta-proof of $(A \wedge B) \Rightarrow (B \wedge A)$ in the new notation becomes

$$\dfrac{\dfrac{\dfrac{\dfrac{}{p : A \wedge B \ \vdash\ p : A \wedge B}\ \textsf{hyp}}{p : A \wedge B \ \vdash\ \textsf{snd}\ p : B}\ \textsf{M}{\wedge}\textsf{E}^{\textsf{R}} \qquad \dfrac{\dfrac{}{p : A \wedge B \ \vdash\ p : A \wedge B}\ \textsf{hyp}}{p : A \wedge B \ \vdash\ \textsf{fst}\ p : A}\ \textsf{M}{\wedge}\textsf{E}^{\textsf{L}}}{p : A \wedge B \ \vdash\ \langle \textsf{snd}\ p\,, \textsf{fst}\ p \rangle : B \wedge A}\ \textsf{M}{\wedge}\textsf{I}}{\vdash\ \lambda p.\, \langle \textsf{snd}\ p\,, \textsf{fst}\ p \rangle : (A \wedge B) \Rightarrow (B \wedge A)}\ \textsf{M}{\Rightarrow}\textsf{I}$$

The λ term we've constructed is of course the familiar *flip* function. Some other familiar function definitions are also well known in logic. Consider, for instance, composition ($\circ$):

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\Gamma \ \vdash\ f : B \Rightarrow A}\ \textsf{hyp} \qquad \dfrac{\dfrac{}{\Gamma \ \vdash\ g : A \Rightarrow B}\ \textsf{hyp} \qquad \dfrac{}{\Gamma \ \vdash\ a : A}\ \textsf{hyp}}{\Gamma \ \vdash\ g\ a : B}\ \textsf{M}{\Rightarrow}\textsf{E}}{\Gamma \ \vdash\ f\ (g\ a) : C}\ \textsf{M}{\Rightarrow}\textsf{E}}{f : B \Rightarrow A,\ g : A \Rightarrow B \ \vdash\ \lambda a.\, f\ (g\ a) : A \Rightarrow C}\ \textsf{M}{\Rightarrow}\textsf{I}}{f : B \Rightarrow A \ \vdash\ \lambda g.\, \lambda a.\, f\ (g\ a) : (A \Rightarrow B) \Rightarrow A \Rightarrow C}\ \textsf{M}{\Rightarrow}\textsf{I}}{\vdash\ \lambda f.\, \lambda g.\, \lambda a.\, f\ (g\ a) : (B \Rightarrow C) \Rightarrow (A \Rightarrow B) \Rightarrow A \Rightarrow C}\ \textsf{M}{\Rightarrow}\textsf{I}$$

where $\Gamma = f : B \Rightarrow C, g : A \Rightarrow B, a : A$. In Natural Deduction, this proof is

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{B \Rightarrow C}\ f \qquad \dfrac{\dfrac{}{A \Rightarrow B}\ g \qquad \dfrac{}{A}\ a}{B}\ {\Rightarrow}\textsf{E}}{C}\ {\Rightarrow}\textsf{E}}{A \Rightarrow C}\ {\Rightarrow}\textsf{I}_a}{(A \Rightarrow B) \Rightarrow A \Rightarrow C}\ {\Rightarrow}\textsf{I}_g}{(B \Rightarrow C) \Rightarrow (A \Rightarrow B) \Rightarrow A \Rightarrow C}\ {\Rightarrow}\textsf{I}_f$$

Another familiar function is the S combinator:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\Gamma\ \vdash\ f : P \Rightarrow A \Rightarrow B}\,\mathsf{hyp} \quad \dfrac{}{\Gamma\ \vdash\ p : P}\,\mathsf{hyp}}{\Gamma\ \vdash\ f\ p : A \Rightarrow B}\,\mathsf{M}{\Rightarrow}\mathsf{E} \quad \dfrac{\dfrac{}{\Gamma\ \vdash\ a : P \Rightarrow A}\,\mathsf{hyp} \quad \dfrac{}{\Gamma\ \vdash\ p : P}\,\mathsf{hyp}}{\Gamma\ \vdash\ a\ p : A}\,\mathsf{M}{\Rightarrow}\mathsf{E}}{\Gamma\ \vdash\ f\ p\ (a\ p) : B}\,\mathsf{M}{\Rightarrow}\mathsf{E}}{f : P \Rightarrow A \Rightarrow B,\ a : P \Rightarrow A\ \vdash\ \lambda p. f\ p\ (a\ p) : P \Rightarrow B}\,\mathsf{M}{\Rightarrow}\mathsf{I}}{f : P \Rightarrow A \Rightarrow B\ \vdash\ \lambda a. \lambda p. f\ p\ (a\ p) : (P \Rightarrow A) \Rightarrow P \Rightarrow B}\,\mathsf{M}{\Rightarrow}\mathsf{I}}{\vdash\ \lambda f. \lambda a. \lambda p. f\ p\ (a\ p) : (P \Rightarrow A \Rightarrow B) \Rightarrow (P \Rightarrow A) \Rightarrow P \Rightarrow B}\,\mathsf{M}{\Rightarrow}\mathsf{I}$$

where $\Gamma = f : P \Rightarrow A \Rightarrow B, a : P \Rightarrow A, p : P$. In Natural Deduction:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{P \Rightarrow A \Rightarrow B}\,f \quad \dfrac{}{P}\,p}{A \Rightarrow B}\,{\Rightarrow}\mathsf{E} \quad \dfrac{\dfrac{}{P \Rightarrow A}\,a \quad \dfrac{}{P}\,p}{A}\,{\Rightarrow}\mathsf{E}}{\dfrac{B}{P \Rightarrow B}\,{\Rightarrow}\mathsf{I}_p}\,{\Rightarrow}\mathsf{E}}{(P \Rightarrow A) \Rightarrow P \Rightarrow B}\,{\Rightarrow}\mathsf{I}_a}{(P \Rightarrow A \Rightarrow B) \Rightarrow (P \Rightarrow A) \Rightarrow P \Rightarrow B}\,{\Rightarrow}\mathsf{I}_f$$

We can see from the various examples and the nature of the definitions that the process of constructing a proof in Natural Deduction is nothing more than the process of checking that the λ term representation of the proof type checks. The type checking λ terms of type P with free variables Γ are precisely the Natural Deduction proofs from hypotheses in Γ that P is true. In a sense, this meta-proof system is a demonstration that the Curry-Howard correspondence isn't so much a correspondence as it is a recognition that proofs and propositions on the one hand, and terms and types on the other, *were secretly the same thing all along.*

3 Proof Simplification and Reduction Rules

What we now want to consider is how to get from merely having proof objects as λ terms to having computation. The λ calculus, after all, has reduction rules — the simply typed λ calculus given above should have five computational rules, one for each of the constructors and eliminators of $\wedge$, $\vee$ and $\Rightarrow$. The $\top$ constructor and $\bot$ eliminator will be left alone.

In the usual λ calculus terms, we have the following reductions:

$$\mathsf{fst}\ \langle \mathcal{A} , \mathcal{B} \rangle \quad \rightsquigarrow \quad \mathcal{A}$$

$$\mathsf{snd}\ \langle \mathcal{A} , \mathcal{B} \rangle \quad \rightsquigarrow \quad \mathcal{B}$$

$$\mathsf{case}\ (\mathsf{left}\ \mathcal{A})\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathcal{C}\ ;\ \mathsf{right}\ b \mapsto \mathcal{C}'\ \} \quad \rightsquigarrow \quad \mathcal{C}[\mathcal{A}/a]$$

$$\mathsf{case}\ (\mathsf{right}\ \mathcal{B})\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathcal{C}\ ;\ \mathsf{right}\ b \mapsto \mathcal{C}'\ \} \quad \rightsquigarrow \quad \mathcal{C}'[\mathcal{B}/b]$$

$$(\lambda a .\, \mathcal{B})\mathcal{A} \quad \rightsquigarrow \quad \mathcal{B}[\mathcal{A}/a]$$

These reduction rules come quite directly from proof simplification/normalization rules. Let's consider some Natural Deduction proofs for each case. Starting with the fst equation, we derive this by consider the equivalence of proofs

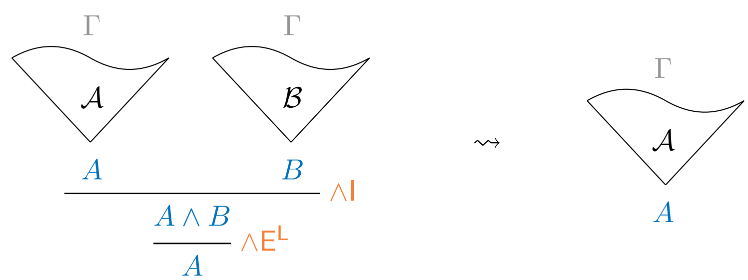

This makes a lot of sense, because on the left, we just use $\mathcal{A} : A$ to build a proof of $A \wedge B$, only to immediately eliminate this back to a proof of A, so why not just sit on A? Similarly for snd:

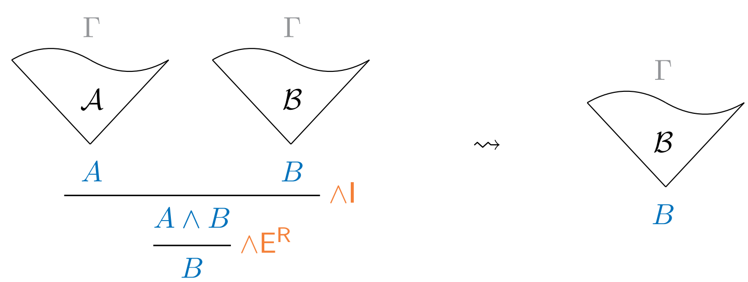

Using the textual representation, we have

$$\Gamma \ \vdash\ \mathsf{fst}\ \langle \mathcal{A} , \mathcal{B} \rangle : A \quad \rightsquigarrow \quad \Gamma \vdash \mathcal{A} : A$$

$$\Gamma \ \vdash\ \mathsf{snd}\ \langle \mathcal{A} , \mathcal{B} \rangle : B \quad \rightsquigarrow \quad \Gamma \vdash \mathcal{B} : B$$

Which, other than some noise, are the reductions we had earlier! Similarly for disjunction:

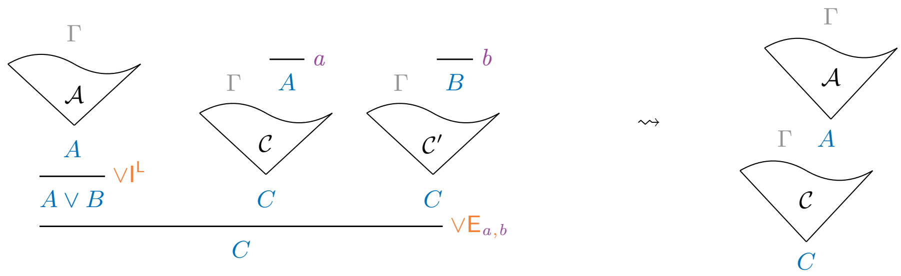

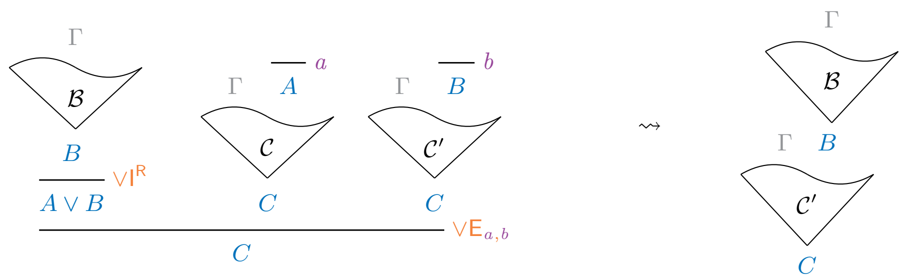

In textual notation:

$$\Gamma \vdash \mathsf{case}\ (\mathsf{left}\ \mathcal{A})\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathcal{C}\ ;\ \mathsf{right}\ b \mapsto \mathcal{C}'\ \} : C \quad \rightsquigarrow \quad \Gamma \vdash \mathcal{C}[\mathcal{A}/a] : C$$

$$\Gamma \vdash \mathsf{case}\ (\mathsf{right}\ \mathcal{B})\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathcal{C}\ ;\ \mathsf{right}\ b \mapsto \mathcal{C}'\ \} : C \quad \rightsquigarrow \quad \Gamma \vdash \mathcal{C}'[\mathcal{B}/b] : C$$

And lastly, for implication:

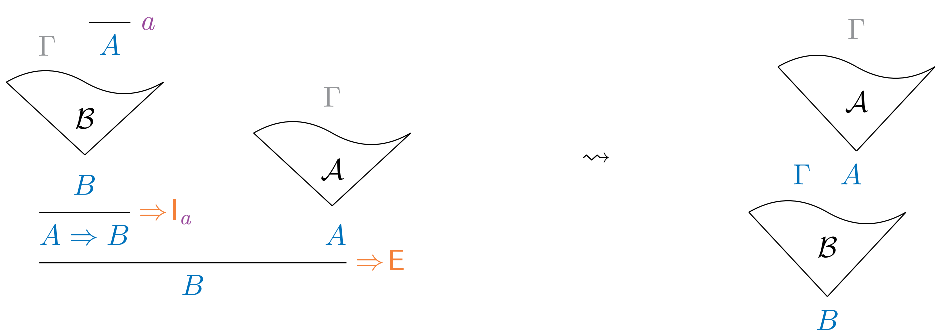

In textual notation:

$$\Gamma \vdash (\lambda a.\, \mathcal{B})\ \mathcal{A} : B \quad \rightsquigarrow \quad \Gamma \vdash \mathcal{B}[\mathcal{A}/a] : B$$

4 Harmony

As we've seen, the proofs in Natural Deduction give rise directly to the type theory of the simply typed λ calculus, and the proof simplification rules give rise directly to the reductions. The proof simplification rules are especially interesting in this regard, because they are precisely the reductions that demonstrate local soundness [LNJP] of the inference rules involved. Some other well known conversions/equivalences from the λ calculus — the η equivalences — are "complication" rules which demonstrate local completeness. For instance, η conversion $\mathcal{F} \leftrightsquigarrow \lambda a.\, \mathcal{F}\, a$ when a does not appear free in $\mathcal{F}$ is

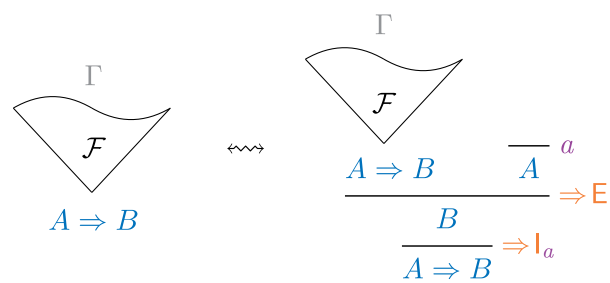

As typing statements this is

$$\Gamma \vdash \mathcal{F} : A \Rightarrow B \qquad \leftrightsquigarrow \qquad \Gamma \vdash \lambda a.\, \mathcal{F}\, a : A \Rightarrow B$$

In the forward direction this is local completeness, in the backwards direction it's simplification. The identity on pairs $\mathcal{P} = \langle \mathsf{fst}\ \mathcal{P}, \mathsf{snd}\ \mathcal{P} \rangle$ is similarly a combination of local completeness and simplification:

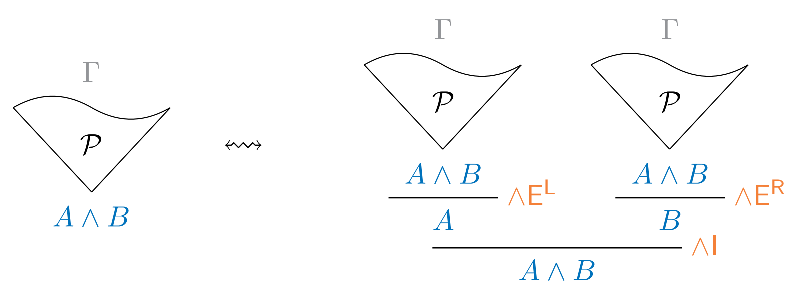

$$\Gamma \vdash \mathcal{P} : A \wedge B \qquad \leftrightsquigarrow \qquad \Gamma \vdash \langle \mathsf{fst}\ \mathcal{P}, \mathsf{snd}\ \mathcal{P} \rangle : A \wedge B$$

And finally, for disjunction we have

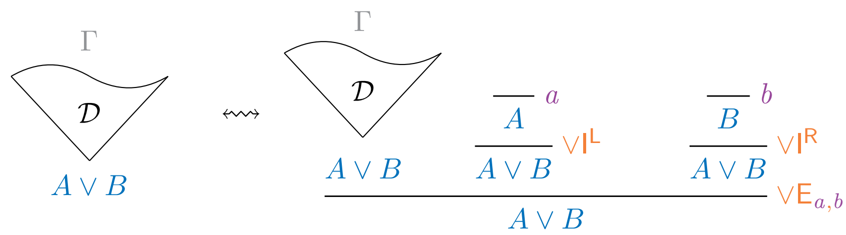

$$\Gamma \vdash \mathcal{D} : A \vee B \quad \rightsquigarrow \quad \Gamma \vdash \mathsf{case}\ \mathcal{D}\ \mathsf{of}\ \{\ \mathsf{left}\ a \mapsto \mathsf{left}\ a\ ;\ \mathsf{right}\ b \mapsto \mathsf{right}\ b\ \} : A \vee B$$

The β and η rules for $\top$ and $\bot$ are a little more subtle in their justification. See Pfenning's notes [LNJP] for an explanation, and an excellent discussion of the concepts of local soundness and local completeness, which together constitute the concept of logical harmony.

5 Bibliography

[GEN] Gentzen, Gerhard. 1934. *Untersuchungen ber das logische Schlieen.* Mathematische Zeitschrift, 39(2): 176–210.

[ITT] Martin-Löf, Per. 1984. *Intuitionistic Type Theory*.

[LNJP] Pfenning, Frank. 2010. *Lecture Notes on Judgments and Propositions*.