Appendix: Soundness and Completeness

In this Appendix, we outline soundness proofs for our PL and QL natural deduction systems, and then go on to outline completeness proofs for both systems. (Extending the proofs to cover the $QL^=$ system doesn't involve any particularly exciting new ideas; so we won't do that here.)

Two comments. First, we will indeed only give *outlines*, explaining the Big Ideas. There is nothing to be gained in an introductory book by giving all the details; doing that would just obscure the overall shapes of the arguments.

Second, the proofs don't presuppose any specific knowledge of mathematics. But they *are*, necessarily, a degree or two more abstract and mathematical in flavour than most proofs earlier in the book. Still, allowing for their abstractness, they are actually quite approachable. Just take things slowly!

A1 Soundness for PL

(a) Let's restate a general definition we've seen before, in §32.3:

> The *live assumptions* at a line of a Fitch-style proof are (i) the wff at that line, if it is a premiss or temporary supposition, plus (ii) any earlier premisses and undischarged (i.e. still available) suppositions.

And now here's another version of a definition we met when first discussing PL soundness in §24.6(c):

> A line in a PL proof is *good* if and only if the wff on that line (if there is one) *is* tautologically entailed by the live assumptions at that line.

However, at the *last* line of a completed proof, no suppositions will still be available. So to say that the last line of a proof is good is just to say that the wff on that line – i.e. the conclusion of the proof – is tautologically entailed by the premisses. Hence the claim that our PL proof system is sound is equivalent to the claim that the *last* line of any complete proof is good.

(b) We will now argue as follows:

(1) The *first* line of any PL proof is good.
(2) If each line in a PL proof before line n is good, so is line n.
So (3) *Every* line of a PL proof must be good.
So (4) The *last* line of a PL proof must be good.

(1) is trivial: if there is a wff on the first line of a proof, it must be a premiss and it is tautologically entailed by the premiss at that line. Next, we of course chose our proof-building rules precisely because they look intuitively truth-preserving: so applying a rule to add another line to a proof ought to take us from some previous good line(s) to another good line. Hence (2) *ought* to be true too. But with the premisses (1) and (2) in place, we know the first line of a proof is good, so the second line is also good, so the third line is also good, and so on. Hence the conclusion (3) follows, giving us the soundness theorem in the form (4).

To fill out this argument, we just(!) have to check the claim underlying (2), i.e. we need to check that applications of a PL proof-building rule always preserves goodness when we add a new line. Which is not difficult, just tedious.

For how can we extend a proof? Let us count the ways. We can add a new premiss or a new temporary supposition. But obviously doing *that* gives us a new good line – for the wff we've added will be entailed by the live assumptions including the newly added one. Or else we can apply one of the eleven rules of inference (iteration, ex falso quodlibet, four pairs of introduction/elimination rules, plus the outlying double negation rule). And now there is nothing for it but to hack through all the different ways of extending a proof, to show that each does indeed keep our proof virtuous.

(c) Before looking at this task, we need a couple of useful facts.

Ask yourself: at a given line, which assumptions (premisses and suppositions) *are* live in the sense defined? Well, start from the given line and walk up the current column. If and when you get to another temporary assumption, step one column left. Go up again. If and when you get to a temporary assumption, step another column left. Keep on going until you find yourself in the home column and proceed to the top of the proof. Then every premiss or supposition you pass en route counts as a live assumption. (Draw a skeletal diagram of a Fitch-style derivation or two, and think about it!)

This implies our two useful facts:

> If the wff at the earlier line j is still available at the later line k, all the assumptions that are live at line j are still live at line k.
>
> If the subproof starting with a temporary supposition at the earlier line i and finishing at line j is still available at the later line k, then all the assumptions that are live at line j except for the supposition at line i are still live at line k.

The first is obvious; the second only takes a little more thought (draw more skeletal diagrams!).

(d) With those useful facts to hand, let's consider just two ways of extending a proof by using a rule of inference. We'll take one 'first level' rule where we infer a wff from two previous available wffs in the proof, and one 'second level' rule where we use an available subproof to infer a new wff.

> (MP) Suppose we have α at the good line i, and $(\alpha \to \beta)$ at the good line j, and suppose both these wffs are still available at the later line k. Because

those earlier lines are good, the live assumptions at i tautologically entail α, and the live assumptions at j tautologically entail $(\alpha \to \beta)$. Combine those earlier live assumptions, and those pooled assumptions must then tautologically entail β. But by our first useful fact, those pooled assumptions are still among the live assumptions at line k. So if we use (MP) to infer β at line k, this will indeed give us another good line, with β entailed by the live assumptions.

(RAA) Suppose we have a subproof starting with α and finishing with $\bot$ at the good line j. And suppose this subproof is still available at line k. By the goodness assumption, the live assumptions at line j, which include α, tautologically entail $\bot$. So the live assumptions at line j apart from α tautologically entail $\neg\alpha$. But by the second useful fact, the live assumptions at line k will be the live assumptions at line j apart from α. So if we use (RAA) to infer $\neg\alpha$ at line k, this will indeed give us another good line.

Another nine arguments like this (exercise: find them!) will complete the demonstration that (2), *every* legitimate way of extending a proof whose lines are good up to now gives us a proof with one more good line. Then we are done.

A2 Soundness for QL

(a) We prove soundness for QL by using an exactly parallel argument. We start by re-defining the appropriate notion of goodness:

A line of a QL proof is *good* if and only if the wff at that line (if there is one) is true on every relevantly expanded q-valuation which makes the live assumptions at that line true.

By a 'relevantly expanded' q-valuation we mean a q-valuation (of the relevant language) that is expanded to assign values to any dummy names that appear in the wff and/or in any assumption which is live at that line.

As with the argument for PL's soundness, we can then argue that every line of a properly constructed QL proof is good, and hence that the last line is good, which establishes QL's soundness. And how do we show that every line of a proof is QL good? Just as before, by noting (1) that a proof (trivially) must start with a good line. And then showing that (2) every legitimate way of extending a QL proof whose lines are good up to now gives us a proof with one more good line.

(b) Again showing (2) is not difficult, though tedious. We just have to go through all the different ways of extending proofs. We will consider one of the new quantifier rules of inference:

($\forall$I) Suppose that we have a wff $\alpha(\delta)$ at the good line j. Suppose the live assumptions at that line are Γ. And to keep things simple to start with, suppose the wffs Γ and the wff $\alpha(\delta)$ involve *no* dummy names except for the δ occurring in $\alpha(\delta)$.

Take some q-valuation q. Since by assumption line j is good, any expansion q_δ which makes Γ true makes $\alpha(\delta)$ true too. But since δ doesn't appear in the wffs Γ, the assignment of a value to that dummy name can't affect the truth values of those wffs. Hence an expanded valuation q_δ makes Γ all true if and only if the original unexpanded q makes Γ all true. So the goodness of line j means: for any q which makes Γ all true, any expansion q_δ will make $\alpha(\delta)$ true, and therefore (by the rule for evaluating universally quantified wffs) q will make $\forall\xi\alpha(\xi)$ true too!

Hence, if we move to a line which still has Γ as live assumptions, and where the wff on that line is $\forall\xi\alpha(\xi)$, this line will be a good one.

We can now easily generalize this line of thought to allow for cases where some other dummy names are being carried along for the ride. And we then get a demonstration that an application of ($\forall$I) universally generalizing on a dummy name which doesn't appear in the live assumptions for a line takes us from a good line to another good line!

We can give similar arguments that the other three ways of extending proofs by applying quantifier rules of inference are virtue-preserving. We also need to check that the eleven propositional rules shared with PL are still virtue-preserving. Dotting every 'i' and crossing every 't' and putting everything together, will take some rather tedious pages. But we can excuse ourselves from the further work. Rounding out our soundness proof for QL can be left as an enterprise for masochistic enthusiasts.

A3 Completeness: what we want to prove

Here are a couple of definitions, from §15.6 and Exercises 22(d*):

The PL wffs Γ are *tautologically consistent* if and only if there is a valuation of the relevant atoms which makes the wffs all true together.

The PL wffs Γ are PL-*consistent* if and only if there no PL proof of $\bot$ using (some of) those wffs as premisses.

And now note the following:

(1) $\Gamma \vDash \gamma$ if and only if $\Gamma, \neg\gamma$ are tautologically inconsistent.

(2) $\Gamma \vdash \gamma$ if and only if $\Gamma, \neg\gamma$ are PL-inconsistent.

Both these should by now look obvious. For the first, see §15.6 again. For the second, note that if Γ proves γ, then Γ plus $\neg\gamma$ will yield absurdity. And if you can get from Γ plus $\neg\gamma$ to absurdity, then you can use a reductio proof to get from Γ to $\neg\neg\gamma$ and hence γ.

So take the PL completeness theorem

(CP) If $\Gamma \vDash \gamma$, then $\Gamma \vdash \gamma$

and use (1) and (2) and contrapose to get the equivalent claim

If $\Gamma, \neg\gamma$ are PL-consistent, then $\Gamma, \neg\gamma$ are tautologically consistent.

But now note that *this* is just a special case of the slightly more general claim that, for any wffs Δ,

(*) If Δ are PL-consistent, then Δ are tautologically consistent.

So in the next section we set out to prove (*), and thereby secure our completeness theorem (CP) for PL. (And similarly, in §A5, we will use an analogue of (*) to secure the completeness theorem for QL: more about that shortly.)

A4 PL completeness proved

(a) We will continue to take the relevant Γ (and so Δ) to be finitely many in number, as first announced in §24.5. Our proof for (*) will take two stages:

(S) We show that the PL-consistent wffs Δ can be beefed up into a bigger collection Δ^+ which is PL-consistent and *saturated*.

(V) We show that there is always a valuation which makes the wffs in a PL-consistent and saturated collection all true.

By (V), then, there is a valuation which makes all the wffs Δ^+ in (S) true. But this valuation will of course make all of the smaller original PL-consistent collection Δ true together. Which proves (*).

(b) So what do we mean by saying that a collection of wffs is saturated? We mean it is *saturated with truth-makers*. And what do we mean by *that*? We mean that for every complex wff in the collection (every wff that isn't an atom or negated atom) there are one or two simpler wffs also in the collection such that if the simpler ones are true, so is the more complex one. Thus:

A collection of wffs Γ is *saturated* if it satisfies the following seven conditions, for any wffs α, β:

(i) If $\neg\neg\alpha$ is one of the wffs Γ, so is α;
(ii) if $(\alpha \wedge \beta)$ is one of the wffs Γ, so are α and β;
(iii) if $\neg(\alpha \wedge \beta)$ is one of the wffs Γ, so is at least one of $\neg\alpha$ and $\neg\beta$;
(iv) if $(\alpha \vee \beta)$ is one of the wffs Γ, so is at least one of α and β;
(v) if $\neg(\alpha \vee \beta)$ is one of the wffs Γ, so are $\neg\alpha$ and $\neg\beta$;
(vi) if $(\alpha \to \beta)$ is one of the wffs Γ, so is at least one of $\neg\alpha$ and β;
(vii) if $\neg(\alpha \to \beta)$ is one of the wffs Γ, so are α and $\neg\beta$.

By inspection, we see in each case that, if the simpler wff(s) on the right of the conditional are true, so is the complex wff on the left. And note that these seven clauses cover every form of wff that isn't an atom or negated atom.

It is more or less immediate from our cunning definition that (V) is true. Suppose wffs Γ are PL-consistent and saturated. Since they are PL-consistent, they can't contain an atom and its negation, so we can consistently choose the valuation which makes all the naked atoms true and makes any other relevant atoms false (and so makes any negated atoms true). This chosen valuation therefore makes the simplest wffs in Γ all true.

We can easily assign a numerical degrees of complexity on which a wff always counts as more complex than its truth-makers (as given by (i) to (vii)). Now look at the wffs in Γ in order of increasing complexity. So after the atoms and negated atoms, the next simplest wffs in Γ are all true on the chosen valuation – since by saturation these wffs have simple truth-makers in Γ, and these simple truth-makers are true on the chosen valuation. Then, going up another level of complexity, the next simplest wffs in Γ are also all true on this valuation – since by saturation these wffs too have simpler truth-makers in Γ, and we've just seen that those are true on the chosen valuation. Keep on going. Truth on the chosen valuation percolates upward through more and more complex wffs in Γ, making them all true on the chosen valuation.

(c) Now to prove (S). Recall the results of Exercises 20(b*), 21(b*) and 22(c*). There are seven of them, neatly matching (i) to (vii) in our definition of saturation. And each result has the same shape: it tells us that *if some wffs are PL-consistent and contain a complex wff of such-and-such form, then we can always add truth-makers for that wff while keeping our collection of wffs PL-consistent!*

So given the PL-consistent wffs Δ, arrange them in some order. Imagine a logical demon walking along the list of wffs.

(1) Every time they encounter a wff of kinds (i), (ii), (v) or (vii), they add the appropriate truth-maker(s) to the end of the list, which we know that they can do while maintaining PL-consistency.

(2) Every time they encounter a wff of kinds (iii), (iv) and (vi), they have to choose a truth-maker to add while maintaining PL-consistency; but again we know that they always *can* do this. Imagine that they use their demonic powers so as always to jump the right way!

The demon's walk will eventually terminate if the initial Δ are finitely many, as the wffs they add at each step will in the end get shorter and shorter. Eventually every complex wff will have had truth-makers added to the list, and these truth-makers (if complex) will have had their truth-makers added, and so on. So the demon will arrive at the desired PL-consistent but fully saturated set Δ^+ where there are no more truth-makers which we need to add. Which establishes (S).

Of course, the demon is just for vividness. The point is the existence of a suitable Δ^+, not its imagined clever constructor! We can put the same argument more abstractly. And – for the record – the argument can be further tweaked to cover the case where we start with an infinite initial collection of wffs Δ. But we needn't pause over that.

A5 QL completeness proved

(a) The QL completeness theorem tells us that

(CQ) If $\Gamma \vDash \gamma$, then $\Gamma \vdash \gamma$,

where Γ are some QL sentences, finitely many. For reasons exactly parallel to those given in §A3, we can prove (CQ) by showing

(**) If Δ are QL-consistent, then Δ are q-consistent,

where again Δ are finitely many QL sentences, and some wffs are QL-consistent just if there is no QL proof of absurdity from them. And the two-stage proof of

(**) is parallel to our proof of (*):

(qS) We show that the QL-consistent sentences Δ can be beefed up into a bigger, perhaps infinite, collection Δ^+ of wffs which is QL-consistent and *q-saturated*.

(qV) We show that there is always an expanded q-valuation which makes the wffs in a QL-consistent and q-saturated collection all true.

By (qV), there is an expanded q-valuation q^+ which makes all the wffs Δ^+ in (qS) true (expanded, because as we will see, Δ^+ will involve dummy names). But q^+ will of course make all of the smaller original QL-consistent collection Δ true. By assumption Δ only contains sentences, so q^+'s assignments of values to dummy names which don't occur in Δ are irrelevant. So ignoring those assignments, we are left with a q-valuation q which makes all of Δ true. Which proves (**).

That, at any rate, is our strategy. We need to start with a definition of q-saturation.

(b) Say that a collection of wffs Γ from a QL language is q-saturated if it satisfies the conditions (i) to (vii) for being saturated, and it also satisfies the following conditions:

(viii) If $\forall\xi\alpha(\xi)$ is one of the wffs Γ, so is $\alpha(\tau)$ for any term τ that appears in any wff from Γ (and there is at least one such term);

(ix) If $\exists\xi\alpha(\xi)$ is one of the wffs Γ, so is $\alpha(\tau)$ for some term τ;

(x) If $\neg\forall\xi\alpha(\xi)$ is one of the wffs Γ, so is the corresponding $\exists\xi\neg\alpha(\xi)$;

(xi) If $\neg\exists\xi\alpha(\xi)$ is one of the wffs Γ, so is the corresponding $\forall\xi\neg\alpha(\xi)$.

In cases (ix) to (xi), if the wff given on the right of the conditional is true, so is the original wff on the left; so again these three clauses give truth-makers. In the special case of (viii), the truth of all the added wffs given on the right will make the quantified wff on the left true if (but only if) the domain contains no more objects than those named by the terms appearing in Γ.

So next we need to show (qV) is true. We use the same basic trick as before, i.e. choose an expanded q-valuation which makes the atomic wffs and negated atomic wffs in a QL-consistent collection of wffs Γ all true. And then show that if Γ are q-saturated, all of them must be true on this chosen q-valuation. How, then, we do choose the needed expanded q-valuation?

(c) Assume Γ are QL-consistent and q-saturated. Remember there are only a finite number (maybe zero) proper names in any QL language, and hence in some wffs Γ constructed in that language. So now proceed as follows:

(1) Take all the terms that appear in Γ, and line them up, proper names first $\tau_0, \tau_1, \ldots, \tau_k$ (finitely many), followed by the dummy names τ_{k+1}, τ_{k+2}, $\tau_{k+3}, \ldots$ (perhaps unlimitedly many of them). (Fine print: if there are no terms at all occurring in Γ, that can only be because there are no quantified wffs in Γ, so we would in effect be back dealing with propositional logic – we can safely ignore that boring case.)

(2) Set the domain for our chosen q-valuation to be the natural numbers

from zero up, one for each of the terms in Γ (so perhaps that's *all* the natural numbers).

(3) Let the reference of the term τ_j be the number j. So note, we *are* ensuring that we are dealing with special case where every object in the domain is the reference of some term on our chosen q-valuation.

(4) And now for the clever trick: suppose φ is a unary predicate. Say that the number j is in the extension of φ if and only if the atomic wff $\varphi\tau_j$ is in Γ! Then that makes all the unnegated wffs of the form $\varphi\tau_j$ in Γ true. And if $\neg\varphi\tau_j$ is in Γ, then by QL-consistency $\varphi\tau_j$ is *not* in Γ, so j is not in the extension of φ, and hence $\varphi\tau_j$ is false and $\neg\varphi\tau_j$ is true. In short, all atoms and negated atoms formed from unary predicates will come out true on our chosen q-valuation.

(5) Say that the pair of numbers $\langle i, j\rangle$ is in the extension of the binary predicate ψ if and only if the atomic wff $\psi\tau_i\tau_j$ is in Γ! Then, by the same reasoning as is in the unary case, all atoms and negated atoms formed from binary predicates will come out true on our chosen q-valuation. Handle predicates of other arities, including zero, similarly.

So, we have arrived at our chosen q-valuation which makes the simplest wffs – i.e. the atomic wffs and negated atomic wffs – in Γ all true together.

Now we just need to show that, as in the PL case, a valuation which makes the simplest wffs in a q-saturated collection of wffs true makes all of them true.

By definition, a q-saturated set, every complex wff has truth-makers in the set. In particular, note that if $\forall\xi\alpha(\xi)$ is in Γ, then so is $\alpha(\tau)$ for every term τ in Γ; but then if those simpler wffs are all true on the chosen q-valuation, then since each number in the domain is named by one of those terms, $\forall\xi\alpha(\xi)$ is true on that valuation too. So, as in the PL case, truth will percolate up from the simplest truth-makers to more and more complex wffs, eventually making *all* the wffs in Γ true. Which establishes (qV) – and moreover, we see that to find a q-valuation which makes all the wffs in the QL-consistent q-saturated set all true we need look no further than valuations constructed out of natural numbers.

(d) It remains to prove (qS). Again, we use the same basic idea as in the proof of (S). So given the QL-consistent wffs Δ, finitely many, arrange them in some order. Imagine again a logical demon walking along the list of wffs. Then:

(1) Every time they encounter a wff of kinds (i), (ii), (v), (vii), (x) or (xii), they add to the end of the list the appropriate truth-maker(s), which we know they can do while maintaining QL-consistency (obvious for the new last two cases).

(2) Every time they encounter a wff of kinds (iii), (iv) and (vi), they have to choose a truth-maker to add, but again we know that they always *can* do this while maintaining QL-consistency. Imagine that they use their demonic powers to jump the right way!

(3) When they encounter a wff of the kind (viii), a universal quantification $\forall\xi\alpha(\xi)$, they add every instance $\alpha(\tau)$ where τ is any term that appears

somewhere in the wffs in their list so far (or if no term appears yet, then use the first dummy name). This is easily seen to maintain QL-consistency, as we in effect noted in Exercises 33(d*).

(4) When they encounter a wff of the kind (ix), an existential quantification $\exists\xi\alpha(\xi)$, they add the instance $\alpha(\tau)$ where τ is the first dummy name that doesn't appear anywhere in the list of wffs so far. Then, with this new term in play, they revisit every universally quantified wff encountered so far, and now also instantiate it with this new term. Again, doing all this maintains QL-consistency, which follows from another point noted in Exercises 33(d*).

And the demon keeps on going.

But note that this time – in contrast with the PL case – the demon's task can be an infinite one, even if they start from a finite collection Δ. Why? Take the case where the original Δ_0 comprises just the single sentence ∀x∃yLxy. First following the instruction (3) and then repeatedly following (4) as new existential quantifications are popped onto the end of the list, the demon's list will start growing in the following stages:

Δ_0: ∀x∃yLxy
Δ_1: ∀x∃yLxy, ∃yLay
Δ_2: ∀x∃yLxy, ∃yLay, Lab, ∃yLby
Δ_3: ∀x∃yLxy, ∃yLay, Lab, ∃yLby, Lbc, ∃yLcy
Δ_4: ...

and there is no end to the process. What to do?

There is a standard mathematical trick we can use. Our logical demon who follows the instructions, starting from a given initial (finite) collection Δ, will build up a perhaps infinite sequence of (still finite) collections Δ_0, Δ_1, Δ_2, ..., each one extending the previous one, and all QL-consistent. *We can now define our desired Δ^+ to be the union of all these Δ_j* (so a wff is in Δ^+ if and only if it is in some Δ_j). Though it is rather more fun to imagine a demon who can go faster and faster and so complete the infinite task of building Δ^+!

So the QL-consistent Δ gets beefed up to a collection Δ^+ which is q-saturated. (Why? Because every complex wff acquires truth-makers along the way). And Δ^+ is still QL-consistent. (Why? Because if it were not, there would be a proof of absurdity using finitely many premisses in Δ^+ – finitely many, since proofs are finite. Hence all these premisses would be in some finite Δ_j: but that would contradict the QL-consistency of Δ_j.) Which establishes (qS).

For the record, this argument too can be tweaked to cover the case where we start with an infinite initial collection of wffs Δ. We needn't pause over that.

A6 A squeezing argument

(a) Let's say that a q-valuation whose domain is some or all of the natural numbers (integers from zero up) is a *tame* valuation. We showed in the last section that, if some wffs are QL-consistent, then there is not just a *q*-valuation which makes them all true, but more particularly there is a *tame* valuation which does the job.

Let's introduce a third turnstile symbol to capture this point in a slightly different way:

> We will use $\Gamma \Vvdash \gamma$ to abbreviate the claim that there is no tame valuation which makes all Γ true and makes γ false.

So what we showed in the last section amounts to this (why?):

(1) If $\Gamma \Vvdash \gamma$ then $\Gamma \vdash \gamma$.

But the soundness theorem

(2) If $\Gamma \vdash \gamma$ then $\Gamma \vDash \gamma$

seems secure. Our rules of inference are absolutely compelling whatever objects we are talking about (however wild the domain might be). And we certainly have

(3) If $\Gamma \vDash \gamma$ then $\Gamma \Vvdash \gamma$.

For if there is no q-valuation at all which makes all Γ true and makes γ false, then there can't in particular be a tame valuation tidily built from the natural numbers which makes Γ true and γ false.

(b) We offered the notion of q-validity as a rational reconstruction of the idea of logical validity for QL arguments. But in §36.2 we worried whether the idea of quantifying over *all* possible q-valuations makes sufficiently clear sense: could it make a difference just which wildly infinite domains for valuations we are prepared to countenance?

We can now see that we needn't worry – a point noted by the logician Georg Kreisel. Going round the circle of implications in (1), (2) and (3), the relations symbolized by the three turnstiles must be equivalent: if any one holds between Γ and γ, so do the other two. So yes, our original definition of q-validity was perhaps somewhat slack; but it characterizes a relation $\vDash$ which is 'squeezed' between the relations $\Vvdash$ and $\vdash$. But those two sharply defined relations are provably equivalent; one holds if and only if the other does. So there is no wriggle room: despite the apparent slackness in our initial less formal definition, that definition is in fact rigorous enough, it says enough to ensure that the relation $\vDash$ has a sharp extension, the same as the extension of $\Vvdash$ and $\vdash$. Which is a relief.