Constructive Logic for All

Greg Restall*
Philosophy Department
Macquarie University

June 14, 2000

Abstract

It is a commonplace in recent metaphysics that one's logical commitments go hand in hand with one's metaphysics. Brouwer, Heyting and Dummett have each championed the move to constructive (intuitionistic) reasoning on the grounds of anti-realism. I hope to break this close connection, to explain why a *realist* ought to reason constructively.

1 Introduction

Let me start by explaining the terms of the discussion. In this area of philosophical logic there seems to be some confusion in the use of terms such as "intuitionistic" and "constructive." It will help to get the use of these terms somewhat fixed before I start to argue my case.

1.1 Logic

First, *logic*. The subject *Logic* is the study of logical consequence. Logical consequence is a matter of the validity and invalidity of arguments. An argument is *valid* just when in any case in which the premises of the argument are true, so is the conclusion. It is often helpful, in discussing logical consequence, to have a *formal language* in which to express the premises and conclusions of arguments. One such language is the language of first-order logic, with primitive expressions for the connectives and quantifiers

$\wedge$	conjunction
$\vee$	disjunction
$\supset$	conditional
$\sim$	negation
$\forall x$	universal quantifier (with x a variable)
$\exists x$	existential quantifier (with x a variable)

For some, logic is at heart the study of the behaviour of these connectives. For others, the singling out of particular parts of the language is an incidental matter. All agree that *formal* logic is about the *forms* of arguments and inferences,

*Research supported the Australian Research Council, Large Grant No. A00000348.

but many disagree about whether the language of first-order logic completely expresses the forms relevant to logical judgements. For this paper, this argument is incidental. All of our discussion of logical consequence will take place comfortably within the language of first-order logic.

We will write validity for arguments as follows:

$$X \vdash A.$$

This states that the argument from premises X to conclusion A is valid. For example, if a negated disjunction follows from both negated disjuncts, we could say the following:

$$\sim A, \sim B \vdash \sim(A \lor B).$$

The premises are listed to the left of the turnstile, and the conclusion to the right.

1.2 Intuitionistic logic and Constructivity

Intuitionistic logic is a particular account of logical consequence, at variance with classical logical consequence, with which I assume familiarity. One way to introduce intuitionistic logic is by means of *constructions*. We give an account of validity of arguments by indicating what it is to *construct* (for the moment, take this to mean something like demonstrate, prove or establish) a statement.

- A construction of $A \land B$ is a construction of A, together with a construction of B.
- A construction of $A \lor B$ is either a construction of A or a construction of B.
- A construction of $A \supset B$ is a technique for converting constructions of A into constructions of B.
- There is no construction of $\perp$.[1]
- A construction of $\forall x A$ is a rule giving, for any object n, a construction of $A(n)$.
- A construction of $\exists x A$ is an object n together with a construction of $A(n)$.

Then we say that an argument is valid if and only if a construction for the premises provides a construction for the conclusion.[2]

With this in hand we can see some of the distinctive behaviour of intuitionistic logic. For example, the inference of *distribution*

$$\forall x(A \lor B) \vdash \exists x A \lor \forall x B \tag{1}$$

is valid in classical logic, but it need not be valid in intuitionistic logic. For example, it is easy to demonstrate that every string of ten digits in the decimal

[1]We define $\sim A$ as $A \supset \perp$, so a construction of $\sim A$ is a technique for converting constructions of A into absurdity. It shows that there are no constructions of A.

[2]This is certainly not a *formalisation* of intuitionistic logic, as I have said nothing of what it is for a construction to 'provide' another, what it is to have a construction 'together with' another, nor what these constructions are, or how rules can convert constructions into other constructions. It is enough, however, to motivate the distinctive features of intuitionistic logic.

expansion of π is either a string of ten zeros, or it is not. A demonstration of this fact, however, does not give us a construction of the claim that either there is a string of ten zeros in π or every string of ten digits in π is not a string of zeros. Any construction of *this* statement must either prove that there is no string of ten zeros in π or to show where one such string is. The constructive content of $\exists xA \vee \forall xB$ is greater than that of $\forall x(A \vee B)$.

So, some constructions for $\forall x(A \vee B)$ are not constructions of $\exists xA \vee \forall xB$. But these constructions are also not constructions of $\sim(\exists xA \vee \forall xB)$, so some constructions must be incomplete. This leads us to more classically valid inferences which fail in intuitionistic logic. Another inference to fail is

$$\vdash A \vee {\sim}A$$

The law of the excluded middle can be thought of as an inference with no premises. It is valid in intuitionistic logic if *any* construction provides a construction for $A \vee {\sim}A$. This fails, for not all constructions will give a construction of A or one of ${\sim}A$. Constructions may well be incomplete. To think otherwise is to take constructions to provide answers to *all* questions, and this is against the spirit of the enterprise. Constructions provide justification for some things and not for others.

As a result of this failure we have the failure of the inference of *double negation elimination*.

$${\sim}{\sim}A \vdash A$$

We may have a construction of ${\sim}{\sim}A$ without this providing a construction of A, for a construction of ${\sim}{\sim}A$ is a technique for converting constructions of ${\sim}A$ into absurdity, and so, is a technique for converting {techniques for converting constructions of A into absurdity} into absurdity. Double negation elimination fails because even though we may not have constructions of $A \vee {\sim}A$, we can always provide a construction of its double negation

$${\sim}{\sim}(A \vee {\sim}A).$$

It is worthwhile seeing how this works. We wish to provide a technique for converting constructions of ${\sim}(A \vee {\sim}A)$ to absurdity. Suppose we do have a construction of ${\sim}(A \vee {\sim}A)$. That would be a technique for converting a construction of $A \vee {\sim}A$ to absurdity. If we had *that*, then we have a technique for converting constructions of A into absurdity (for any construction of A gives you a construction for $A \vee {\sim}A$). But this is a construction for ${\sim}A$, and so we get a construction of $A \vee {\sim}A$, which we can convert into absurdity. So, we have indeed gone from a construction of ${\sim}(A \vee {\sim}A)$ to absurdity as desired. We have a construction of ${\sim}{\sim}(A \vee {\sim}A)$.[3] But it does not follow that we have a construction of $A \vee {\sim}A$, lest constructions be omniscient.

This is enough to provide a *taste* of intuitionistic logic. We will see more as the argument continues. Intuitionistic logic is motivated by the idea of the *construction*. As a result, intuitionistic reasoning is often described as *constructive* reasoning. An intuitionistically valid argument may be described as a *constructive* argument.

[3] Phew!

1.3 Intuitionism

Intuitionistic logic, therefore, is not inherently wedded to considerations of *intuition*. Intuition*ism* is a philosophical view of the foundations of mathematics, introduced by Brouwer [14] formalised by Heyting [13], and generally applied to philosophy by Dummett [9, 10].[4] For intuitionists, mathematical reasoning is a function of the intuition of the creating subject. Mathematical proofs are correct to the extent that they encode the constructions of a creating mathematical reasoner. To this extent, intuitionism is a variety of constructivism. However, intuitionism goes further than other varieties of constructivism, in maintaining that constructive reasoning is appropriate because of the nature of mathematical entities themselves. The entities are the constructions of the reasoner and furthermore, mathematical entities only have the properties bestowed upon them by their construction.

Intuitionism about mathematical objects is a kind of anti-realism, of a piece with Dummett's semantic anti-realism. Truths about mathematical entities cannot outstrip what we can truly say about them, and these cannot outstrip our capacities to describe those entities. Constructivity seems a fitting partner of this kind of anti-realism.

This much is fairly straightforward, and is true. However, the marriage of anti-realism and constructive reasoning is taken to be much more intimate than what I have expressed so far. Many have taken it that the two go hand in hand everywhere, and are inseparable partners. In this paper I will argue that this is not the case. Anyone, regardless of metaphysical commitments, can and ought reason constructively. Intuitionistic logic is to be freed from its ties to intuitionism and other anti-realist philosophical views.

2 Everyone *can be* a Constructivist

I will start by examining closely a particular instance of mathematical reasoning which will help us understand the interrelationship between constructive considerations and the metaphysics of mathematics. The lessons are more general than simply mathematical, and we will then go on to see how constructive considerations may apply globally, wherever reasoning is applicable. In doing so, this section will show that everyone *can* be constructivist.

2.1 Construction and Proof

Here is a scenario from Shapiro's *Philosophy of Mathematics* [17]. It is a fairly simple piece of reasoning, which will illustrate the role of construction in reasoning and proof.

> PROFESSOR A: Next I will prove the Bolzano–Weierstrass Theorem: every bounded infinite set has at least one cluster point. Let C_0 be an arbitrary, bounded infinite set. To prove the theorem, we must produce a point p with the property that every neighbourhood of p contains infinitely many points in C_0 ... We divide C_0 into four equal squares by intersecting lines. One of these smaller squares must contain infinitely many points of C_0 ... choosing such a subsquare, label it C_1. We have $C_0 \subseteq C_1$ and both are closed and bounded. Now repeat this process. Divide C_1 into four squares

[4] However, it is unclear to what extent Dummett is an *intuitionist* over and above being a *constructivist*. It is better, I think, to defer to his own label for his position: he is a *semantic anti-realist* [10].

> ... By continuing this, we generate a sequence of closed squares C_n ... appealing to the nested set property, there must be a point p that lies in all the sets C_n. This is the point that will turn out to be a cluster point for S.
>
> At this moment, a student with a double major in mathematics and philosophy raises her hand and is recognised.
>
> STUDENT: You are using a constructional language in this lecture. You do not actually mean that you or some ideal mathematician has done this construction, do you? How can anyone do an infinite number of things, and then after *all* of them — on the basis of them — do something else, in this case pick the point p?
>
> PROFESSOR A: Do not take this lecture literally. Of course, there is no such constructional process. I am *describing* a property of *the plane*. From principles of cardinality, I *infer* the existence of infinitely many points in some square C_1, and then in C_2. The axiom of replacement implies the existence of the whole sequence $\langle C_n \rangle$. Finally, from the nested-set property, I deduce the existence of a point in all of the C_is. I let p be the name of one such point. Now this is what I *mean* by this lecture; the constructional language makes it easier for you to see. [17, page 184][5]

There are a number of things going on in this example. Firstly, it shows how images of 'construction' are used to clarify the processes of reasoning leading to a mathematical proof. As Professor A indicates, this language does not have to be taken literally; it may be used as a heuristic device.

However, the language is not *merely* heuristic. If you are presented with an infinite set of points and are asked to *find* a cluster point for that set, the process you are given in this proof is one you might think to use. However, to use it, you must, as the student points out, perform an infinite task. The Professor is wrong in thinking that the infinite task is impossible. Sometimes, you can actually *do* it, depending on which infinite set you are given. For example, if I am given the unit square $[0, 1] \times [0, 1]$, then I can choose each C_n as $[0, 1/n] \times [0, 1/n]$ and my limit point will be $(0, 0)$. I *can* perform infinitely many choices if they choices are "easy" enough. However, the choices are not always that simple, and indeed, there is no algorithm providing for each infinite set of points a cluster point for that set.

Shapiro uses this example to motivate an understanding of the difference between intuitionist and classical mathematics. His explanation is a good example of the orthodox view of the relationship between classical and intuitionistic mathematics, so it is worth quoting at some length.

> One crucial difference between the classical constructive mode of thought [by this, Shapiro means the "constructing" language used in mathematical demonstration] and the intuitionistic mode is that the former seems to presuppose that there is a (static) external mathematical world that mirrors the constructs. A traditional Platonist (such as Proclus) might claim that the existence of the mathematical world is what *justifies* or grounds the constructs ...
>
> Whatever its metaphysical status, the supposition of an external world suggests certain inferences, some of which are the nonconstructive parts of mathematical practice rejected by intuitionism. For example, if a classical

[5] In this quote I have changed Shapiro's 'S' to 'C_0' to clarify the reasoning. Shapiro starts with the set S then calls it C_0 without explaining the transition.

mathematician proves that *not all* natural numbers *lack* a certain property (i.e., she proves a sentence in the form $\sim\forall x \sim\Phi$), she can infer the *existence* of a natural number with this property ($\exists x\Phi$). Following existential instantiation, she can give a "name" to some such number and do further constructional operations on it. In the proof of the Bolzano–Weierstrass theorem ... we have a similar instance of excluded middle at work. At each stage, the constructor knows that at least one of four squares has infinitely many points from the given set, but he may not know which one. Nevertheless the constructor can pick one such square, and go on from there.

Intuitionists demur from such inferences and constructions because they understand the principles as relying on the independent, objective existence of the domain of discourse. For them, every assertion must report (or correspond to) a construction. In the present example, an intuitionist cannot assert the existence of a natural number with the said property, because such a number was not constructed. The intuitionist cannot choose a square with infinitely many points from the given sets, because such a square was not identified. Bishop [5] understands the law of the excluded middle as a principle of omniscience. [17, page 187]

Here the orthodox view is clear. Classical reasoners accept such infinitary constructions because they do not step from truth to untruth. If not everything is not Φ then something is Φ, regardless of the availability of a construction of this object. For intuitionists, this is unacceptable, for there may well *be* no such object.

What about the non-intuitionist? What if we are happy with the idea of an objective mind-independent mathematical realm? (We may not think that there *is* such an objective mind-independent realm; we may think, however, that mathematics proceeds *as if* there is such a realm. The distinction is irrelevant here [1, 11, 12, 17].) Can constructivist considerations have any force for us? I think that the answer to this may still be positive. Although as non-intuitionists we may well be happy that there actually *is* a cluster point of the infinite set, we may agree that it may well be impossible to *find*. Being presented with an infinite set does not in and of itself give one the means to find a cluster point for that set. All are agreed on this point, at least in cases where we are unable to complete the infinite division-and-choice task discussed in the proof. This consideration has force *whatever* your view of the nature of the mathematical universe, and this distinction, between the constructible and the non-constructible, is modelled in intuitionistic reasoning. In this way, intuitionistic reasoning can appeal to the non-intuitionist. You need not be an anti-realist for those considerations to have force. Shapiro describes the situation in the following way:

The supposition that there is a static mathematical universe that mirrors the dynamic language sanctions these infinitary procedures. Consider, once more, the aforementioned proof of the Bolzano–Weierstrass theorem ... On the classical view, each construct refers to a *fact* in the static mathematical world. The construction of the first square C_0 reflects the *existence* of a corresponding square in the plane. A more neutral way to put it is that the construction proceeds as if there were a corresponding square in the plane. Similarly for C_1 and C_2. From the discourse, it becomes clear that the construction *could* be continued as far as one wished. This reflects the existence of a square that corresponds to C_n for any particular natural number n. The crucial supposition is that the corresponding squares exist whether

> or not the construction is actually performed. The existence of the entire sequence of squares is then deduced. By the nested-set theorem, there is a point p in all of the C_ns. The proof concludes with the demonstration of some facts about p.
>
> In this example at least, the supposition allows one to proceed as if an infinite process had been completed — more precisely, the supposition of a static universe suggests that one can infer the existence of what would be the result of such a process, in this case a point in *every* element of the infinite sequence of squares. In the jargon of mathematics, we pass to the limit.
>
> Because, as noted, intuitionistic construction proceeds without this supposition, its ideal constructors are not endowed with the ability to finish infinite tasks. [17, page 188]

Note here that all of the considerations Shapiro brings in favour of an constructive understanding of proof apply independently of the nature of the objects being reasoned about. What we take our ideal constructors to be able to do is a feature of our understanding of construction; it is not a feature determined by the nature of the things constructed. We can impose constructive limits in our reasoning about anything we like.

2.2 A Semantics

At this point, it will serve us well to clarify how this is done. To do this we need to give a more detailed constructive analysis of the logical particles. Our interpretation will be a version of the Brouwer–Heyting–Kolmogorov interpretation ('BHK' for short) of the logical constants. To show that our approach follows the constructive tradition I will introduce this by quoting Errett Bishop:

> To prove the statement (P and Q) we must prove the statement P and prove the statement Q, just as in classical mathematics. To prove the statement (P or Q) we must either prove the statement P or the statement Q, whereas in classical mathematics it is possible to prove (P or Q) without proving either the statement P or the statement Q.
>
> The connective "implies" is more complicated. To prove (P implies Q) we must show that P necessarily entails Q, or that Q is true whenever P is true. The validity of the computational facts implicit in the statement P must insure the validity of the computational facts implicit in the statement Q, but the way this actually happens can only be seen by looking at the proof of the statement (P implies Q). Statements formed with this connective, for example, statements of the type ((P implies Q) implies R), have a less immediate meaning than the statements from which they are formed, although in actual practice this does not seem to lead to difficulties in interpretation.
>
> The negation (not P) of a statement P is the statement (P implies $0 = 1$). Classical mathematics makes no distinction between the content of the statements P and not (not P), whereas constructively the latter is a weaker statement. [5, pages 7,8]

Note here that we are recursively defining the behaviour of complex mathematical statements in terms of the behaviour of simple ones. You can understand this as giving the meanings of the connectives in terms of the meanings of something more fundamental or prior to them, in particular, the notion of proof. However, this is not compulsory. You can read this analysis as giving the sense of the notion of constructive proof. Here is what it is to give a proof of a conjunction: it is to give a proof of each conjunct. Here is what it is to give a proof of a disjunction: it is to give a proof of either disjunct. Here is what it is to give a proof of a conditional: it is to convert proofs of the antecedent into proofs of the consequent. And so on. Seeing the BHK interpretation in this light makes it open to all, *even* those who think that classical reasoning is applicable in the domain in question. The BHK interpretation gives an analysis of *constructive reasoning* about objects.

Here is a formal account expanding on the BHK interpretation which both does justice to the practice of constructive mathematics and opens the way for a reading of that practice for non-intuitionists. This truth conditional semantics is simply Kripke's semantics for intuitionistic logic. This semantics makes use of *constructions* which are partially ordered by *strength* (written '$\sqsupseteq$'). We abbreviate "c is a construction of A" as "$c \Vdash A$".

- $c \Vdash A \wedge B$ if and only if $c \Vdash A$ and $c \Vdash B$.
- $c \Vdash A \vee B$ if and only if $c \Vdash A$ or $c \Vdash B$
- $c \Vdash A \supset B$ if and only if for any $d \sqsupseteq c$, if $d \Vdash A$ then $d \Vdash B$.
- $c \Vdash {\sim}A$ if and only if for any $d \sqsupseteq c$, $d \nVdash A$.
- $c \Vdash \exists x A(x)$ if and only if for some a constructed at c, $c \Vdash A(a)$.
- $c \Vdash \forall x A(x)$ if and only if for any $d \sqsubseteq c$ and for some a constructed at d, $d \Vdash A(a)$.

The clauses for conjunction and disjunction are straightforward transcriptions of our pre-formal notion of constructions. The rules for implication and negation are motivated by the BHK interpretation in the following way. A construction proves $A \supset B$ if and only if when *combined with* any construction for A you have a construction for B. The assumption guiding Kripke models is that a construction for $A \supset B$ combined with one for A will be a *stronger* construction.[6] So, a construction c proves $A \supset B$ if and only if any stronger construction d for A is also a construction for B.

Constructions are *incomplete* and hence should not be expected to construct, for every claim A, either it or its negation ${\sim}A$. Constructions have computational content, so a construction of $A \vee B$ should be a construction of A or a construction of B. This jointly ensures that $A \vee {\sim}A$ ought fail. This can not necessarily be constructed.

The rules for the quantifiers are also straightforward. A construction for $\exists x A(x)$ must first construct an object, and then construct the claim that this object has property A. (We presume no special facts about which objects are constructed at any stage, except that every construction constructs *some* object, and if a construction c constructs some object and $d \sqsupseteq c$ then d constructs this object too.) For the universal quantifier, we must be more circumspect. To construct $\forall x A(x)$ we need more than just a construction for $A(x)$ for each object x available at this stage. For there may well be stronger constructions which

[6] This is the assumption challenged by relevant accounts of implication. In constructive mathematics, where relevance is not at issue, this account is appropriate.

construct *more* objects. If we have a construction for $\forall x A(x)$ and a stronger construction providing a new object a, this stronger construction must construct $A(a)$ too. So, for $c \Vdash \forall x A(x)$ we need $d \Vdash A(a)$ for each $d \sqsupseteq c$ where d constructs a.

Now we can formalise and clarify the failure of the inference of distribution. $\forall x(A \vee B) \vdash \exists x A \vee \forall x B$ fails because we may well have $c \Vdash \forall x(A \vee B)$ without having $c \Vdash \exists x A$ or $c \Vdash \forall x B$. A construction may be enough to ensure that *every* object constructed at any later stage is either an A or a B. However, we may neither have every future constructed object an A *nor* an object constructed now which is constructed as a B. Consider this simple model with two stages $c_1 \sqsubset c_2$, at which c_1 constructs one object 1 and c_2 constructs two objects 1 and 2. Let us further suppose that $c_1 \Vdash B(1)$ and $c_2 \Vdash B(1)$ and $c_2 \Vdash A(2)$. We do have $c_1 \Vdash \forall x(A \vee B)$ but we do not have $c_1 \Vdash \exists x A$ (as no object constructed at c_1 is constructed as having property A) but neither do we have $c_1 \Vdash \forall x B$, for c_1 is extended by c_2 at which not every object is indeed a B. This is a model for the reasoning situation discussed earlier. We may well have an construction which provides for us assurance that every digit in the decimal expansion of π is either even or odd. Suppose construction c_n provides the first n digits of π, and that *every* construction tells us that every digit of π is either even or odd. (That is straightforward, because for this you need know only that π has a decimal expansion.) So, $c_1 \Vdash \forall x(\mathsf{Even}(x) \vee \mathsf{Odd}(x))$ but $c_1 \nVdash \exists x \mathsf{Even}(x) \vee \forall x \mathsf{Odd}(x)$ because c_1 does not provide us an even digit for π (it gives us only the first digit — 3 — which is odd) and it certainly does not prove that *all* digits of π are even. So, this model of the reasoning gives the same account of the failure of distribution as before. A construction of $\forall x(A \vee B)$ need not itself be a construction of $\exists x A \vee \forall x B$.

So, this model coheres well with the considerations underpinning intuitionistic logic, yet it is available to the realist. The central feature is the elaboration of the notion of a construction. Nothing in this notion requires that we be anti-realists about objects in question. Nothing requires us to believe that all of the digits of π are not yet "there" to "make" mathematical claims true or false. We simply need a notion of construction which models our access to mathematical truth and mathematical demonstration. Intuitionistic logic demands at most a modest, realistic account of our access to mathematical objects and mathematical truths. We need not be anti-realists about those objects to appreciate this.

2.3 Where is Truth?

But what of mathematical truth? If we endorse intuitionistic logic, what can we say about the truth of $\exists x A \vee \forall x B$, if we know that $\forall x(A \vee B)$ is true? The inference from the former to the latter fails in intuitionistic logic. Suppose the premise is true. Is the conclusion true? What does our appeal to intuitionistic logic say about mathematical truths?

The first thing to admit is this: our appeal to intuitionistic logic doesn't say very much about truth at all. The failure of non-intuitionistic inferences such as distribution means only this: there is a construction of the premises which is not a construction of the conclusion too. This says nothing in and of itself about whether it may be that the premise is true and the conclusion is not. Here is a simple example: in constructive mathematics you cannot *prove* that every real number is either greater than zero, equal to zero, or less than zero (the trichotomy law). The reason is straightforward. A real number is an infinite series of successive approximations to a point on the real line. For example, it may be presented as a descending sequence of *intervals* (p_n, q_n) bounded by rational numbers p_n and q_n, descending so tightly that the distance $|p_n - q_n|$ between the endpoints is no more than $1/2^n$. The real number given by this sequence can be thought of as the limiting point of that sequence, the point on which the intervals converge. Now, suppose that someone is presenting you with such a sequence, one step at a time, and each interval they have given you so far includes 0. (Say, the first n intervals they give you are $(-1/2^{n+1}, 1/2^{n+1})$.) A construction which gives us the first n intervals making up the number does not in and of itself give us any assurance as to whether the number is greater than zero, less than zero, or zero itself. The only way to do that is to be presented with the *whole* sequence. It need not follow that not all real numbers are greater than or equal to or less than zero. We do not have a constructive proof the trichotomy law It does not follow that the law is not true. Such a conclusion only follows if all that the only truths about numbers are those that may be constructed.

The situation is the same with the law of the excluded middle. The fact that not all constructions are constructions of $A \vee {\sim}A$ is one thing. It does not follow that $A \vee {\sim}A$ is not true, or even, not *necessarily* true. It is consistent to maintain that all of the truths of classical logic hold, and that all of the arguments of classical logic preserve truth, with the use of constructive mathematical reasoning, and the rejection of certain classical inferences. The crucial fact which makes this position consistent is the shift in context. Classical inferences are valid, *classically*, in the sense that they never step from truth to falsity. They are not *constructively* valid. If we use a classical inference step, say the inference from $\forall x(A \vee B)$ to $\exists xA \vee \forall xB$, then we have not (we think) moved from truth to falsity, and we cannot move from truth to falsity. It is impossible for $\forall x(A \vee B)$ to be true and for $\exists xA \vee \forall xB$ to be false. However, such an inference *can* take one from a truth which can be constructed to one which cannot, as we have seen. So, the inference, despite being classically valid, can be rejected on the grounds of non-constructivity. Truth is one thing, and what a construction may construct is another.

3 Everyone *ought to be* Constructivists

In constructive mathematics the goal is to gain understanding of mathematical structures, and to prove theorems about them (just as in classical mathematics). However, the goal is to prove mathematical theorems with constructive, or computational content. If a statement asserting the existence of some mathematical object is proved in a constructive manner (using the rules of intuitionistic logic) then this proof will contain the means of specifying the object or structure in question. Wittgenstein illustrates the advantages of constructive proof over its classical cousin by drawing out its implications for our *understanding*.

> A proof convinces you that there is a root of an equation (without giving you any idea *where*) — how do you know that you understand the proposition that there is a root? [22, page 146]

This feature of constructive mathematics is guaranteed by the structure of constructive proofs. Any intuitionistically valid proof of a disjunction (from no premises, of course) will prove a disjunct. Any intuitionistically valid proof of an existentially quantified statement (from no premises, of course) will prove an instance of that quantifier. These proofs track mathematical constructions.

It is one thing for a proof to encode mathematical construction. It is another for this notion of proof to be properly in the domain of logic. Is there any call to consider constructively invalid arguments as genuinely *invalid*? Or is it more appropriate to consider constructive considerations as in the domain of the theory of computation? To put the matter succinctly: *Is* intuitionistic consequence genuinely *logical* consequence?

Recall my introduction to the notion of validity and logical consequence. I said that an argument is valid if and only if in any case in which the premises are true, so is the conclusion. Intuitionistic logic fits this scheme if we take cases to be constructions. An argument is intuitionistically valid if and only if any construction for the premises is also a construction for the conclusion. There is no doubt that intuitionistic logic, as modelled by the Kripke semantics of Section 2.2, fits the general scheme of logical consequence described here.

Do intuitionistically valid arguments ever step from truth to untruth? Can one have an intuitionistically valid argument in which the premises are true and the conclusion is not? If there is such, then intuitionistic logic will be useless in providing truth. Thankfully, there is no need to think that this will ever be the case. We may be assured that intuitionistic reasoning preserves truth if *one* construction is so expansive that it constructs *all* truths. If this is the case, then this construction will provide a counterexample to all arguments with true premises and false conclusion. This construction, if you like, is a God's eye view of truth. It may not be a construction which is at all feasible for an earth-bound mathematical reasoner, but its admission does not make intuitionistic reasoning any less constructive.[7] On the contrary, admitting more constructions provides *more* opportunity for arguments to fail, not less, so its admission, if adding any restriction on the repertoire of proof, is not in the direction of allowing non-constructive forms of reasoning.[8]

So, intuitionistically valid arguments preserve truth-in-constructions. Any construction in which the premises are true is a construction in which the conclusion is true. Intuitionistically valid arguments never step from outright truth to outright untruth. For any intuitionistically invalid argument, there is a construction in which the premises are true and in which the conclusion is not true. This seems to justify the appellation "logic" for intuitionistic logic. People *ought* reason constructively if they have any care for what may be constructed or demonstrated or shown. This is not an idle claim. In the computational representation of mathematical reasoning, intuitionistic logic is completely natural, and is widespread.[9]

[7]It does make the reasoning less *intuitionistic* but, as we have already seen, constructivism is not to be identified with intuitionism.

[8]This strategy will not be available to the dialetheist, for whom some contradictions $A \wedge {\sim}A$ are true. Given dialetheism, the "world" w cannot be a construction (or be modelled by a construction) for $w \Vdash A$ and $w \Vdash {\sim}A$, which conflicts with the assumption that $w \sqsupseteq w$ and the clause for negation.

[9]Lloyd Humberstone points out a useful analogy. In saying that people ought to be constructivists, it is like I am commending that people eat vegetables. The analogy is not with vegetarianism. The analogue of vegetarianism is the *exclusive* endorsement of constructive reasoning, which

4 Everyone ought to be *more than* Constructivists

Nothing we have said so far that $A \vee {\sim}A$ is not true, or even, not *necessarily* true. One can agree with everything we have said so far and still say that every (truth-apt) claim is either true or false, that all of the theses of classical logic are true, and that all of the arguments of classical logic preserve truth. The way is open for the constructivist to agree that classical inferences are 'valid' in the (admittedly more restricted) sense that they never step from truth to falsity. Of course, they are are not *constructively* valid. If we use a classical inference step, say the inference from $\forall x(A \vee B)$ to $\exists xA \vee \forall xB$, then we have not (we think) moved from truth to falsity, and we cannot move from truth to falsity. It is impossible for $\forall x(A \vee B)$ to be true and for $\exists xA \vee \forall xB$ to be false. However, such an inference *can* take one from a truth which can be constructed to one which cannot, as we have seen. So, the inference, despite being classically valid, can be rejected on the grounds of non-constructivity. The constructive reasoner is free to agree that classical logic is an organon for the preservation of truth. Its failing is not because of any step from truth to falsity *tout court*. Its failing is a failing of constructivity.

Can we push the case further and argue that not only ought we reason constructively, but that we ought reason classically[10] *as well*? Here the ground is much weaker. There is a nearly century-old tradition of constructive reasoning, and there are no universally convincing knock-down arguments to the effect that more than that is needed. I will provide just two swift reasons why one might find a constructivist-only policy for inference unacceptable. (It is hopefully clear that this paper advocates a *pluralism* about logical consequence. For more of this, see my papers with Beall [2, 3].)[11]

First, a constructivist-only policy, according to which the *only* deductively acceptable arguments are intuitionistic, faces problems of expressivity. Typically, the proponent of such a view will want to express reasons for the *failure* of inferences such as the law of the excluded middle or distribution. But there is no straightforward way of doing this.[12] You cannot express your rejection of $A \vee {\sim}A$ (for some A) in a way which is rules out iys truth. For to say something which rules out the truth of $A \vee {\sim}A$ is to be committed to ${\sim}(A \vee {\sim}A)$, which is, intuitionistically speaking, absurd. You cannot say that for some instance of A and B, $\forall x(A \vee B)$ is true and $\exists xA \vee \forall xB$ is *not* true, because $\forall x(A \vee B) \wedge {\sim}(\exists xA \vee \forall xB)$ is likewise a contradiction. The intuitionist who *rejects* the use of classical inference such as ${\sim}{\sim}A \vdash A$ cannot express that rejection by saying that the inferences step from truth to untruth, for to say that would be to be committed to ${\sim}{\sim}A$ and ${\sim}A$, a contradiction.

It follows that intuitionistic demurral from classical inferences cannot be readily expressed in the positive language provided by intuitionistic logic.[13]

I reject.

[10] Or at least, non-constructively. I have said nothing of considerations of paraconsistency or relevance.

[11] This section is at most a sketch. To make this case at all convincingly, I must provide an elaboration of the normativity of "ought" in play here.

[12] You can indeed say that, for example, $A \vee {\sim}A$ is not provable, but an intuitionist wants to say more than this. The classical logician will agree, for example, that p is not provable, but be happy to assert p.

[13] Proponents of non-classical logics may indeed avail themselves of *denial* or *rejection* as speech acts independent of assertion (but no doubt connected to it). A proponent of intuitionistic logic

This brings us to the *second* swift argument for the "acceptability" of classical propositional logic for constructivists. Constructivists themselves are committed to the conclusion that we never make a *mistake* while stepping from premises to conclusions of classically valid arguments, at least if by a "mistake" you mean a step in which the premises are true and the conclusion false. If $X \vdash A$ is classically propositionally valid, then we have a construction of

$$\sim(\bigwedge X \wedge \sim A).$$

It is *never* the case that the premises of a classically valid argument are true and the conclusion is false. If we make the step from premises to conclusion, we do *not* make a *mistake*. So, the acceptability of classical reasoning, at least in this mild form, is represented already in intuitionistic logic. The extent to which it is appealing to avoid stepping from truth to falsity is the extent to which classical reasoning ought to appeal to someone for whom constructivist considerations are paramount.

These arguments are certainly not decisive. Intuitionistic reasoners are not *compelled* to join a classical bandwagon. However, the case of this paper is cumulative. We have shown how the virtues of constructive reasoning are available to all. The anti-realist metaphysics of Brouwer or Dummett may help motivate intuitionistic logic, but it is not the only possible motivation. The tight connection between intuitionism and anti-realism is at most a historical connection. Realists too can avail themselves of intuitionistic reasoning and constructivist distinctions. If the case for anti-realism is partly based on the case for intuitionistic logic (as it seems to be in the work of Dummett [10] and Tennant [20]) then this case is undermined. It is one thing to justify intuitionistic logic; it is altogether another to justify an anti-realist metaphysics.

5 Inconsistency?

I must complete this paper by confronting an important issue for anyone who employs both constructive and classical reasoning. Intuitionistic mathematics is famous for not only being weaker than classical mathematics, by not taking certain classical results to be valid: it differs from classical mathematics by taking certain classical mathematical results to be outright *false*. This seems to commit me to an inconsistency. If I as a classical reasoner can prove A and as an intuitionistic reasoner can prove $\sim A$, then as a pluralist I seem to be committed to $A \wedge \sim A$. What can I do?

Here is the general strategy: I must examine closely the intuitionistic proofs in question. If there is a genuine intuitionistic proof of $\sim A$, where A is provable in classical mathematics, then there must be a premise used which is not true in classical mathematics. It is certainly not the *logic* which gives us $\sim A$, as intuitionistic logic is weaker than classical logic. So, if we genuinely have a proof of $\sim A$, we have used premises which are false in classical mathematics.

Of course, upon examination we may see that we have not really *proved* $\sim A$, but we have proved that we cannot prove A. And this is another way to resolve the conflict. For the statement that we cannot (constructively) prove A is not in conflict with the claim that we can (classically) prove A. The situation here

who wishes to provide a formal understanding of the *rejection* of some instance of $A \vee \sim A$ is invited to develop a formal account featuring a notion of *rejection*. I know of no attempt in this direction.

differs no more than that with such trivialities as $A \vee {\sim}A$ (where A is constructively undecidable). This is classically provable but constructively unprovable. It does not follow that ${\sim}(A \vee {\sim}A)$!

However, there are intuitionistic theories in radical variance with classical theories. There are theories in which you can prove the negations of classical theorems. We must understand what to say about these if our position is to be consistent.

The first line of defence then is a deference to an important tradition in constructive mathematics. The constructivism of Errett Bishop [5, 6], Douglas Bridges [7], Fred Richman [15, 16] and others can best be described as mathematics *pursued in the context of intuitionistic logic*.[14] This brand of constructive mathematics is explicitly consistent with classical mathematics. Bishop-style constructivists reject any inference in conflict with classical reasoning [16]. This is the approach we must take also.

However, to see what this means in practice, we must examine cases of conflict in some detail. There are *two* major ways intuitionistic theories conflict with classical theories. These conflicts arise from notions of *choice sequences* and *realisability*. We will consider each in turn.

The remainder of this section is quite technical. It may be ommitted without cost if you believe that we can develop constructive theories as consistent subtheories of classical theories.

5.1 The First Route to Inconsistency: Choice Sequences

First choice sequences.[15] A choice sequence is $\alpha(0)$, $\alpha(1)$, $\alpha(2)$, ... of natural numbers. We let '$\overline{\alpha}(k)$' stand for the sequence $\langle \alpha(0), \alpha(1), \ldots, \alpha(k-1) \rangle$, the initial segment of length k of α. These sequences are taken to encode the choices of a creating mathematical subject. A choice sequence may be completely freely constructed (by the analogue of tossing a coin at each stage) or it may be completely determined by law (such as the law defining the decimal expansion of π) or it may be somewhat constrained but somewhat free (each step may be free within certain constraints, such as taking $\alpha(k)$ to be the k-th digit in the expansion of π if the toss is heads, and 9 minus that digit if the toss is tails).

A typical intuitionistic thesis about choice sequences is this: If F is a function on choice sequences, then given that choice sequences may be *free* creations, F must depend on some initial segment of the choice sequences accepted as inputs. There is no way to assume more about the choice sequence, as at any stage of reasoning not all of the sequence has been constructed. So, there is another function f on segments such that $\forall\alpha\exists k\big(F(\alpha) = f(\overline{\alpha}(k))\big)$. A consequence of this is Brouwer's *continuity principle for functions*:

$$\forall\alpha\exists k\forall\beta\big(\overline{\alpha}(k) = \overline{\beta}(k) \supset F(\alpha) = F(\beta)\big) \tag{2}$$

Given the continuity principle we have conflict with classical analysis. We must reject one classical principle, the *limited principle of omniscience*.

$$\forall\alpha\big(\exists x(\alpha(x) = 0) \vee \forall x(\alpha(x) \neq 0)\big) \tag{3}$$

[14] Tait provides more explicitly *philosophical* account which draws very similar distinctions to the work of constructive mathematicians [18, 19].

[15] My presentation closely follows van Dalen and Troelstra's helpful short expositions [8, 21].

It is instructive to see why we cannot constructively prove (3). It is *not* a problem with the law of the excluded middle. Identity for natural numbers is decidable — when given two natural numbers, I have a simple routine for determining if they are identical. So, we can prove

$$\forall\alpha\big(\forall x(\alpha(x) = 0 \vee \alpha(x) \neq 0)\big)$$

as for any value the choice sequence presents to us, we can determine if it is a zero or not. The crucial move not allowed in intuitionistic logic is the following inference of *distribution* referred to before (1).

$$\forall x(A \vee B) \vdash \exists xA \vee \forall xB$$

As we have seen, distribution is constructively undesirable because a routine showing that every x is either A or B does not necessarily provide a routine to find an A or to show that *all* objects in the domain are B. For this we need to survey the domain. So, our routine for verifying $\exists xA \vee \forall xB$ outstrips routines for verifying $\forall x(A \vee B)$.

Now, if (3) were *true* (as I take it to be, as a pluralist) we would have a function F such that

$$F(\alpha) = \begin{cases} 0 & \text{if } \exists x(\alpha(x) = 0) \\ 1 & \text{if } \forall x(\alpha(x) \neq 0) \end{cases}$$

Applying the continuity condition, F must be determined by an initial segment of its input. In particular, since F applied to the constant 1 choice sequence (β where $\beta(x) = 1$ for each x) gives 1, continuity tells us that there is some sequence $\langle 1, 1, \ldots, 1\rangle$ such that every continuation γ yields 1: that is $F(\gamma) = 1$. However, there are many continuations of the series $\langle 1, 1, \ldots, 1\rangle$ which contain zeros.

It follows that such a function F cannot exist. But the existence of F is a consequence of the classical tautology:

$$\forall\alpha\big(\exists x(\alpha(x) = 0) \vee \forall x(\alpha(x) \neq 0)\big)$$

What do we do? The intuitionistic response is to assert the negation of the thesis:

$$\sim\forall\alpha\big(\exists x(\alpha(x) = 0) \vee \forall x(\alpha(x) \neq 0)\big)$$

A pluralist response cannot follow the orthodoxy of intuitionism. We must look elsewhere if we are to maintain consistency. Is there any well motivated option open to us?

One option is this: Reject the continuity principle. Once we reject intuitionism we have no reason to agree that functions on choice sequences must be determined by initial segments of those sequences. The function F is a case in point. For a pluralist, (3) is *true* without being *constructively provable*. Functions such as F may *exist* without being *constructed*. Constructive considerations give us no reason to endorse (2).

So, one route to inconsistency fails. Choice sequences are unacceptable to the pluralist, for they make illegitimate assumptions. They rule out of existence functions like F which are classically demonstrable. If we have reason to allow the existence of such functions (as I think we have) then we have reason to reject choice sequences.

This means that we reject certain *branches* of constructive mathematics, not the whole study. The constructive mathematics of Bishop, Bridges and Richman [5, 6, 7, 15, 16] makes no use of choice sequences, and it makes no counter-classical claims.

5.2 The Second Route to Inconsistency: Realisability

Take an enumeration of all partial recursive functions, with $\{n\}$ the function with index n, so $\{m\}(n)$ is the function with index m applied to the natural number m. (Details of how we deal with *partiality* and undefined results I leave to elsewhere [4].) Similarly, we encode *pairing*, so that $(n)_0$ and $(n)_1$ are the first and second item in the pair n. so $(\langle n, m\rangle)_0 = n$ and $(\langle n, m\rangle)_1 = m$, and $\langle (n)_0, (n)_1\rangle = n$. With this technology we define a relation between (codes of) functions and sentences of the language of arithmetic in the following way:

Realisability Fact	*Condition*
$n \mathrel{\mathsf{r}} A$, if A is atomic	A is true.
$n \mathrel{\mathsf{r}} A \wedge B$	$(n)_0 \mathrel{\mathsf{r}} A$ and $(n)_1 \mathrel{\mathsf{r}} B$.
$n \mathrel{\mathsf{r}} A \vee B$	If $(n)_0 = 0$ then $(n)_1 \mathrel{\mathsf{r}} A$
	and if $n_1 \neq 0$ then $(n)_1 \mathrel{\mathsf{r}} B$.
$n \mathrel{\mathsf{r}} A \supset B$	For all m, if $m \mathrel{\mathsf{r}} A$ then $\{n\}m \mathrel{\mathsf{r}} B$.
$n \mathrel{\mathsf{r}} \exists x A(x)$	$(n)_1 \mathrel{\mathsf{r}} A((n)_0)$.
$n \mathrel{\mathsf{r}} \forall x A(x)$	For all m, $\{n\}(m) \mathrel{\mathsf{r}} A(m)$.

The justifications for these clauses are straightforward:

- Atomic sentences are self-justifying. We take *everything* to be a realisation of an atomic sentence.
- A realisation for a conjunction is a pair of realisations for each conjunct.
- A realisation for a disjunction is a realisation for a disjunct, combined with an indication of *which* disjunct has been realised.
- A realisation for a conditional is a function transforming realisations for the antecedent to realisations for the consequent.
- A realisation for an existential quantifier is an object together with the realisation that *that* object satisfies the formula under the quantifier.
- A realisation for a universal quantifier is a function sending objects to realisations that the object satisfies the formula under quantifier.

A nice result is that every thesis of Heyting Arithmetic (Peano Arithmetic using intuitionistic predicate logic: we write this as 'HA') is realisable. That is, if $\mathsf{HA} \vdash A$ then for some n, $n \mathrel{\mathsf{r}} A$. However, *more* is realised than simply the theses of HA. Consider what counts as a realisation of a $\forall\exists$ formula. If $n \mathrel{\mathsf{r}} \forall x \exists y A(x, y)$ then for each m, $\{n\}(m) \mathrel{\mathsf{r}} \exists y A(m, y)$, and given the definition of a realiser for an existentially quantified formula, we have that for each m,

$$(\{n\}(m))_1 \mathrel{\mathsf{r}} A\big(m, (\{n\}(m))_0\big)$$

But *this* has consequences of its own. Abstracting out the m, we have a realiser l such that

$$l \mathrel{\mathsf{r}} \forall x A\big(x, (\{n\}(x))_0\big)$$

where l is the code of the recursive function sending m to $(\{n\}(m))_1$. As a result, we have

$$\langle n, l\rangle \;\mathsf{r}\; \exists e \forall x A\big(x, \{e\}(x)\big)$$

Now, the function which sends n (the realiser of $\forall x \exists y A(x, y)$) to $\langle n, l\rangle$ is itself recursive, so the code of this function is a realiser for the following claim:

$$\forall x \exists y A(x, y) \supset \exists e \forall x A\big(x, \{e\}(x)\big) \tag{4}$$

This has become known as *Church's Thesis*.[16] This thesis states that given any true $\forall\exists$ formula, there is a recursive choice function choosing the appropriate instance of the existential quantifier for each input into the universal quantifier.

This thesis is *false*. Given that there is one non-recursive function, f then we have $\forall x \exists y\big(f(x) = y\big)$, but we do *not* have $\exists e \forall x\big(f(x) = \{e\}(x)\big)$, as this states that f is recursive (as it is identical in extension to the recursive function $\{e\}$).

If Church's thesis is false, then why is it realisable? Look back to the verification that (4) is realisable. We argued that if n was a realiser for $\forall x \exists y A(x, y)$ then there is a realiser (recursively constructible from n) for $\exists e \forall x A\big(x, \{e\}(x)\big)$. Does it follow that if $\forall x \exists y A(x, y)$ is *true*, so is $\exists e \forall x A\big(x, \{e\}(x)\big)$? We can safely deny this. After all, there may be a sentence of the form $\forall x \exists y A(x, y)$ which is true without having a realiser. It is true that all theorems of HA have realisers, but it may not be the case that all *arithmetic truths* have realisers. If truth in arithmetic outstrips truth in HA (with Church's thesis) then we have no reason to think that simply because the *realisability* of $\forall x \exists y A(x, y)$ transforms into the realisability of $\exists e \forall x A\big(x, \{e\}(x)\big)$ that in addition, the truth of the former gives us the truth of the latter.

Realisability semantics have only *recursive* realisations are in play. It is little surprise that in this 'universe' all functions are recursive. If we wish to reason constructively about the mathematical universe studied by classical mathematicians, *this* realisability semantics will not do. We will need more constructions than those provided by recursive functions. There is no force in the argument that realisability semantics motivates a departure from classical arithmetic. We can reason constructively without fear of contradiction with classical theories.[17]

References

[1] MARK BALAGUER. *Platonism and Anti-Platonism in Mathematics*. Oxford University Press, 1998.

[2] JC BEALL AND GREG RESTALL. "Logical Pluralism". Under consideration, Year 2000 special issue on Logic of the *Australasian Journal of Philosophy*, 2000.

[3] JC BEALL AND GREG RESTALL. "Defending Logical Pluralism". In B. BROWN AND J. WOODS, editors, *Logical Consequences*. Kluwer Academic Publishers, to appear.

[4] MICHAEL BEESON. *Foundations of Constructive Mathematics: Metamathematical Studies*. Springer Verlag, Berlin, 1985.

[5] ERRETT BISHOP. *Foundations of Constructive Analysis*. McGraw-Hill, 1967. Out of print. A revised and extended version of this volume has appeared [6].

[16] Beware: This is *not* the Church–Turing thesis to the effect that every computable function is recursive. *This* Church's thesis is much stronger, to the effect that *every* function is recursive.

[17] Thanks to JC Beall, and an audience at Monash University, especially Lloyd Humberstone and Dirk Baltzy, for helpful comments on this paper.

[6] ERRETT BISHOP AND DOUGLAS BRIDGES. *Constructive Analysis*. Springer-Verlag, 1985.

[7] DOUGLAS S. BRIDGES. *Constructive Functional Analysis*, volume 28 of *Research Notes in Mathematics*. Pitman, 1979.

[8] DIRK VAN DALEN. "The Intuitionistic Conception of Logic". In *The Nature of Logic*, volume 5 of *The European Review of Philosophy*. CSLI Publications, 1999.

[9] MICHAEL DUMMETT. *Elements of Intuitionism*. Oxford University Press, Oxford, 1977.

[10] MICHAEL DUMMETT. *The Logical Basis of Metaphysics*. Harvard University Press, 1991.

[11] HARTRY FIELD. *Science without numbers : a defence of nominalism*. Blackwell, 1980.

[12] HARTRY FIELD. *Realism, Mathematics and Modality*. Blackwell, 1991.

[13] AREND HEYTING. *Intuitionism: An Introduction*. North Holland, Amsterdam, 1956.

[14] AREND HEYTING. *Brouwer Collected Works I*. North Holland, Amsterdam, 1975.

[15] RAY MINES, FRED RICHMAN, AND WIM RUITENBURG. *A Course in Constructive Algebra*. Springer-Verlag, 1988.

[16] FRED RICHMAN. "Interview with a constructive mathematician". *Modern Logic*, 6:247–271, 1996.

[17] STEWART SHAPIRO. *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, 1997.

[18] W. W. TAIT. "Against Intuitionism: Constructive Mathematics is Part of Classical Mathematics". *Journal of Philosophical Logic*, 12:173–195, 1983.

[19] W. W. TAIT. "Truth and Proof: The Platonism of Mathematics". *Synthese*, 69:314–370, 1986.

[20] NEIL TENNANT. *The Taming of the True*. Clarendon Press, Oxford, 1997.

[21] A. S. TROELSTRA. "Concepts and Axioms". Technical Report ML-1998-02, University of Amsterdam, 1998.

[22] LUDWIG WITTGENSTEIN. *Remarks on the Foundations of Mathematics*. MIT Press, 1967. Edited by G. H. von Wright, R. Rees and G. E. M. Anscome.

Greg Restall ~ Department of Philosophy, Macquarie University, Sydney, NSW 2109, AUSTRALIA.
`Greg.Restall@mq.edu.au` ~ `http://www.phil.mq.edu.au/staff/grestall/`