INVENTION IS THE MOTHER OF NECESSITY: MODAL LOGIC, MODAL SEMANTICS, AND MODAL METAPHYSICS

Greg Restall
restall@unimelb.edu.au*

VERSION 0.8
October 11, 2005

ABSTRACT: Modal logic is a well-established field, and the possible worlds semantics of modal logics is essential to our understanding of the logical features of modal concepts such as possibility and necessity. However, the significance of possible worlds models for a genuine theory of meaning—let alone for metaphysics—is less clear. In this paper I shall explain how and why the use of the concepts of necessity and possibility could arise (and why they have the logical behaviour charted out by standard modal logics) without either taking the notions of necessity or possibility as primitive, and without starting with possible worlds. Once modal logic is *explained*, we can give an account of possible worlds, and explain why possible worlds semantics is a natural fit for modal logic without being the source of modal concepts.

♦ ◊ ♦

1 ON AN EXPLANATORY CIRCLE

In this section, we see (1) that to take possibility to explain possible worlds, *and* possible worlds to explain possibility, is to endorse a tight circle with no explanation of what ought to be explained. (2) That to take possible worlds as primitive is nice because it gives you the logical structure, but unacceptable because it does not tell you what is *modal* about possibility. (3) That to take possibility as primitive is nice because it is lighter on the ontology and it does seem to get the order of explanation correct. However it needs supplementation to explain why possibility has the logical features that it does have. (4) That the fact that the very tight circle allows a lot of variability (and not much grip on anything *outside* that tight circle) should let us know that there is something more to be said in the analysis of possibility.

It is *possible* that everything is physical if and only if there is some possible world in which everything is physical. If every possible world contains something non-physical, then and only then is it impossible for everything to be physical. It is *necessary* that everything is determined if and only if there is no possible world in which not everything is determined. If three is a possible world in which things happen undetermined, then and only then is it not necessary that everything be determined.

This way of connecting possibility and necessity on the one hand, and the notion of a "possible world" is widespread among philosophers, and with good

*This is an early and rough draft, posted for comments. Please check the website for a later version, and do not cite before checking with me. ¶ Comments posted at http://consequently.org/writing/invention are most welcome. ¶ Thanks to audiences at ESSLLI 2005 and LOGIC COLLOQUIUM 2005 for helpful feedback on some of this material and to Conrad Asmus, Allen Hazen, Mark Lance and Øystein Linnebo for fruitful and enjoyable conversations on these topics. Thanks too, to Dan Isaacson for hospitality during my visit to Wolfson College and the Philosophy Faculty at Oxford University. This research is supported by the Australian Research Council, through grant DP0343388.

reason. It does a great deal of work. Connecting the *intensional* notions of possibility and necessity with the *extensional* notion of truth-in-an-x seems like a good idea for many different reasons.

On the one hand, if we agree that something is necessary if and only if it is true in every possible world, then we can derive a great many distinctive logical features of necessity. The universal quantifier (the *every*, in "every possible world") is well understood, logically speaking, and so we can derive a great many logical principles concerning necessity, by way of this connection between necessity and truth-in-every-possible-world. For example, if two statements are necessary, their conjunction is necessary; if something is necessary, it is true; if something is necessary, it is necessarily necessary; something is true then it is necessary that it is not necessary that its negation is true, and many more things besides. The logic of necessity (and of possibility, too) becomes tractable and understandable if we connect necessity with truth-in-each-possible-world.

On the other hand, if we observe this reading of truth-in-a-possible-world, then we have a powerful principle concerning what possible worlds there are. Frankly speaking, the notion of a possible world is mysterious (we never *see* any possible world, except perhaps for the one we are in) and if we connect possibility to possible worlds in this way, we have a tremendous principle for explaining what possible worlds there in fact are. If some statement is possibly true (no matter how outlandish) then there is *some* possible world in which it *is* true. On the other hand, if something is *necessary* then it is true in *every* possible world. The unruly notion of a 'possible circumstance' is brought under some measure of control by the simple constraint that the goings on in such a possible circumstance must be *possible*. It is constrained by whatever is *necessary*.

So, this connection between possibility and necessity on the one hand, and possible worlds on the other, is fruitful. However, there is a sense in which we cannot have it both ways. If we want this connection to provide an *explanation*, then it seems that we must choose: either the notions of possibility and necessity may be used to explain possible worlds, or the notion of truth-in-a-possible-world may be used to explain possibility and necessity. To attempt to have both (on face value, at least) is to commit yourself to an all-too-tight explanatory circle.

The choice of a direction of explanation is not straightforward. Both directions have an appeal, and both also have a distinct shortcoming. If we *explain* the behaviour of necessity and of possibility by reference to the notion of truth-in-a-possible world, then we are free to do so. The overwhelmingly *striking* logical fact, then "possibly" behaves like an existential quantifier, and "necessarily" like a universal quantifier, is given a straightforward explanation: That at bottom, they *are* quantifiers. However, this explanation comes at a great cost. We cannot then use possibility and necessity as a prior constraint on truth-in-a-possible world, since we use this very notion to explain possibility and necessity. We must either take the notion of a possible world as primitive and unexplained, or explain it in terms of something else.[1] Furthermore, if

[1] There are attempts to do so. David Lewis's *On the Plurality of Worlds* [10] charts many options and comes down on a criterion of spatio-temporal connectedness. A possible world is a maximally spatio-temporally connected region of Reality. David Armstrong's *A Combinatorial Theory of Possibility* [1] analyses a possible world as a recombination of objects and universals.

possibility and necessity are disguised ways of talking about truth-somewhere-else, then we are given the extra burden of attempting to explain what makes this kind of truth *modal*, and how we might come to know anything modal in the first place.[2] There is no doubt that we come to know modal truths by some means *other* than bumping into other possible worlds and taking account of what is true in any of them.[3] Arthur Prior seems to have had it right that we come to know what is true in a possible world by knowing what is *possible* but, not *vice versa*.

> . . . possile worlds, in the sense of possible states of affairs are not *really* individuals (just as numbers are not *really*) individuals. To say that a state of affairs obtains is just to say that something is the case; to say that something is a possible state of affairs is just to say that something could be the case; and to say that something is the case 'in' a possible state of affairs is just to say that the thing in question would necessarily be the case if that state of affairs obtained, i. e. if something else were the case . . . We understand 'truth in states of affairs' because we understand 'necessarily'; not *vice versa*. [12, pages 243–244].

It should be clear where my sympathies lie. I am much more inclined to think that possibility and necessity can be used to give an account of possible worlds than the converse. This seems to get the order of explanation correct, and it seems to put the ontology of possible worlds in the right place.[4] However, this leaves *me* with just as significant a problem. If I use possibility to explain what can count as true in some possible world, then I have to explain why it is that possibility has the features it would have had, were it the kind of thing that was explained in terms of truth-in-some-possible world. Just why is it that necessity has the features of universal quantification? Why is it that we can connect these modal notions with these extensional ones? Where does this notion of possibility come from, and why does it have the features that it does? Taking modal notions as *primitive* seems to ask more questions than it answers.

[2]The image that God might take place all of these possible worlds in a platonic incinerator and do away with them, without this making *any* difference to us or what is necessary or possible is not an *argument*, but it does point to something worth explaining if you are to take possible worlds the source of modal truth. I first learned of this image from Graham Priest, who used it in relation to the platonic conception of number [11].

[3]In fact, possible worlds are such slippery creatures that it is hard to see how any particular possible world does any work on its own. After all, we never specify a single possible world by specifying what is true in it. For any claim taken to be true in some possible world (a world in which England retains the Ashes after a victorious tour of Australia 2006–7) — we don't specify a *world* by considering a world in which that happens. We consider a whole raft of them. It seems that we can further specify things by adding another, claim independent of the specification so far (say, that there'll be a song by a white rapper in the top 10 of the Australian charts in December 2006). A world is specified by a consistent, complete set of sentences, not in some restricted vocabulary, but at the very least, in the vocabulary you and I are speaking. I know of no way of specifying a particular consistent and complete theory in English.

[4]There is a *reason* that 'ersatz' accounts of possible worlds are more popular than mad-dog modal realism, than it is not merely the failure of nerve among modal metaphysicians. It's that to a great many of us, possible worlds seem like maximal descriptions, propositions, theories, or whatever else, and not universes just like ours but elsewhere.

This problem is quite severe. The fact that there is such a tight circle between modality and possible worlds should tell us that something is wrong. If all that we can say about possible worlds is what we know about possibility, and all we can say about possibility is in terms of truth-in-some-possible-world, then we have said very little about what possibility consists in. For all we have said, there is only one world and the only possibly true things are the actually true things. On the other hand, for all we have said, there is a possible world for absolutely every formally consistent set of sentences.[5] Surely we can say more about possibility than that. If taking possibility as primitive and taking possible worlds as the starting point are both unsatisfactory, then we must look somewhere else for fresh insight, and for a tighter constraint on the notions in we have before us.

This paper to open up another line of attack on the problem: I will sketch a way in which we can explain possible worlds in terms of possibility and necessity, and in which the modal notions of possibility and necessity are themselves explained in other terms.

To explain the meaning of a term like "possible" or "necessary", you ought to look to a general theory of meaning. In the present day and age we do this either by looking for an account of the *truth conditions* of sentences using the term, or you look for a *semantically anti-realist* account, often couched in terms of the "use" of the vocabulary. (Though the Wittgensteinian slogan is worn out from overuse and I will leave it here.) In our case, the only plausible account of the truth conditions of modal claims seems to be the possible worlds one, which we wish to avoid. If we are to stay with a truth-conditional account of meaning we either invent our own new account for modality or become an eliminativist about modal notions and agree with Quine that modality is hopeless. Constructing your own truth-conditional semantics for modal notions seems—frankly—destined for failure. Success in the enterprise would involve the construction of a new truth-conditional meaning theory for possibility and necessity, and there is such an established market leader (possible worlds semantics) that it doesn't seem worth the bother. Possible worlds semantics looms so large that it is hard to imagine any other kind of truth conditional semantics for modality. Given that error-theory about modality is not an option (the logic of modal notions is so well-behaved[6] as to rule that out). The other option is to be semantically *anti*-realist, and to say that in an explanation of the meaning of modal vocabulary you look for something other than truth conditions.[7] This is both easier and more difficult than looking for a truth-conditional meaning theory. It is easier to attempt because we do not have a

[5]I am not arguing her that necessitarianism (that everything that is true is necessarily true) or the identification of necessity with derivability in in some formal logic are *mistaken*. (However, they both *are* mistaken.) The point is merely that they are substantive views of possibility and necessity, and that your account of the meaning of the vocabulary should have something more to say about these limiting cases. It can do this by saying something more about what things are possible worlds, or saying something more about possibility and necessity.

[6]Well, the *propositional* logic is well-behaved. One would hope that at some stage *predicate* modal logics would be well-understood.

[7]This is not to say that you must be an anti-realist about possible worlds or necessity in any other sense. To be anti-realist in *semantics*, in the vocabulary of Michael Dummett [6] is to give an account of meaning not in terms of truth conditions. That is the kind of anti-realism in play at this point. It is true that semantic anti-realism makes *open* the possibility of other kinds of anti-realism concerning the subject matter of the discourse given a non-truth-conditional meaning theory. However, it does not

market-leading non-truth-conditional account of modal vocabulary. It is more difficult, therefore, because it means that we must construct one ourselves. It is not enough to wave our hands and say (possibility claims mean ...) without checking either that the result is one of the modal logics we already know and understand, or by explaining and justifying the departures from best practice up to now. The logical features of modality are well-understood, and it will be a cost to any explanation of the meaning of modal discourse if it cannot explain these features. When we look at the traditions in logic best suited to playing a role in a theory of meaning, two candidates stand out. *Model* theory, broadly construed, provides the raw materials for a truth-conditional theory of meaning, and *proof* theory as doing the same for a semantically anti-realist theory of meaning. So, it is in the direction of proof theory that I shall turn.

2 ON ASSERTING (AND DENYING)

In which the reading of logical constants in terms of coherence of *positions* of assertions and denials is explained and defended. It is shown how a well-understood notion of *logical consequence* falls out, without starting out with a truth-conditional reading. (This will be useful because we don't want to explain things in terms of truth-in-a-possible-world, either.)

If we start with proof theory, the deliverances of a 'logic' will at least be a collection of *derivations* or *proofs* from premises to conclusions. It is one thing to have such a collection, and it is another to apply them in a theorem of meaning. The most well-worn path is to either say that a proof delivers you an assurance that if the premises are *true* the conclusion is *true*, and to keep truth at the heart of our semantics. If we want to avoid this, the obvious alternative is to read a proof as something that preserves something other than truth, like *warrant* [6]. I will not consider this option at present because its major exponents take this to justify *intuitionistic* logic rather than *classical* logic, and my target here is to give a semantically anti-realist account of traditional, *classical* modal reasoning. Instead, I wish to start with an argument from a premise P to a conclusion C, and I want to consider what it *does* tell us, if we abstract away from talk of *truth* or even talk of *warrant*.[8] Consider our options if we are in possession of a proof which leads from a premise P to a conclusion C. We do not have to *conclude* C, if we do not adhere to the starting point P. I have argued elsewhere that a fruitful way to understand the significance of a proof from premise to conclusion is that it *precludes* accepting the premise and rejecting the conclusion [13]. It does *not* mean that if you accept the premise you ought to accept the conclusion (take a case where you accept P but you ought not. There is a trivial proof from P to P. It does not follow that you *ought* to accept P). I think a case could be made that it does not even mean that it ought to be the case that if you accept the premise you accept the conclusion. (Take a case where there is an argument from a premise P to a conclusion that is the disjunction of P with something completely irrelevant, hard to remember and not worth your time thinking about. *Ought* you accept this conclusion? If so, why? For what good purpose?) The clear edict, on the basis of the proof, is that if you accept the premise then rejecting the conclusion is ruled out. Equivalently, if you reject the conclusion, then accepting the premise is ruled out. If you accept the premise and reject the conclusion, then you are in a self-defeating position. Similarly, if you *assert* the premise, and *deny* the conclusion, your acts are in conflict. A proof with a premise P and a conclusion C provides a way to

necessarily follow that this option must be taken.

[8] Let me underline at this point: we do *not* need to say that truth or warrant have no place. It is rather that our reading of the building blocks of semantic theory will start elsewhere and later *incorporate* talk of truth and of warrant.

evaluate *positions*. A proof is normative insofar as it provides a constraint on the linguistic practice of asserting and denying, and on the cognitive states of accepting and rejecting. (In this paper I will focus on the public linguistic practice of asserting and denying. Much of what I will say will apply to the cognitive states as well.)

The picture, then, is one in which the practice of making assertions and denials is normative [4, 5]. From this very thin starting point, a great deal of proof theory can be interpreted and taken up into a theory of meaning. We will consider what norms might be thought to underwrite any practice of assertion and denial. Consider the very practice of making the kinds of normative evaluations of *positions* in a discourse in which a certain number of assertions and denials have been made. At this stage it does not matter what form these assertions and denials might take – all we need is that they may be identified, and that it makes sense to say that something has been denied here, and asserted there (perhaps by the same person, perhaps by some other agent). We assume, in other words, that we are evaluating *positions* of the form $[X : Y]$, which consist of a collection X of *things asserted* and Y of *things denied* [13]. You can think of positions as keeping score in a discussion as claims are made, hypotheses are considered, suppositions entered into, and retracted, and conclusions are attacked and defended.[9] Now think of the kind of *constraint* made on positions by proof. This constraint is a *deductive* one, which me might call 'coherence.' Consider the structural features of coherence; the kinds of features of coherence we would expect to be delivered by any notion of proof. Two notions on constraint are straightforward:

IDENTITY: Asserting and denying the same thing (in the one position) is a failure of coherence. That is, the position $[A : A]$ is incoherent.

STRENGTHENING: If a position $[X : Y]$ is incoherent, then so is $[X, A : Y]$, and so is $[X : A, Y]$. In other words, you cannot get out of a hole (incoherence) by further digging (by adding more commitments).

Identity is the simplest connection between assertion and denial. (They are opposed notions.) The *Strengthening* condition tells us that incoherence is a deductive matter. If we know that X and Y clash, then adding more information does not make the clash go away. It is not that the position $[X : Y]$ is *unclear* and would become better if we added some A (either as an assertion or denial) that would make the position palateable. If there is a clash in $[X : Y]$, it remains in the presence of A.[10]

[9] Why assertions *and* denials? Can't we encode the denial of something as the assertion of its negation? Perhaps we can, but it seems like logical vocabulary can have a grip in a language where the proponents cannot use a negation as a fully fledged composable operator and judgements, and we wish to use logical consequence in our task to *explain* negation. We do not want to *assume* a composable negation operator in order to explain coherence. This point is taken up at greater length in "Multiple Conclusions" [13].

[10] The position $[\text{Tweety is a bird}, \text{Tweety can't fly} : \ ]$ is *odd* but not incoherent. The position $[\text{Tweety is a bird}, \text{Tweety can't fly}, \text{Tweety is a penguin} : \ ]$ is less odd. Adding extra information might make a position less strange, but it will not make it less incoherent. The point is that the notion of proof at hand in deductive logic is *monotonic*.

The most interesting constraint seems to be the next one, which also connects assertion and denial:

NO DILEMMA: If $[X : Y]$ is coherent, then either $[X, A : Y]$ is coherent or $[X : A, Y]$ is coherent (or perhaps both).

The motivation is straightforward. If $[X : Y]$ is a coherent position, then if it is $[X, A : Y]$ is incoherent, then asserting A is ruled out by $[X : Y]$. If $[X : Y]$ rules out the assertion of A, then it (implicitly) *denies* A, so the denial of A is coherent with $[X : Y]$. Coherence does not necessarily involve *warrant* or any rich notion. Even if we have no evidence for or against A, then either asserting or denying A will be coherent.[11]

These rules governing positions of assertions/denials suffice to define a *consequence* relation, $X \vdash Y$, where we have $X \vdash Y$ if and only if $[X : Y]$ is incoherent. This is the standard picture, at least when Y has only one member. For example, we have $A \vdash A$ (that is *identity*) and if $X \vdash A$, and $Y, A \vdash B$, then $X, Y \vdash B$ (that is *no dilemma*, or the transitivity of consequence). In general, $X \vdash Y$ means that the position of asserting each member of X and denying each member of Y is incoherent. If you assert each member of X, then don't deny each member of Y.[12] This relationship is worth calling "consequence", since if $A \vdash B$ (that is, if $[A : B]$ is incoherent), then if we *assert* A (in some position, for example $[X, A : Y]$, where X and A are asserted, and Y is denied) then since $[X, A : B, Y]$ is *incoherent* (by $A \vdash B$ and the *stregthening* rule), it follows that $[X, A, B : Y]$ is coherent if $[X, A : Y]$ is. (Note, it is not necessarily coherent, since $[X, A : Y]$ need not be coherent). It follows that any story you tell, *if* it involves A as an assertion, and if it is coherent, will not deny B. To make the story as *full* as possible, it must include B as another assertion. So, if you want to express an opinion about B, it must be *asserted*, and not denied. If $[X, A : Y]$ is a position, then $[X, A, B : Y]$ is no-cost expansion of it. This is a good thing to mean by *entails*.[13,14]

This bare notion of coherence is enough to expose the behaviour of logical connectives such as conjunction and negation. For example we can make the following *definition*, introducing the term '$\neg$' into the vocabulary:

---

[11]As of course it must: we may take up both positions, by adding A as an assertion, or adding A as a denial, and considering the consequences. We wouldn't want the fact that we have no warrant for either position to mean that we cannot coherently inquire as to their consequences.

[12]Notice, this does not mean that you should assert some member of Y. We have $A \vee B \vdash A, B$. It is incoherent to accept the disjunction $A \vee B$ while at the same time denying both A and B. However, if I assert $A \vee B$, I need not assert A or assert B—I might not know which of A or B is true.

[13]Though not, of course, the *only* thing to mean by "entails." See Beall and Restall, *Logical Pluralism* [2].

[14]Furthermore, if $[X : B]$ is incoherent (if $X \vdash B$), then any context in which X is asserted may be—at no cost to coherence—expanded with the assertion of A. In this case, we can think of X entailing B. Furthermore, if $[A : Y]$ is incoherent, any context in which Y is denied may be expanded with the further denial of A. The interesting case is $[X : Y]$ where both X and Y contain more than one statement. In that case, if we want *determinate* advice you single out a particular formula, and think of that as the "conclusion." So, the incoherence of $[X, A : Y]$ tells us that if we at least assert X and deny Y, then asserting A is incoherent: if we are to express an opinion on it, we must deny. $[X : B, Y]$ tells us, similarly, that if we assert X and deny Y then, if we are to express an opinion on B (coherently) we must assert it.

RULE 1 [NEGATION] The one-place operator $\neg$ on statements is determined by the following rules:

Negation assertion: $[X, \neg A : Y]$ is coherent if and only if $[X : A, Y]$ is coherent.

Negation denial: $[X : \neg A : Y]$ is coherent if and only if $[X, A, Y]$ is coherent.[15]

(For this definition to succeed, we need to show that the relation of coherence defined on the new vocabulary satisfies the conditions of identity, strengthening and no dilemma. It does.) We can view these rules as introducing some vocabulary that enables us to *make explicit* what is implicit in the behaviour of assertion and denial. Conjunction is another cconnective which has behaviour that may be characterised in terms of assertion and denial conditions:

RULE 2 [CONJUNCTION] The two-place operator $\wedge$ on statements is determined by the following rules:

Conjunction assertion: $[X, A \wedge B : Y]$ is coherent if and only if $[X, A, B : Y]$ is coherent.

Conjunction denial: $[X : A \wedge B, Y]$ is coherent if and only if either $[X : A, Y]$ is coherent, or $[X : B, Y]$ is coherent.

More might be said concerning whether this is a definition of a single connective or a class of connectives (this is what it takes to be a 'conjunction'). It depends on how you want to individuate things, of course.

The rules for conjunction may be understood as follows. The upshot of asserting a conjunction is just the same as for asserting the conjuncts individually. The upshot ofdenying a conjunction is more involved. I can deny $A \wedge B$ (I hereby deny the conjuction: "Australia will regain the Ashes in 2006–7 and England will retain the Ashes in 2006–7") without implicitly denying either A or B (I think that one of these conjuncts is more likely than the other but this will not lead me to *deny* either). Thankfully, we do not need at this point to say what else is asserted or denied when a conjunction is denied: we need merely to say when it is coherent to make such a denial. This is simple enough. Denying $A \wedge B$ is certainly coherent if denying A is coherent, and it is coherent if denying B is coherent. Could it be coherent in other circumstances too? Not given the *No Dilemma* rule. Suppose $[X : A \wedge B, Y]$ is coherent while $[X, A : Y]$ and $[X, B : Y]$ are incoherent. Now, since $[X : A, Y]$ is incoherent $[X : A \wedge B, A, Y]$ must be too (by *Strengthening*) and so $[X, A : A \wedge B, Y]$ is coherent (by *No Dilemma*). Similarly, $[X, A, B : A \wedge B, Y]$ is coherent too (since $[X : A \wedge B, A, B, Y]$ is a strengthening of the supposedly coherent $[X : B, Y]$). But this cannot be right, since $[A, B : A \wedge B]$ is *not* coherent (by *conjunction assertion* and the incoherence of $[A \wedge B : A \wedge B]$).[16]

The same kind of story may be told for the other connectives of classical logic. This is a roundabout way to understand the traditional sequent system for classical logic [7, 13, 15].[17]

[15]In other words, we could accept these rules:

$$\frac{X \vdash A, Y}{X, \neg A \vdash Y} \qquad \frac{X, A \vdash Y}{X \vdash \neg A, Y}$$

which are familiar from the proof theory of classical logic.

[16]This is a reading in terms of *coherence* of the derivation of the conjunction-on-the-right rule in the sequent calculus from the conjunction-on-the-left rule, using the structural rules of the system.

[17]One way to carve out the logical vocabulary is to say that it is the vocabulary for which the norms of assertion/denial can be given just in terms of coherence. We

We taken what might seem to be a roundabout route to arrive at an account of logical consequence. We have used only the minimal vocabulary of the coherence of collections of assertions and denials, in order to give an account of norms governing of logical vocabulary.[18] A *formal* notion of coherence will stay at this level of generality, giving accounts of coherence-notions on an unstructured vocabulary satisfying the conditions of *identity*, *strengthening* and *no dilemma*. Other, *non*-formal, notions will say more, by including other incoherence claims concerning the basic vocabulary. A rich theory of meaning would explain the behaviour of other units of the vocabulary either in terms of coherence, or in other terms which would have consequences for coherence. You could think, for example, of the meanings of colour claims as delivering collections of incompatibilities ($Rx, Gx \vdash$ – don't say that x is red all over and that x is green all over) and entailments ($Rx \vdash Cx$ – don't say that x is red and deny that it's coloured). It doesn't matter for our purposes here whether coherence will form a core part of the explication of meaning for other parts of the vocabulary. It suffices that whatever else the theory does, it will have consequences for the coherence of assertions and denials, which will then play a role in the account of logical vocabulary.

3 ON SUPPOSING

To begin to address issues of *modality*, we must pay particular attention to our practices of assertion and denial when we engage in modal reasoning. These practices are quite complicated. Consider what it means to *suppose* some claim P. To *suppose* is to enter the claim into the 'asserted' side of the ledger, even when it might not have any evidential right to be there. We are simply 'trying it on for size,' to see what follows.[19] Typically, supposings are made and then *discharged* once we have made deductions. For example, we conclude something about the conditional $A \supset B$ on the basis of the *supposition* of the antecedent A.

In which we see that to understand a "dialectical position" appropriate to modal reasoning, we must do more. It is not merely a collection of assertions/denials. It is a collection of assertions/denials partitioned into different *contexts*. The point is that we can make use of *positions* which are not our own (we do this already with material reasoning when we at the very least take up positions that are *stronger* than our own), and then conditional reasoning that is non-material (when the addition of an assumption or hypothesis yields a *removal* of commitments). *Supposing* provides the contexts we require. Sometimes we then combine our *actual* commitments with hypothetical reasoning (the H_2O/XYZ case is a stark example of this), and so we need the full gamut of an arbitrary colelctions. Given a context partition, we see necessity and possibility working straightforwardly: they have simple rules [which validate standard modal conditions quite readily] and they are beholden to a simple additional structure (just as one would expect given the truth-conditional models.)

When we make suppositions, this can occur in a variety of ways. We may add the claim P to our basket of assertions and proceed from there. However, sometimes when we make the addition, we do something else as well. Consider a reasoning context in which the sun is beating down on us, and there is not a cloud in the sky. And I ask you to *suppose it's raining* ... You *could* go on

may say that the act of asserting $\neg A$ has exactly the coherence significance of denying A. And denying $\neg A$ has the significance (when it comes to coherence) of asserting A. We *could* say this. The nice thing is that the addition of such a rule to our vocabulary is conservative. It can count as a *definition*. It is for "free". And to use Brandomian jargon, it makes *explicit* what is *implicit* in the language of assertion and denial.

[18] This account, as far as I can make out, fits quite nicely with the aims of Simon Blackburn's *quasi-realist* project. (I have not checked the details.) This is close to Blackburn's own account of logical consequence as inconsistent commitments. At the very least, the account presented here avoids the kinds of problems noticed with Blackburn's own account in generating the classically valid inferences [14]. This account *explicitly* underwrites the classical sequent calculus, unlike Blackburn's own, which is incomplete for classical logic as it stands. However, nothing in the account given here *requires* a quasi-realist meaning theory, let alone a distinction between statements whose assertions are *truth stating* and those which are *expressive*.

[19] Note, supposing P is not necessarily supposing that we have a *warrant* for P. Take P to be 'Q but there is no warrant for Q', for example. It is just to place under consideration a *position* in which P is asserted.

to infer from this that it's raining *and* not raining, by *adding* the claim that it is raining to your assertion basket. Typically, you don't do that. You 'consider an alternative situation'. This isn't a matter of your getting in touch with some other circumstance by means of a possible-worlds-o-scope. You delete "it's raining" from the basket of assertions, and make other modifications as necessary.

I shall not consider (in this paper at least) the kind of counterfactual reasoning in which we consider what *would* have happened were things a little different. I want to consider, instead, canonical cases of reasoning in which the notions of metaphyiscal necessity and possibility play a role in philosophical discourse. Consider the kind of reasoning that might lead one towards thinking that if water is H_2O then water is *necessarily* H_2O.

> Suppose that water is H_2O. Now suppose things are different. There are lots of oceans, rivers and lakes of XYZ on the 'Earth'. Then the oceans, rivers and lakes aren't filled with water, the lakes are filled with other stuff. Water is only a trace compound on 'Earth' if it's present at all. So, in that case, water is H_2O, it's not XYZ. It follows that water is *necessarily* H_2O.

Consider the interplay of assertions made in this little piece of reasoning. When the reasoner says "there are lots of oceans, rivers and lakes of XYZ on the 'Earth'" she is not thereby *committed* to the prevalence of XYZ around her. She is opening up a conversational context, relative to which she *is* so committed. She continues to reason with this premise (using the *commitment* in the *original* context to water being H_2O) to the conclusion that XYZ is not water, but H_2O is, and this conclusion is drawn in the 'hypothetical' context. Then she moves back to the original context, to draw the conclusion that water is *necessarily* H_2O.

How can we understand this reasoning? I think that it is clear that it is not *explicitly* committing the reasoner first-and-foremost to 'alternative worlds' which are being quantified over. (We don't read "suppose things are different, suppose that ..." as "consider a possible world in which ..."). Instead, consider the 'dialectical score' as one in which the reasoner has asserted "water is H_2O" in *one* context and "the oceans, rivers and lakes of "Earth" are filled with XYZ" *in another*. Her *score* (or commitment, or normative status) in this fragment of reasoning involves two kinds of assertions, but they are not held together in the same way that assertions in the one context are held together. (To see this more starkly, consider a piece of reasoning in which I say "... P is the case. But now consider what happens if $\neg P$ ..." I have not contradicted myself. The assertion of P and the assertion of its negation $\neg P$ are fenced off from one another when I say "but now consider what happens if".) We do not need to discuss the *detail* of this reasoning, let aloneendorse or take issue with it, because there are a number of difficult issues of names, and natural kinds which are beyond the scope of this paper. I merely wish to introduce the consideration that contexts of reasoning (flagged by markers such as "suppose" and "on the other hand" etc.) can give rise to collections of assertions and denials that are *distributed* among *contexts*. Let me make, then, a proposal:

DEFINITION 3 [POSITIONS] A *position* is then a *collection* of the form

$$[X_0 : Y_0]\ [X_1 : Y_1]\ [X_2; Y_2]\ \cdots$$

where in context i, each of X_i are asserted and Y_i are denied.

(If you *like*, you can mark one of these as the *actual* context, where taking that package on board is to take X_0 to be true, Y_0 to be false, and to take it to be possible that X_1 to be true and Y_1 to be false, and so on.) Now, just as we can see the behaviour of Boolean connectives as making explicit some of what is implicit in the practice of making assertions and denials, so we can make explicit some of the implicit behaviour of contexts. Take necessity: To assert $\Box A$ in some context would thereby commit you to A in an arbitrary context. At the very least, any position like this

$$[\Box A : \ ]\ [\ : A]$$

in which $\Box A$ is asserted in one context and A is denied in a context (a different one, or the same one, it does not matter which) is incoherent. If I take $\Box A$ to be necessary, then if I go on to consider how things could be different, if I conclude that here A fails, then this is inconsistent with my claim for A to be *necessary*.

More generally, we can connect the coherence of an arbitrary position in which $\Box A$ is asserted, to the coherence of a position in which it is A is asserted. This is managed by shifting A into any context you care to choose. On the other hand, if I *deny* that A is necessary, this denial is coherent if and only if A can be coherently denied in some other context. In this case, the context may be *new*. So, coherence rules for necessity look like this:

RULE 4 [NECESSITY] *Necessity assertion*: $[X, \Box A : Y]\ [X' : Y'] \cdots$ is coherent if and only if $[X : Y]\ [X', A : Y'] \cdots$ is coherent (for each choice of context $[X' : Y']$).

Necessity denial: $[X : \Box A, Y]\ \cdots$ is coherent if and only if $[\ : A]\ [X : Y]\ \cdots$ is coherent

In other words, it is incoherent to assert $\Box A$ in some context (relative to your other assertions and denials) whenever it is incoherent to assert A in *any* of the contexts in play (given those same assertions and denials). On the the *denial* $\Box A$ is coherent if and only if denying A is coherent some other context.

Similarly, we can see the behaviour of the notion $\Diamond$ of possibility. You could think of $\Diamond A$ as $\neg\Box\neg A$ (A is possible if and only if its negation is not necessary), but it is instructive to consider its behaviour directly. The assertion of the claim that A is possible is incoherent just when it would be incoherent to add a new context in which A is asserted. On the other hand, it is incoherent to deny that A is possible when the denial of A is incoherent in *some* context or other. We have

{That is hard to parse!}

RULE 5 [POSSIBILITY] *Possibility assertion*: $[X, \Diamond A : Y]\ \cdots$ is coherent if and only if $[A : \ ]\ [X : Y]\ \cdots$ is coherent.

$[X : \Diamond A, Y]\ [X' : Y']\ \cdots$ is coherent if $[X : Y]\ [X' : A, Y'] \cdots$ is coherent, for each choice of context $[X' : Y']$.

From these straightforward rules, standard modal logic follows. Let's consider how some of the reasoning goes.

The detail of this, and an exposition of a 'direct' form of modal reasoning, in which proofs lead from premises to conclusions, is contained in a companion paper "Proofnets for s5," in preparation.

EXAMPLE 6 It is incoherent to assert $\Box A \wedge \Box B$ and deny $\Box(A \wedge B)$. That is, $[\Box A \wedge \Box B : \Box(A \wedge B)]$ is incoherent. This is because the position

$$[\Box A \wedge \Box B : \ ]\ [\ : A \wedge B]$$

is incoherent. This is incoherent since the position

$$[\Box A \wedge \Box B : \;] [\; : A]$$

is incoherent (since $[\Box A : \;] [\; : A]$ is), and since the similar position $[\Box A \wedge \Box B : \;] [\; : B]$ is also incoherent.

Similarly, we can show that $\Box A \vdash A$ – that is, $[\Box A : A]$ is incoherent (which is straightforward), and that $\Box A \vdash \Box\Box A$. More interesting is the characteristic axiom for the strong modal logic s5, $A \vdash \Box\Diamond A$.

EXAMPLE 7 Suppose $[A : \Box\Diamond A]$ is coherent. This means that

$$[A : \;] [\; : \Diamond A]$$

is coherent. But in that *other* context, $\Diamond A$ is denied. So, this means denying A in *every* context (including our starting context), so

$$[A : A]$$

must be coherent. But it isn't. So we have a contradiction. $[A : \Box\Diamond A]$ is not coherent.[20]

We will not detain ourselves with examining the other formal properties if possibility and necessity as delivered up by this account. Suffice it to say that the 'logic' defined here is just the simple modal logic s5, which is given the simplest possible worlds semantics, described at the beginning of this paper. What we have found is a semantically anti-realist account of the modal operators that delivers up for us the kind of logical properties we were looking for.

It should also be clear that while our account is bears more than a passing affinity to possible worlds semantics, contexts are not *themselves* possible worlds, but are merely ways to regiment coherence in a way appropriate to modal discourse. It is coherent to assert A somewhere and deny A somewhere else, provided that the contexts differ. If you like, contexts are a way to regiment this kind of plurality in our practices of assertion and denial.

4 ON WORLDS

So, where have the possible worlds gone? Have we thrown the baby out with the bathwater? I started this paper concerned to *explain* possibilit in a way that did not appeal to possible worlds. This was not because I do not like possible worlds. We never considered *rejecting* the claim that a statement is possibly true if and only if it is is true in some possible world. But our semantics tells us nothing about what is true, and it tells us nothing – explicitly at least – about what is true in a world, or what a world might be.

That last claim is not quite right. The semantics and the logic as provided doesn't *quite* tell you what is true, but comes very close: It does tell you when something is *undeniable*. If $[\; : B]$ is incoherent, then we can go on to assert

In which the scope for talking about *truth* in this context is made. Not adding a truth predicate, but the concept of the 'actual world' or context relative to which truth is evaluated. It's not completely trivial, for these *positions* are not worldlike entities at all (and neither are the separate contexts inside a position). But there is a sense in which a context in a position 'points to' a world in which it is embedded. We can think of the 'world' as the idealisation or completion of what holds in a context. From *a* world it is a simple step to *all* worlds by doing the completion thing with all contexts, with contexts spewing out more contexts with each possibility claim. The limit of this process is a collection of worlds. Something is possible iff it is true in some possible world. We have our reconstruction of the biconditional, but modal notions explain worlds, rather than *vice versa*.

[20] It follows, I think, that if you *reject* the inference from A to $\Box\Diamond A$ (in that you think it's coherent to accept A while denying $\Box\Diamond A$) then you're using contexts in a different way than I have described them here. This will not mean that there is no context-first definition of your favourite modal logic. It just means that contexts will have a richer structure that we have countenanced here.

B at no cost to coherence in any context. Does it follow that B *true*? To say that B is true requires a little more, it seems, than to accept that B is undeniable. To say that B is true is to play the game—it is to stop talking *about* positions, and to take one for yourself. Given that I am playing the game (if the other options are to express ignorance about B or to deny it, instead of eschewing the vocabulary altogether) then the incoherence of [: B] suffices for me to take B to be true.[21] The position of *refraining* from taking B to be true (if it is not an *eschewing* of the vocabulary, and is a *denial*) is itself incoherent. Given the vocabulary of the language in which assertions and denials are expressed, we can think of a process of investigation as attempting to *maximise* the context. The constraint *No Dilemma* tells us that wherever we are, provided the context is coherent, then expanding it to include either asserting A or to include denying A will maintain that coherence. A sufficiently *comprehensive* story will decide for A one way or another. (Note well: it may do this in the complete absence of evidence or warrant for A.) Does it follow that A either true or false? There is an important sense in which it is. Relative to a starting context, the fullest coherent extensions decide A one way or another. If A is not decided already on the basis of what is asserted and what is denied in the starting context, then it has extensions in which A is asserted, and others in which A is denied.

Any coherent context is surrounded, then, by a family of "maximal" extensions which decide every statement. You can think of these as the "worlds" in which that context would be "correct." If the language contains negation as I have defined it above, each world-context asserts one exactly of a statement and its negation. There is a sense, of course, in which this fact is implicit in each context at its core. Relative to any context $[X : Y]$ it is not coherent to deny both A and $\neg A$ (that is, $[X : A, \neg A, Y]$ is incoherent). Similarly, it is not coherent to assert both A and $\neg A$ (that is, $[X, A, \neg A : Y]$ is incoherent). So in every context, if we extend the story far enough, we will add one, and only one, of A and its negation $\neg A$ as an assertion, or equivalently, we will add A either as an assertion, or as a denial.

So, we can think of the "actual world" from the point of view of a given context, as one if its maximal coherent extensions: one of these maximally opinionated stories. You can turn this picture on its head, of course, and think of yourself as sitting not at a context, but at a world. Unless you are omniscient, you do not know which world you are at, of course. But now, instead of thinking of what is asserted or denied at a context, you can think of what is true or not at your *world* (which is one of those worlds extending the context). A conjunction is true in this world if and only if both conjuncts are true, a disjunction is true if and only if one disjunct is true, etc. We have truth conditions. These worlds work just as Boole (or Tarski, or Davidson) intended. But here we take the soundness and completeness theorem (connecting worlds on the one hand, with proof theory on the other) as a result justifying the *truth tables*, not as justifying the *proof system*. We did not explain the meaning of

[21] This hedging about taking up the game is in order to leave open cases of pejorative discourse: perhaps the package [: x is Boche $\vee$ x is not German] is incoherent, but I need not enter the discourse according to which it's true that (x is Boche or x is not German) if I wish to eschew the use of the vocabulary. On the other hand, it may well be that to judge the coherence and incoherence of claims like these is *already* to use them. You might like to say the same sort of thing about moral or aesthetic discourse, if you think that the use of such vocabulary requires certain attitudes or capacities.

the logical vocabulary (or any vocabulary) in terms of the conditions under which something is true, but on the conditions under which something is coherently asserted or denied.

Given that we have gone this far, it is not too much more work to give an account of the class of possible worlds, too. The idealisation process continues. Suppose we have some coherent context (a context to which we're committed).

$$[X_0 : Y_0]\ [X_1 : Y_1]\ [X_2; Y_2]\ \cdots$$

Now we can think of the 'filling out' process as not only filling out one context $[X_0, Y_0]$, but also as filling out each of the other contexts at the same time. Occassionally in the filling out process, a context will find itself asserting $\Diamond A$ (or denying $\Box A$). In this case, part of the 'filling out' process will involve checking if any context is already asserting A (denying A), or if it can be coherently added. If we can't, we add a new context in which A is asserted (or A denied), and keep on filling them out. This process has a limit. And it's an *actual* world and a family of other *possible worlds*. And indeed, if $\Diamond A$ is true at a world then A is true at some world. And if $\Diamond A$ is false at a world, then A is false at every world. (Otherwise we would be incoherent, but we maintain coherence at every step.)

A system of possible worlds is a then a kind of maximal filling out of a particular position. Given a such a 'filling out' each context is a kind of 'possible world', a maximal coherent context. In this limit, the worlds are consistent and complete, and $\Diamond A$ is true in one of these worlds if and only if then A is true in *at least one*. $\Box A$ is true in one of these worlds if and only if A is true in *all* of them. They act just as you expect worlds to act.

This, again, is the soundness and completeness for the proof system relative to possible worlds models. If you *like* the completeness theorem (and it is basically a consequence of the *No Dilemma* rule that we can do this, for any formula – throw it in one side or the other, together with enough idealisation to consider the result of completing an infinite process) then the possible worlds account is *constructed* or even *vindicated* by this process out of a prior account of the logical structure of modality. The fact that the worlds constructed out of this process are so obviously *idealisations* seems to me to fit nicely with the epistemic inaccessibility of each individual world.

5 ON QUANTIFICATION

We have explained modality and its logical structure in terms other than (and prior to) possible worlds talk, and we have explained how possible worlds talk can be vindicated on this account, with worlds being *constructed* out of contexts, constrained by the modal notions of possibility and necessity. We have not yet explained *why* possible worlds talk seems so compelling. Why is it that the modal notions of possibility and necessity behave, logically, like quantifiers over objects? Isn't it so much easier to maintain that they *are* quantifiers, instead of treating them differently? In this section I want to explain why we should resist this move.[22]

It is easy to lose track of just *why* that construction works, and why $\Diamond$ is *so much like* an existential quantifier, and why $\Box$ is *so much like* a universal quantifier. I'll explain why this is the way it is, and why they are *like* quantifiers without *being* quantifiers in the language at hand.

[22] After drafting this section, I came across Stephen Yablo's "How in the World?" [16], which contains an illuminating discussion of why the move to ontological committing talk should be resisted, together with an interesting claim about

We can explain much more directly (and not through the completeness proof sketched in the previous section) that it is no accident that modality has the inferential structure of the quantifiers. It is a simple consequence of the way that coherence of modal claims relates to contexts. Look at positions, and *label* each context. For example, we might have

$$[P : Q]_a \ [Q : R]_b \ [R : P]_c$$

Now, move the labels *inside*, and feel free to shuffle the formulas around, so that all of the assertions are to the left and all of the denials are to the right.

$$[P_a, Q_b, R_c : P_c, Q_a, R_b]$$

The idea is that we have asserted P in context a, Q in context b and R in context c, and denied P in context c, Q in context a and R in context b. Now, this is very nearly *predication*. If we make the move from thinking of P true *at* a context, to thinking of P *categorising* the context, we might be more comfortable in writing the position as follows:

$$[Pa, Qb, Rc : Pc, Qa, Rb]$$

Then, the rules for the modalities become versions of rules for *quantifiers*, more or less. The universal quantifier rules, saying that

$$[(\forall x)Px : Pa]$$

is incoherent, is only a hop and a skip and a jump from saying that

$$[\Box Pc : Pa]$$

is incoherent, for any context c you like.[23] Similarly, the rule for *denial* of a universal quantifier, for which

> $[X : \forall x Px, Y]$ is coherent if and only if $[X : Pa, Y]$ is coherent for a name a not present in X and Y.

is obviously analagous to the denial rule for $\Box$, in which the denial of $\Box A$ is coherent just when we can coherently deny A in a new context.

So, the structure of assertion and denial rules modal formulas bears an obvious family resemblance to the rules for the quantifiers. This explains why they share so many logical features. It does not mean, however, that we should think of worlds as disguised quantifiers. There are differences as well as similarities which mean that the treatment of modal operators as quantifiers

possibility quite congenial to that discussed here – the idea is that an important constraint about possibility is that if $\Diamond S$ is possible, then either $\Diamond(S \wedge T)$ or $\Diamond(S \wedge \neg T)$. This is clearly explainable on our account in terms of the context shift of possibility (move to a context in which S is asserted) and the *No Dilemma* constraint on that context (here it is coherent to assert T or to deny it).

[23]We do have a distinction here, in that the context in which a modal claim is made still hangs around. The analogy is quite tight with the Tarskian semantics for predicate logic, in which the variable x is assigned a value in each assignment α with respect to which we evaluate the sentence $(\forall x)Px$. It is just that do not consult that assignment of a value to the variable while evaluating $(\forall x)Px$, since the quantifier looks after that spot. In just the same way, we do not need to look at the world in which $\Box P$ is evaluated–its value is constant over each world.

is a significant extra theoretical commitment which may well be sensible to avoid.

The difference in the original rules for modalities with their corresponding rules for quantifiers is in the presence of *names*. In the proof-first or semantically anti-realist account of quantifiers, names feature prominently. Given a categorisation of propositions or judgement contents into predicate and name, there is a raw material for defining quantifiers.[24] Given the presence of names, it is a simple matter to define whether or not a name occurs in a statement, and when the same name occurs within two different statements. In the case of contexts in a position, these features are less clear. In a single position,

$$[\ : A]\ [A, B : C]\ [C : B]$$

for example, we know that we have three contexts, in which the statements are distributed as assertions and denials. Suppose I apply the rule of necessity denial to step to the context

$$[\ : \Box A]\ [A, B : C]\ [C : B]$$

Is the context $[\ : \Box A]$ the *same* context as the original context in which A was the sole statement denied? Nothing in what we have said says that it is the same context, and nothing in what we have said says that it is different. In our discussion, the only theoretical 'work' done by contexts was in the partition of assertions and denials. Nothing in what we have done has assumed anything about the identity conditions of such 'contexts.' If we are to treat modal operators as disguised *quantifiers*, then we nominalise contexts by bringing them up to the surface for scrutiny, and in so doing, we need to decide on identity conditions. As far as I can see, there is no *need* for us to do that. The fact that there are no terms in our language that seem to *require* a nominalised treatment of contexts[25] might make us think that treating operators as quantifiers is a conceptual shift rather than an analysis.[26]

6 ON INVENTION

I will briefly conclude with an explanation of why this construction is something we might actually *use*. Why might we embark on partitioning our assertions and denials into different contexts, and *modalise*? If the account here is to be believed, then the answer will be found by examining why we *suppose* and why we suppose in ways that do not merely *add* (at least temporarily) to our stock of commitments, but also hypothetically *subtract* from them. Let me attempt a just-so story as to why even in everyday discourse this kind of practice might arise.

---

[24]The best account of the treatment of quantifiers on these lines is Mark Lance's "Quantification, Substitution, and Conceptual Content" [9].

[25]Perhaps there is one counterexample: the *actuality* operator, which requires a distinction between the *actual* world and other worlds. It seems to me that this could be handled by distinguishing the assertions and denials that are unconditionally expressed and those that are the result of *supposings*. However, that is conjecture.

[26]On the other hand, it may well be *interesting* to make this shift. There are innovations in modal logic that take nominalisation *seriously*. The result is not simply predicate logic in another guise, but something distinctive, with its own new operators: *Hybrid Logic* [3]. Hints of it can be seen in Prior's *Worlds, Times and Selves* [12].

Consider the kind of discussion between two believers engaged in practical reasoning in an environment. The point of making assertions and denials, in this situation, is at the very least, to form shared opinions and engage in joint action. Imagine two agents (call them 'you' and 'me') who are both on the look-out for predators, and one (say, you) takes there to be a leopard in the vicinity, and the other (say, me) does not, and takes the environment to be quite safe. Experience has shown us that the judgements "there's a leopard in the vicinity" (L) and "the environment is quite safe" (S) are incompatible.[27] Instead of arguing about this while there might be a leopard in the vicinity, we go back home and discuss the matter in the safety of the camp, so we can more quickly come to agreement when it matters. You point out that the bushes were rustling. I didn't notice that at the time, but I say that it could have been the monkeys. You point out that monkeys make much more noise than we heard. I agree, and come to the conclusion that there was, indeed, a leopard in the vicinity.

What do we do now? It seems that we could have the discussion all over again next time we hear the bushes rustling. On the other hand, I could go on as follows: I could, here in the safety of the camp, *suppose* that the bushes are rustling (R) and that there are no monkeys in the vicinity ($\neg M$) and conclude from this that there's a leopard around (L). Then I conclude a conditional statement, of the form "if S and $\neg M$ then L." We can agree on this *now* and keep it in mind, as it were, when we're out and about. But look at how this reasoning worked: all of it is hypothetical. When we are talking in the camp, there are no leopards about (I deny L). The bushes aren't rustling (I deny R). I am *manifestly* safe (I affirm S, incompatible with L). The reasoning is *hypothetical*, because we revise our stock of assertions and denials in the safety of the camp, in order to make a decision now about something that concerns us later.

When we make this kind of hypothetical update, we have the raw materials for modal reasoning and the modal concepts of possibility and necessity. Our capacity to engage in *invention*, making things up, and supposing things are thus-and-so, is the source of our grasp of necessity and possibility. This truism can be grasped by everyone—whether semantic realist or anti-realist. But the semantic anti-realist—who takes the practice of asserting and hypothesising to feature in semantic explanation—can use this fact to explain modal vocabulary.

♦ ◊ ♦

REFERENCES

[1] D. M. ARMSTRONG. *A Combinatorial Theory of Possibility*. Cambridge University Press, Cambridge, 1989.

[2] JC BEALL AND GREG RESTALL. *Logical Pluralism*. Oxford University Press, Oxford, forthcoming.

[3] PATRICK BLACKBURN. "Representation, Reasoning and Relational Structure: A Hybrid Logic Manifesto. 1999". *Logic Journal of the IGPL*, 8(3):339–365, 2000.

[27] I owe this particular example to Mark Lance, in discussion. We both take it to demonstrate the importance of updating.

[4] ROBERT B. BRANDOM. *Making It Explicit*. Harvard University Press, 1994.

[5] ROBERT B. BRANDOM. *Articulating Reasons: an introduction to inferentialism*. Harvard University Press, 2000.

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

[7] GERHARD GENTZEN. "Untersuchungen über das logische Schliessen". *Math. Zeitschrift*, 39:176–210 and 405–431, 1934. Translated in *The Collected Papers of Gerhard Gentzen* [8].

[8] GERHARD GENTZEN. *The Collected Papers of Gerhard Gentzen*. North Holland, 1969. Edited by M. E. Szabo.

[9] MARK LANCE. "Quantification, Substitution, and Conceptual Content". *Noûs*, 30(4):481–507, 1996.

[10] DAVID K. LEWIS. *On the Plurality of Worlds*. Blackwell, Oxford, 1986.

[11] GRAHAM PRIEST. "An Anti-Realist Account of Mathematical Truth". *Synthese*, 57:49–65, 1983.

[12] ARTHUR N. PRIOR. "Worlds, Times and Selves". In PER HASLE, PETER ØHRSTRØM, TORBEN BRAÜNER, AND JACK COPELAND, editors, *Papers on Time and Tense*, pages 241–256. Oxford University Press, 2003. New edition.

[13] GREG RESTALL. "Multiple Conclusions". to appear in the proceedings volume of the 12th International Congress of Logic, Methodology and Philosophy of Science, Oviedo 2003. Online preprint available at `http://consequently.org/writing/multipleconclusions`, 2004.

[14] JORN SONDERHOLM. "Why an Expressivist should not Commit to Commitment-Semantics". *Proceedings of the Aristotelian Society*, ???:403–409, 2005.

[15] A. S. TROELSTRA AND H. SCHWICHTENBERG. *Basic Proof Theory*, volume 43 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, second edition, 2000.

[16] STEPHEN YABLO. "How in the World?". typescript, available at `http://www.mit.edu/%7Eyablo/hiw%3f.pdf`.