4

Is the Ternary R Depraved?

Graham Priest

4.1 Introduction: Routley–Meyer Semantics

When modal logic was reinvented by C. I. Lewis early in the twentieth century, it was formulated simply as a number of axiom systems. There were certainly intuitions which drove thoughts about what should be an axiom, and what should not. But the axiom systems had no formal semantics. In modern logic, a system of proof with no such semantics appears distinctly naked. In some sense, without a semantics, we would not seem to know what the system is *about*. It was therefore a happy event when world-semantics for modal logics were invented by Kripke (and others).

In a very similar way, relevant logics were also invented (some years later) as purely axiomatic systems. It was therefore a happy event when Routley and Meyer (and others) produced a world-semantics for them. As they put it in the first of a ground-breaking series of papers:[1]

> Word that Anderson and Belnap had made a logic without semantics leaked out. Some thought it wondrous and rejoiced, that the One True Logic should make its appearance among us in the Form of Pure Syntax, unencumbered by all that set-theoretic garbage. Others said that relevant logics were Mere Syntax. Surveying the situation Routley . . . found an explication of the key concept of relevant implication. Building on [this], and with help from our friends . . . we use these insights to present here a formal semantics for the system R of relevant implication.

But critics were not impressed. The world-semantics of modal logic were certainly contentious, but no one could deny that the notion of a possible world, and of a binary relation of relative possibility, clearly had intuitive and relevant content.

[1] Routley and Meyer (1973, 199).

The Routley–Meyer semantics, by contrast, said the critics, employed devices of a purely technical nature. As one put it:[2]

> [T]he Routley–Meyer semantics . . . fails to satisfy those requirements which distinguish an illuminating and philosophically significant semantics from a merely formal model theory.

Or, as another critic put it more bluntly, commenting on the promises held out by those who took this kind of semantics to explain why contradictions do not entail everything: 'What else can one do but ask for one's money back?'[3]

The devices in the semantics which drew the ire of critics were two:[4] a function, * (the Routley Star), from worlds to worlds, employed in giving the truth conditions of negation:[5]

$$\nu_w(\neg\alpha) = 1 \text{ iff } \nu_{w^*}(\alpha) = 0$$

and a ternary relation, R, on worlds, employed in giving the truth conditions of the conditional:

$$\nu_w(\alpha \to \beta) = 1 \text{ iff for all } x, y, \text{ such that } Rwxy, \text{ when } \nu_x(\alpha) = 1, \nu_y(\beta) = 1$$

The first of these is not now in such bad shape. A plausible understanding of the Routley * may be given in terms of a primitive notion of incompatibility.[6] Matters with the ternary relation are in a less happy state. Some interpretations have certainly been put forward. The most successful so far, I think, is in terms of the notion of information flow, suggested by a connection between relevant semantics and situation semantics.[7] But even this appears somewhat tenuous.[8] The following question is still, therefore, a pressing one: what, exactly, does the ternary relation mean, and why is it reasonable to employ it in giving the truth conditions of a conditional? In what follows, I will attempt an answer to the question.

4.2 Validity

Before I turn to this matter, however, it will be useful to take a step back, and put the matter in perspective, so that we can see what is being asked of the semantics and why. Let's start at the beginning.

[2] Copeland (1979, 400). He repeats the charge in Copeland (1983), citing others who have made the similar claims: Scott, van Bentham, Hintikka, Keilkopf, and (David) Lewis.

[3] Smiley (1993, 19).

[4] It it worth noting that there are semantics for relevant logics which avoid both of these techniques. See Priest (2008, chs. 8, 9).

[5] $\nu_w(\alpha) = 1[0]$ means that the value of α at world w is true[false].

[6] As in Restall (1999).

[7] See, for example, Restall (1995), Mares (1997).

[8] See Priest (2008, 10.6.)

When we reason, we deploy premises and conclusions. Successful reasoning requires that the premises really do support that conclusion: that is, that the argument is valid. Logic is essentially the study of validity. We need a logical theory to tell us which inferences are valid and which are not. But the theory should do more than just give us two washing lists: logic, like, arguably, any science, should explain the *why* of things. Without understanding why things are valid, we are in no position, for example, to evaluate inferences that do not, as yet, appear on either list, or to adjudicate disputes about whether something should be on one or other of the lists.

Now, again, when we reason, we reason in a natural language. The language may be augmented by technical notions, such as those of physics or mathematics; but it is a natural language nonetheless. The notion of validity which logic investigates must apply to such arguments. However, in the methodology of modern logic, an account of validity is given, not for natural languages, but for various formal languages. Of course, there must be a connection. Some, such as Montague, have taken English itself to be a formal language. This is somewhat implausible. No one ever used a formal language to write poetry or make jokes. And even if English is a formal language, it is not one of those which is in standard use in logic (e.g. the first-order predicate calculus). Perhaps more plausibly, the formal languages we use can be thought of as providing reliable models (in the scientists' sense) of certain aspects of natural language. Thus, a correlation is made between certain formal symbols and certain worlds of natural language. Standardly, '$\wedge$' is paired with 'and', '$\rightarrow$' with 'if', '$\forall$' with 'for all', and so on. No one would suppose that the formal symbol behaves exactly as does its natural-language counterpart. For example, 'and' is much more versatile, often in an idiosyncratic way, than is '$\wedge$'. Nonetheless the behaviour of '$\wedge$' provides a model of certain key aspects of the behavior of 'and'; and so on for the other paired symbols. How exactly to understand this is a hard matter, which, fortunately, we may bypass here. The important point is that once we have an account of the validity of inferences in the formal language, the correlation will provide, transitively, an account for the natural-language inferences.

We are not at the why of the matter yet, though. We still want for an account of why the inferences of the squiggle language are valid or invalid. In particular, any appeal to what is valid or invalid in natural language will not provide the required explanation of why certain inferences in the squiggle language are valid or invalid; for understanding validity in the formal language was meant to deliver an understanding of natural language validity, not the other way around.

So how is this to be provided? In modern logical methodology, two very different strategies for solving the problem are commonly espoused: one is proof-theoretic; the other is model-theoretic. The proof-theoretic strategy applies, as far as I am aware, only to deductive validity. The thought is that an inference is valid in virtue of the meanings of some of the symbols involved. The meanings are constituted, in turn, by the rules of inference governing these constituents. There are many obstacles to pursuing this strategy successfully. However, this approach is not my concern here, and I mention it only to set it aside. Our focus is on the other strategy.[9]

4.3 Pure and Applied Semantics

The model-theoretic strategy is quite different, and can be applied to both deductive and inductive inference.[10] We define certain set-theoretic structures called *interpretations*, and also what it means for a sentence of our formal language to *hold* in an interpretation. A valid inference is one the conclusion of which holds in every interpretation in which all the premises hold.[11] So much is easy. But we hardly have something which gives us the why of validly yet. An arbitrary semantics of this kind will not provide it. Thus, for example, we can give a semantics for intuitionist logic in which interpretations have as a component a topological space; sentences are assigned open subsets of the space, and the sentences that hold in the interpretation are the ones that get assigned the whole space. There is no reason whatever (at least without a *very* much longer story) as to why the fact that an inference preserves taking as a value the whole of a topological space should explain how it is that the premises of a valid argument provide any rational ground for the conclusion.

If this is not clear, just recall that given any set of rules that is closed under uniform substitution, we can construct a many-valued semantics in a purely formulaic way. The values are the formulas themselves; the designated values are the theorems, and the sentences which hold in the interpretation are those which get designated values. Given some fairly minimal conditions, it can be shown that an inference is vouchsafed by the rules iff it is valid in this model theory.[12] A semantic construction which can be made to fit virtually any set of inferences whatever, does not have the discrimination required to justify any one of them.

We are forced to distinguish, then, between a model theory with explanatory grunt, and one without. The distinction is a well-acknowledged one, though

[9] In Priest (2006, ch. 11), the proof-theoretic strategy is discussed at length, and I argue that it cannot, in the end, be made to work.

[10] See Priest (2006, ch. 11).

[11] Or every interpretation of a certain kind, in the inductive case.

[12] See Priest (2008, 7.10.9–7.10.10).

terminology and a precise characterization vary. One fairly standard account of it is given by Haack, who calls the distinction one between a formal semantics and an applied semantics—or, following Plantinga (1979), a pure and a depraved semantics. For her, any construction of the kind I have described is a formal (pure) semantics. An applied (depraved) semantics is a pure semantics which has a suitable interpretation. She describes the matter as follows:[13]

I distinguished . . . four aspects relevant to one's understanding of ordinary, non-modal sentence logic; the distinction applies, equally, to modal logic. One has:

(i) the syntax of the formal language
(ii) informal readings of (i)
(iii) formal semantics for (i) (pure semantics)
(iv) informal account of (iii) ('depraved semantics')

In the case of the sentence calculus, the formal semantics (iii) supplies a mathematical construction in which one of t, f is assigned to wffs of the calculus, and in terms of which (semantic) validity is defined and consistency and completeness results proved. For all the formal semantics tells one, however, the calculus could be a notation representing electrical circuits, with 't' standing for 'on', and 'f' for 'off' . . . But the claim of the calculus to be a sentence logic, to represent arguments the validity of which depends upon their molecular sentential structure, depends on one's understanding the formal semantics in such a way that 't' represents truth and 'f' falsehood; it depends, in other words, on the informal account of the formal semantics—level (iv).

Dummett characterizes the distinction as one between a semantic notion of logical consequence, properly so called, and a merely algebraic one. For him, the difference is that, in the former, the notions involved must themselves be semantic ones, having an appropriate connection to meaning:[14]

We have examples of purely algebraic [interpretations]. For instance, the topological interpretations of intuitionist logic were developed before any connection was made between them and the intended meanings of the intuitionistic logical constants. Thus, intuitionistic sentential and predicate logic is complete with respect to the usual topology on the real line, under a suitable interpretation, relative to that topology, of the sentential operators and quantifiers. No one would think of this as in any sense giving the meanings of the intuitionistic logical constants, because we have no idea what it would mean to assign an actual statement, framed within first-order logic, a 'value' consisting of an open subset of the real line.

Semantic [interpretations] are framed in terms of concepts which are taken to have a direct relation to the use which is made of the sentences of a language; to take the most obvious example, the concepts of truth and falsity. It is for this reason that the definition of the semantic valuation of the formula under a given interpretation of its schematic letters is

[13] Haack (1978, 188f). [14] Dummett (1975, 293 of reprint).

thought of as giving the meanings of the logical constants. Corresponding algebraic notions define a valuation as a purely mathematical object—an open set, or a natural number—which has no intrinsic connection with the uses of sentences.

Though Haack and Dummett characterize the distinction differently, the difference between them is, I think, somewhat superficial. Both agree on the fundamental point: to have a model-theory with philosophical grunt, the notions used in the model-theoretic construction must be ones which either are, or may be interpreted as, intrinsically semantic ones—something, to put it in Fregean terms, to do with sense and reference. That Routley–Meyer semantics are not of this kind is essentially the critics' complaint. Thus Copeland again:[15]

The key semantical function of [Routley and Meyer's] theory [truth at a point in an interpretation] receives no more than a bare formal treatment in their writings, and we are offered no explanation of how the formal account of the logical constants given in the theory is to be related to the body of linguistic practices within which the logical constants receive their meaning. The Routley–Meyer 'semantics' as it stands, then, is merely a formal model-theoretic characterisation of the set of sentences provable in NR, no connection being exhibited between the assessment of validity and the intended meanings of the logical constants . . . it is totally unclear what account of the meanings of the logical constants is given in the Routley–Meyer 'semantics'.

4.4 Model-Theoretic Validity

None of this is to disparage the usefulness of pure semantics. Clearly, such semantics are very useful in proving various metatheoretic results concerning independence, and so on. But for a semantics to give us an account of the why of validity, its notions must be (interpretable as) semantic in an appropriate way. But what way?

When we reason, we reason about all sorts of situations: actual, merely possible, and maybe impossible as well. Deductive reasoning is useful because a valid argument is one which gives us a guarantee that whatever situation we are reasoning about, if the premises are true of that situation, so is the conclusion. This is exactly what an applied semantics is all about. Its interpretations are not literally situations—at least, certainly not in the case of actual situations: real situations are

[15] Copeland (1979, 406). Copeland also requires that the semantics deliver the classical meanings of the connectives, and especially negation. This is a much more contentious point, and he is taken to task over the claim in Routley et al. (1982b). He replies in Copeland (1983), but I think that this particular criticism sticks. The very issue here is whether classical semantics or relevant semantics get the meaning of negation in the vernacular right. See Priest (2006, ch. 10, esp. 10.9).

not sets. But the set-theoretic constructions represent situations. They do this by sharing with them a certain structure, in virtue of which a sentence holding in an interpretation faithfully represents being true in the situation represented.[16]

It should be noted, as an aside, that it would seem to be necessary on this conception of model-theoretic validity that every situation about which we reason has an appropriate set-theoretic representation, or we have no reason to suppose that valid arguments will do the job—or better, if they do, we still lack an explanation of why they do. This is not a toothless requirement. In fact, it is not even satisfied by standard model theory, couched in ZF set theory. One situation about which we reason—indeed, about which we reason when we do model theory—is set theory. And there is no interpretation the domain of which contains all sets. This is not a problem that arises if we use a set theory not so limited, such as a naive set theory.[17] But it is a problem which must be faced if we wish to conduct our model theory in ZF. To discuss strategies for doing this and their adequacies,[18] would, however, take us away from the matter at hand.

Part of what is normally involved in an account of truth in an interpretation is the provision of recursive truth-in-a-model conditions for the logical operators. The thought here is that these spell out the meanings of the operators. That the meaning of a sentence is provided by its truth conditions is one which is widely subscribed to in logic—though people may disagree about which notion of truth should be deployed here: for example, whether or not it should be epistemically constrained. The meanings of the logical operators are then naturally thought of as being their contributions to the truth conditions of sentences in which they occur. And this, in turn, is naturally thought of as provided by their recursive truth conditions. Model theory adds a twist to this picture. In model theory, we are not dealing with truth *simpliciter*, but with truth-in-an-interpretation—or, at one remove, truth in a situation. The extension is a natural one, however. To know the meaning of a sentence is not just to know what it is for it to be true or false of some particular situation. The understanding has to be one which can be ported to any situation with which one is presented. If someone knows what it would be for a sentence of the form $\alpha \wedge \beta$ to be true in some situations but not others, we would be disinclined to say that they knew what '$\wedge$' meant. Thus,

[16] Thus, to reason about a situation in which there is a desk and two books (with various properties), take an interpretation where the domain is the set containing those three objects, the extension of one monadic predicate, P, is the set containing the desk, the extension of another, Q, is the set containing the books, etc. With this interpretation, sentences of the language express things about the situation, and may be used to reason about it.

[17] See Priest (2006, 11.5, 11.6).

[18] Such as that in Kreisel (1967).

the truth-in-an-interpretation conditions of a logical operator can be thought of as specifying the meaning of that operator, and so transitively, of the vernacular notion with which is coordinated. Of course, given any particular formal language with its semantics, there is an issue of how good a model it is. There can be legitimate philosophical disagreement about this. Thus, both intuitionists and dialetheists, for example, will insist that the account of negation given in classical semantics provides a bad model for the meaning of ordinary language negation. Nonetheless, when the model is right, the truth conditions of the operator give its meaning.

But what is it for a model to be right? Hard issues in the philosophy of language lie here, such as those that divide realists and Dummettian anti-realists. But all can agree, as Haack and Dummett insist, that for the model to be an adequate one, the notions deployed in stating the recursive truth conditions must be ones which are plausibly thought of as connected with the meaning of the natural language correlate of the formal operator, and are deployed in an appropriate fashion. Presumably, there can be no objection to stating the truth conditions "homophonically", as one would normally do for conjunction. In any interpretation:

$$\nu_w(\alpha \wedge \beta) = 1 \text{ iff } \nu_w(\alpha) = 1 \text{ and } \nu_w(\beta) = 1$$

That the connection is met here is patent. (Though it should be remembered that the natural language 'and' is already to be understood as regimented in a certain way.) But appropriate truth conditions are not necessarily homophonic. Thus, in the world-semantics for modal logic, the truth conditions for $\Box$ are the familiar:

$$\nu_w(\Box\alpha) = 1 \text{ iff for all } w' \text{ such that } wRw', \nu_{w'}(\alpha) = 1$$

Indeed, such non-homophonic truth conditions may well be highly desirable; for example, if our grasp of the behavior of the operator for which truth conditions are being given is insecure, and the truth conditions are given in terms of notions our grasp of which is more secure. However, if we are giving truth conditions for some form of necessity, the following must be considered failures:

$$\nu_w(\Box\alpha) = 1 \text{ iff } f(\alpha, w) = 37$$

where f is some function from sentences and worlds to numbers, and:

$$\nu_w(\Box\alpha) = 1 \text{ iff } \nu_w(\alpha) = 1$$

The one does not even get to first base because, even if the definition gets the extension of $\Box\alpha$ correct, having value 37 is not a semantic notion at all. (Just being 37 has no connection with meaning.) In the other, truth (having value 1) is certainly a semantically relevant notion, but the posited connection between it and necessity is all wrong.

What we have seen, in summary, then, is this: applied 'semantics do not come free. The notions deployed must be intelligible [as semantic notions] in their own right, and their deployment in the framing of a semantics similarly so'.[19] We have seen some of the things this means. There is surely more to be said about matters. But we at least have enough to turn, at last, to Routley–Meyer semantics.

4.5 The Ternary Relation

If Routley–Meyer semantics are to be more that merely algebraic, the ternary relation employed must have an intuitive meaning, and one, moreover, that is plausibly connected with conditionality. How is this to be done? I claim no originality for the answer I will offer. In some sense, it is part of the folklore of relevant logic.[20] If what follows has any originality, it is in dragging the idea from the subconscious of relevant logicians into the full light of day. In the rest of this section I will describe the main idea. In the next, we will see how it is implemented.

The idea that a proposition is a function is a familiar one in modern logic. For example, in intensional logics one can take a proposition as a function from worlds to truth values. One can think of this as something like the sense of the sentence: given a world/situation, it takes us to its truth value there. The idea that the propositional content of a conditional is a particular sort of function is also familiar. In intuitionist logic, the semantic content of a conditional, $\alpha \to \beta$, is a construction which applies to any proof of α to give a proof of β. This construction is obviously a function. I want to suggest that the conditional in relevant logic is also best thought of as a function. Clearly, a conditional is something which, in some sense, takes you from its antecedent to its consequent. It is therefore natural to think of the proposition expressed by the conditional $\alpha \to \beta$ as a function which, when applied to the proposition expressed by α, gives the proposition expressed by β. The ternary relation can be understood in these terms. Let us look at the details.[21]

[19] Priest (2008, 585).

[20] For example, the thought that applying a conditional is something like functional application is found in the motivating remarks of Slaney (1990); the idea that $Rxyz$ means that $x(y) = z$ can be found in Fine (1974); and Restall (2000, 12–13 and 246–8), notes the connection between relevant conditionals and functions.

[21] After this paper was written, Beall et al. (2012) was written. This offers three understandings of the ternary relation. The one offered in this paper is the second of these.

4.6 Interpreting Routley–Meyer Semantics

In what follows, I take it that we are dealing with a positive propositional relevant logic, so avoiding issues to do with the Routley Star. Let us consider the most fundamental of these, the relevant logic B^+. I will discuss its extensions in the next section.

A Routley–Meyer interpretation for this[22] is a tuple $\langle @, N, W, R, \nu \rangle$. W is a set of worlds (situations); N (the normal worlds) is a non-empty subset of W; $@$ is a distinguished member of W; R is a ternary relation on W; ν is a function which assigns a truth value (1 or 0), $\nu_w(p)$, to every parameter at each world, w.

For $x, y \in W$, the relation $x \leqslant y$ is defined as follows: $\exists n \in N$, $Rnxy$. The worlds of a structure must satisfy the following conditions:

R0 $@ \in N$
R1 $x \leqslant x$
R2 If $x \leqslant y$ and $Ryzw$ then $Rxzw$
R3 If $x \leqslant y$ and $\nu_x(p){=}1$ then $\nu_y(p){=}1$

where, in R3, p is any propositional parameter. R3 is called the *heredity condition*, and, employing the truth conditions of the connectives and R2, it can be shown to extend to all formulas, not just propositional parameters.

The truth conditions for the logical constants of the language are as follows:

T1 $\nu_w(\alpha \wedge \beta) = 1$ iff $\nu_w(\alpha) = 1$ and $\nu_w(\beta) = 1$
T2 $\nu_w(\alpha \vee \beta) = 1$ iff $\nu_w(\alpha) = 1$ or $\nu_w(\beta) = 1$
T3 $\nu_w(\alpha \to \beta) = 1$ iff for all x, y, such that $Rwxy$, when $\nu_x(\alpha) = 1$, $\nu_y(\beta) = 1$

α holds in an interpretation if $\nu_@(\alpha) = 1$, and an inference is valid if in every interpretation in which all the premises hold, the conclusion holds.[23]

We may understand the meanings of the various notions as follows. Sentences express propositions. We do not need to worry too much about what these are; they are just whatever it is that sentences express. I write a, b, etc. for propositions. If α and β express the propositions a and b, I write the propositions expressed by $\alpha \wedge \beta$, $\alpha \to \beta$, as $a \wedge b$, $a \to b$, etc. We do not need to worry too much, either, about what, exactly, worlds are. It will suffice that they are the sort of thing

[22] As found, for example, in Routley et al. (1982a, 4.1–4.6). They do not include a base world, @. Including one makes no difference to what is valid in the semantics, and brings the semantics into line with the abstract characterization I gave.

[23] Note, then, that the situation about which we reason is (represented by) @, which is to be thought of as coming with its own raft of alternative worlds.

characterized by a set of propositions. In fact, as a matter of convenience, we may simply identify a world with a set of propositions. Each world is closed under under conjunction, and is prime (that is, whenever a disjunction is a member, so is at least one disjunct). In particular, we have, for all $w \in W$:

P1 $\mathfrak{a} \wedge \mathfrak{b} \in w$ iff $\mathfrak{a} \in w$ and $\mathfrak{b} \in w$
P2 $\mathfrak{a} \vee \mathfrak{b} \in w$ iff $\mathfrak{a} \in w$ or $\mathfrak{b} \in w$

Further, say that $\mathfrak{a}$ *entails* $\mathfrak{b}$ just if every world that $\mathfrak{a}$ is in, $\mathfrak{b}$ is in. It follows that each world is also closed under entailment.

The propositions expressed by conditionals are functions. Specifically, the proposition $\mathfrak{a} \to \mathfrak{b}$ is a function which maps $\mathfrak{a}$ to $\mathfrak{b}$. One can think of the function as a procedure which takes certain propositions into others: one which, when applied to the proposition expressed by α, gives one expressed by β. Given the conceptual connection between conditions and inference, it is natural to take this procedure to be one grounded in inference. Thus, it might take things of a logical form of α into things of a logical form of β.[24]

Now, if $x, y \in W$, let $x[y]$ be:

$\{\mathfrak{b}$: for some $\mathfrak{a} \in y$, $\mathfrak{a} \to \mathfrak{b} \in x\}$

Thus, $x[y]$ is the result of taking any $\mathfrak{a}$ in y, and applying any function of the form $\mathfrak{a} \to \mathfrak{b}$ in x. Note that $x[y]$ may not be a world. For example, there is no reason to suppose it to be prime. However, we can use it to define the relationship R on worlds as follows:

$Rxyz$ is $x[y] \subseteq z$

In other words, $Rxyz$ iff whenever the result of applying any function, $\mathfrak{a} \to \mathfrak{b}$, in x to a proposition, $\mathfrak{a}$, in y is in z.[25] Given that a function and its application are involved, why there should be a three-place relation is obvious: one place is for the function; one is for its argument; and one is for its value. The third place in the relation simply records the propositions one gets by applying the relevant functions to the relevant arguments.

It remains to say what @, ν, and N are. @ is (represents) the situation about which we are reasoning. $\nu_w(\alpha) = 1$ means that α is true at a w: that is, if α

[24] There is a well-known connection between the Routley–Meyer semantics and the λ-calculus (and combinatory logic)—see, for example, Dunn and Meyer (1997). In the λ-calculus *all* objects are functions. This suggests that an investigation of the present proposal in that context might be fruitful.

[25] The semantics given here are the non-simplified semantics. In the simplified semantics, to give truth conditions for $\to$ uniformly in terms of R, we need the condition: $Rnyz$ iff $y = z$ (where n is a normal world). See Priest (2008, 10.2). For this condition to hold on the present account, we would need: $n[y] \subseteq z$ iff $y = z$. This clearly fails, since we can have distinct z_1 and z_2 for which $n[y] \subseteq z_1$ and $n[y] \subseteq z_2$.

expresses the proposition a, then $a \in w$. The members of N are exactly those worlds, n, such that for any a and b, $a \to b \in n$ iff a entails b. Thus, since worlds are closed under entailment, we have, for any w:

P3 If $a \to b \in n$ then if $a \in w$, $b \in w$

Given these explanations of the semantic notions, both the conditions R0–R3, and the truth conditions of the connectives make perfectly good sense. That is, they are justified by these understandings.

R0 says that, in the situation about which we are reasoning, $\to$ really represents the entailment relation. That is, we may interpret sentences of the form $\alpha \to \beta$ as saying that α entails β. In other words, $\to$ gets the right meaning. (This is not a vacuous constraint, since at non-normal worlds, $\to$ may represent a different relation.)

For R1: We need to show that, for some $n \in N$, $Rnxx$: that is, $n[x] \subseteq x$. Choose any $n \in N$, and let $b \in n[x]$. Then for some $a \in x$, there is an $a \to b \in n$. By P3, $b \in x$.

For R3: Suppose that $x \leqslant y$. Then for some $n \in N$, $Rnxy$: that is, $n[x] \subseteq y$. If $a \in x$ then, since $a \to a \in n$, $a \in n[x]$. So $a \in y$. That is, $x \subseteq y$. Now suppose that $\nu_x(p) = 1$. Then if a is the proposition expressed by p, $a \in x$. Hence, $a \in y$: that is, $\nu_y(p) = 1$.

For R2: Suppose that $x \leqslant y$. Then, as we have just seen, $x \subseteq y$. It follows that $x[z] \subseteq y[z]$. (For if $b \in x[z]$, then for some $a \in z$, $a \to b \in x$. Since $x \subseteq y$, $a \to b \in y$, so $b \in y[z]$.) Thus, if $Ryzw$, that is, $y[z] \subseteq w$, it follows that $x[z] \subseteq w$: that is, $Rxzw$.

Turning to the truth conditions: P1 and P2 obviously deliver T1 and T2. For T3: Suppose that the function $a \to b \in w$, and that $Rwxy$: that is, $w[x] \subseteq y$. Then if $a \in x$, $b \in y$. Conversely, if the function $a \to b \notin w$, then it is natural to suppose that there are worlds, x and y, such that $Rwxy$, with $a \in x$ and $b \notin y$. For take a world, x, which contains just a and what it entails; we can form y by applying all the functions $c \to d \in w$ to it. Because $a \to b \notin w$, b will not be in y. (More precisely, if $b \in y$, then there is something, c, entailed by a, such that $c \to b \in w$. But in that case, $a \to b$ would be in w, since $c \to b$ entails $a \to b$.)[26]

4.7 Extensions

Positive logics stronger than B^+ are obtained, in a standard way, by adding further constraints on the ternary relation. For example, the relevant logic R^+—

[26] Essentially, this is the heuristic which is implemented in constructing the canonical model in the completeness proof for relevant logics. See Routley et al. (1982a, 4.6).

the strongest standard-relevant logic—is obtained by imposing the following constraints:

R4 If $Rxyz$ then $Ryxz$
R5 If $Rxyz$ then $\exists w(Rxyw$ and $Rwyz)$
R6 If $\exists w(Rxyw$ and $Rwuv)$ then $\exists w(Rxuw$ and $Rywv)$

The functional interpretation of R does not, in itself, make these constraints plausible. Indeed, it makes them implausible. Consider R4. This says that for any x, y, z, if $x[y] \subseteq z$ then $y[x] \subseteq z$. $x[y]$ is obtained by applying functions in x to arguments in y; $y[x]$ is obtained by applying functions in y to arguments in x. These are not, in general, the same—functional application is not commutative! Similarly, R5 tells us that for all x, y, z, if $x[y] \subseteq z$ then, for some w, $x[y] \subseteq w$ and $w[y] \subseteq z$. But the things guaranteed to be in w are the results of applying functions in x to arguments in y. There is no reason to suppose that, if these are functions, applying them to arguments in y *again* will give the same things: the application of a second function could take us anywhere. Similar comments apply to R6.

In the canonical model construction, used in the completeness proofs for these stronger logics, one invokes the appropriate axiom to show that the corresponding constraint holds in the model. (Thus, in the case of R4, one invokes the axiom $\alpha \to (\beta \to \gamma) \vdash \beta \to (\alpha \to \gamma)$). One might therefore appeal to the plausibility of such inferences about conditionality to justify the corresponding constraint. However, in this case, the semantics cannot be used to justify the properties of the conditional, on pain of circularity. In other words, they fail to explain the *why* of things, as is required for a genuine applied semantics.[27] (See the discussion in section 4.3 of appealing to the behaviour of vernacular notions.)

The interpretation of the semantics we have been looking at cannot, therefore, be used to justify the stronger relevant logics—or if they can, this requires a much more complicated version of the story than the one told here. Some might see this as a vice. Personally, I see it as a virtue. One of the embarrassments of relevant logics is their multiplicity. A challenge has always been to single out one of the multiple as the correct relevant logic. From this perspective, the discriminating nature of the interpretation of the Routley–Meyer semantics which I have offered is an advantage. It is true that it does not justify the stronger relevant logics, like R, which have been the favourites of North American relevant logicians. Instead, it justifies weaker (depth-relevant) logics, like B. These have always been preferred

[27] In a similar way, one cannot invoke the plausibility of the inference $\Box\alpha \vdash \alpha$ to justify the reflexivity of the binary R in modal logic. If justification is to be forthcoming, this has to be in terms of the *meaning* of the binary R.

by Australian relevant logicians (such as myself), because of their applications to naive truth theory and set theory.

4.8 Conclusion

Let me summarize. As we saw, for a semantics to provide a satisfactory model-theoretic account of validity, it must be possible to understand it as an applied (depraved) semantics. Routley–Meyer semantics, and especially its ternary relation, have always had a problem being seen in this light. We now see that they can be. Specifically:

1. It is perfectly natural to understand the meaning of a conditional as a function.
2. If one does this, then an intelligible meaning for the semantic ternary relation is straightforward. Essentially, it records the results of applying the function which the conditional expresses.
3. Understanding the meaning of the conditional in this way motivates the relevant logic B^+, though not stronger relevant logics, like R^+.

Critics of the semantics of relevant logic have had a tendency to see them as perverted—a turning away from true semantics. In particular, the ternary relation has been taken to be depraved: that is, debased, corrupt (*OED*). What we have seen is that it is not: its depravity is not of the vicious kind, but of the virtuous.[28]

References

Beall, Jc, Brady, Ross, Dunn, J. Michael, Hazen, A.P., Mares, Edwin, Meyer, Robert K., Priest, Graham, Restall, Greg, Ripley, David, Slaney, John, and Sylvan, Richard. 2012. "On the ternary relation and conditionality." *Journal of Philosophical Logic* 41:595–612.

Copeland, B.J. 1979. "On when a semantics is not a semantics: some reasons for disliking the Routley–Meyer semantics for relevance logic." *Journal of Philosophical Logic* 8:399–413.

Copeland, B.J. 1983. "Pure and applied semantics." *Topoi* 2:197–204.

Dummett, Michael. 1975. "The justification of deduction." *Proceedings of the British Academy* 59:201–32. Reprinted as ch. 17 of *Truth and Other Enigmas*, London: Duckworth, 1978.

[28] A version of this paper was given at a workshop of the Foundations of Logical Consequence at the University of St Andrews, January 2009. It was also given to the meeting of the Australasian Association for Logic, Melbourne University, July 2009. I am grateful to the audiences for their comments, and particularly to Jc Beall, Colin Caret, Ira Kiourti, Stephen Read, Greg Restall, Crispin Wright, and Elia Zardini.

Dunn, J. Michael and Meyer, Robert K. 1997. “Combinators in structurally free logic.” *Logic Journal of the IGPL* 5:505–37.

Fine, Kit. 1974. “Models for entailment.” *Journal of Philosophical Logic* 3:347–72.

Haack, Susan. 1978. *Philosophy of Logics*. Cambridge: Cambridge University Press.

Kreisel, Georg. 1967. “Informal rigour and completeness proofs.” In I. Lakatos (ed.), *Problems in the Philosophy of Mathematics*, 138–86. Amsterdam and London: North-Holland.

Mares, Edwin. 1997. “Relevant logic and the theory of information.” *Synthese* 109:345–60.

Plantinga, Alvin. 1979. *The Nature of Necessity*. Oxford: Oxford University Press.

Priest, Graham. 2006. *Doubt Truth to Be a Liar*. Oxford: Oxford University Press.

Priest, Graham. 2008. *An Introduction to Non-classical Logic: From Ifs to Is*. Cambridge: Cambridge University Press, 2nd edition.

Restall, Greg. 1995. “Information flow and relevant logics.” In J. Seligman and D. Westerståhl (eds.), *Logic, Language, and Computation: The 1994 Moraga Proceedings*, 463–77. Stanford, CA: CSLI Publications.

Restall, Greg. 1999. “Negation in relevant logics: how I stopped worrying and learned to love the Routley star.” In D. Gabbay and H. Wansing (eds.), *What is Negation?* 53–76. Dordrecht: Kluwer.

Restall, Greg. 2000. *An Introduction to Substructural Logics*. London: Routledge.

Routley, Richard and Meyer, Robert K. 1973. “The semantics of entailment.” In H. LeBlanc (ed.), *Truth, Syntax, and Modality*, 199–243. Amsterdam and London: North-Holland.

Routley, Richard, Plumwood, Val, Meyer, Robert K., and Brady, Ross. 1982a. *Relevant Logics and Their Rivals*, volume I. Atascadero, CA: Ridgeview.

Routley, Richard, Routley, Valerie, Meyer, Robert K., and Martin, Erroll. 1982b. “On the philosophical bases of relevant logic.” *Journal of Non-classical Logic* 1:71–102.

Slaney, John K. 1990. “A general logic.” *Australasian Journal of Philosophy* 68:74–88.

Smiley, Timothy J. 1993. “Can contradictions be true?, I.” *Aristotelian Society Supplementary Volume* 68:17–33.