Technical Report TR-ARP-3-95

Automated Reasoning Project
Research School of Information Sciences and Engineering
and Centre for Information Science Research
Australian National University

February 20, 1995

NEGATION IN RELEVANT LOGICS (How I stopped worrying and learned to love the Routley Star)

Greg Restall

Abstract Negation raises three thorny problems for anyone seeking to interpret relevant logics. The frame semantics for negation in relevant logics involves a 'point shift' operator *. Problem number one is the interpretation of this operator. Relevant logics commonly interpreted take the inference from A and $\sim A \vee B$ to B to be invalid, because the corresponding relevant conditional $A \wedge (\sim A \vee B) \to B$ is not a theorem. Yet we often make the inference from A and $\sim A \vee B$ to B, and we seem to be reasoning validly when we do so. Problem number two is explaining what is really going on here. Finally, we can add an operation which Meyer has called *Boolean negation* to our logic, which is evaluated in the traditional way: $x \models -A$ if and only if $x \not\models A$. Problem number three involves deciding which is the 'real' negation. How can we decide between orthodox negation and the new, 'Boolean' negation. In this paper, I present a new interpretation of the frame semantics for relevant logics which will allow us to give principled answers to each of these questions.

Negation in Relevant Logics

(How I stopped worrying and learned to love the Routley Star)

Greg Restall
Automated Reasoning Project, Australian National University
Email: `Greg.Restall@anu.edu.au` WWW: `http://arp.anu.edu.au/arp/gar/gar.html`

Dedicated to Richard Sylvan on the occasion of his 60th Birthday.

Relevant logic was born out of a desire to formalise an account of conditionality and of entailment which respected relevance. The central idea is that for $A \to B$ to be a true conditional (or to record a valid entailment), there must be some kind of connection between A and B. This has consequences for the semantics of negation, for it at least seems that on this policy, $A \wedge \sim A \to B$, and $A \to B \vee \sim B$ should fail. This in turn has consequences for any kind of frame semantics of the logic — if we take the failure of $A \wedge \sim A \to B$ to mean that there is some point x is a model for which $x \models A \wedge \sim A$, and $x \not\models B$, then we seem committed to points which allow inconsistency. Further, we seem to need points y at which $y \models A$ but $y \not\models B \vee \sim B$. Some points must reject commonly accepted logical truths.

The traditional way of ensuring this is in the semantics of relevant logics is to evaluate negations in what appears to be a non-standard fashion. We say that $x \models \sim A$ not when $x \not\models A$, but rather, when $x^* \not\models A$, where x^* is a point somehow related to x. The operator * was introduced to relevant logic by Routley and Roultey [23]. If $x \neq x^*$, then certainly we can get both $A \wedge \sim A \to B$ and $A \to B \vee \sim B$ to fail, but there is a price. The price is the obligation to explain the meaning of the operator *. And some have argued that this price is too high to pay [6].

However, the troubles for the interpretation of negation do not end there. Traditional relevant logics not only reject the paradoxes of implication like $A \wedge \sim A \to B$, and $A \to B \vee \sim B$, but they also reject what is often called the *disjunctive syllogism*: $A \wedge (\sim A \vee B) \to B$. This is certainly a problem, for without a doubt we do use the inference from A and $\sim A \vee B$ to B (at least occasionally) in our reasoning. And it certainly does seem that we are reasoning validly when we do so. How can we understand this?

Finally, negation gives us one more problem for the interpretation of relevant logics. Given that 'orthodox' relevant negation invokes a 'point-shift' in its evaluation, it seems that there is another operator, which Meyer has dubbed *Boolean negation*, and which has the more traditional interpretation: $x \models -A$ if and only if $x \not\models A$. Given this clause, we must ask ourselves: Is this a legitimate connective for relevant logics? If so, how does it relate to 'orthodox' negation?

In this paper I sketch an interpretation of the frame semantics of relevant logics which provides principled answers to each of these questions. I will argue for the following answers to our questions.

1. Against Copeland [6], I argue that the frame semantics for relevant logics is (when properly interpreted) an *applied semantics*, not merely a *pure* one. In particular, sense can be made of the interpretation of negation, and the Routley star function.

2. Against the relevantist of Belnap and Dunn [11], I argue that in an important sense, disjunctive syllogism is valid. Not that it is valid *sometimes*, but rather, that it is valid, in a sense to be explained later.

3. Finally, against Meyer [13], I argue that under the intended interpretation of relevant logics, *real* negation turns out to be de Morgan negation (or something very much like it) and *not* Boolean negation, and that under the intended interpretation, Boolean negation is a *senseless* connective.

The answers when taken together define a new position in the philosophy of relevant logic. These answers stem from a reading of the Routley semantics, partly inspired by Dunn's work on the semantics of negation [8], and by my own work connecting the semantics of relevant logic with work on information flow and situation semantics [19]. The guiding considerations will be semantic, not syntactic. That is, we will not take proof theoretical considerations as given and only then construct a 'semantics' which fits the proof theory. Instead, we will take our semantic considerations as primary, and see what this means for validity. As a result, the results vary from original work in relevant logics. This paper is not an apologetic for the particular formal systems **E** (the focus of the magisterial *Entailment* [1], [2]) or **R** (the focus of Dunn's excellent survey article on relevant logics [7]). The situation is similar to that of classical modal logic given the appearance of possible worlds semantics. Considerations about the structure of possible worlds could dictate logical investigations, and as a result, systems like **S4** and **S5** which had perspicuous possible worlds semantics were favoured above systems which did not. Furthermore, different modal logics were discovered and found a permanent place in the study of modal logics.

The situation is similar here. Our semantic considerations will not necessarily point to the relevant systems **R** and **E**, and in fact they may point away to a whole range of alternative systems. However, the results of this study do serve to support work on relevant logic, just as various interpretations of the frame semantics of modal logics support Lewis' work on strict implication.

In what follows I will first introduce the notions required to interpret the frame semantics for relevant logics. Then, with these notions in hand, we will see how they naturally give rise to the frame semantics for relevant logics. After that, we will consider the issues of disjunctive syllogism and Boolean negation. But first we will look at the semantic primitives of our study.

1 Ways

There are ways the world could be. There is a way the world is. There are also ways parts of this world are, and ways parts of this world could be. I take it that these are relatively uncontroversial claims, when taken at their face value. Of course, when taken as a part of a philosophical system, they can be much more controversial. For example, the *modal realism* of David Lewis [12] is a controversial philosophical view which identifies ways worlds could be with worlds, each the same sort of entity as the actual world. I do not advocate that approach of fleshing out the concept of a 'way,' for reasons which should become fairly obvious later in the paper. However, I do not wish to sketch a rival ontology. That task is best left for another occasion. For now, I am satisfied with the uncontroversial claim that there are ways the world could be, and I will use these entities (whatever they turn out to 'be') in the applied semantics which will be the focus of our attention. However, in the case of relevant logics, ways the world could be are not enough. We must broaden our horizons a little.

Just as there are ways the world could be, there are ways the world couldn't be. For example, the world couldn't be such that there are square circles. It also couldn't be such that angles are trisected with ruler and compass. Such things are impossibilities. And here we have examples of *different* ways the world couldn't be, since inconsistent ways can be inconsistent differently. Classical possible worlds formalisms have no place for inconsistencies because they are purely a systematisation of the possible. But that isn't to say that there aren't any ways the world couldn't be. (After all, if there weren't any ways the world can't be, then it would surely follow that anything is possible; and while some of my friends and colleagues take this to be true [15], I don't.) So, let us agree that there are ways the world can't be. It ought to be clear that this is incompatible with Lewis-style modal realism, as a simple example will suffice to show. I take it that the world can't include a square circle. So, there is a way the world can't be which involves a square circle. So, by Lewis-style modal realism, and treating ways the world can't be on a par with ways the world can be, there is some world in which there is a square circle. But then, there is a square circle. Which we have assumed is impossible.[1] But it must be emphasized that this is not a shortcoming of our view of ways the world couldn't be. Rather, it is a shortcoming of Lewis' account of ways the world could be that it cannot also uniformly treat ways the world couldn't be.

For one final generalisation, we ought to note that classical modal logic deals only with *total* ways the world could be, and it has no place for merely

[1] The problem is the quantifier shift involved in modal realism of the form exhibited by Lewis. We can move from $(\exists w)\big(\text{In } w\text{: } (\exists y)Fy\big)$, to $(\exists y)Fy$ for many predicates F (not all). If we allowed this kind of shift, we could move from the truth of impossibilities in ways the world couldn't be, to the truth of a impossibilities simpliciter.

partial ways, which describe limited parts of the world. This is another shortcoming when it comes to conditionality and entailment, as many have noticed.[2] I have argued elsewhere that when discussing entailment and conditionality we ought take account of partial ways [18].

So we will consider both ways that parts of our world could be, and ways that parts of our world couldn't be, together with ways our whole world could be, and ways our whole world couldn't be. All of these entities will be called *states*.[3] States admit of both incompleteness (ways need not answer every issue with a 'yes' or a 'no') and inconsistency (some ways might answer both 'yes' and 'no' to some issues). Note that we are speaking of 'partial states' in the sense of states which are incomplete — not in the sense of parts *of* a state. I have no idea what mereological relations states enter into.

It is clear that there is need for inconsistent and incomplete states, if we are to discriminate between conditionals with inconsistent antecedents or necessarily true succedents. As an example, we ought to distinguish the following two conditionals.

If I trisected an angle with ruler and compass, I would become famous.
If I trisected an angle with ruler and compass, grass would be purple.

The first seems true. If I trisected an angle with ruler and compass, I would publish the result and there would be a great deal of rethinking of Euclidean geometry to be done. The second conditional does not, at least at first glance, seem to be true. It appears that there is no connection between trisecting an angle and the colour of grass. My becoming famous follows from my trisecting an angle in a way that grass being purple doesn't. One way to understand this distinction is to note that there there are (undoubtedly inconsistent) states involving my trisecting an angle and my fame, without also involving grass being purple.

Much more could be said here, but it would only rework old discussions [1], [7], [22], which are much better than anything I could attempt here. The main point is that even if A and A' have the same truth value in all possible worlds we do *not* necessarily have $A \to B$ if and only if $A' \to B$, and nor do we have $B \to A$ if and only if $B \to A'$. To evaluate conditionals, we need more than just possible worlds. And it is my contention in this paper that *states* are just what is needed.

It would be interesting to chart the connections between states as we have sketched them and other entities like *events*, *times*, *locations*, *objects*,

[2] Situation semantics [4] as an industry is born largely out of the need to move from total possible worlds to more partial objects such as situations in order to model various phenomena which require the 'informational sensitivity' of situations over and above mere possible worlds.

[3] In an earlier draft of this paper I tried calling them 'ways.' However, as Pragati Jain pointed out to me, it is a lot harder to construct grammatically correct sentences this way. Calling these things 'states' is a lot easier.

states of affairs, *propositions*, and many other things besides. However, this is neither the time nor the place for that kind of metaphysics. Suffice it to say that a coherent comprehensive view of states ought also tell us how these things fit together. For now, we will use *states* as the points in our frames for relevant logics.

2 Positive Logic

We have a collection W of states. There is a relationship of *involvement* $x \leq y$ between states. To say that $x \leq y$ is just to say that being y includes being x, or that y involves x. We can assume that $\leq$ is transitive and reflexive, and it might in fact be a lattice ordering (for any two states there is a least state which involves them both) though I will not assume that here.

Fundamental to our semantics is a relation $\models$ between states and claims in a particular language. We read '$x \models A$' as 'according to x, A is true,' or equivalently that x (by itself) carries the information that A. The details of the language in question are irrelevant, except that it include the usual binary connectives '$\wedge$', '$\vee$' and '$\rightarrow$', and the unary connective '$\sim$'. Given the nature of involvement, it follows that if $x \leq y$, and if $x \models A$ then $y \models A$. By expanding a state to include more, you don't lose any information you were already given. If being y involves being x, then information gleaned from x is certainly given by y as well, because y includes x. As an example the state of this paper is included in the state of all of my, and this is included in the state of the whole world. So, there are non-trivial inclusion relations between states.[4] This makes them differ from possible worlds.[5] If a world x is included in a world y in this sense, then x must include y as well. For worlds are complete and consistent, so if $x \leq y$, and $y \models A$, then $x \not\models \sim A$ (lest y be inconsistent, since whatever x carries is also carried by y) so by the completeness of x, $x \models A$. This argument will become important later, when we come to discuss the place of Boolean negation.

For now we will consider how complex claims (like conjunctions, disjunctions and conditionals) are related to their components. Conjunctions and disjunctions are simplest. We simply require the following equivalences.

- $x \models A \wedge B$ if and only if $x \models A$ and $x \models B$
- $x \models A \vee B$ if and only if $x \models A$ or $x \models B$

The condition for conjunction is unproblematic. But as is often the case, disjunction needs a more extended discussion. There is no problem with reasoning from $x \models A$ or $x \models B$ to deduce $x \models A \vee B$ (provided that we

[4] It is important to note that this inclusion relation need not be the mereological 'part of' relation on states, though it perhaps *could* be that relation.

[5] Only if worlds are consistent and complete. But that is what most people take a 'world' to be in this context.

read '$\vee$' as inclusive disjunction). If a state carries the information that A, then it also carries $A \vee B$. The problems occur in the other direction. Perhaps a state could support a disjunction without supporting either disjunct. We do not have space for a complete discussion of this point here, as this is a paper primarily about negation, and not about disjunction.[6] Instead, I will just sketch the way ahead. The crux of the matter is this: We can simply construe 'states' so that the only way a state makes a disjunction true is to make one disjunct or the other true. On this construal, states are maximally specific about their own subject matter. For example, the way my papers are seems to be maximally specific in just this sense. This distinguishes states from pieces of information or states of affairs as they have been traditionally construed. Certainly there is a state of affairs of 'my car being either green or blue,' which will not satisfy the disjunction condition (on a suitable reading), but this state of affairs does not count as a 'state' on our account. The state of my car, counts as a state, but it is plausible that this does satisfy the disjunction condition. For my car is green or blue just when my car is green or my car is blue.

We must be careful here, for there is an argument against the disjunction condition, which goes as follows. Suppose the disjunction condition holds, and that $x \models A$. Then we have $x \models (A \wedge B) \vee (A \wedge \sim B)$ and hence, $x \models A \wedge B$ or $x \models A \wedge \sim B$, by the disjunction condition. It follows that if $x \models A$, then either $x \models B$ or $x \models \sim B$, putting paid to our claim that states can be incomplete. What goes wrong here? There are two plausible places to stop the deduction. One, at the disjunction property, and the other at the inference from $x \models A$ to $x \models (A \wedge B) \vee (A \wedge \sim B)$. And on our account of states, this latter step is quite unwarranted, since we may not have $x \models B \vee \sim B$, as not all states need give all necessary truths — states can be incomplete. So, we are well within our rights to construe states as maximally specific about their subject matter, as this argument need not deter us.[7] As a final reassurance, we may note that we haven't lost anything with our treatment of states as satisfying the disjunction condition, for we can model states of affairs (or 'non-disjunctive states') as sets of states if we so choose. The state of affairs of my car being either green or blue can be modelled by the set of all states in which my car is either green or blue. This will not satisfy the disjunction condition (if we say a state of affairs makes something true just when all of 'its states' makes that thing true), and it has the sorts of properties we would require of 'non-disjunctive states.'

That is enough of disjunction and conjunction. For a conditional, we rely

[6] I argue the details elsewhere [18]. The argument is couched in terms of truthmakers instead of 'states,' but the points are similar.

[7] This makes states analagous to the situations of Barwise and Perry's situation theory [4]. In fact, Barwise and Perry's *abstract situations* are one possible account of states. However, I do not want to identify states with abstract situations, not least because I do not think that states are set theoretical constructions.

on a ternary relation R on states. $Rxyz$ holds when all of the information given by x, after applying the information in y, holds in z. You are to think of the information supported by x as 'data,' and that supported by y as information to be applied to the data. If the results are all supported by z, then $Rxyz$ holds. For example, y might be a state which includes the laws of physics, and x might be an initial state of a system. Then if z includes all consequent states of the system, we would take it that $Rxyz$ holds. However, y might be much more partial than the comphrehensive choice we made, and z might contain a lot less. Given this account of R we have the following sensible condition for the conditional.

- $x \models A \to B$ if and only if for each y, z where $Ryxz$, if $y \models A$ then $z \models B$

Let's see whether this condition makes sense. Firstly, if $x \models A \to B$, and we have y and z where $Ryxz$ and $y \models A$, then since $x \models A \to B$, and all information got by applying that in x to y is given by z, then we must have $z \models B$ as required. Conversely, if for every y and z where $Ryxz$, if $y \models A$ then $z \models B$, it certainly appears that x licences the inference from A to B. So, the clause does, in fact, make sense.

To ensure that if $x \models A \to B$ and $x \leq x'$ then $x' \models A \to B$, we need to relate $\leq$ and R. A sufficient condition for preservation of conditionals goes as follows — if $Ryxz$, $x' \leq x$, $y' \leq y$ and $z' \geq z$, then $Rx'y'z'$. I leave it to the reader to confirm that this condition coheres well with our interpretation of R.

Of course, there will be certain conditions that R satisfies beyond these, but it is not our place to consider those here — because this is a paper on negation, and not one on conditionals. However, we can simply note that there's an interpretation of R in terms of 'application' which works, and which provides an applied semantics. We need not flesh out all the properties of R in order to provide an interpretation. Just as in the possible worlds semantics explanation of physical possibility in terms of an accessibility relation P where $w_1 P w_2$ is interpreted as "all of the physical laws true in w_1 are true in w_2." We need not pin down all of the behaviour of P to agree that it provides an applied semantics. Similarly, for our interpretation of R in terms of application. We can agree that this is an applied semantics, without worrying about the specific behaviour of R.

So, what we have is an applied semantics. We have reasoned about the relationship between states and certain sorts of claims, so far conjunctions, disjunctions and conditionals. Leaving aside the exact behaviour of R as a matter for discussion elsewhere[8] it is sufficient to note that the evaluation clause for the arrow captures a meaningful conditional. The important factor is that there are some conditional forming operators (like those we discussed

[8]Check, for example, Routley et. al.'s [22], or my [19].

earlier) which must appeal to the 'finely-grained' resources of states (or something like them) in their semantics, because they discriminate between propositions which are true in all the same worlds. These conditionals are 'state functional' but not 'world functional,' in the sense that we need the finer structure of states in order to discriminate between conditionals with necessarily truth-functionally equivalent antecedents. Restricting ourselves to worlds results in the ironing out of too many differences.

We must leave it up to further research to decide whether all aspects of natural language conditionals are captured in this kind of scheme, or whether you need to appeal to neighbourhood semantics of a Lewis-Stalnaker style-conditional. Whatever we say about natural language conditionals, a ternary relational semantics is clearly applicable for entailment.[9]

That is enough of positive logic. We must move on to negation, for this is where most of the problems in the interpretation of relevant logic have arisen. We will find that the detour through the interpretation of the postive logic will stand us in good stead for dealing with the interpretation of negation.

3 Negation and Compatibility

The first point to note is an important one. We have been considering states, and the relationship $\models$ between states and claims in some language. When it comes to negation, it should be obvious that negation ought *not* respect the 'classical' clause

- $x \models \sim A$ if and only if $x \not\models A$.

This goes wrong in just too many ways. First, as states may be incomplete, we ought not have $x \models \sim A$ whenever $x \not\models A$. The state x might only be a state of a proper part of the world, so it might 'have nothing to say' about A. Further more, as states may be inconsistent, we ought allow $x \models A$ and $x \models \sim A$. But this is not compatible with the traditional clause either. So, something must change. Instead of forcing the issue of whether or not $x \models \sim A$ to depend on whether or not $x \models A$ alone, we should allow it to depend on more.

Consider what it is for $x \models \sim A$ and $y \models A$ to hold. Then x and y are incompatible, because according to x, A is false, while according to y, A is true. (This is more than the case where $x \not\models A$ and $y \models A$, for then x and y may still be compatible, for x may be incomplete 'about' A — x may neither

[9]The problem here is the fact that we can prove that if $\vdash A \to A'$ (that is, if $x \models A$ then $x \models A'$ for every x) then $\vdash (A' \to B) \to (A \to B)$. That is, we can strengthen antecedents, which is a move often said to fail for certain natural language conditionals. There is a serious question of how to handle this in our context. I have discussed a little of this elsewhere [19], and Richard Sylvan has unpublished work on relevant conditional logics which deny these kinds of moves. However, there is much more to be said on this matter.

support A nor $\sim A$.) Conversely, if $x \not\models \sim A$, then there is some state (say y) that is compatible with x such that $y \models A$. Otherwise, if there's no y that is compatible with x where $y \models A$, then it does seem that x 'rules A out' after all, for there is no state that is compatible with x which supports A. But then, if x rules A out, we should take it that $x \models \sim A$. We should be clear here that this doesn't conflict with the possible incompleteness of states — we can't infer that a state x 'rules A out' just from $x \not\models A$, for x may be incomplete, and it may rule out neither A nor $\sim A$.

So, given a relation C of compatibility between states, we may evaluate negations in a simple manner. We say xCy if and only if y is compatible with x. Then we have the following evaluation condition for negations.

- $x \models \sim A$ if and only if for every y where xCy, $y \not\models A$

We would expect C to interact with $\leq$ in the following way. If xCy and $x' \leq x$ and $y' \leq y$ then $x'Cy'$. That is, if y is compatible with x, then any part of y is compatible with any part of x. This is quite a plausible condition, and if C satisfies this condition, then if $x \models \sim A$ and $y \geq x$ then $y \models \sim A$ too, as we would hope to see. There is little new in the formalism. The semantics for negation in terms of a binary accessibility relation has been known for some time. I point the reader to a recent example [8].

This provides us with an applied semantics for negation. It important to note that it is not *queer*, even though the evaluation clause is not the classical one. Despite this, the clause is exactly the kind of semantics we would expect to see, given that states are sometimes inconsistent and sometimes incomplete.

This is my 'baseline committment' about negation and its relationship with states. More can be said about the kinds of conditions C will satisfy, but I am less certain of these than of the conditions we have seen so far. The discussion ahead is not intended to be 'the complete definititive story' about negation, but only one way that our account of negation can be developed.

So, let us consider the kinds of conditions C might satisfy. Some decisions are simple. Given our interpretation of states, we do not want C to be reflexive, since some states are inconsistent. We would expect there to be states x such that xCx fails, for if xCx, then if $x \models \sim A$ we cannot have $x \models A$, as you would expect, since x is 'self-compatible,' or as we will say from now on, x is *consistent.*

Another possible constraint on C which turns out to be inappropriate is as follows. We could say that xCy if and only if there is a consistent z where $x, y \leq z$. That is, x is compatible with y just when there is a consistent state combining them. This conception of defining 'consistency-with' as the consistency of the union of the objects in question is appropraite in some situations but it turns out to be mistaken here. For we take xCy to say that nothing given by x is rejected by y. We may well have xCy even when x is not compatible with itself — for even if x is inconsistent (suppose

$x \models A \wedge \sim A$) x may still be compatible with y because y may support neither A nor $\sim A$. Then x doesn't contradict y because y says nothing about A, one way or another. Given this interpretation, we may have states which are themselves inconsistent but which can be compatible with other states. This situation is readily understandable — I may not contradict you, even if I have contradicted myself on some topic, for you may not make any assertions about the topic in which I contradicted myself.

So much for the *via negativa*. Now we can consider some properties which it would be plausible to assume that C has. For example, compatibility certainly does *seem* to be symmetric. That is, if xCy then yCx. Given that this is the case, it follows that whenever $x \models A$, then $x \models \sim\sim A$ (which we abbreviate to $A \vdash \sim\sim A$) as the reader is encouraged to verify.

Another possibility for C is that it be directed — that is, for every x, there is a y where xCy. This makes sense if we ignore the absolutely inconsistent state, (or if we have both the absolutely inconsistent state and the absolutely vacuous state). I prefer the former route, and given that we do that, and if we allow our language to have two propositional constants, $\top$ and $\bot$, where $x \models \top$ alstates, and $x \models \bot$ never, then directedness amounts to the validity of $\sim \top \vdash \bot$.

Finally, we might want C to be what we will call 'convergent.' The idea is that if anything is consistent with x, then there will be a maximal point consistent with x. So, if $(\exists y)(xCy)$, then $(\exists y)(xCy \wedge (\forall z)(xCz \Rightarrow z \leq y))$. This condition amounts to the validity of $\sim(A \vee B) \vdash \sim A \wedge \sim B$, as you can verify. This condition is perhaps not as plausible as the others. If x is consistent, and incomplete (say x is a state that a small part of this world is) then the condition provides us with a state y which collects together everything consistent with x. And this state will be wildly inconsistent about everything on which x does not dictate. If $x \not\models A \vee \sim A$, then $y \models A \wedge \sim A$. Perhaps this generates too many odd states for some people's comfort. But then, the step from $x \models \sim(A \vee B)$ to $x \models \sim A \wedge \sim B$ is appealing, and if we do think that $\sim(A \vee B)$ entails $\sim A \wedge \sim B$ in the strong sense that any state involving the former also involves the latter, then we are committed to such states (given the evaluation clause for negation we have already seen). I am not going to argue here for the absolute necessity of this condition. Rather, I want to simply note that it is a condition for which one might find a place.

Now, all of these conditions taken together gives us a way to understand the Routley star. For if every state x has a corresponding state with respect to which it is compatible, and is maximal among all such, we can call that point x^*. Then symmetry of C means that $x \leq x^{**}$, and if we impose $x^{**} \leq x$ (which is necessary, in this context, for $\sim\sim A \vdash A$) we have all of the conditions of the Routley star from traditional relevant logics. In addition, to that, the clause for negation familiar to relevant logicians

- $x \models \sim A$ if and only if $x^* \not\models A$

is equivalent to our clause involving C. This is because $x^* \not\models A$ if and only if for each y where xCy, $y \not\models A$, for we have xCy if and only if $y \leq x^*$. In other words, x^* is a 'cover all' for each state y compatible with x. The Routley star is a simplification of our compatibility clause for negation when we assume that C is symmetric, directed and convergent. Now this formal result in and of itself is not new. The details were worked out by Dunn [8]. What is new is the use we will make of it.

The significant result is that we have a reading of the Routley star which makes a great deal of sense. Given that we take the compatibility relation to be symmetric and directed (which it certainly seems to be) and convergent (which is more controversial, but certainly not out of the question, when it comes to 'states') then the star semantics for negation is a simple retelling of the compatibility semantics. Given that the compatibility semantics makes sense and is an applied semantics, it follows that its simple retelling, involving the Routley star, also makes sense, and it too is an applied semantics. So, our first question is answered. The semantics for negation involving star is not a 'hack' designed simply to evaluate negation in a state congenial to the strictures of relevant logic. It is a special case of a general applied semantics, which we have seen is the appropriate evaluation for negations in our context.

It ought to be clear that this semantic structure does invalidate the disjunctive syllogism. There are states x such that $x \models A \wedge ({\sim} A \vee B)$ but $x \not\models B$, because $x \models A \wedge {\sim} A$. If x is an inconsistent state, then if A is true only at x and the states containing x, we have $x \models A$. But take any point y where xCy. We cannot have $x \leq y$, lest we have xCx. So, $y \not\models A$, and hence $x \models {\sim} A$. But we need not have $x \models B$, so disjunctive syllogism is invalid. This is just as we would expect, given that states may be inconsistent. However, disjunctive syllogism does seem to be valid in some sense. Our task is to determine if our semantics can provide for us a sense in which it is valid.

4 Worlds and Disjunctive Syllogism

The literature in relevant logic is equivocal when it comes to the interpretation of points in frames for relevant logics. You see them referred to as worlds (whether possible or impossible ones) [21], theories [14], or setups [22], [23]. I will not enter the discussion of the relative merits of these proposals, except for the first, for the identification of points in the frames of relevant logics has done the interpretation of relevant logics a disservice. Whether the worlds are said to be possible or not, they are not what are modelled by the points in frames of relevant logics. Points in our model structures can be thoroughly incomplete. They can be in non-trivial inclusion relations. One would expect that truths in a world would satisfy more closure conditions than the few that are imposed on points in our model structure. For what is true at a state is only closed under (relevant) entailment, and the laws for conjunction and disjunction. Given any atomic proposition p, the set $\{A : \vdash p \to A\}$ satisfies these conditions, and it is far from plausible that this set picks out a world. Even if you don't believe that every world must be consistent and complete with respect to negation, they must satisfy some sort of closure condition above and beyond those needed for states. So the move to take all points in our model structure as *worlds* seems to give a mistaken impression.

However, the use of the term is understandable, for two reasons. Firstly, possible worlds semantics are very popular. We use 'worlds' all the time, so it is no surprise that when you get a model structure which looks familiar (except for the ternary relation, and the odd clause for negation) you think of the points as worlds. Secondly, the actual world at least ought to play some part in our logic. We would like there to be a way we could read off the 'standard model' what is true simpliciter. That is, what is true in the actual world. So, this leads us to an important point — we ought to consider how we can fit worlds (or at the very least, the actual world) into our model structure. We have part of our answer already — not every state is a world, or as it is better to say, not every state picks out a unique world. But some states might pick out worlds. Or more generally, some collections of states might pick out unique worlds. Given our collection W of all of the states, we can consider $\mathcal{W}$ the collection of all of the states of parts of this world. This collection $\mathcal{W}$ is closed *downwards*, in the sense that if $x \in \mathcal{W}$, and $y \leq x$, then $y \in \mathcal{W}$ too, obviously. And this collection $\mathcal{W}$ is an obvious surrogate for the actual world in our model structure.

There is more that we can say about $\mathcal{W}$ than it being closed downwards. If we make the (commonly held) assumption that the world is self-compatible, then all of the elements of $\mathcal{W}$ will be pairwise compatible. To see what work this does, let $\mathcal{W} \models A$ be the claim that for some $x \in \mathcal{W}$, $x \models A$. Then $\mathcal{W} \models A$ if and only if A is true in the actual world. The condition that all elements of $\mathcal{W}$ be pairwise compatible ensures that we never have $\mathcal{W} \models A \wedge {\sim} A$, as is simple to verify. For if $\mathcal{W} \models {\sim} A$, then $x \models {\sim} A$ for some $x \in \mathcal{W}$. So for any $y \in \mathcal{W}$, we have xCy, and so, $y \not\models A$ by the evaluation condition for negation. Thus, $\mathcal{W} \not\models A$.

That is enough to encode the assumption that the world is consistent. To encode the completeness of the world we must do more work. A plausible first thought[10] is to require that $\mathcal{W}$ be maximal. That is, if yCx for every $y \in \mathcal{W}$, then $x \in \mathcal{W}$. If $x \notin \mathcal{W}$, there is some $y \in \mathcal{W}$ such that ${\sim}(yCx)$. But this is mistaken. Consider the simple frame pictured below. Containment is read upwards (so if x is under y in the picture, $x \leq y$) and we set xCy to be simply $(\exists z)(x \leq z \wedge y \leq z)$. That is, we assume all points are self-compatible, and compatibility generally is equivalent to having a common

[10] I call it plausible mainly because it was *my* first thought.

'ancestor.'

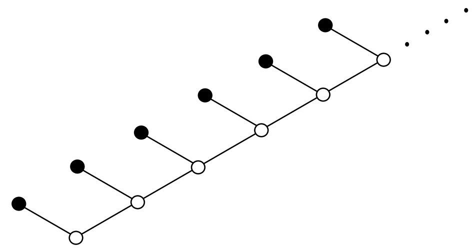

The set $\mathcal{W}$ of all white points is a maximal pairwise compatible set, as is simple to check. No black point is compatible with all of the white points, yet all of the white points are pairwise compatible. If A is true at the black points but nowhere else, then $\mathcal{W} \not\models A$, and $\mathcal{W} \not\models \sim A$, since for any $x \in \mathcal{W}$ there is a (black) y where xCy and $y \models A$. As a result, not every maximal consistent set counts as a world. So, what we need for $\mathcal{W}$ to be complete with respect to negation is a more complex condition.

To get that more complex condition, we can reason as follows: $\mathcal{W} \models A \vee \sim A$ if and only if there's some $x \in \mathcal{W}$ where $x \models A$, or there's an $x \in \mathcal{W}$ where $x \models \sim A$. But $x \models \sim A$ only when for each y where xCy, we have $y \not\models A$. Now, this holds for any proposition A we choose. In particular, we can let A be the proposition that is true at only the points not in $\mathcal{W}$ (as $\mathcal{W}$ is closed downwards, its complement is closed upwards). As a result, we must have by the second disjunct, an $x \in \mathcal{W}$ where for all y where xCy, $y \in \mathcal{W}$. That is, there is a point x in $\mathcal{W}$ such that *everything* compatible with x is in $\mathcal{W}$. We can then argue backwards to show that if this condition is satisfied, then $\mathcal{W}$ must be consistent. (I leave the details of that for you to verify.) We will call this the *witnessing condition* and the point x the *witnessing point* of the world $\mathcal{W}$.

Note that this condition ensures maximal consistency (if $\mathcal{W}$ is consistent) since *every* point compatible with the witnessing point appears in $\mathcal{W}$, so adding anything else would introduce an incompatibility within $\mathcal{W}$. The witnessing condition is very close to saying that the witnessing point does duty for the whole world, that any point in $\mathcal{W}$ is under the witnessing point. But that would be a mistake, as we can see from this next example.

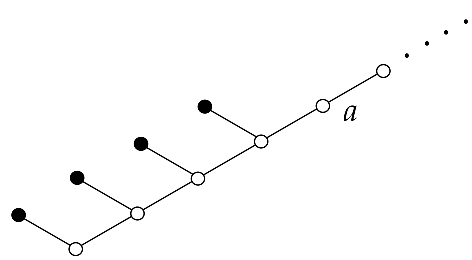

Again, let $\mathcal{W}$ be the set of white points, and let xCy be as before. In this frame a is a witnessing point for $\mathcal{W}$, yet it is not the top element in $\mathcal{W}$

(there is no top element in $\mathcal{W}$). So, for arbitrary worlds $\mathcal{W}$, there may be a maximal element, but there need not be.

So what are our options in finding worlds, such as the actual world, modelled in our frames? The first, and the simplest option, is to take worlds to be (modelled as) maximal consistent states. That is, a state w stands for a world if it is consistent (if wCw) and it is maximally so (if wCz then $z \leq w$). Then the set $\mathcal{W}$ of all states z where $z \leq w$ will certainly model a world in the sense we have discussed above. This simple approach takes states the whole world could be to be the objects which model worlds.

However, some have worried about the existence of such total entities as worlds.[11] For people with these sorts of qualms, we can offer two sorts of comfort. Firstly, they could take worlds to be modelled by a set $\mathcal{W}$ of states, satisfying our two conditions we've discussed above. However, it does seem that the condition that in any world $\mathcal{W}$ there be a witnessing state w is a little too much like accepting a 'totality' for people with worries about totalities. But I leave it to those who have these worries to decide whether or not this approach suits. There is another approach which might help people who do not like total worlds, and that is to restrict the propositions available for evaluation on the frame. In our first example frame, we showed $\mathcal{W}$ to be incomplete by considering the proposition true at all points not in $\mathcal{W}$. If we had some independent reason for thinking this is not a proposition, and if we could restrict our attention to a smaller class of propositions on this frame for which worlds (in the sense of maximal consistent sets) were complete, then we could leave worlds as maximal consistent sets of states, without needing worlds to satisfy any more conditions. However, until there is a principled reason for restricting propositions in this state, it does seem that this is an unsatisfactory way of dealing with worlds.

However we resolve this issue (and I lean most heavily to the first option), we have an interesting phenomenon. For any world $\mathcal{W}$, $\mathcal{W} \models \sim A$ if and only if $\mathcal{W} \not\models A$. That is, negation is completely truth functional when it comes to worlds. This is exactly what we would expect, and it is not really a surprise, seeing that we designed worlds to be complete and consistent. However, what is more of a surprise is that this holds *no matter what the 'underlying' logic of negation happens to be.* We need not have any special property of negation as it relates to states in order to have a very classical negation when it comes to worlds. This is what we would hope to see, given the motivations behind relevant logics. They are not designed to be a different account of negation. They are designed to give us 'finer' notions of validity and of conditionality, and they do this through the use of the finely individuated states, over and above the coarse worlds. (Recall our discussion in the first part of this paper which pointed to the need for considering

[11] I am thinking primarily of Barwise and Etchemendy, and their discussion of the Liar paradox in the context of situation theory [3], but others will no doubt also fit the bill.

states when evaluating conditionals.) However, worlds can reappear in our semantic structures, and it is pleasing to see that negation interacts with worlds in exactly the manner you would expect.[12] But once we have worlds, we have another notion of validity. We can define $\Sigma \vdash_{\mathcal{W}} A$ to mean that for all worlds $\mathcal{W}$ at which every element of Σ is true, then so is A. This is akin to the original definition of validity, setting $\Sigma \vdash A$ to mean that for every state x at which every element of Σ is true, so is A. For $\vdash_{\mathcal{W}}$ we restrict our attention to worlds. Given this definition, it is clear that $\vdash_{\mathcal{W}}$ is classical propositional validity on the fragment of the language including $\wedge$, $\vee$ and $\sim$.[13] (Again, this holds *whatever* the 'underlying' properties of negation.) As a corollary of this fact, disjunctive syllogism is valid in the form $A, \sim A \vee B \vdash_{\mathcal{W}} B$. We have the general result that $\Sigma, A \vdash_{\mathcal{W}} B$ if and only if $\Sigma \vdash_{\mathcal{W}} \sim A \vee B$, reinforcing the view that the material 'conditional' is, in a certain sense, a conditional, as it satisfies a kind of deduction theorem. Of course, the material conditional does not respect any canons of relevance, as we have known for years. This does not mean that it is not a conditional of some sort. It only means that the notion of validity with respect to which it satisfies the deduction theorem is itself rather coarse, as it also does not respect any canons of relevance.

This explains the validity of disjunctive syllogism. It is valid in the weak sense that whenever a world makes the premises true, then that world (of necessity) also makes the conclusion true. This does not mean that the corresponding conditional is relevantly valid, for relevant validity is a more robust kind of validity, designed to discriminate even in inconsistent or incomplete contexts. But both sorts of validity have their place, and both can coexist, as can be seen by our models. The more fundamental sort requires truth preservation across all states. The second, less discriminating kind of validity requires only truth preservation across all worlds. This traditional, classical account of validity cannot help but identify all classical tautologies since they are true in all worlds. But, they are not identified in all states. So in a semantic scheme which includes states, relevant validity has its place and its uses.

[12] You might think that this need not be enlightening, because this kind of move is possible with *any* kind of 'deviant' negation. But the same kind of objection is possible to any kind of worlds semantics. We know that almost *anything* can be given a worlds semantics — so it would seem to follow that it isn't enlightening at all. But that is clearly wrong. Enlightenment comes when you have an interpretation of the formal semantic structures. Given our interpretation of states, and given that we take worlds to be consistent and complete, then this analysis makes sense. The fact that anyone can do the formal work is neither here nor there.

[13] It is clearly *at least* classical propositional validity, as every world induces a propositional evaluation. However, for any propositional evaluation, we can define a one point frame which echoes that evaluation precisely (In that frame on one point x, we have $Rxxx$, xCx, and so $\{x\}$ is a world, and we set $x \models A$ if and only if A is true under that evaluation) so there are no extra constraints on $\vdash_{\mathcal{W}}$.

This is approach is new option in the debate about the validity of the disjunctive syllogism. (Though it is hinted at in a number of places in the literature, such as the Routleys' original paper [23].) The extended discussion given by Belnap and Dunn [11] (reprinted as Section 80 in [2]) is an analysis of different positions possible in regards to disjunctive syllogism, but it does not consider this position. The closest position to ours is the option that Belnap and Dunn dub 'I'm all right Jack.' According to this option, we can reason as follows. The argument from A and $\sim A \vee B$ to B is fine when the premises are not inconsistent. (For then it must be B which is true, and not $\sim A$ in the disjunction.) In the cases in which we generally use disjunctive syllogism, the premises are not inconsistent, so we can validly deduce B. But as Belnap and Dunn point out, this is no way to solve the problem, because adding another premise (that the original premises are not inconsistent) is not going to restore consistency if the premises are inconsistent to start with. It only adds a falsehood to what is already a problem.

However, there is a sense in which the offending reasoning is valid. We do know that if the premises are consistent, then it must be B that is true among A and $\sim A \vee B$.[14] Of course, adding a premise C for consistency is not going to provide us with an entailment of the form $A, \sim A \vee B, C \vdash B$, since we could find an (inconsistent) state x for which $x \models A \wedge \sim A \wedge C$, while $x \not\models B$ (unless $C \vdash B$, I suppose, but that would be cheating). However, we have seen that there is another way of encoding this. The restriction to consistency is not simply another premise which we add in order to find a *relevantly* valid argument. There is no premise which will do that, in general. Instead, we have an argument which is valid by the canons of $\vdash_{\mathcal{W}}$. We have $A, \sim A \vee B \vdash_{\mathcal{W}} B$, and there is no need for any extra premises. The restriction to consistency is 'wired in' to the definition of $\mathcal{W}$, which considers worlds as the only salient units of semantic evaluation.

It is important to note that this analysis of disjunctive syllogism does not amount to saying that $A \wedge (\sim A \vee B) \vdash B$ is valid *sometimes*. (Though clearly we can relevantly deduce B from $A \wedge (\sim A \vee B)$ sometimes, such as when $A \vdash B$ or $\sim A \vdash B$.) Unlike Routley [20] this position doesn't posit that disjunctive syllogism is 'locally valid' or that it is 'valid in the context of local consistency,' even though Routley's analysis and ours have some obvious similarities. This analysis is that disjunctive syllogism is valid, but valid in a sense which differs from the validity encoded by $\vdash$. Truth preservation across all states differs from truth preservation across all worlds. Truth is preserved in the step from $A \wedge (\sim A \vee B)$ to B when we evaluate the formulae at worlds, but not when we evaluate the formulae at all states. Both are fine notions for logical study, and both have some correlation to our pretheoretic

[14] Yes, in deducing this, I do use disjunctive syllogism. The extent to which you follow me in this reasoning only goes to show that you use disjunctive syllogism too.

notion of validity. A conditional which respects relevance will interact with $\vdash$ (the relevance respecting validity) more smoothly than it will with $\vdash_{\mathcal{W}}$. The reverse is true for the material 'conditional.' This interacts quite smoothly with $\vdash_{\mathcal{W}}$ (witness the deduction theorem), but not so well with $\vdash$.

This whole account is not so much a rival to classical logic as an extension to it. It extends classical logic with a new relation of validity, which respects the canons of relevance by considering states. According to this conception, there are (at least) two notions of validity, and each have their own uses. As an example of these uses, consider the example of a database, made popular by Belnap's discussion [5]. Suppose you have a database with information from many different sources, not all of which are completely reliable. As a result, it may contain inconsistencies. There are a number of questions we can ask about the content of such a database. The first is the question of the content of the database. This is determined by the atomic facts in the database, but these facts do not exhaust that content. For example, if the database contains an entry which states that Rajeev Goré works in the Automated Reasoning Project, and it also contains an entry stating that Rajeev is a vegetarian, then it follows that the database itself dictates that someone in the ARP is a vegetarian. This is one of the claims made by the database, even though it is not an atomic fact of the database. It is a consequence of the explicit content of the database, using a deduction of the form $Fa \wedge Ga \Rightarrow \exists x(Fx \wedge Gx)$. It seems quite plausible to suppose that the content of the database is found by closing the collection of atomic facts under some sort of consequence relation. But it is just as clear that the kind of consequence relation needed cannot be classical, because the content of an inconsistent database is not the trivial collection of all claims in the language used. If a database is inconsistent, it has some inconsistencies as a part of its content — but it does not follow that the database dictates on anything at all. If a database claims that A and it claims that $\sim A$, then it also claims that $A \wedge \sim A$, but it does not follow that it claims that B for any B we like. Similarly, if a database claims that A and it claims that $\sim A \vee B$, then it does not necessarily follow that it claims that B. (After all, its claim that $\sim A \vee B$ may rest on $\sim A$). The content of a database need not be closed under disjunctive syllogism. It is quite plausible to suppose that the content of a database is closed under a relation like the $\vdash$ of our models.

However, we can ask another question about the database. Instead of simply contenting ourselves with the issue of whether the database makes the claim that A or not, we may assume that the content of the database is correct (that the world is as the database says) and ask ourselves what else we can say about the world, given that assumption. Then something much more like $\vdash_{\mathcal{W}}$ is the relation to consider. We consider the collection of all of the ways the world could be, we pick out those which validate all of the facts in the database, and the consequences are the claims true in all of those worlds. If there are no worlds in this set, then the database is inconsistent, and the world will not (and cannot) be like the database says it is. If there are worlds in this collection, then the database is consistent, and we can deduce things about the world, given the hypothesis. Of course, some things will 'follow' from the database not because of any fact in the database at all, but simply because of the structure of worlds. For example, every necessary truth will be true in all worlds, and so, every necessary truth will feature in all worlds in our restricted set of worlds matching the database's criteria. These truths are not necessarily a part of the database's *content* (they do not all follow relevantly from the content of the database), but they are facts about the world which are true, given that the database is correct.[15]

Before leaving our discussion of worlds we ought at least mention the interaction of worlds and the ternary relation R. For a world $\mathcal{W}$ we ought have at least that if $x, y \in \mathcal{W}$, then for some $z \in \mathcal{W}$, $Rxyz$. Or more simply, that $Rwww$. This ensures the validity of $A \wedge (A \rightarrow B) \vdash_{\mathcal{W}} B$, which certainly looks right. I add this because I am quite partial to accounts of entailment for which $Rxxx$ might fail. After all, if $x \models A$ and $x \models A \rightarrow B$, we may not have $x \models B$, but only that some state larger than x gives me B. If this is in fact the case (and I argue for it elsewhere [17]), we would still like worlds to be closed under *modus ponens*, even if individual states may not be.

I must spare a thought for those who have not been able to follow this section because of their views about negation. If you think that the world could be incomplete [10], or that the world could be inconsistent [16], then you will disagree with the account of worlds we have seen. However, you need only slightly modify this account of worlds in one direction or the other to make it palatable for your view of negation. Provided that you can distinguish between worlds and the more general 'states,' by positing some kind of condition on a state (or a set of states) for it to count as (modelling) a world, then you can still have a 'coarse' view of validity which requires only preservation across worlds (whatever worlds turn out to be), and a finer, relevant validity, which requires preservation across all states. The only difference is that your 'worldy' validity will not be classical validity. So, the insights of this section still apply, save for the discussion of disjunctive syllogism (if you take the world to be inconsistent).[16] But even there, if you take it that disjunctive syllogism is valid in a more limited sense, then perhaps this account will reappear in your treatment of $\vdash_{\mathcal{W}}$.

[15]So, the 'given that' in this clause is to be analysed much more like the material conditional than a relevant conditional. This is fine in our formalism, for we can use the material conditional to analyse it if we so choose.

[16]If you think that the world is truly inconsistent, then of course you must take disjunctive syllogism to be invalid in general, beacuse there are *counterexamples*.

5 Against Boolean Negation

Consider our standard model, with the collection of all states W. Can we not add a connective $-$ to our set of connectives, which is evaluated as follows?

- $x \models -A$ if and only if $x \not\models A$

The surprise is that Boolean negation satisfies conditions like $A \wedge -A \vdash B$, and $A \vdash B \vee -B$, without corrupting the behaviour of $\wedge, \vee, \rightarrow, \sim$ and $\vdash$. The heartache that this fact has caused has made many in the community of relevant logicians worry long and hard. These worries have been well expressed by Belnap and Dunn.

> ... what was all the fuss about "fallacies of relevance"? What were the complaints lodged against contradictions' implying everything and against the disjunctive syllogism? Boolean negation trivially satisfies these principles; so what can be the interest of De Morgan negation's failing to satisfy them? Will the real negation please stand up? ([2], page 174)

It is clear that fallacies of relevance by themselves have nothing much to add to the debate, for it could be argued that $A \wedge -A$ is truly relevant to any B you like.[17] No, simply arguing about relevance will not extend the debate in any particularly fruitful way. Fortunately, there are ways to extend the debate. The first involves our semantic structures, and the second, the concept of conservative extension.

First, note what Boolean negation is doing in our semantic structures. It is clear that under interpretation, it is functioning as a 'negation as failure,' and clearly it's not 'not' as anyone means it, as we have all agreed that states can be both incomplete and inconsistent. Since states can be incomplete, the mere fact that $x \models -A$ only tells us that $x \not\models A$. And in our models, there is no guarantee from $x \not\models A$ to $y \not\models A$ where $x \leq y$. In fact, if x is incomplete, and it is extended by y, we can expect that some 'gaps' in x are filled in y. This means either that Boolean 'facts' are not facts that are carried by states (for those facts are preserved at states are filled out) or that the inclusion relation $\leq$ is trivialised to identity. This makes states collapse into possible worlds, and we do not want that, for the reasons we have already seen. For the former, there is no reason for our semantic formalism to encode any and every fact *about* states. The mere fact that a state is mentioned in this paper or that some state has not been considered by anyone do not appear to be the kinds of facts that are supported by

[17] If we accept $\perp$, and if we accept $\circ$ as the binary connective residuated by the conditional, then we are already committed to $A \circ (A \rightarrow \perp) \rightarrow B$ by dint of $(A \rightarrow \perp) \rightarrow (A \rightarrow B)$, and it is right that $A \circ (A \rightarrow \perp)$ is relevant to B by the $\perp$ ('everything is true') lurking there. Perhaps it could be argued that there is a similar implicit $\perp$ in $-A$.

those states themselves. (By the world as a whole — of course — but not necessarily by that state itself.) There is a difference between claims about states, and claims supported by states. We have an independent criterion for information supported by a state, and that is that the information be preserved as the state is increased. Boolean negation does not pass that test, and so, it does not appear to be a meaningful connective in this context.

This point needs to be explained further, for claims of 'meaninglessness' are particularly hard to defend. The claim here is simply that there is no connective $-$ on propositions such that $A \wedge {\sim} A \vdash B$ and $A \vdash B \vee -B$ come out as valid for all choices of A and B. For if there were, then either the relation $\leq$ would collapse into identity, or propositions would fail (in general) to be persistent. Of course we can *define* the locution "$x \models -A$" to be read as saying that $x \not\models A$. But that does not mean that claims of that form express a relationship between a state and a new kind of proposition. For $-A$ is not a proposition (on pain of failing to be preserved upward by $\leq$). Talk of $-A$ can be said to be meaningful, but when you engage in that kind of talk, you are no longer simply talking of what is supported *by* states — you are also considering other things which are true *of* states.[18]

At this point it is worthwhile to step back a little, for there is another argument for the extension of relevant logics with Boolean negation which deserves consideration. Meyer [13] argues that since Boolean negation is a conservative extension of the relevant logic **R**, it is a worthwhile extension to that logic.[19] There is more to be done with this argument before it proves its point, for the following reason. Two can play the game of conservative extensions. The opponent of Boolean negation can play it like this: Any relevant logic with a frame semantics can be conservatively extended with a binary connective $\supset$, which residuates ordinary conjunction. That is, it will satisfy $A \wedge B \vdash C$ if and only if $A \vdash B \supset C$.[20] This is a conservative extension of the relevant logic **R**, and almost every other logic in the relevant family. So, we can add it if we like (by the argument that any conservative extension is a good extension). Call the system $\mathbf{R}^{\supset}$. But what do we find? The addition of Boolean negation to $\mathbf{R}^{\supset}$ is not conservative. In $\mathbf{R}^{\supset}$ we have $\not\vdash ((A \supset B) \supset A) \supset A$, but once we have Boolean negation, $A \supset B$ is identified with $-A \vee B$, and we have $\vdash ((A \supset B) \supset A) \supset A$ by simple moves. So, what can the defender of Boolean negation say at this point? She must at least say that the extension $\mathbf{R}^{\supset}$ was incomplete. But the only plausible way to argue that $\mathbf{R}^{\supset}$ as presented is incomplete is to appeal to Boolean negation. (How else do you argue for Peirce's law?) So,

[18] And you may ask what how the proposition $x \not\models A$ is to be cashed out on our terms. Well, a plausible choice is that $\mathcal{W} \models {\sim}(x \models A)$, where $\mathcal{W}$ is the actual world.

[19] **Y** is a conservative extension of **X** just when **Y** is an extension of **X**, and **Y** restricted to the language of **X** agrees with **X**.

[20] Give $\supset$ the standard intuitionistic clause: $x \models B \supset C$ if and only if for each $y \geq x$, if $y \models B$ then $y \models C$.

what the defender of Boolean negation needs at this point is an independent argument for Boolean negation, and not just an argument which relies on the conservative extension result. Two can play that game, and it gets neither of them anywhere.

But what are the independent reasons for wanting Boolean negation? I can think of only one. And that is that the logic without Boolean negation is not expressive enough to do certain things we want of it — and in particular, disjunctive syllogism. But we have seen how to do that, in a way which is more in harmony with our interpertation of frames for relevant logics. There is no need to expand our language to recapture classical validity. We simply need to be aware that our structures are rich enough to encode two different sorts of validity, one of which has all of the power (and, of course, all of the weaknesses) of classical validity.

Of course, one can use the model structures of relevant logics for purposes other than giving an account of what information is carried by particular states. If that is the case (say, if you are interested in reasoning *about* states, or if the points in your frame are something other than states) then boolean negation may be a sensible connective on that frame. The argument here only applies to this particular understanding of frames for relevant logics.

6 Logic and Truth

There is one last issue which ought to be faced, for there is a potential confusion when it comes to interpreting what counts as truth in a model. Given the 'standard model' of all states, we have the obvious choice of truth in that model, as truth in the actual world in that model. However, a model may not come with a distinguished actual world, and there are a number of choices for what counts as truth in that model. One choice which is not going to be too useful is to take truth in a model as truth at *every* point in the model. But this is next to useless as an account of truth in a model, for hardly anything of interest turns out to be true at every point in a model. (Ways can be very incomplete. Logical truths can fail at states.) So, truth at every point is not the notion for us. Another notion is truth at every world. This is simply recorded for us by $\vdash_{\mathcal{W}}$, in that A is true at every world just when $\vdash_{\mathcal{W}} A$.

However, there is yet another notion of truth in a model which is worth our notice, and it is *this* notion which has appeared in the literature on relevant logic. Relative to a particular model, we can say that A (relevantly) entails B just when for every x, if $x \models A$ we have $x \models B$. We can say that a state x in a model *respects relevant entailment* just when whenever A entails B in that model, $x \models A \to B$. It is clear that any state x such that $Ryxz \Rightarrow y \leq z$ will respect relevant entailment. Let's call the class of all of these states $\mathcal{L}$, for these are states which respect *logic* (in the sense encoded by '$\vdash$'). In the relevant logic literature, it has been truth in all of the states in $\mathcal{L}$ which has counted as truth in a model — or more particularly, a 'reduced model' is one in which we can make $\mathcal{L}$ a set of the form $\{y : 0 \leq y\}$ for some point y, and it has been truth at 0 which counts as truth in the model. And this is yet another notion of truth in a model.

Once we have split apart these notions of truth in a model, the question of the relationship between them is open as well. If we let t be true at all and only those states in $\mathcal{L}$, then the truth of A at $\mathcal{L}$ amounts to $t \vdash A$. One issue is whether $t \vdash A$ ensures that $\vdash_{\mathcal{W}} A$, or more simply, whether t is true in all worlds (is t necessarily true?). That t be necessarily true seems like an obviously desirable state of affairs, but it does not follow from the conditions we have seen so far. So, we need to posit the condition that for every world $\mathcal{W}$, $\mathcal{L} \cap \mathcal{W} \neq \emptyset$.

It is clear from the way we have constructed worlds that $A \vee \sim A$ is true in all worlds. It is not so clear that $t \vdash A \vee \sim A$ is desirable. For a plea to reject $t \vdash A \vee \sim A$, consider this quote, from Slaney.

> This "law of the excluded middle" suffices in the context of elementary De Morgan lattice logic to establish all the classical tautologies in the connectives $\wedge$, $\vee$, $\sim$. So much is well known and is usually claimed by supporters of relevant logic as a mark in their favour, as showing their non-deviance with respect to the classical connectives. Yet if it is seriously to be maintained that pure implication is the Heart of Logic then to have *logic* confer full honours on a formula which is not constructed from any record of an inference is somewhat anomalous at best. Presumably the motivation for excluded middle is semantic, that it comes out true "no matter what"; but logic is supposed to sort out what follows from what, and as such has surely no place for these material tautologies which just sit around being true and are no inference tickets at all. [24]

While we need not agree with all of Slaney's claims here — after all, I have just argued that there is a place in logic for what is true in all worlds — the distinction he draws is a valid one. First, there are the inference tickets, those relevant conditionals which record valid entailments. These are true in $\mathcal{L}$, and are 'theorems' in this sense. Excluded middles, on the other hand, are not inference tickets, and there seems to be no reason to dictate that they must be valid in $\mathcal{L}$. Of course, they are true no matter what, and this is best recorded by the fact that $\vdash_{\mathcal{W}} A \vee \sim A$, not by requiring that $t \vdash A \vee \sim A$. It is a mark of the naturalness of our semantic proposal that it can make sense of the distinction on which Slaney draws. Our formal picture shows the difference between being an inference ticket (being a valid entailment, or being true in $\mathcal{L}$) and being true 'no matter what' (being true in all worlds).

7 Future Work

So what we have seen is a general account of the semantics of relevant logics, which makes sense of its treatment of negation, and which coheres with classical accounts of validity. There is still, obviously, much to be done. First and foremost we need a proof theory for $\vdash_{\mathcal{W}}$. Our characterisation has been purely semantic. A proof theoretical characterisation would certainly be nice.

Further food for thought is the addition of modality and quantification. Modalities seem easier, though there are many options. For example, modalities need not respect necessary material equivalence. For many kinds of modal operator (not just conditionals) this seems quite right. But of course, there are some modal operators which do seem to respect necessary material equivalence (like old fashioned necessity). It is an open question as to how all of these modalities should be captured in our framework. Quantification, on the other hand, is a very thorny issue. The only adequate semantics for the traditional quantified relevant logics to date [9] does not seem well suited to our enterprise. Yet there is no obvious alternative. All of this requires much more thought, and much more hard work. But that, thankfully, must be left for another time and another place.[21]

References

[1] Alan Ross Anderson and Nuel D. Belnap. *Entailment: The Logic of Relevance and Necessity, Volume 1.* Princeton University Press, Princeton, 1975.

[2] Alan Ross Anderson, Nuel D. Belnap, and J. Michael Dunn. *Entailment: The Logic of Relevance and Necessity, Volume 2.* Princeton University Press, Princeton, 1992.

[3] Jon Barwise and John Etchemendy. *The Liar.* Oxford University Press, Oxford, 1987.

[4] Jon Barwise and John Perry. *Situations and Attitudes.* MIT Press, Bradford Books, 1982.

[5] Nuel D. Belnap. A useful four-valued logic. In J. Michael Dunn and George Epstein, editors, *Modern Uses of Multiple-Valued Logics*, pages 8–37. D. Reidel, Dordrecht, 1977.

[21]Many people have contributed in one way or another to this work. I must first of all thank Bob Meyer, Nuel Belnap and Mike Dunn, for their encouragement, their example, and their hard work on these issues. Thanks too go to Graham Priest, John Slaney and Richard Sylvan, for helping me to form my views on negation, logic and many other issues over the last few years. Finally, helpful discussion or encouragement has come from as far and wide as Mic Detlefsen, Rajeev Goré, Allen Hazen, Ian Hinckfuss, Pragati Jain, Frank Jackson, Gary Malinas, Errol Martin, Chris Mortensen, Daniel Nolan, John Perry, Tim Surendonk, all the students at the 1995 Logic Summer School at the ARP, and an audience at the 1995 joint session of the AAL/AAP conferences in Armidale. Thanks to you all!

[6] B. J. Copeland. 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, 1979.

[7] J. Michael Dunn. Relevance logic and entailment. In D. Gabbay and D. Guenther, editors, *Handbook of Philosophical Logic*, volume III, pages 117–229. D. Reidel, 1986.

[8] J. Michael Dunn. Star and perp: Two treatments of negation. In James E. Tomberlin, editor, *Philosophical Perspectives*, volume 7. Ridgeview Publishing Company, Atascadero, California, 1994.

[9] Kit Fine. Semantics for quantified relevance logic. *Journal of Philosophical Logic*, 17:27–59, 1988.

[10] Patrick Grim. *The Incomplete Universe.* The MIT Press, Cambridge, MA, 1991.

[11] N. D. Belnap Jr. and J. M. Dunn. Entailment and the disjunctive syllogism. In F. Fløistad and G. H. von Wright, editors, *Philosophy of Language/Philosophical Logic*, pages 337–366. Martinus Nijhoff, The Hague, 1981.

[12] David K. Lewis. *On the Plurality of Worlds.* Blackwell, Oxford, 1986.

[13] Robert K. Meyer. New axiomatics for relevant logics — I. *Journal of Philosophical Logic*, 3:53–86, 1974.

[14] Robert K. Meyer and Errol P. Martin. Logic on the australian plan. *Journal of Philosophical Logic*, 15:305–332, 1986.

[15] Chris Mortensen. Anything is possible. *Erkenntnis*, 30:319–337, 1989.

[16] Graham Priest. *In Contradiction: A Study of the Transconsistent.* Martinus Nijhoff, The Hague, 1987.

[17] Greg Restall. *On Logics Without Contraction.* PhD thesis, The University of Queensland, January 1994.

[18] Greg Restall. Truthmakers, entailment and necessity. Technical Report TR-ARP-10-94, Automated Reasoning Project, Australian National University, 1994. Available from `ftp://arp.anu.edu.au/pub/techreports/TR-10-94.ps`.

[19] Greg Restall. Information flow and relevant logics. In Jerry Seligman and Dag Westerstahl, editors, *Logic, Language and Computation: The 1994 Moraga Proceedings.* CSLI, 1995.

[20] Richard Routley. Relevantism, material detachment, and the disjunctive syllogism argument. *Canadian Journal of Philosophy*, 14:167–188, 1984.

[21] Richard Routley and Robert K. Meyer. Semantics of entailment 1. In Hugues Leblanc, editor, *Truth Syntax and Modality*, pages 199–243. North-Holland, Amsterdam, 1973. Proceedings of the Temple University Conference on Alternative Semantics.

[22] Richard Routley, Val Plumwood, Robert K. Meyer, and Ross T. Brady. *Relevant Logics and their Rivals.* Ridgeview, 1982.

[23] Richard Routley and Valerie Routley. Semantics of first-degree entailment. *Noûs*, 3:335–359, 1972.

[24] John K. Slaney. A metacompleteness theorem for contraction-free relevant logics. *Studia Logica*, 43:159–168, 1983.