Theoria 37 (1971) 74-85

Completeness and decidability of three logics of counterfactual conditionals[1]

by

DAVID LEWIS

(Princeton University)

Language

Our language will be that of the ordinary propositional calculus supplemented with the counterfactual conditional connective $\Box\!\!\rightarrow$. The sentence $\varphi \Box\!\!\rightarrow \psi$ may be read as 'If it were the case that φ, then it would be the case that ψ'.

Intented interpretation

A sentence $\varphi \Box\!\!\rightarrow \psi$ is intended to mean, roughly, that ψ holds in certain of the possible worlds in which φ holds: those of them that are most closely similar to our actual world. We could capture this intention most straightforwardly by positing a function f which selects, for any sentence φ and world i, a set $f(\varphi, i)$ of worlds regarded as the set of worlds most closely similar to i out of the worlds in which φ holds.

But this approach is open to objection. Just as no real number greater than 1 is closest to 0, so it may be that none of the worlds in which φ holds is most closely similar to i. It may be that for each of them, there is another still closer. To meet this difficulty we could introduce the notion of degrees of similarity between worlds, and take $\varphi \Box\!\!\rightarrow \psi$ to mean that, unless no worlds in which φ holds are similar to any degree to our actual world, there is some degree of similarity to our actual world within which there are

[1] This research was supported by a fellowship from the American Council of Learned Societies. I am grateful to Kit Fine for valuable suggestions and comments.

some worlds in which φ holds, and within which ψ holds in all worlds in which φ holds.

To introduce the notion of degrees of similarity, it is fortunately not necessary to suppose that the similarity of worlds admits of numerical measurement. We could posit a family, indexed by worlds, of sets $\$_i$ of sets of worlds; each S in $\$_i$ is regarded as the set of all worlds similar to at least a certain degree to the world i. Or we could posit a family, again indexed by worlds, of comparative similarity relations $\leqslant_i$ over sets of all or some worlds; $j \leqslant_i k$ is regarded as meaning that the world j is at least as similar to the world k to the world i.

A further account of the philosophical motivation and consequences of such an interpretation of counterfactual conditionals will be given in [2].

Model theory

Corresponding to these various approaches to the interpretation of counterfactuals, we will consider three different versions of the model theory for our language. In each version, the intended models are those in which I is the set of all possible worlds; $[\varphi]$ is the set of those worlds in which the sentence φ holds; and f, $\$$, or $\leqslant$ is as described above.

An *α-model* is any triple $\langle I, [\ ], f\rangle$ such that:

(α.0.1) I is a nonempty set,
(α.0.2) $[\ ]$ assigns to each sentence φ a subset $[\varphi]$ of I,
(α.0.3) $[\sim\varphi] = I - [\varphi]$, $[\varphi \,\&\, \psi] = [\varphi] \cap [\psi]$, and so on for the other truth-functional connectives,
(α.0.4) f assigns to each i in I and sentence φ a subset $f(\varphi, i)$ of I,
(α.0.5) $[\varphi \Box\!\!\rightarrow \psi] = \{i \in I : f(\varphi, i) \subseteq [\psi]\}$.

The intended α-models also meet at least three further conditions:

(α.0.6) $f(\varphi, i) \subseteq [\varphi]$,
(α.0.7) if $f(\varphi, i) \subseteq [\psi]$ and $f(\psi, i) \subseteq [\varphi]$, then $f(\varphi, i) = f(\psi, i)$,

(α.0.8) either $f(\varphi \vee \psi, i) \subseteq [\varphi]$ or $f(\varphi \vee \psi, i) \subseteq [\psi]$ or $f(\varphi \vee \psi, i) = f(\varphi, i) \cup f(\psi, i)$.

Let us call any α-model that meets conditions (α.0.6—8) *standard*.

A *β-model* is any triple $\langle I, [\ ], \$\rangle$ such that:

(β.0.1—3) (same as (α.0.1—3) above),
(β.0.4) $\$$ assigns to each i in I a nonempty set $\$_i$ of subsets of I,
(β.0.5) $[\varphi \;\Box\!\!\rightarrow \psi] = \{i \in I: [\varphi] \cap \bigcup \$_i = \Lambda \text{ or } \exists S \in \$_i\, [\Lambda \neq [\varphi] \cap S \subseteq [\psi]]\}$.

All intended β-models also meet at least this further condition:

(β.0.6) $\$_i$ is nested (that is, if $S, T \in \$_i$ then $S \subseteq T$ or $T \subseteq S$).

Let us call any β-model that meets condition (β.0.6) *standard*.

A *γ-model* is any triple $\langle I, [\ ], \leqslant\rangle$ such that:

(γ.0.1—3) (same as (α.0.1—3) above),
(γ.0.4) $\leqslant$ assigns to each i in I a 2-place relation $\leqslant_i$ over a subset S_i of I,
(γ.0.5) $[\varphi \;\Box\!\!\rightarrow \psi] = \{i \in I: [\varphi] \cap S_i = \Lambda$ or $\exists k \in [\varphi] \cap S_i\ \forall j \in [\varphi]\ [\text{if } j \leqslant_i k \text{ then } j \in [\psi]]\}$

All intended γ-models also meet at least this further condition:

(γ.0.6) $\leqslant_i$ is a total preordering of S_i (that is, it is transitive and strongly connected in S_i).

Let us call any γ-model that meets condition (γ.0.6) *standard*.

It is highly plausible that all intended models meet two further conditions: roughly, that a world is at least as similar to itself as any other world is to it, and that no other world is as similar to a world as that world itself is. Let us call any standard model that meets these two conditions *1-standard*. More precisely, an α-model is 1-standard iff it is standard and:

(α.1.1) if $i \in [\varphi]$ then $i \in f(\varphi, i)$,
(α.1.2) if $i \in [\varphi]$ then $j \in f(\varphi, i)$ only if $j = i$.

A β-model is 1-standard iff it is standard and:

(β.1.1) $i \in \bigcap \$_i$,
(β.1.2) unless $\bigcup \$_i = \Lambda$, $\{i\} \in \$_i$.

A γ-model is 1-standard iff it is standard and:

(γ.1.1) $i \in S_i$, and if $j \in S_i$ then $i \leqslant_i j$,
(γ.1.2) unless $S_i = \Lambda$, $j \leqslant_i i$ iff $j = i$.

It has been suggested by Stalnaker and Thomason, in [4] and [5], that all intended models meet another condition: roughly, that for each world i and sentence φ, unless no world in which φ holds is at all similar to i, there is a unique closest world to i in which φ holds. Let us call any 1-standard model that meets this condition *2-standard*. More precisely, an α-model is 2-standard iff it is 1-standard and,

(α.2) $f(\varphi, i)$ contains at most one member.

A β-model is 2-standard iff it is 1-standard and:

(β.2) if $[\varphi] \cap \bigcup \$_i \neq \Lambda$, there are $S \in \$_i$ and $j \in I$ such that $[\varphi] \cap S = \{j\}$.

A γ-model is 2-standard iff it is 1-standard and:

(γ.2) if $[\varphi] \cap S_i \neq \Lambda$, there is $k \in [\varphi] \cap S_i$ such that for any $j \in [\varphi]$, $j \leqslant_i k$ only if $j = k$.

Apart from inessential technical differences, the models for sentential conditional logic considered by Stalnaker and Thomason are exactly our 2-standard α-models.

A sentence φ is *true at* i in I *in* a model $\langle I, [\ ], \ldots\rangle$ iff $i \in [\varphi]$; φ is *valid in* such a model iff $I = [\varphi]$, so that φ is true at every i in I.

Equivalent models

We call two models (of the same or different sorts) *equivalent* iff they have the same first and second components—the same I and the same $[\ ]$—and differ only in the third component: the f, $\$$, or $\leqslant$ as the case may be. We shall see that if we are given a standard α-model, we can convert it into an equivalent β-model or into an equivalent γ-model.

Given a standard α-model $\langle I, [\ ], f\rangle$, we begin by defining a family, indexed by I, of relations $\preccurlyeq_i$ over all sentences:

$$\varphi \preccurlyeq_i \psi =_{df} \Lambda \neq f(\varphi, i) \subseteq f(\varphi \vee \psi, i) \text{ or } f(\psi, i) = \Lambda.$$

In intended models, $\varphi \preccurlyeq_i \psi$ iff the closest worlds to i in which φ holds are at least as close to i as are the closest worlds to i in which ψ holds.

Note that if $[\eta] \subseteq [\theta]$ and $[\eta] \cap f(\theta, i) \neq \Lambda$ then, by (α.0.6—8), $f(\eta, i) = [\eta] \cap f(\theta, i)$. Note also that if $[\eta] \subseteq [\theta]$ and $f(\theta, i) = \Lambda$ then by (α.0.6—7), $f(\eta, i) = \Lambda$. Using these two observations, we can prove three useful lemmas about the relations $\preccurlyeq_i$.

First lemma: $\preccurlyeq_i$ is a total preordering of all sentences. To prove that $\preccurlyeq_i$ is strongly connected, note that if $[\varphi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, then $\varphi \preccurlyeq_i \psi$; if $[\psi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, then $\psi \preccurlyeq_i \varphi$; otherwise $f(\varphi \vee \psi, i) = \Lambda$, so $f(\varphi, i) = f(\psi, i) = \Lambda$, so $\varphi \preccurlyeq_i \psi$ and conversely. To prove that $\preccurlyeq_i$ is transitive, note that if $[\varphi] \cap f(\varphi \vee \psi \vee \chi, i) \neq \Lambda$, then $\varphi \preccurlyeq_i \psi$; if $[\varphi] \cap f(\varphi \vee \psi \vee \chi, i) = \Lambda$ and $[\chi \,\&\, {\sim}\varphi] \cap f(\varphi \vee \psi \vee \chi, i) \neq \Lambda$ then not $\varphi \preccurlyeq_i \chi$; if $[\varphi] \cap f(\varphi \vee \psi \vee \chi, i) = \Lambda$ and $[\chi \,\&\, {\sim}\varphi] \cap f(\varphi \vee \psi \vee \chi, i) = \Lambda$ and $[\psi \,\&\, {\sim}\chi \,\&\, {\sim}\varphi] \cap f(\varphi \vee \psi \vee \chi, i) \neq \Lambda$, then not $\chi \preccurlyeq_i \psi$; otherwise, $f(\varphi \vee \psi \vee \chi, i) = \Lambda$, so $f(\psi, i) = \Lambda$ and again $\varphi \preccurlyeq_i \psi$. So in all cases if $\varphi \preccurlyeq_i \chi \preccurlyeq_i \psi$ then $\varphi \preccurlyeq_i \psi$.

Second lemma: if $\varphi \preccurlyeq_i \psi$, then $[\psi] \cap f(\varphi, i) \subseteq f(\psi, i)$. If $[\varphi] \cap f(\varphi \vee \psi, i) \neq \Lambda$ and $[\psi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, the consequent of the lemma holds; if $[\varphi] \cap f(\varphi \vee \psi, i) \neq \Lambda$ and $[\psi] \cap f(\varphi \vee \psi, i) = \Lambda$, $[\psi] \cap f(\varphi, i) = \Lambda$ and again the consequent holds; if $[\varphi] \cap f(\varphi \vee \psi, i) = \Lambda$ and $[\psi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, then not $\varphi \preccurlyeq_i \psi$; otherwise $f(\varphi \vee \psi, i) = \Lambda = f(\varphi, i) = f(\psi, i)$, so again the consequent holds.

Third lemma: if $[\varphi] \cap f(\psi, i) \neq \Lambda$, then $\varphi \preccurlyeq_i \psi$. If $[\varphi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, the consequent of the lemma holds; if $[\varphi] \cap f(\varphi \vee \psi, i) = \Lambda$ and $[\psi] \cap f(\varphi \vee \psi, i) \neq \Lambda$, $f(\psi, i) = [\psi] \cap f(\varphi \vee \psi, i)$ so the antecedent of the lemma fails; otherwise $f(\varphi \vee \psi, i) = \Lambda = f(\psi, i)$ so again the antecedent fails.

Now let $\$$ assign to each i in I the set of all subsets S of I such that, for some sentence φ, $S = \cup\{f(\psi, i) : \psi \preccurlyeq_i \varphi\}$. We call $\langle I, [\ ], \$\rangle$ the *β-conversion* of our original α-model $\langle I, [\ ], f\rangle$. Using our

three lemmas about the relations $\preccurlyeq_i$, it is easily verified that the β-conversion of a standard α-model is a standard β-model; that the β-conversion of a standard α-model meeting (α.1.1) meets (β.1.1); that the β-conversion of a standard α-model meeting (α.1.2) meets (β.1.2); and that the β-conversion of a standard α-model meeting (α.2) meets (β.2).

Alternatively, let $\leqslant$ assign to each i in I the relation $\leqslant_i$ such that $j \leqslant_i k$ iff there are sentences φ and ψ such that $j \in f(\varphi, i)$, $k \in f(\psi, i)$, and $\varphi \preccurlyeq_i \psi$. We call $\langle I, [\ ], \leqslant\rangle$ the *γ-conversion* of our original α-model. It is easily verified that the γ-conversion of a standard α-model is a standard γ-model; that the γ-conversion of a standard α-model meeting (α.1.1) meets (γ.1.1); that the γ-conversion of a standard α-model meeting (α.1.2) meets (γ.1.2); and that the γ-conversion of a standard α-model meeting (α.2) meets (γ.2).

We went from α-models to β-models and γ-models for the sake of generality; and we did indeed gain generality. Although every standard α-model can be converted to an equivalent β-model or γ-model, the opposite is not the case. Consider, for instance, the 1-standard β-model $\langle I, [\ ], \$\rangle$ in which I is the set of real numbers; in which $[\sigma]$ is $\{i \in I\colon 0 < i\}$, $[\sigma_1]$ is $\{i \in I\colon i < \frac{1}{2}\}$, $[\sigma_2]$ is $\{i \in I\colon i < \frac{1}{4}\}$, and in general $[\sigma_n]$ is $\{i \in I\colon i < 2^{-n}\}$ throughout a countable sequence of sentence letters; and in which $\$_i$ is the set of all closed intervals $[i, x]$ with $i \leqq x$. Or consider the equivalent 1-standard γ-model in which I and $[\ ]$ are as just specified and $\leqslant_i$ is the usual ordering of the real numbers greater than or equal to i. There can be no α-model equivalent to these two models. For in the two models, $\sigma \mathrel{\Box\!\!\rightarrow} \sigma$, $\sigma \mathrel{\Box\!\!\rightarrow} \sigma_1$, $\sigma \mathrel{\Box\!\!\rightarrow} \sigma_2, \ldots, \sigma \mathrel{\Box\!\!\rightarrow} \sigma_n, \ldots$ are all true at 0. But there is no real number at which all of $\sigma, \sigma_1, \sigma_2, \ldots, \sigma_n, \ldots$ are true in the original models. So there is no i in I that could serve as a member of $f(\sigma, 0)$ in an equivalent α-model. Nor could $f(\sigma, 0)$ be empty, since $\sigma \mathrel{\Box\!\!\rightarrow} {\sim}\sigma$ is not true at 0 in the two original models.

We might expect that as a result of the greater generality gained by moving to β-models and γ-models, some valid sentences would be lost. As will be seen, this is not correct. Exactly the same sentences are valid in all standard α-models, in all standard β-

models, and in all standard γ-models; and likewise for 1-standard and 2-standard models of the three sorts.

Deductive systems

We may identify a deductive system with the set of its theorems. This will be the smallest set of sentences closed under certain rules of inference and containing certain axioms (more precisely, all instances of certain axiom-schemata). Our rules of inference are, first, the rule of *tautological implication* (*TI*):

$$\frac{\quad}{\psi} \quad \text{when } \psi \text{ is a tautology,}$$

$$\frac{\chi_1, \ldots, \chi_n}{\psi} \quad \text{when } (\chi_1 \,\&\, \ldots \,\&\, \chi_n) \supset \psi \text{ is a tautology;}$$

and, second, the rule of *deduction within conditionals* (*DWC*):

$$\frac{\psi}{\varphi \mathrel{\Box\!\!\rightarrow} \psi} \qquad \frac{(\chi_1 \,\&\, \ldots \,\&\, \chi_n) \supset \psi}{((\varphi \mathrel{\Box\!\!\rightarrow} \chi_1) \,\&\, \ldots \,\&\, (\varphi \mathrel{\Box\!\!\rightarrow} \chi_n)) \supset (\varphi \mathrel{\Box\!\!\rightarrow} \psi)}$$

The deductive system **C0** is generated by the rules *TI* and *DWC* and three basic axioms:

Axiom A. $\varphi \mathrel{\Box\!\!\rightarrow} \varphi$,
Axiom B. $((\varphi \mathrel{\Box\!\!\rightarrow} \psi) \,\&\, (\psi \mathrel{\Box\!\!\rightarrow} \varphi)) \supset ((\varphi \mathrel{\Box\!\!\rightarrow} \chi) \equiv (\psi \mathrel{\Box\!\!\rightarrow} \chi))$,
Axiom C. $(\varphi \vee \psi \mathrel{\Box\!\!\rightarrow} \varphi) \vee (\varphi \vee \psi \mathrel{\Box\!\!\rightarrow} \psi) \vee ((\varphi \vee \psi \mathrel{\Box\!\!\rightarrow} \chi) \equiv (\varphi \mathrel{\Box\!\!\rightarrow} \chi) \,\&\, (\psi \mathrel{\Box\!\!\rightarrow} \chi))$.

C0 is the weakest system that has any claim to be called a logic of *conditionals*; a system missing some of Axioms *A*—*C* might better be called a logic of sententially indexed modalities.

The system **C1** is generated by the rules *TI* and *DWC*, Axioms *A*—*C*, and two further axioms:

Axiom D. $(\varphi \mathrel{\Box\!\!\rightarrow} \psi) \supset (\varphi \supset \psi)$,
Axiom E. $\varphi \,\&\, \psi \supset (\varphi \mathrel{\Box\!\!\rightarrow} \psi)$.

In my opinion, **C1** is the correct logic of counterfactual conditionals as we ordinarily understand them.

The system **C2** is generated by the rules *TI* and *DWC*, Axioms *A*—*C*, Axiom *D*, and one further axiom:

Axiom F. $(\varphi \mathrel{\Box\!\!\rightarrow} \psi) \vee (\varphi \mathrel{\Box\!\!\rightarrow} {\sim}\psi)$.

C2 contains **C1**, since Axiom *E* follows by *TI* from Axioms *D* and *F*. **C2** is the same as Stalnaker's differently axiomatized system of the same name in [4]. Stalnaker and Thomason [5] contains a proof, in effect, that **C2** is sound and complete for the class of 2-standard α-models. We shall proceed to obtain similar results for the weaker systems **C0** and **C1**, with respect to models of all three sorts.

Soundness results

Each of the following observations is easily verified. The rule *TI* preserves truth at any i in I in any model, and hence preserves validity. The rule *DWC*, though not truth-preserving, preserves validity in any α-model, in any standard β-model, and in any standard γ-model. Axiom *A* (more precisely, any instance thereof) is valid in any α-model that meets condition (α.0.6), in any β-model, and in any γ-model. Axiom *B* is valid in any α-model that meets condition (α.0.7), in any standard β-model, and in any standard γ-model. Axiom C is valid in any α-model that meets condition (α.0.8), in any standard β-model, and in any standard γ-model. Therefore any theorem of **C0** is valid in any standard model.

Further, Axiom *D* is valid in any α-model that meets condition (α.1.1), in any β-model that meets condition (β.1.1), and in any γ-model that meets condition (γ.1.1). Axiom *E* is valid in any α-model that meets condition (α.1.2), in any β-model that meets condition (β.1.2), and in any γ-model that meets condition (γ.1.2). Therefore any theorem of **C1** is valid in any 1-standard model.

Further, Axiom *F* is valid in any α-model that meets condition (α.2), in any β-model that meets condition (β.2), and in any γ-model that meets condition (γ.2). Therefore any theorem of **C2** is valid in any 2-standard model.

Deducibility and consistency

A sentence φ is *deducible from* a set Σ of sentences *in* a deductive system L iff φ belongs to the smallest set of sentences closed under the rule TI and including both Σ and L. (Note that it is inappropriate to require closure under the non-truth-preserving rule DWC.) A set of sentences is *L-consistent* iff not every sentence is deducible from it in L. A set of sentences is *maximal L-consistent* iff it is L-consistent but not a subset of any larger L-consistent set. Note that if a sentence φ belongs to every maximal L-consistent set that includes a set Σ of sentences, then φ is deducible in L from some finite subset of Σ.

Canonical models

The methods employed in the remainder of this paper are adapted from those developed for modal logic by Kaplan in [1] and by Lemmon and Scott in unpublished work described by Segerberg in [3].

If L is any deductive system closed under the rules TI and DWC, the *canonical α-model for L* is the triple $\langle I, [\ ], f\rangle$ where I is the set of all maximal L-consistent sets of sentences, $[\varphi]$ is $\{i \in I: \varphi \in i\}$ and $f(\varphi, i)$ is $\cap\{[\psi]: (\varphi \mathrel{\Box\!\!\rightarrow} \psi) \in i\}$. It is easily verified that this triple is indeed an α-model. (It is not an intended model, since possible worlds are not really sets of sentences.) Part of this verification merits closer examination: given that $f(\varphi, i) \subseteq [\psi]$ show that $(\varphi \mathrel{\Box\!\!\rightarrow} \psi) \in i$ as follows. By hypothesis, ψ belongs to every member of $f(\varphi, i)$; that is, to every maximal L-consistent set of sentences that includes the set Σ of all those sentences χ such that $(\varphi \mathrel{\Box\!\!\rightarrow} \chi) \in i$. Therefore ψ is deducible in L from a finite subset of Σ. Either ψ itself is a theorem of L or there are $\chi_1, \ldots, \chi_n$ in Σ such that $(\chi_1 \,\&\, \ldots \,\&\, \chi_n) \supset \psi$ is a theorem of L. Then since L is closed under DWC, either $\varphi \mathrel{\Box\!\!\rightarrow} \psi$ or $((\varphi \mathrel{\Box\!\!\rightarrow} \chi_1) \,\&\, \ldots \,\&\, (\varphi \mathrel{\Box\!\!\rightarrow} \chi_n)) \supset (\varphi \mathrel{\Box\!\!\rightarrow} \psi)$ is a theorem of L; so, since i is maximal L-consistent, $(\varphi \mathrel{\Box\!\!\rightarrow} \psi) \in i$.

It is also easily verified that if L contains Axiom A, the canonical α-model for L meets condition (α.0.6); if L contains Axiom B, it meets (α.0.7); if L contains Axiom C, it meets (α.0.8); if L contains Axiom D, it meets (α.1.1); if L contains Axiom E, it meets (α.1.2); and if L contains Axiom F, it meets (α.2). Thus the canonical α-model for C0 is standard, the canonical α-model for C1 is 1-standard, and the canonical α-model for C2 is 2-standard.

If L is any deductive system closed under TI and DWC and containing Axioms A—C, so that the canonical α-model for L is standard, then let us call the β-conversion of the canonical model for L the *canonical β-model for L*, and let us call the conversion of the canonical α-model for L the *canonical γ-model for L*. The three canonical models for L are equivalent; and if the canonical α-model for L is 1-standard or 2-standard, then so are the other two canonical models for L.

Completeness results

If any sentence is valid in every member of some class of models that includes one of the canonical models for a deductive system L, then that sentence belongs to every maximal L-consistent set of sentences, so it must be a theorem of L. By this argument together with our results about standardness of canonical models and our soundness results, we have now proved the following theorem. (The third part restates a result of Stalnaker and Thomason [5].)

THEOREM

All and only theorems of C0 *are valid in all standard* $\left\{\begin{matrix}\alpha\\ \beta\\ \gamma\end{matrix}\right\}$*-models;*

All and only theorems of C1 *are valid in all 1-standard* $\left\{\begin{matrix}\alpha\\ \beta\\ \gamma\end{matrix}\right\}$*-models;*

All and only theorems of C2 *are valid in all 2-standard* $\left\{\begin{matrix}\alpha\\ \beta\\ \gamma\end{matrix}\right\}$*-models.*

Decidability results

We can now also show that the systems C0, C1, and C2 are decidable. If L is one of these systems and φ is any sentence, proceed as follows to decide whether φ is a theorem of L. First choose a set J with exactly 2^n members, where n is the number of subsentences of φ (including φ itself). Then consider the set M of all β-models $\langle I, [\ ], \$\rangle$ such that: (1) $I \subseteq J$; (2) $[\sigma] = \Lambda$ whenever

σ is a sentence letter that does not occur in φ; and (3) $\langle I, [\,], \$\rangle$ is standard (if L is C0), 1-standard (if L is C1), or 2-standard (if L is C2). There are finitely many models in M, and each of them is such that we can decide whether φ is valid in it. If φ is valid in every model in M, decide that φ is a theorem of L; if φ is invalid in some model in M, decide that φ is not a theorem of L.

If φ is a theorem of L, it follows from our soundness results that φ is valid in every model in M, so in this case the procedure decides correctly.

If φ is not a theorem of L, it follows from our completeness results that φ is invalid in some β-model $\langle I, [\,], \$\rangle$ which is standard (if L is C0), 1-standard (if L is C1), or 2-standard (if L is C2). Call i and j in I *φ-indistinguishable* iff every subsentence of φ is true in that model at both or neither of i and j. Call $\langle {}^*I, {}^*[\,], {}^*\$\rangle$ a *φ-filtration* of $\langle I, [\,], \$\rangle$ iff it is a β-model and there is a function $*$ from I into J such that: (1) ${}^*I = \{*i\colon i \in I\}$; (2) $*i = *j$ iff i and j are φ-indistinguishable; (3) if σ is any sentence letter that occurs in φ, ${}^*[\sigma] = \{*i\colon i \in [\sigma]\}$; (4) if σ is any sentence letter that does not occur in φ, ${}^*[\sigma] = \Lambda$; and (5) for each j in ${}^*I$ there is i in I such that $j = *i$ and ${}^*\$_j = \{\{*k\colon k \in S\}\colon S \in \$_i\}$. We can easily verify that there do exist φ-filtrations of $\langle I, [\,], \$\rangle$. We can verify by induction on subsentences of φ that if ψ is any subsentence of φ (in particular if ψ is φ itself) and $\langle {}^*I, {}^*[\,], {}^*\$\rangle$ is any φ-filtration of $\langle I, [\,], \$\rangle$ then ${}^*[\psi] = \{*i\colon i \in [\psi]\}$. It follows that φ is invalid in any φ-filtration of our original model $\langle I, [\,], \$\rangle$. We can also verify that if $\langle I, [\,], \$\rangle$ is standard, 1-standard, or 2-standard, then so is any φ-filtration of it. (To verify that 2-standardness is preserved, we should first note that if $\langle {}^*I, {}^*[\,], {}^*\$\rangle$ is any φ-filtration and ψ is any sentence there is a truth-functional compound χ of subsentences of φ such that ${}^*[\psi] = {}^*[\chi]$; and for any such χ, ${}^*[\chi] = \{*i\colon i \in [\chi]\}$.) It follows that the φ-filtrations of our original model $\langle I, [\,], \$\rangle$ belong to M. Therefore the non-theorem φ is invalid in some model in M; so in this case also the procedure decides correctly.

This completes the proof of the following theorem.

THEOREM

C0, C1, *and* C2 *are decidable.*

References

DAVID KAPLAN. Review of Saul Kripke, "Semantical analysis of modal logic I." *The journal of symbolic logic*, vol. 31 (1966), pp. 120—122.

DAVID LEWIS. *Counterfactuals*. In preparation.

KRISTER SEGERBERG. "Decidability of S4.1." *Theoria*, vol. 34 (1968), pp. 7—20.

ROBERT STALNAKER. "A theory of conditionals." In Nicholas Rescher (ed.), *Studies in logical theory. American philosophical quarterly* supplementary monograph series (Oxford: Blackwell, 1968).

ROBERT STALNAKER AND RICHMOND THOMASON. "A semantic analysis of conditional logic." *Theoria*, vol. 36 (1970), pp. 23—42.

Received on December 21, 1970.