introduced. A tableau version of consequence—deduction from premises—is easy. One says X follows from a set S of formulas if there is a closed tableau starting with $F\,X$, allowing the additional rule that for any $Z \in S$, $T\,Z$ can be added to the end of any branch.

Finally, I have used signed formulas, but one could just as well work with an unsigned version. Instead of $F\,X$ use $\neg X$, and instead of $T\,X$, just use X. I'll leave a full formulation to you (or see [20,49]). The use of signs brings some additional power, however. It is easier to establish a connection with the Gentzen sequent calculus, as we will do in Section 7. There is a simple signed tableau system for intuitionistic logic, something that is not possible without signs. And one can add extra signs to create proof systems for many-valued logics.

5.2 Destructive Modal Tableaus

Modal tableaus come in more than one version. Some logics have *destructive* tableau systems, [17,31]. These will be presented in this section—a different approach is given in Section 6. The terminology comes from the fact that some destructive tableau rule applications lose information. Destructive tableau proofs tend to be more useful metatheoretically than other kinds of tableaus—for example, one can devise a simple proof of interpolation theorems using such tableau systems.

To continue the uniform treatment begun with the *alpha/beta* grouping, two new categories, *nu* and *pi*, and their components are introduced in Table 9, to take care of the modal operators—both $\Box$ and $\Diamond$ are taken as primitive now.

ν	ν_0	π	π_0
$T\,\Box X$	$T\,X$	$T\,\Diamond X$	$T\,X$
$F\,\Diamond X$	$F\,X$	$F\,\Box X$	$F\,X$

Table 9
Nu and Pi Formulas

In Table 5 I gave a definition of $S^\sharp$ for several modal logics. As it happens, logics whose semantics involve symmetry don't have simple (or any) destructive tableau systems, so these must be dropped. And I'm now allowing more connectives and modal operators as primitive than before. So a definition appropriate for this section is given in Table 10—connections with the earlier version should be clear. A few observations about this definition. First, for all six of the logics we have monotonicity: $S_1 \subseteq S_2$ implies $S_1^\sharp \subseteq S_2^\sharp$. And second, for the **K4**, **S4**, **D4** group we have $S^\sharp \subseteq S^{\sharp\sharp}$. Both of these are easily checked, and both play a role in later soundness and completeness proofs.

Logic	$S^\sharp$
$\mathbf{K}, \mathbf{T}, \mathbf{D}$	$\{\nu_0 \mid \nu \in S\}$
$\mathbf{K4}, \mathbf{S4}, \mathbf{D4}$	$\{\nu_0, \nu \mid \nu \in S\}$

Table 10
Revised Definition of $S^\sharp$

Destructive tableau rules for the logics of Table 10 are as follows. First, all the classical rules of the previous section continue to apply. And in addition there are the rules given in Table 11. These rules require some comment. The second, from ν to get ν_0, is of the same general kind as earlier tableau rules: a branch containing ν can have ν_0 added to the end. The third is slightly different since it is premiseless: at any point on a tableau branch we can add $T\,\Diamond\top$. The first, however, is of a very different nature. Let us call it the π *rule*, though technically what is displayed is actually several rules, depending on the definition of $S^\sharp$. The π rule says that, given a branch containing π, and with S as the entire set of (other) signed formulas on it, that branch can be *replaced* with a new one containing the members of $S^\sharp$, and π_0. The π rule is the reason for the terminology *destructive*—application of this rule removes formulas.

For all logics: $\dfrac{S, \pi}{S^{\sharp}, \pi_0}$ For **T** and **S4**: $\dfrac{\nu}{\nu_0}$ For **D** and **D4**: $\dfrac{}{T\,\Diamond\top}$

Table 11
Destructive Modal Rules

An example of a destructive tableau can be found in Figure 5. It provides a tableau proof, in the **K4** system, of $(\Box\Diamond X \wedge \Box Y) \supset \Box\Diamond(X \wedge Y)$. In it, 2 and 3 are from 1, and 4 and 5 are from 2 by α. Then a destructive π rule applies, with 3 as the π formula. The original branch is replaced by a new one, shown below the line, with 6 from 3, 7 and 8 from 4, and 9 and 10 from 5; formulas 1 and 2 disappear entirely. Another π rule application now happens, with 8 as the π formula, producing the new branch shown below the second line. Item 11 is from 8; 12 and 13 are from 6; 14 and 15 are from 7; 16 and 17 are from 9. Finally β applied to 13 produces 18 and 19, and both branches are closed.

$F\,(\Box\Diamond X \wedge \Box Y) \supset \Box\Diamond(X \wedge Y)$ 1.

$T\,\Box\Diamond X \wedge \Box Y$ 2.

$F\,\Box\Diamond(X \wedge Y)$ 3.

$T\,\Box\Diamond X$ 4.

$T\,\Box Y$ 5.

$F\,\Diamond(X \wedge Y)$ 6.

$T\,\Box\Diamond X$ 7.

$T\,\Diamond X$ 8.

$T\,\Box Y$ 9.

$T\,Y$ 10.

$T\,X$ 11.

$F\,\Diamond(X \wedge Y)$ 12.

$F\,X \wedge Y$ 13.

$T\,\Box\Diamond X$ 14.

$T\,\Diamond X$ 15.

$T\,\Box Y$ 16.

$T\,Y$ 17.

$F\,X$ 18. $\quad$ $F\,Y$ 19.

Fig. 5. **K4** Destructive Tableau Example

Destructive rules add a level of complexity to tableaus. Tableau rules are non-deterministic—they say what can be done, but the order of rule application is not specified. It can be shown that, for the classical system of Section 5.1, a kind of Church-Rosser property applies. If a formula X is a tautology, any attempt to provide a closed tableau for $F\,X$ will succeed, no matter in what order the rules are applied, provided only that on each branch, every non-atomic formula eventually has a rule applied to it. The π rule changes things. We might have a tableau branch containing, among other things, both $T\,\Diamond X$ and $T\,\Diamond Y$. Either could be used as π in a π-rule application, but when so used, it will cause the deletion of the other formula. It can happen that a proof is obtainable when one choice is made, but not the other—we can choose badly. This means that a systematic proof search must allow backtracking, and so will be inherently more time-consuming than a systematic search in the classical system.

5.3 Soundness and Completeness

A proof of soundness can also serve to motivate the rules of Table 11. We'll say a signed formula is *realized* at a possible world w of a model $\mathcal{M}$ if the formula is $T\,X$ and $\mathcal{M}, w \Vdash x$, or the formula is $F\,X$ and $\mathcal{M}, w \not\Vdash X$. Let **L** be one of the six logics for which tableau rules have been provided. Call a set S of signed formulas **L**-*satisfiable* if there is some **L** model $\mathcal{M}$, and some possible world, w, of it that realizes all the members of S. Call a tableau branch **L**-satisfiable if the set of signed formulas on it is **L**-satisfiable. And call a tableau **L**-satisfiable if some branch is. The key fact is that satisfiability is an invariant for tableau construction. That is, each tableau rule preserves **L**-satisfiability. Let us call a rule *sound* if it has this satisfiability preserving feature. It is the need to have sound rules that dictates some of the features of the systems we have seen—for instance, this is why the version of $S^\sharp$ that works for **K4** will not serve for **K**.

Proposition 5.1 *Suppose $\mathcal{T}$ is an* **L** *tableau that is* **L***-satisfiable. If any* **L** *tableau rule is applied to $\mathcal{T}$ the resulting tableau is still* **L***-satisfiable.*

Proof. Suppose branch θ of $\mathcal{T}$ is **L**-satisfiable, say its members are realized at world w of model $\mathcal{M}$. And say a tableau rule is applied to $\mathcal{T}$. If it is applied on a branch other than θ, the resulting tableau is trivially **L**-satisfiable, so now assume the tableau rule has been applied on θ.

If the applied rule was a β rule θ branches, technically it is replaced with two new branches which we'll call θ, β_1 and θ, β_2, using the obvious notation. Since β was realized at w, a check of each case in the definition of β shows that either β_1 is realized at w, or β_2 is. Consequently, either the branch θ, β_1 is satisfied, at w, or the branch θ, β_2 is. Either way the resulting tableau is **L**-satisfiable. The argument if the rule application was an α or a negation rule is even simpler, and is omitted.

Now suppose the applied rule was the π rule from Table 11. The key thing we need is this. For each of the logics under consideration, if members of S are realized at world w of an **L** model $\mathcal{M}$, and if w' is any world of $\mathcal{M}$ that is accessible from w, then members of $S^\sharp$ are realized at w'. Verification of this is left to you. Now, suppose θ consists of the members of S, and the signed formula π, and w realizes all the signed formulas on θ. Since π is realized at w, there must be an alternate world w' at which π_0 is realized. As we just noted, at w' all members of $S^\sharp$ are realized. Then in the resulting tableau there is still a satisfiable branch, though its members are realized at a different world than the one realizing the members of the original branch.

The other rules from Table 11 are straightforward. □

Proposition 5.2 *If X has a proof using the tableau rules for* **L***, then X is* **L** *valid.*

Proof. I'll show the contrapositive. Suppose X is not **L** valid; so there is some world of some **L**-model at which X is false, Then $\{F\,X\}$ is **L**-satisfiable. Any tableau proof of X must start with the tree with only $F\,X$, at its root. This is an **L**-satisfiable tableau, so Proposition 5.1 says only **L**-satisfiable tableaus will be produced. A **L**-satisfiable tableau cannot be closed. Hence X can have no **L** tableau proof. □

Next we turn to completeness. A common way of showing completeness for tableau systems involves devising a systematic way of applying tableau rules. Such a systematic approach is presented for classical logic in [49], for instance. While such a method has utility when computer implementations are involved, it is often hard work. Fortunately the method used to show completeness for axiom systems in Section 2.2 can also be applied, and is much simpler.

First we need a small generalization of the notion of tableau. So far we have started tableau constructions with a single signed formula. From now on, if S is a finite set of signed formulas, a tableau for S will be any tableau starting with a single branch containing the members of S, and continuing using the usual tableau rules. Then, a tableau proof of a formula X is a closed tableau for the set $\{F\,X\}$.

Let **L** be one of the logics for which tableau rules have been provided in Table 11. Call a set S of *signed* formulas **L**-*inconsistent* if there is a closed **L** tableau for some finite subset of S, and call S **L**-*consistent* if it is not **L**-inconsistent. Clearly this notion of **L**-consistency is of finite character, and so the Lindenbaum construction applies, see (1). Every **L**-consistent set of signed formulas can be extended to a maximal **L**-consistent set.

We construct a model $\mathcal{M} = \langle \mathcal{G}, \mathcal{R}, \mathcal{V} \rangle$ much like we did in Section 2.2. $\mathcal{G}$ is the collection of all maximal **L**-consistent sets of signed formulas. For $w_1, w_2 \in \mathcal{G}$, $w_1 \mathcal{R} w_2$ provided $w_1^\sharp \subseteq w_2$. And finally, $w \in \mathcal{V}(P)$ provided $T\,P \in w$. One cannot, at this point, show an exact counterpart of (2) (though it is, in fact, true). But one can show the following, involving an implication instead of an equivalence. For every *signed* formula $\mathcal{X}$ and possible world $w \in \mathcal{G}$

$$\mathcal{X} \in w \Longrightarrow w \text{ realizes } \mathcal{X} \text{ in the model } \mathcal{M} \tag{4}$$

The proof is by induction on the complexity of signed formulas. Since we have several connectives as primitive now, I'll make use of uniform notation. Here are the cases needed to establish (4).

Suppose P is atomic. If $T\,P \in w$ then $\mathcal{M}, w \Vdash P$ by definition of $\mathcal{V}$, so $T\,P$ is realized at w. Likewise if $F\,P \in w$, since w is **L**-consistent, $T\,P \notin w$, and so $\mathcal{M}, w \not\Vdash P$, and again $F\,P$ is realized at w.

The negation cases are straightforward, and are omitted.

Suppose we have a β signed formula, $\beta \in w$, and (4) and is known for β_1 and β_2. Since w is **L**-consistent, it follows from the tableau rules that one of $w \cup \{\beta_1\}$ or $w \cup \{\beta_2\}$ is **L**-consistent. Then it follows by maximality of w that either $\beta_1 \in w$ or $\beta_2 \in w$. By the induction hypothesis, either β_1 or β_2 is realized at world w of $\mathcal{M}$. And an examination of the cases in the definition of β formulas shows this is enough for β to be realized at w as well.

The α case is similar, and is omitted.

Suppose we have a ν formula, $\nu \in w$, and (4) is known for ν_0. Let w' be any member of $\mathcal{G}$ with $w \mathcal{R} w'$. For each choice of **L**, $\nu_0 \in w^\sharp$ and since $w^\sharp \subseteq w'$ by definition of $\mathcal{R}$, $\nu_0 \in w'$. By the induction hypothesis, w' realizes ν_0. A check of cases in the definition of ν shows that, since w' was arbitrary, ν is realized at w.

Finally, suppose we have a π formula, $\pi \in w$, and (4) is known for π_0. Using the π rule from Table 11, it follows that $w^\sharp \cup \{\pi_0\}$ is consistent. Let w' be a maximal **L**-consistent extension of this. Then $w' \in \mathcal{G}$ and since $\pi_0 \in w'$, the induction hypothesis gives us that π_0 is realized at w'. Finally, since $w \mathcal{R} w'$, it follows, for each case in the definition of π, that π is realized at w.

Now that (4) has been established, putting the final pieces together is easy. For **L** being any of the six logics from Tables 10 and 11, one can easily check that the construction described above produces a model $\mathcal{M}$ that is, in fact, an **L**-model. Here is part of such a verification. For any of the six logics being treated, $S_1 \subseteq S_2$ implies $S_1^\sharp \subseteq S_2^\sharp$. For **L** being one of **K4**, **S4**, or **D4**, $S^\sharp \subseteq S^{\sharp\sharp}$. Now, suppose $w_1 \mathcal{R} w_2$ and $w_2 \mathcal{R} w_3$. Then $w_1^\sharp \subseteq w_2$ and $w_2^\sharp \subseteq w_3$, so $w_1^{\sharp\sharp} \subseteq w_2^\sharp \subseteq w_3$. So if **L** is one of **K4**, **S4**, or **D4**, $w_1^\sharp \subseteq w_3$, and hence $w_1 \mathcal{R} w_3$. Thus for these three logics, the model is transitive. I'll leave the other conditions to you.

Now, if X is not provable by **L**-tableaus, $\{F\,X\}$ is **L**-consistent. Extend it to a maximal **L**-consistent set w, which will be a world of the **L** model constructed above. It is a world at which X is false, by (4). Thus X is not **L**-valid. This establishes the following.

Proposition 5.3 *For* **L** *being one of* **K**, **T**, **D**, **K4**, **S4**, **D4**, *the* **L** *tableau rules are complete.*

5.4 The Logic **GL**

In Section 2.3 we saw that, while a straight canonical model argument was not able to prove completeness for **GL**, still a completeness argument could be given. A variation of that argument works for an appropriate destructive tableau system for **GL** as well. I'll sketch this here.