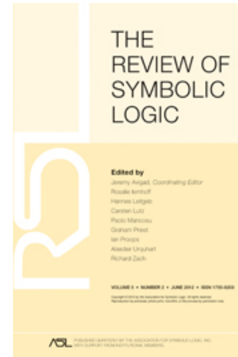

CONSERVATIVELY EXTENDING CLASSICAL LOGIC WITH TRANSPARENT TRUTH

DAVID RIPLEY

The Review of Symbolic Logic / Volume 5 / Issue 02 / June 2012, pp 354 - 378
DOI: 10.1017/S1755020312000056, Published online: 26 March 2012

THE REVIEW OF SYMBOLIC LOGIC
Volume 5, Number 2, June 2012

CONSERVATIVELY EXTENDING CLASSICAL LOGIC WITH TRANSPARENT TRUTH

DAVID RIPLEY

Department of Philosophy, University of Connecticut and School of Historical and Philosophical Studies, University of Melbourne

Abstract. This paper shows how to conservatively extend a classical logic with a transparent truth predicate, in the face of the paradoxes that arise as a consequence. All classical inferences are preserved, and indeed extended to the full (truth-involving) vocabulary. However, not all classical metainferences are preserved; in particular, the resulting logical system is *nontransitive*. Some limits on this nontransitivity are adumbrated, and two proof systems are presented and shown to be sound and complete. (One proof system features admissible Cut, but the other does not.)

§1. Introduction.

Adding a truth predicate to a language governed by classical logic is not easy. It is particularly tricky when the truth predicate is intended to be *transparent*—such that $T\langle A\rangle$ and A are everywhere intersubstitutable. The trouble, as is well-known, comes from such paradoxes as the liar; because of them, theories of truth typically either use a nontransparent truth predicate (Halbach, 2011; Maudlin, 2004) or a logic weaker than classical (Beall, 2009; Field, 2008). (Others—e.g., Priest, 2006—have used both subclassical logics and nontransparent truth.) This paper will take a different avenue, showing how to conservatively extend a classical logic with a fully transparent truth predicate in the face of the paradoxes.

Of course, there is a reason that the above-cited approaches have not taken this route. They have held to *transitivity* of consequence: in its simplest form, they have held that if $A \vdash B$ and $B \vdash C$, then $A \vdash C$. The logical system to be outlined in this paper is nontransitive; this is the escape route that allows transparent truth to mesh with classical logic. Despite this, the system to be considered presently validates all classically valid inferences; as a result, the nontransitivity is relatively contained.[1] (Transitivity itself is a metainference, a closure property on the set of valid arguments. Not all metainferences usually associated with classical logic will be preserved here; the target system is thus *weakly classical* in the sense of Field, 2008.)

Here's the plan: in §2, I'll lay out the models for the classical logic that will be in play. This logic extends classical logic with distinguished names for formulas, allowing the language to talk about itself. (This extension itself is nonconservative, as will be discussed presently.) The models used are not the usual models for classical logics; they involve three values instead of two. In §3, I'll show how to conservatively extend the language with a transparent truth predicate. §4 and §5 present two distinct proof systems for the logic arrived at, and explore the relations between them, as well as proving them both to be sound and complete for the model-theoretically defined logic. §6 concludes.

Received: January 19, 2011.

[1] This paper builds on Cobreros *et al*. (2011), which applies a similar approach to providing a nontransitive theory of vagueness that validates both full tolerance and all classical inferences.

doi:10.1017/S1755020312000056

§2. Three-valued models for classical logic. I'll work with two languages: $\mathcal{L}$ and $\mathcal{L}^+$. The base language $\mathcal{L}$ is an ordinary first-order language with identity, including constants $\top$ and $\bot$; $\mathcal{L}^+$ is $\mathcal{L}$ plus a distinguished unary predicate T—a truth predicate. The eventual target is a theory of truth in $\mathcal{L}^+$. To ensure that there will be paradoxes around, the interpretation of $\mathcal{L}$ and $\mathcal{L}^+$ will be partially constrained. In particular, the individual constants will come divided into two countable sets. The members of one set—call them "ordinary names"—will function as the usual sort of individual constant, receiving their denotation from each model. The members of the other set—call them "distinguished names"—will receive their denotations independently. We fix a 1-1 function τ from distinguished names onto formulas of $\mathcal{L}^+$, and require that a distinguished name n denote $\tau(n)$ in every model. This way of handling reference to $\mathcal{L}^+$ is borrowed from Barwise & Etchemendy (1987). It means that $\mathcal{L}^+$ must be a subset of the domain of every model; I thus work only with infinite models. (This is usual for theories of truth; see, e.g., Kremer, 1988.)

This approach allows for paradoxes of all sorts. For example, if m and n are distinguished names, then by letting $\tau(n) = \neg Tn$, we can ensure that n denotes a liar sentence; by letting $\tau(m) = Tn$ and $\tau(n) = \neg Tm$, we can create a liar pair; by letting $\tau(n) = (Tn \supset \bot) \wedge p$, we can create a contingent Curry sentence (paradoxical iff p); and so on. Throughout the paper, I'll suppose some fixed such τ; everything goes through no matter what τ is, and so goes through in the presence of paradoxes of all sorts. Given a formula A of $\mathcal{L}^+$, I'll write $\langle A \rangle$ for the distinguished name n such that $\tau(n) = A$; thus, a liar sentence λ is such that $\lambda = \tau(\langle \lambda \rangle) = \neg T \langle \lambda \rangle$.[2]

This approach to naming renders the connection between A and $\langle A \rangle$ not purely syntactic, as it depends on τ. If this is seen as a drawback (as it is by two anonymous referees), there are a few options. The most thorough would be to consider this approach in a setting including arithmetic, or some other setting rich enough to contain the syntactic theory of $\mathcal{L}^+$. Such a setting, however, is too rich to allow for a proof of completeness; the present setting, on the other hand, allows for completeness, as will be shown in §5.

Another option would be to follow, for example, Kremer (1988) and include in the language a quote-name-forming device Q such that $Q(A)$, for any formula A, is a singular term required to denote A in any model. This is very similar to the present approach in its upshots, but there are some differences. Most of these differences stem from the actual syntactic occurrence in $Q(A)$ of the formula A (by contrast, A doesn't occur in $\langle A \rangle$; $\langle A \rangle$ is a syntactic simple); as a result, no formula can contain its own quote-name, while on the present approach, formulas can contain their own distinguished names. The difference extends to loops: no formula A can contain the quote-name of a formula B which itself contains the quote-name of A; but A can contain the distinguished name of a formula B which itself contains the distinguished name of A, and so on to bigger loops.

On the other hand, the τ-function approach can completely simulate the quote-name approach, via a particular choice of τ. We can simply let τ work in "layers": first, assign distinguished names to all the formulas of $\mathcal{L}^+$ that do not themselves contain distinguished names; second, assign distinguished names to all the formulas that contain only distinguished names that have already been assigned; third, repeat the second step until every distinguished name has been assigned, and every formula has received a distinguished name. Not every possible τ can be reached in this way, but some can. The distinguished

---

[2] Given some choices of τ, there may be no such λ, and given other choices, there may be many such; I'll assume throughout that we've chosen τ so that there is at least one.

names together with a τ resulting from this process will behave relevantly like quote-names; no formula will be able to contain its own distinguished name, nor will there be any loops of distinguished names. I think we may as well work with the τ-function approach for its increased generality. If you nonetheless prefer the quote-name approach, only minor modifications will be necessary to what follows.

DEFINITION 2.1. *An ST-model for $\mathcal{L}$ is a structure $\langle D, I \rangle$ such that:*

- *D is a domain such that $\mathcal{L}^+ \subseteq D$, and*
- *I is an interpretation function:*
    - *For an ordinary name or variable a, $I(a) \in D$*
    - *For a distinguished name n, $I(n) = \tau(n)$; thus, $I(\langle A \rangle) = A$*
    - *For an n-ary predicate P, $I(P) \in \{0, \frac{1}{2}, 1\}^{D^n}$*
    - *For an atomic formula $A = P(t_1, t_2, \dots, t_n)$,* $I(A) = I(P)(\langle I(t_1), I(t_2), \dots, I(t_n) \rangle)$
    - $I(s = t) = 1$ *iff* $I(s) = I(t)$, *and* 0 *otherwise*
    - $I(\top) = 1$; $I(\bot) = 0$
    - $I(\neg A) = 1 - I(A)$
    - $I(A \wedge B) = \min(I(A), I(B))$
    - $I(A \vee B) = \max(I(A), I(B))$
    - $I(A \supset B) = \max(1 - I(A), I(B))$
    - $I(\forall x A) = \min(\{I'(A) : I' \text{ is an } x\text{-variant of } I\})$
    - $I(\exists x A) = \max(\{I'(A) : I' \text{ is an } x\text{-variant of } I\})$

Equivalently, we can take $\wedge$, $\neg$, and $\forall$ as primitive, and use them to define $\vee$, $\supset$, and $\exists$ in the usual way. (Because of these equivalences, proofs by induction will only consider $\neg$, $\wedge$, and $\forall$ cases.)

ST-models, except for their adjustment to handle distinguished names, are the usual sort of model for the logic LP or Strong Kleene (K3) logic (for details, see, e.g., Beall & van Fraassen, 2003; Priest, 2008)—two nonclassical logics well-suited to transparent truth. Both of those logics define a valid argument as one which preserves designated value, but they differ in which values they take to be designated: LP designates both $\frac{1}{2}$ and 1, while K3 designates only 1. Here, we'll define validity differently, having use for both sorts of designation. Following Cobreros *et al*. (2011), I'll call them *tolerant* and *strict satisfaction* respectively:

DEFINITION 2.2. *An ST-model $M = \langle D, I \rangle$ strictly satisfies a formula A ($M \Vdash^s A$) iff $I(A) = 1$. It tolerantly satisfies A ($M \Vdash^t A$) iff $I(A) > 0$.*

We can now use these two notions of satisfaction to define four notions of consequence, again as in Cobreros *et al*. (2011):

DEFINITION 2.3. *Where Γ, Δ are sets of formulas in $\mathcal{L}$, and $m, n \in \{s, t\}$: $\Gamma \vDash^{mn} \Delta$ iff every model M such that $M \Vdash^m \gamma$ for every $\gamma \in \Gamma$ is such that $M \Vdash^n \delta$ for some $\delta \in \Delta$.*

Where $m = n = t$, this amounts to preservation of tolerant satisfaction—LP—and where $m = n = s$, this amounts to preservation of strict satisfaction—K3. But when $m \neq n$,

$\vDash^{mn}$ is neither of these.[3] As it happens, though, $\vDash^{st}$ is a familiar consequence relation nonetheless: a classical consequence relation.[4,5] The rest of this section proves this claim.

DEFINITION 2.4. *A classical model with distinguished names for $\mathcal{L}$ is a structure $\langle D, I\rangle$ such that:*

- *$\langle D, I\rangle$ is an ST-model for $\mathcal{L}$, and*
- *For each n-ary predicate P, $I(P) \in \{0, 1\}^{D^n}$*

In other words, classical models with distinguished names are just a special kind of ST-model: the kind that never uses the value $\frac{1}{2}$ for atomic formulas. Induction on formula length will show that classical models never assign the value $\frac{1}{2}$, to any formula of any length. These are familiar models for classical logic, restricted to handle distinguished names.

DEFINITION 2.5. *A set Δ of formulas is a CLDN consequence of a set Γ of formulas iff every classical model with distinguished names $M = \langle D, I\rangle$ such that $I(\gamma) = 1$ for every $\gamma \in \Gamma$ is such that $I(\delta) = 1$ for some $\delta \in \Delta$.*

CLDN consequence is the familiar notion of classical consequence, extended to handle distinguished names.

DEFINITION 2.6. *A model $M' = \langle D', I'\rangle$ is a classicization of a model $M = \langle D, I\rangle$ (written $M \preceq M'$) iff:*

- *$D' = D$,*
- *for any term t, $I(t) = I'(t)$, and*
- *for any ntuple $x \in D^n$ and any n-ary predicate P, if $I(P)(x) = 0$ or 1, then $I'(P)(x) = I(P)(x)$.*

FACT 2.7. *Take any ST-models $M = \langle D, I\rangle$ and $M' = \langle D', I'\rangle$ such that $M \preceq M'$. For any formula A, if $I(A) = 0$ or 1, then $I'(A) = I(A)$.*

Proof. Induction on A's construction. □

When $M \preceq M'$, then, M' assigns value $\frac{1}{2}$ in a subset of the cases that M does; and when this condition holds for predicates, that suffices for it to extend to all formulas.

[3] The relation $\vDash^{ts}$ will not enter much in what follows. It's almost empty—if we remove $=$, $\top$, and $\bot$ from the language, it *is* empty—and it won't do any work here. Note as well that these 'mixed' consequence relations are special cases of ideas that have been explored before. For example, Malinowski's (Malinowski, 1990) notion of q-consequence encompasses $\vDash^{ts}$, and Frankowski's (Frankowski, 2004) notion of p-consequence encompasses $\vDash^{st}$. The application to semantic paradox, however, is new. For applications of a closely related system to paradoxes of vagueness, see Cobreros *et al*. (2011).

[4] Due to the presence of $\mathcal{L}^+$ in the domain of every model, it is a nonconservative extension of basic classical logic: it amounts to classical logic over infinite domains. Thanks to an anonymous referee for calling attention to this.

[5] Note that Definition 2.3 defines our consequence relations to hold between sets of premises and sets of conclusions. As is explored in, for example, Shoesmith & Smiley (1978), this allows for finer distinctions between logics. For example, supervaluationist logics, while sometimes presented as 'classical', are classical only in their single-conclusion fragments; a multiple-conclusion framework reveals their subclassicality, as is pointed out in Hyde (1997); Ripley (2011b). Nothing like this happens in the present case. The relation $\vDash^{st}$ remains classical even in a multiple-conclusion setting.

LEMMA 2.8. $\vDash^{st}$ *is CLDN consequence.*

Proof. Suppose $\Gamma \vDash^{st} \Delta$. Then there is no ST-model that provides a counterexample. Since every classical model with distinguished names is an ST-model, and being a CLDN counterexample is the same as being an st-counterexample (assigning 1 to all the premises and 0 to all the conclusions), no classical model with distinguished names is a counterexample. So Δ is a CLDN consequence of Γ.

On the other hand, suppose $\Gamma \nvDash^{st} \Delta$. Then there is an ST-model $M = \langle D, I \rangle$ such that $I(\gamma) = 1$ for every $\gamma \in \Gamma$ and $I(\delta) = 0$ for every $\delta \in \Delta$. Consider any model $M' = \langle D, I' \rangle$ where I' is obtained from I by assigning 0 or 1 (it doesn't matter which) wherever $I(P)$ assigns $\frac{1}{2}$, for any predicate P. It follows that $M \preceq M'$, and so $I'(\gamma) = 1$ for every $\gamma \in \Gamma$ and $I'(\delta) = 0$ for every $\delta \in \Delta$. M' is a classical model with distinguished names; thus, M' is a CLDN counterexample, and Δ is not a CLDN consequence of Γ. □

So far, this is just a way to recover a classical logic from ST-models. But we know from, for example, Kripke (1975) that these models play nice with transparent truth. So we can use this model theory to see a way in which this classical logic can be made to play nice with transparent truth as well.

§3. Adding truth.

This section shows how to conservatively add a transparent truth predicate (in a sense to be made precise presently) to our language.

3.1. A conservative extension.

DEFINITION 3.1. *An ST-model for $\mathcal{L}^+$ must meet all the conditions in Definition 2.1, plus one more: that for all formulas A, $I(T\langle A \rangle) = I(A)$.*

This suffices for the real goal: that for all terms a, whenever $I(a) \in \mathcal{L}^+$, $I(T)(I(a)) = I(I(a))$. We could in addition require for terms a that if $I(a) \notin \mathcal{L}^+$, then $I(Ta) = 0$. This, though, results in noncompact consequence relations (in the presence of our other assumptions), as is shown in Kremer (1988). I won't consider it further: for our purposes here, when $I(a) \notin \mathcal{L}^+$, $I(Ta)$ can take any value.

DEFINITION 3.2. *Where Γ, Δ are sets of formulas in $\mathcal{L}^+$, and $m, n \in \{s, t\}$: $\Gamma \vDash^{mn}_{+} \Delta$ iff every model M such that $M \Vdash^m \gamma$ for every $\gamma \in \Gamma$ is such that $M \Vdash^n \delta$ for some $\delta \in \Delta$.*

This simply extends our earlier definitions of consequence to models for the expanded language. To state the precise sense in which truth will be transparent, we need the notion of a T-transform.

DEFINITION 3.3. *We define T-transform as follows: for any formulas A, B, C, B', C',*

- *A is a T-transform of A.*
- *A is a T-transform of $T\langle A \rangle$.*
- *$T\langle A \rangle$ is a T-transform of A.*
- *If B, C are T-transforms of B', C':*
    - *$\neg B$ is a T-transform of $\neg B'$.*
    - *$B \wedge C$ is a T-transform of $B' \wedge C'$.*
    - *$B \vee C$ is a T-transform of $B' \vee C'$.*
    - *$B \supset C$ is a T-transform of $B' \supset C'$.*
    - *$\forall x B$ is a T-transform of $\forall x B'$.*
    - *$\exists x B$ is a T-transform of $\exists x B'$.*

- *If A is a T-transform of B and B is a T-transform of C, then A is a T-transform of C.*
- *A set of formulas Γ is a T-transform of a set of formulas Γ' iff every $\gamma \in \Gamma$ is a T-transform of some $\gamma' \in \Gamma'$, and every $\gamma' \in \Gamma'$ is a T-transform of some $\gamma \in \Gamma$.*
- *These are the only T-transforms.*

FACT 3.4. *On any model $M = \langle D, I \rangle$, if A is a T-transform of A', then $I(A) = I(A')$.*

Proof. Induction on the definition of T-transform. □

COROLLARY 3.5 (Transparency). *Where Γ, Δ are T-transforms of Γ', Δ', respectively: $\Gamma \models^{mn}_{+} \Delta$ iff $\Gamma' \models^{mn}_{+} \Delta'$.*

This gives the precise sense in which T is *transparent*. No amount of adding Ts or removing them can make a valid argument invalid, or vice versa. This is so even if the Ts are used to modify subformulas of premises or conclusions, and it is so no matter which of our four notions of validity is in play.

This tells us we have transparent truth, but it's not yet a guarantee that it's playing nice. After all, the universal consequence relation on $\mathcal{L}^+$ (the relation $\models$ according to which $\Gamma \models \Delta$ for every Γ, Δ) also features transparent truth, in this sense.

In fact, this universal relation is the relation we arrive at if we try to add transparent truth to ordinary two-valued classical models. If we require that $I(\lambda) = I(T\langle\lambda\rangle)$ (transparent truth), that $I(\lambda) = I(\neg T\langle\lambda\rangle)$ (because λ is $\neg T\langle\lambda\rangle$), and that $I(T\langle\lambda\rangle) \neq I(\neg T\langle\lambda\rangle)$ (by the usual classical clause for $\neg$), then we've imposed inconsistent requirements on our models. Supposing (as I will throughout this paper) a classical metalanguage, this means there are no models, and thus no countermodels; every argument is valid. Is $\models^{st}_{+}$ any better than this? Yes, it is.

FACT 3.6. *For any model $M = \langle D, I \rangle$ of $\mathcal{L}$, there is a model $M^+ = \langle D, I^+ \rangle$ of $\mathcal{L}^+$ such that, for any $A \in \mathcal{L}$, $I(A) = I^+(A)$.*

Proof. The locus classicus is Kripke (1975). Kripke assumes that M only assigns values 0 and 1, never $\frac{1}{2}$, but this assumption plays no role in his proof, and is dispensed with here. For a much more general proof, from which the present fact follows, see Leitgeb (1999). □

COROLLARY 3.7 (Conservative Extension). *For any $\Gamma, \Delta \subseteq \mathcal{L}$, $\Gamma \models^{mn} \Delta$ iff $\Gamma \models^{mn}_{+} \Delta$.*

Proof. If a model for $\mathcal{L}^+$ counterexamples the $\models^{mn}_{+}$ claim, it already is a model for $\mathcal{L}$ that counterexamples the $\models^{mn}$ claim. Conversely, if there is a model for $\mathcal{L}$ that counterexamples the $\models^{mn}$ claim, then (by Fact 3.6) there is a model for $\mathcal{L}^+$ that counterexamples the $\models^{mn}_{+}$ claim. □

Thus, $\models^{st}_{+}$ conservatively extends CLDN with a fully transparent truth predicate. But we can go farther. So far, we've guaranteed that classical inferences *drawn from* $\mathcal{L}$ are unchanged by the addition of T. But it's worth wondering, how far do these inferences extend into the expanded language? For example, do we have excluded middle ($\top \models^{st}_{+} A \vee \neg A$) in full generality, or only applied to As from $\mathcal{L}$? How about explosion ($A \wedge \neg A \models^{st}_{+} \bot$)? The next fact answers these questions:

FACT 3.8 (Inference Extension). *If $\Gamma \models^{mn} \Delta$ and $\Gamma \cup \Delta \subseteq \mathcal{L}$, then $\Gamma^* \models^{mn}_{+} \Delta^*$, where * is any uniform substitution (of open formulas for predicates, avoiding bound-variable conflict in the usual ways) on $\mathcal{L}^+$.*

Proof. Suppose there is a counterexample to the claim that $\Gamma^* \vDash_+^{mn} \Delta^*$; that is, suppose there is some model $M = \langle D, I \rangle$ such that $M \Vdash^m \gamma^*$ for every $\gamma^* \in \Gamma^*$, but $M \nVdash^n \delta^*$ for any $\delta^* \in \Delta^*$.

Consider a model $M' = \langle D', I' \rangle$ such that:

- $D' = D$,
- I' matches I on terms, and
- For each n-ary predicate P occurring in $\Gamma \cup \Delta$, $I'(P)$ is the function from D^n to values that assigns to each $n$tuple the same value that $I(P^\star)$ does.[6]

Since there are no restrictions on the ranges of values that predicates in $\mathcal{L}$ can take, there is such an M'. Since all connectives are value-functional, and all quantifiers depend only on the ranges of values of their arguments, it follows that $M' \Vdash^m \gamma$ for every $\gamma \in \Gamma$, but $M \nVdash^n \delta$ for any $\delta \in \Delta$. Thus $\Gamma \nvDash^{mn} \Delta$. □

No similar proof will work to show that the $\vDash_+^{mn}$s themselves are closed under uniform substitution; they are not. The proof would break down in its final paragraph: some ranges of values *are* barred to T. For example, where λ is $\neg T\langle\lambda\rangle$, $I(T)$ can only assign the value $\frac{1}{2}$ to λ. More prosaically, we get conditional restrictions: if Fa gets a certain value, that constrains $T\langle Fa\rangle$ to take that same value. For this reason, we don't quite have unrestricted uniform substitution on $\mathcal{L}^+$; but then, we never do when trying to provide a logic of truth, as uniform substitution neglects the connection between A and $\langle A\rangle$.

Facts 3.4, 3.6, and 3.8 get us:

THEOREM 3.9. $\vDash_+^{st}$ *is a conservative extension of CLDN with a fully transparent truth predicate, and all valid inferences of CLDN extend to the full vocabulary.*

3.2. More properties of $\vDash_+^{st}$***.*** $\vDash_+^{st}$ is thus a very classical consequence relation. But there are, of course, some surprises, as there must be to accommodate fully transparent truth. In particular, some familiar metainferences fail. ($\vDash_+^{st}$ is thus *weakly classical*, in the terminology of Field, 2008.) The purpose of this section is, first, to highlight some important metainferences that still hold for $\vDash_+^{st}$, and second, to show some examples of metainferences that fail. We'll see that the noteworthy feature here is a certain sort of *nontransitivity*. I'll go on to outline a broad class of areas where transitivity is 'safe'; it may not hold everywhere, but it holds in most places you'd want to go.

First, the familiar properties:

FACT 3.10. $\vDash_+^{st}$ *is reflexive and monotonic. That is:*

- *For any formula* A, $A \vDash_+^{st} A$.
- *For any* $\Gamma \subseteq \Gamma'$, $\Delta \subseteq \Delta'$, *if* $\Gamma \vDash_+^{st} \Delta$, *then* $\Gamma' \vDash_+^{st} \Delta'$.

Proof. The definitions of $\vDash_+^{st}$, $\Vdash^s$, and $\Vdash^t$ suffice. □

FACT 3.11. $\Gamma, A \vDash_+^{st} \Delta$ *iff* $\Gamma \vDash_+^{st} \neg A, \Delta$, *and* $\Gamma \vDash_+^{st} A, \Delta$ *iff* $\Gamma, \neg A \vDash_+^{st} \Delta$.

[6] That is, suppose P is an $n$ary predicate and $P^* = A(x_1, \ldots, x_n)$. Then $I'(P)(\langle d_1, \ldots, d_n\rangle) = J(P^*)$, where J is the variant of I such that $J(x_i) = d_i$ for $1 \leq i \leq n$. Thanks to an anonymous referee for this clarification.

Proof. For any model M, $M \Vdash^s A$ iff $M \nVdash^t \neg A$, and $M \Vdash^s \neg A$ iff $M \nVdash^t A$. Thus, for either of these inference pairs, a counterexample to one of its members is a counterexample to the other. □

FACT 3.12. *Where Γ' comes from Γ by possibly conjoining some of its members, and Δ' comes from Δ by possibly disjoining some of its members,* $\Gamma \vDash^{st}_{+} \Delta$ *iff* $\Gamma' \vDash^{st}_{+} \Delta'$.

Proof. The definition of $\vDash^{st}_{+}$ and the clauses for $\wedge$ and $\vee$ suffice. □

These two together give us:

FACT 3.13 (Deduction Theorem). $\Gamma, A \vDash^{st}_{+} B, \Delta$ *iff* $\Gamma \vDash^{st}_{+} A \supset B, \Delta$.

We also have:

FACT 3.14 (Proof by Cases). *If* $\Gamma, A \vDash^{st}_{+} \Delta$ *and* $\Gamma, B \vDash^{st}_{+} \Delta$, *then* $\Gamma, A \vee B \vDash^{st}_{+} \Delta$.

Proof. Again, the definition of $\vDash^{st}_{+}$ and the clause for $\vee$ suffice. □

So far, so familiar. But of course *something* had to give to enable transparent truth to appear in classical logic. Here it is:

FACT 3.15. *There are formulas $A, B, C \in \mathcal{L}^+$ such that $A \vDash^{st}_{+} B$ and $B \vDash^{st}_{+} C$ but $A \nvDash^{st}_{+} C$.*

Proof. Return to our liar sentence λ. $\top \vDash^{st}_{+} \lambda \wedge \neg\lambda$, and $\lambda \wedge \neg\lambda \vDash^{st}_{+} \bot$, but $\top \nvDash^{st}_{+} \bot$. □

This failure of transitivity, though, occurs *exactly where we want it!* Theorem 3.9 shows that the loss of transitivity does not result in the loss of any classically valid inferences; it simply prevents transparent T from bringing new unwanted inferences with it. After all, if $\bot$ were a consequence of $\top$, that would be a Bad Thing—every inference would be valid. This Bad Thing is exactly what happens to classical logic when transparent truth is added to the ordinary models, as described in §3.1.

The most natural way to understand this nontransitivity, I think, is in terms of a difference between *strict assertion* and *tolerant assertion*—two types of assertion related so that everything strictly asserted is tolerantly asserted but not vice versa. (One can see these two types of assertion reflected in the notions of strict and tolerant satisfaction.) An argument is $\vDash^{st}_{+}$-valid when it allows for safe passage from strictly asserted premises to tolerantly asserted conclusions. This explains why it is not always safe to chain two valid arguments together. Let's look at the case above. If $\top$ is strictly asserted (as it always is), we may pass to tolerant assertion of $\lambda \wedge \neg\lambda$, but only with strict assertion of $\lambda \wedge \neg\lambda$ could we pass to tolerant assertion of $\bot$; this strict assertion is not guaranteed. In fact, it is forbidden, since $\bot$ can never be asserted, even tolerantly. (For more in-depth discussion of this issue, see Ripley, 2011a; for more on how to understand the corresponding issue in the case of vagueness, see Cobreros *et al.*, 2011.)

Given the distinction between strict and tolerant assertion, one might wonder about more than just the constraints that strict assertions place on tolerant assertions ($\vDash^{st}_{+}$). One might wonder as well about the constraints that strict assertions place on other strict assertions ($\vDash^{ss}_{+}$), the constraints that tolerant assertions place on other tolerant assertions ($\vDash^{tt}_{+}$), or the constraints that tolerant assertions place on strict assertions ($\vDash^{ts}_{+}$). This provides a way to make philosophical sense of all of the consequence relations in play. Given this picture, it is $\vDash^{ss}_{+}$ and $\vDash^{tt}_{+}$ that best fit the familiar picture of logical consequence as *preserving* some favored status. It is this preservation-based picture that feeds a cumulative use of logical reasoning, where past conclusions are taken as future premises. $\models^{st}_{+}$, despite its classicality, does not fit this preservation-based picture, due to its nontransitivity, and so cannot be used for cumulative reasoning in this way.[7]

Despite its failure to hold in full generality, there is a wide variety of cases where transitivity *does* hold. First, and most directly, Theorem 3.9 gives us a large class of safe instances of transitivity. We can expand our grip on when transitivity holds even further by using our other notions of consequence: an inference that preserves tolerant satisfaction can be added to the *end* of an $\models^{st}_{+}$-valid argument to produce another $\models^{st}_{+}$-valid argument, and an inference that preserves strict satisfaction can be added to the *beginning*. Since $\models^{ss}_{+}$-validity and $\models^{tt}_{+}$-validity are special cases of $\models^{st}_{+}$-validity, these too get us more transitivity.

FACT 3.16 (Limited Transitivity).

- *If* $\Gamma \models^{st} A, \Delta$ *and* $\Gamma', A \models^{st} \Delta'$ *(where* $\Gamma \cup \Gamma' \cup \Delta \cup \Delta' \subseteq \mathcal{L}$*), then* $\Gamma^*, \Gamma'^* \models^{st}_{+} \Delta^*, \Delta'^*$*, for any uniform substitution* * *on* $\mathcal{L}^+$.
- *If* $\Gamma \models^{st}_{+} A, \Delta$ *and* $\Gamma', A \models^{tt}_{+} \Delta'$*, then* $\Gamma, \Gamma' \models^{st}_{+} \Delta, \Delta'$.
- *If* $\Gamma, A \models^{st}_{+} \Delta$ *and* $\Gamma' \models^{ss}_{+} A, \Delta'$*, then* $\Gamma, \Gamma' \models^{st}_{+} \Delta, \Delta'$.

Proof. The first claim follows from Fact 3.8 and the fact that transitivity holds for classical logic. The latter two claims follow from the definition of $\models^{mn}_{+}$. □

In fact, we can exactly characterize the cases where transitivity holds and where it fails:

FACT 3.17 (The Extent of Transitivity). $\Gamma \models^{st}_{+} A, \Delta$ *and* $\Gamma', A \models^{st}_{+} \Delta'$ *but* $\Gamma, \Gamma' \not\models^{st}_{+} \Delta, \Delta'$ *iff the following conditions are met:*

- *There is at least one model* M *such that* $M \Vdash^{s} \gamma$ *for every* $\gamma \in \Gamma$ *and* $M \nVdash^{t} \delta$ *for any* $\delta \in \Delta$*; call the class of all such models* $\mathfrak{M}$*,*
- *There is at least one model* M *such that* $M \Vdash^{s} \gamma$ *for every* $\gamma \in \Gamma'$ *and* $M \nVdash^{t} \delta$ *for any* $\delta \in \Delta'$*; call the class of all such models* $\mathfrak{M}'$*,*
- *A takes value 0 nowhere in* $\mathfrak{M}$*,*
- *A takes value 1 nowhere in* $\mathfrak{M}'$*, and*
- $\mathfrak{M} \cap \mathfrak{M}'$ *is nonempty.*

Proof. The definition of $\models^{st}_{+}$ suffices. □

Fact 3.17 is relevant for the treatment of so-called ‘contingent’ paradoxes. Note that a failure of transitivity occurs iff A takes value $\frac{1}{2}$ in all models in $\mathfrak{M} \cap \mathfrak{M}'$: this is the model-theoretic sign of A’s being paradoxical, relative to the side premises and conclusions involved. This relative paradoxicality is just what contingent paradoxes exhibit.

The simplest case is the case in which $\Gamma = \Gamma'$ and $\Delta = \Delta'$; then $\mathfrak{M} = \mathfrak{M}'$. We have a failure of transitivity for A in this context iff assigning 1 to everything in Γ and 0 to everything in Δ forces assigning $\frac{1}{2}$ to A. This can happen (to use an anonymous referee’s (slightly tweaked) example), when A is $\neg Tn$, n is an ordinary name, Γ is $\{n = \langle \neg Tn \rangle\}$, and Δ is $\{\bot\}$. A behaves like a liar paradox here, but only contingently; it is only in light of the premise in Γ that it takes on its paradoxical flavor. The conditions for Fact 3.17 are met: there are plenty of models M such that $M \Vdash^{s} n = \langle \neg Tn \rangle$ and $M \nVdash^{t} \bot$, and in all of them A takes value $\frac{1}{2}$.

[7] I assume here that preservation is itself transitive. Thanks to an anonymous referee for fruitful remarks leading to this paragraph.

The conditions given in Fact 3.17 for a failure of transitivity are reminiscent of Kripke (1975)'s definition of the notion of *paradoxicality* relative to a model of $\mathcal{L}$. For Kripke, A is paradoxical relative to a model M iff A takes value $\frac{1}{2}$ in all fixed points over M. In the present setting, the focus is on paradoxicality relative to sets of side-premises and -conclusions, so the approach turns on classes of models rather than individual models, but the idea is quite similar. As Kripke is at pains to point out, it can be arbitrarily difficult to figure out whether a particular formula is or isn't paradoxical; as a result, Fact 3.17 can be arbitrarily difficult to apply in any particular case. In tricky situations, Fact 3.16 can thus still be of considerable use; the conditions it specifies for transitivity to hold are much easier to recognize, although they do not exhaust the safe applications of transitivity.

§4. Two proof theories. In this section, I present and consider two distinct proof theories for these models. Both proof theories are based on three-sided sequents, and each can be shown to be sound and complete for all four notions of consequence ($\vDash^{st}_{+}$, $\vDash^{tt}_{+}$, $\vDash^{ss}_{+}$, and $\vDash^{ts}_{+}$). The difference between the two different sequent systems corresponds to a difference in ways to interpret two-sided sequent systems, which I briefly outline in §4.1. Once we consider sequents with more than two sides, however, the difference becomes more than a difference in interpretation; it affects the sequent calculi themselves. In §§4.2 and 4.3, I present the sequent systems. The completeness proofs are in §5.

4.1. Interpreting two-sided sequents. For much of this subsection, I'll consider standard (two-valued) classical models and standard (2-sided) sequents for them. The details of a particular sequent system will be irrelevant here; I assume only that it is sound and complete. I'll use $\rhd$ as a sequent separator.

When a sequent $\Gamma \rhd \Delta$ is provable, that can be read as telling us two things about the space of models. On the one hand, it might tell us that there is *no* model $\mathcal{M} = \langle D, I\rangle$ such that $I(\gamma) = 1$ for *every* $\gamma \in \Gamma$ *and* $I(\delta) = 0$ for *every* $\delta \in \Delta$. Call this the *negated-conjunctive reading*: it tells us that no model meets a particular conjunctive (and universal) condition. (This reading of two-sided sequents is intimately connected to that in Restall, 2005.) Think about Identity sequents on this reading: that $A \rhd A$ is provable guarantees that there is no model on which A receives both the value 1 and the value 0. Identity, on the negated-conjunctive reading, is an *exclusivity* condition on values. On the other hand, think about the rule of Cut (here in its additive form):

$$\frac{\Gamma, A \rhd \Delta \qquad \Gamma \rhd A, \Delta}{\Gamma \rhd \Delta}$$

This tells us that if there is no model that assigns all of Γ along with A to 1 and all of Δ to 0, and there is no model that assigns all of Γ to 1 and A along with all of Δ to 0, then there is no model that assigns all of Γ to 1 and all of Δ to 0. In other words: if, given a background assignment of 1 to all of Γ and 0 to all of Δ, there is no model on which A receives value 1, and there is no model on which A receives value 0, then it's not A's fault—the trouble is with the background assumptions. In other words yet again: so long as we have a model, A has to be able either to take value 1 or to take value 0. Cut, on the negated-conjunctive reading, is an *exhaustivity* condition on values.

On the other hand, we might read $\Gamma \rhd \Delta$ as telling us something different: that *every* model $\mathcal{M} = \langle D, I\rangle$ is such that either $I(\gamma) = 0$ for *some* $\gamma \in \Gamma$, *or* $I(\delta) = 1$ for *some* $\delta \in \Delta$. Call this the *disjunctive reading*: it tells us that every model meets a particular disjunctive (and existential) condition. (Note that, on the negated-conjunctive reading of $\Gamma \rhd \Delta$, Γ's place is associated with value 1 and Δ's with 0, whereas on the disjunctive reading, Γ's place is associated with value 0 and Δ's with 1.)

On the disjunctive reading, Identity and Cut swap roles: it is now Identity that expresses exhaustivity of values, and Cut that expresses exclusion. Consider a sequent of the form $A \rhd A$, and read it disjunctively. It tells us that every model either assigns A value 0 or assigns A value 1: an exhaustion constraint. And return to Cut. Now it tells us that if every model either assigns 0 to something in $\Gamma \cup \{A\}$ or assigns 1 to something in Δ, and every model either assigns 0 to something in Γ or assigns 1 to something in $\{A\} \cup \Delta$, then every model either assigns 0 to something in Γ or assigns 1 to something in Δ. In other words, if, given that nothing in Γ gets value 0 and nothing in Δ gets value 1, A must receive value 1, and, given those same conditions, A must receive value 0, then the conditions are already unmeetable on their own: A cannot be forced to receive both values. This is an exclusion constraint.

Exclusion and exhaustion are features that extend to the three-valued models of interest in this paper. No formula can receive more than one of the three values on any model, and every formula must receive at least one of the three values on every model. But the symmetry between exclusion and exhaustion that holds in the two-valued case is, to some extent, broken in our three-valued models.

Exclusion constraints are about values taken *pairwise*. In an ST-model, no formula can receive values 1 and $\frac{1}{2}$, and no formula can receive values $\frac{1}{2}$ and 0, and no formula can receive values 1 and 0: there are three pairwise exclusion constraints. On the other hand, exhaustion constraints are about values taken *all together*. In an ST-model, every formula must receive either value 1 or value $\frac{1}{2}$ or value 0: there is a single three-way exhaustion constraint.

In two-valued models, there is no difference between considering the values pairwise and considering them all together: there is only a pair of values. So Identity and Cut look the same, even as we swap between interpretations. But that will not be the case for our sequent systems here. One of them—the negated-conjunctive one—will feature three Identity axioms, expressing the three pairwise exclusion constraints, and a single Cut rule, expressing the one three-way exhaustion constraint. The other system—the disjunctive one—will feature a single Identity axiom, expressing the one three-way exhaustion constraint, and three Cut rules, expressing the three pairwise exclusion constraints. What is just a difference in interpretation in the two-sided case becomes more substantive in the three-sided case.[8]

4.2. Disjunctive sequents. Our disjunctive sequents are of the form $\wr\ \Xi \parallel \Theta \parallel \Sigma\ \int$, where Ξ, Θ, Σ are sets of formulas.

DEFINITION 4.1. *A disjunctive sequent* $\wr\ \Xi \parallel \Theta \parallel \Sigma\ \int$ *is satisfied by a model* $M = \langle D, I\rangle$ *iff either* $I(\xi) = 0$ *for some* $\xi \in \Xi$ *or* $I(\theta) = \frac{1}{2}$ *for some* $\theta \in \Theta$ *or* $I(\sigma) = 1$ *for some* $\sigma \in \Sigma$. *A sequent is valid iff it is satisfied by every model. A model is a counterexample to a sequent iff it does not satisfy the sequent.*

Note that each place in the sequent—left, middle, right—has an associated value—0, $\frac{1}{2}$, 1. This is an extension to three-valued models of the familiar two-valued/two-sided strategy.

[8] To see how this difference plays out in the n-sided case, albeit without T or $=$, see Baaz *et al.* (1992).

The proof system consists of a number of axioms and a number of rules. A sequent is provable iff it follows from the axioms by some number (possibly 0) of rule-applications. Later, in §5, we will see that a disjunctive sequent is provable iff it is valid.

4.2.1. Axioms.

- For atomic A (including atomic T-predications): $\wr\ \Xi, A \mid\mid \Theta, A \mid\mid \Sigma, A\ \int$ is an axiom.
- For s, t any terms, A, B any *distinct* formulas, all of $\wr\ \Xi, s = t \mid\mid \Theta \mid\mid \Sigma, s = t\ \int$, $\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, t = t\ \int$, and $\wr\ \Xi, \langle A\rangle = \langle B\rangle \mid\mid \Theta \mid\mid \Sigma\ \int$ are axioms.

4.2.2. Rules. First, three rules of Cut. Remember, on the disjunctive reading, Cut is related to exclusion between values, and exclusion is three pairwise constraints when three values are in play. (Since our sequents work with sets, the effects of Contraction and Exchange are built in; and Weakening is built in to the axioms, so there are no other structural rules needed.)

- $$\frac{\wr\ \Xi, A \mid\mid \Theta \mid\mid \Sigma\ \int \qquad \wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int}$$
- $$\frac{\wr\ \Xi, A \mid\mid \Theta \mid\mid \Sigma\ \int \qquad \wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int}$$
- $$\frac{\wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma\ \int \qquad \wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int}$$

There is also a single three-premise rule, derived from these, that will be convenient to have around for the completeness proof in §5. The rule, which I'll call Derived-Cut, is:

- $$\frac{\wr\ \Xi, A \mid\mid \Theta, A \mid\mid \Sigma\ \int \qquad \wr\ \Xi, A \mid\mid \Theta \mid\mid \Sigma, A\ \int \qquad \wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma, A\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int}$$

and it can be derived as in Figure 1.

Next, we need operational rules. For each connective, there is a rule to introduce it in each place. In the rules for $\forall$, t is any term, and a is an eigenvariable (a variable not occurring in the rule's conclusion sequent).

- $$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A\ \int}{\wr\ \Xi, \neg A \mid\mid \Theta \mid\mid \Sigma\ \int}$$
  $$\frac{\wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma\ \int}{\wr\ \Xi \mid\mid \Theta, \neg A \mid\mid \Sigma\ \int}$$
  $$\frac{\wr\ \Xi, A \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, \neg A\ \int}$$
- $$\frac{\wr\ \Xi, A, B \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi, A \wedge B \mid\mid \Theta \mid\mid \Sigma\ \int}$$
  $$\frac{\wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma, A\ \int \qquad \wr\ \Xi' \mid\mid \Theta', B \mid\mid \Sigma', B\ \int \qquad \wr\ \Xi'' \mid\mid \Theta'', A, B \mid\mid \Sigma''\ \int}{\wr\ \Xi, \Xi', \Xi'' \mid\mid \Theta, \Theta', \Theta'', A \wedge B \mid\mid \Sigma, \Sigma', \Sigma''\ \int}$$

$$
\dfrac{\dfrac{\dfrac{\wr\; \Xi, A \mid\mid \Theta, A \mid\mid \Sigma \;\int \qquad \wr\; \Xi, A \mid\mid \Theta \mid\mid \Sigma, A \;\int}{\wr\; \Xi, A \mid\mid \Theta \mid\mid \Sigma \;\int} \qquad \wr\; \Xi \mid\mid \Theta, A \mid\mid \Sigma, A \;\int}{\wr\; \Xi \mid\mid \Theta, A \mid\mid \Sigma \;\int} \qquad \dfrac{\wr\; \Xi, A \mid\mid \Theta, A \mid\mid \Sigma \;\int \qquad \wr\; \Xi, A \mid\mid \Theta \mid\mid \Sigma, A \;\int}{\wr\; \Xi, A \mid\mid \Theta \mid\mid \Sigma \;\int}}{\wr\; \Xi \mid\mid \Theta \mid\mid \Sigma \;\int}
$$

Fig. 1. Deriving Derived-Cut.

$$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A\ \int \qquad \wr\ \Xi' \mid\mid \Theta' \mid\mid \Sigma', B\ \int}{\wr\ \Xi, \Xi' \mid\mid \Theta, \Theta' \mid\mid \Sigma, \Sigma', A \wedge B\ \int}$$

- $$\frac{\wr\ \Xi, A(t) \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi, \forall x A(x) \mid\mid \Theta \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta, A(a) \mid\mid \Sigma, A(a)\ \int \qquad \wr\ \Xi' \mid\mid \Theta', A(t) \mid\mid \Sigma'\ \int}{\wr\ \Xi, \Xi' \mid\mid \Theta, \Theta', \forall x A(x) \mid\mid \Sigma, \Sigma'\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A(a)\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, \forall x A(x)\ \int}$$

Dual rules can be constructed for $\vee$ and $\exists$, and a related rule for $\supset$; I omit them here, officially considering these to be defined connectives. Finally, we move on to rules for $=$ and T:

- $$\frac{\wr\ \Xi, A(t) \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi, A(s), s = t \mid\mid \Theta \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta, A(t) \mid\mid \Sigma\ \int}{\wr\ \Xi, s = t \mid\mid \Theta, A(s) \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A(t)\ \int}{\wr\ \Xi, s = t \mid\mid \Theta \mid\mid \Sigma, A(s)\ \int}$$

- $$\frac{\wr\ \Xi, A(t) \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi, A(s), t = s \mid\mid \Theta \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta, A(t) \mid\mid \Sigma\ \int}{\wr\ \Xi, t = s \mid\mid \Theta, A(s) \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A(t)\ \int}{\wr\ \Xi, t = s \mid\mid \Theta \mid\mid \Sigma, A(s)\ \int}$$

- $$\frac{\wr\ \Xi, A \mid\mid \Theta \mid\mid \Sigma\ \int}{\wr\ \Xi, T\langle A\rangle \mid\mid \Theta \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta, A \mid\mid \Sigma\ \int}{\wr\ \Xi \mid\mid \Theta, T\langle A\rangle \mid\mid \Sigma\ \int}$$

$$\frac{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, A\ \int}{\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma, T\langle A\rangle\ \int}$$

It's worth noting that the rules for T are not purely syntactic, as the connection between A and $\langle A\rangle$ is itself not purely syntactic, but instead depends on the function τ. Given a different way of handling self-reference (recall the discussion in §2), this can come out differently.

FACT 4.2. *Cut is not admissible in the above proof system.*

Proof. The sequent $\wr \; || \; \lambda \; || \; \int$ is provable with Cut:

$$\dfrac{\dfrac{\dfrac{\dfrac{\wr\, T\langle\lambda\rangle \;||\; T\langle\lambda\rangle \;||\; T\langle\lambda\rangle \,\int}{\wr\, \lambda \;||\; \lambda \;||\; \lambda \,\int}}{\wr\, T\langle\lambda\rangle \;||\; \lambda \;||\; \lambda \,\int}}{\wr \;||\; \lambda \;||\; \lambda \,\int} \qquad \dfrac{\dfrac{\dfrac{\wr\, T\langle\lambda\rangle \;||\; T\langle\lambda\rangle \;||\; T\langle\lambda\rangle \,\int}{\wr\, \lambda \;||\; \lambda \;||\; \lambda \,\int}}{\wr\, \lambda \;||\; \lambda \;||\; T\langle\lambda\rangle \,\int}}{\wr\, \lambda \;||\; \lambda \;||\; \int}}{\wr \;||\; \lambda \;||\; \int}$$

Recall that $\lambda = \neg T\langle\lambda\rangle$; then the first step in each subproof above is three applications of the $\neg$ rules, one in each place; the middle step in each subproof is an application of a T-rule; and the last step in each subproof is an application of a single $\neg$ rule, collapsing the two occurrences of λ that end up in the same position. The subproofs are then combined by an application of Cut.

The sequent $\wr \; || \; \lambda \; || \; \int$ is not provable without Cut, however. Call a sequent $\wr\, \Xi \;||\; \Theta \;||\; \Sigma \,\int$ *middle-only* iff $\Xi = \Sigma = \emptyset$. All of the axioms have the property of being *not* middle-only, and all of the rules but Cut preserve this property. Since $\wr \; || \; \lambda \; || \; \int$ is middle-only, it is not provable without Cut.[9] □

Despite the nonadmissibility of Cut, however, the proof system is sound and complete for ST-models for $\mathcal{L}^+$. This claim will be proved in §5. First, though, let's look at this sequent system's negated-conjunctive twin.

4.3. Negated-conjunctive sequents. Here, our sequents are of the form $\int \Xi \;||\; \Theta \;||\; \Sigma \wr$, where Ξ, Θ, Σ are sets of formulas. We use a different notion of satisfaction, as befits the negated-conjunctive reading:

DEFINITION 4.3. *A negated-conjunctive sequent* $\int \Xi \;||\; \Theta \;||\; \Sigma \wr$ *is satisfied by a model* $M = \langle D, I\rangle$ *iff it is not the case that:* $I(\xi) = 0$ *for every* $\xi \in \Xi$ *and* $I(\theta) = \frac{1}{2}$ *for every* $\theta \in \Theta$ *and* $I(\sigma) = 1$ *for every* $\sigma \in \Sigma$. *A sequent is valid iff it is satisfied by every model.*

Note that each place in the sequent—left, middle, right—has an associated value—0, $\frac{1}{2}$, 1. (Unlike the classical case, where the value associated with a place changes depending on the reading, here both disjunctive and negated-conjunctive sequents use the same associations between places and values, for ease of reading.)

Again, the proof system consists of a number of axioms and a number of rules; a sequent is provable iff it follows from the axioms via the rules; and the system is sound and complete, although proof of this will await §5.

4.3.1. Axioms.

- For atomic A (including atomic T-predications): all of $\int \Xi, A \;||\; \Theta, A \;||\; \Sigma \wr$, $\int \Xi, A \;||\; \Theta \;||\; \Sigma, A \wr$, and $\int \Xi \;||\; \Theta, A \;||\; \Sigma, A \wr$ are axioms.
- For s, t any terms, A, B any *distinct* formulas, all of $\int \Xi, \;||\; \Theta, s = t \;||\; \Sigma \wr$, $\int \Xi, t = t \;||\; \Theta \;||\; \Sigma \wr$, and $\int \Xi \;||\; \Theta \;||\; \Sigma, \langle A\rangle = \langle B\rangle \wr$ are axioms.

4.3.2. Rules. First, structural rules. This time, there is a single rule of Cut. Remember, on the negated-conjunctive reading, Cut is related to exhaustion amongst the values: a single constraint no matter how many values are in play. (Since our sequents still work

[9] Many thanks to Greg Restall for this efficient demonstration that $\wr \; || \; \lambda \; || \; \int$ isn't provable without Cut.

with sets, the effects of Contraction and Exchange are again built in, and Weakening is again taken care of by the axioms.)

• $$\frac{\int \Xi, A \mid\mid \Theta \mid\mid \Sigma \wr \qquad \int \Xi \mid\mid \Theta, A \mid\mid \Sigma \wr \qquad \int \Xi \mid\mid \Theta \mid\mid \Sigma, A \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma \wr}$$

On to operational rules. Again, there is one rule to introduce each connective in each place; again, for the $\forall$ rules, t is any term and a an eigenvariable.

• $$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A \wr}{\int \Xi, \neg A \mid\mid \Theta \mid\mid \Sigma \wr}$$

$$\frac{\int \Xi \mid\mid \Theta, A \mid\mid \Sigma \wr}{\int \Xi \mid\mid \Theta, \neg A \mid\mid \Sigma \wr}$$

$$\frac{\int \Xi, A \mid\mid \Theta \mid\mid \Sigma \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, \neg A \wr}$$

• $$\frac{\int \Xi, A \mid\mid \Theta \mid\mid \Sigma \wr \qquad \int \Xi', B \mid\mid \Theta' \mid\mid \Sigma' \wr}{\int \Xi, \Xi', A \wedge B \mid\mid \Theta, \Theta' \mid\mid \Sigma, \Sigma' \wr}$$

$$\frac{\int \Xi \mid\mid \Theta, A \mid\mid \Sigma, B \wr \qquad \int \Xi' \mid\mid \Theta', B \mid\mid \Sigma', A \wr \qquad \int \Xi'' \mid\mid \Theta'', A, B \mid\mid \Sigma'' \wr}{\int \Xi, \Xi', \Xi'' \mid\mid \Theta, \Theta', \Theta'', A \wedge B \mid\mid \Sigma, \Sigma', \Sigma'' \wr}$$

$$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A, B \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A \wedge B \wr}$$

• $$\frac{\int \Xi, A(a) \mid\mid \Theta \mid\mid \Sigma \wr}{\int \Xi, \forall x A(x) \mid\mid \Theta \mid\mid \Sigma \wr}$$

• $$\frac{\int \Xi \mid\mid \Theta, A(a), A(t) \mid\mid \Sigma \wr \qquad \int \Xi \mid\mid \Theta, A(a) \mid\mid \Sigma, A(t) \wr}{\int \Xi \mid\mid \Theta, \forall x A(x) \mid\mid \Sigma \wr}$$

$$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A(t) \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, \forall x A(x) \wr}$$

We again take $\vee$, $\exists$, and $\supset$ to be defined connectives. Finally, we move on to rules for $=$ and T:

• $$\frac{\int \Xi, A(t) \mid\mid \Theta \mid\mid \Sigma \wr}{\int \Xi, A(s) \mid\mid \Theta \mid\mid \Sigma, s = t \wr}$$

$$\frac{\int \Xi \mid\mid \Theta, A(t) \mid\mid \Sigma \wr}{\int \Xi \mid\mid \Theta, A(s) \mid\mid \Sigma, s = t \wr}$$

$$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A(t) \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A(s), s = t \wr}$$

• $$\frac{\int \Xi, A(t) \mid\mid \Theta \mid\mid \Sigma \wr}{\int \Xi, A(s) \mid\mid \Theta \mid\mid \Sigma, t = s \wr}$$

$$\frac{\int \Xi \mid\mid \Theta, A(t) \mid\mid \Sigma \wr}{\int \Xi \mid\mid \Theta, A(s) \mid\mid \Sigma, t = s \wr}$$

$$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A(t) \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A(s), t = s \wr}$$

- $$\frac{\int \Xi, A \mid\mid \Theta \mid\mid \Sigma \wr}{\int \Xi, T\langle A \rangle \mid\mid \Theta \mid\mid \Sigma \wr}$$

  $$\frac{\int \Xi \mid\mid \Theta, A \mid\mid \Sigma \wr}{\int \Xi \mid\mid \Theta, T\langle A \rangle \mid\mid \Sigma \wr}$$

  $$\frac{\int \Xi \mid\mid \Theta \mid\mid \Sigma, A \wr}{\int \Xi \mid\mid \Theta \mid\mid \Sigma, T\langle A \rangle \wr}$$

Again, note that the T rules depend on τ, and so are not purely syntactic.

FACT 4.4. *Cut is admissible in this proof system.*

This Fact will be proven in §5, via completeness.

It might at first seem that the negated-conjunctive proof system is superior to the disjunctive proof system, since it features admissible Cut. However, there are two things worth noticing in this connection. The first is that neither proof system enjoys a subformula property, admissible Cut or no. This is because the rules for T (which are the same in both systems) allow us to introduce atomic formulas from formulas of any complexity. So one of the usual reasons to prefer systems with admissible Cut does not apply here. The second is that the disjunctive proof system enjoys a simpler connection to the two-sided consequence relations that are a driving concern, as we will see presently.

4.4. Two-sided and three-sided consequence. Both the disjunctive and the negated-conjunctive proof systems presented above are sound and complete (given their respective readings) for ST-models for $\mathcal{L}^+$. However, our initial concern was with two-sided consequence relations, like $\vDash^{tt}_{+}$, $\vDash^{ss}_{+}$, and, especially, $\vDash^{st}_{+}$. The connections are easiest to see for disjunctive sequents:

FACT 4.5. *Connections between disjunctive sequents and two-sided consequences:*

- $\Gamma \vDash^{st}_{+} \Delta$ *iff* $\wr\ \Gamma \mid\mid \Gamma, \Delta \mid\mid \Delta \int$ *is valid.*
- $\Gamma \vDash^{ss}_{+} \Delta$ *iff* $\wr\ \Gamma \mid\mid \Gamma \mid\mid \Delta \int$ *is valid.*
- $\Gamma \vDash^{tt}_{+} \Delta$ *iff* $\wr\ \Gamma \mid\mid \Delta \mid\mid \Delta \int$ *is valid.*
- $\Gamma \vDash^{ts}_{+} \Delta$ *iff* $\wr\ \Gamma \mid\mid\ \mid\mid \Delta \int$ *is valid.*

Proof. These follow directly from the definitions of $\vDash^{mn}_{+}$ and validity for disjunctive sequents. □

The connections to negated-conjunctive sequents are a bit hairier, for the most part:

FACT 4.6. *Connections between negated-conjunctive sequents and two-sided consequences:*

- $\Gamma \vDash^{st}_{+} \Delta$ *iff* $\int \Delta \mid\mid\ \mid\mid \Gamma \wr$ *is valid.*
- $\Gamma \vDash^{ss}_{+} \Delta$ *iff all sequents of the form* $\int \Delta_1 \mid\mid \Delta_2 \mid\mid \Gamma \wr$ *are valid, where* $\Delta_1 \cup \Delta_2 = \Delta$.

- $\Gamma \models^{tt}_{+} \Delta$ *iff all sequents of the form* $\int \Delta \,||\, \Gamma_1 \,||\, \Gamma_2 \wr$ *are valid, where* $\Gamma_1 \cup \Gamma_2 = \Gamma$.
- $\Gamma \models^{ts}_{+} \Delta$ *iff all sequents of the form* $\int \Delta_1 \,||\, \Delta_2, \Gamma_1 \,||\, \Gamma_2 \wr$ *are valid, where* $\Gamma_1 \cup \Gamma_2 = \Gamma$ *and* $\Delta_1 \cup \Delta_2 = \Delta$.

Proof. These follow directly from the definitions of $\models^{mn}_{+}$ and validity for negated-conjunctive sequents. □

Thus, to prove, say, that $p, p \supset q \models^{ss}_{+} q, r$ via negated-conjunctive sequents, one must prove the following four sequents (other sequents of the relevant form are all weakenings of these four, and so are provable if these four are):

- $\int q, r \,||\, \,||\, p, p \supset q \wr$
- $\int q \,||\, r \,||\, p, p \supset q \wr$
- $\int r \,||\, q \,||\, p, p \supset q \wr$
- $\int \,||\, q, r \,||\, p, p \supset q \wr$

Of course, since $p, p \supset q \models^{ss}_{+} q, r$, all of these are indeed valid, and since our proof system is complete, it is possible to prove all four.[10] But it is inconvenient to have to do so; this is an advantage of the disjunctive system, in which it suffices to prove the single sequent $\wr\; p, p \supset q \,||\, p, p \supset q \,||\, q, r \int$. The inconvenience is only multiplied as we look to two-sided arguments with more premises and more conclusions.

This difference does not directly affect the connection to $\models^{st}_{+}$, our primary concern here. However, as we've seen above, keeping track of which instances of transitivity are safe in $\models^{st}_{+}$ requires attention to the weaker logics based on $\models^{ss}_{+}$ and $\models^{tt}_{+}$ as well. In addition, since $\models^{ss}_{+}$ and $\models^{tt}_{+}$, respectively, are the familiar logics K3 and LP with transparent truth, they are of independent interest.

In two-sided sequent systems, transitivity is carried by the rule of Cut, so it's tempting to look for a connection between three-sided Cut and transitivity of the two-sided systems that result; but as far as I can see no such connection plays any interesting role here. In an "ordinary" (two-sided) sequent calculus for $\models^{st}_{+}$, Cut fails, since it amounts to transitivity.[11] Similarly, Cut in the three-sided systems is something like "transitivity", but generalized to hold of three-place relations rather than two-place ones. However, Cut in the three-sided systems does not guarantee Cut in the two-sided systems.

The action is in the way two-place consequence relates to three-place consequence. Consider the case of the disjunctive three-sided system. Because the places where Γ occurs overlap the places where Δ occurs in the relation for $\models^{st}_{+}$, $\models^{st}_{+}$ can be nontransitive.

[10] For example, here is a proof of the first (recall that $p \supset q$ is officially an abbreviation for $\neg(p \wedge \neg q)$):

$$\dfrac{\dfrac{\int p \,||\, \,||\, p \wr \qquad \dfrac{\int q, r \,||\, \,||\, q \wr}{\int q, r, \neg q \,||\, \,||\, \wr}}{\int p \wedge \neg q, q, r \,||\, \,||\, p \wr}}{\int q, r \,||\, \,||\, p, \neg(p \wedge \neg q) \wr}$$

Variations on this proof will produce proofs of the other three sequents.

[11] For presentation and philosophical discussion of such a sequent calculus, see Ripley (2011a).

(Similarly, it's because there is a place where neither Γ nor Δ occurs in the reduction for $\vDash^{ts}_{+}$ that $\vDash^{ts}_{+}$ can be nonreflexive.) In some ways, three-sided consequence is the "real" target system, revealing the full structure of the constraints on strict and tolerant assertion, and the two-sided consequences are just slices that reveal various aspects of this structure. So the whole thing can be "transitive" (in its way) even when some part of it is not transitive (in the usual way). This all holds whether or not Cut is admissible in the three-sided system; Cut's admissibility seems not to matter.

§5. Completeness. In this section, the above proof systems are shown to be sound and complete for ST-models for $\mathcal{L}^+$. That is, it is shown that a sequent is provable iff it is valid. (Remember that the definitions of satisfaction, and so validity, differ from disjunctive to negated-conjunctive sequents.) In §5.1, this is shown for the proof system involving disjunctive sequents, and in §5.2, this is shown for the proof system involving negated-conjunctive sequents.

5.1. Disjunctive sequents. In this section, discussion is restricted to disjunctive sequents. Thus, *sequent* should be taken to mean sequent of the form $\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int$, and all talk of axioms, rules, validity, provability, etc, should be taken relative to the proof system for these sequents presented in §4.2.

LEMMA 5.1 (Soundness). *If a sequent* $\wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int$ *is provable, then it is valid.*

Proof. The axioms are valid, and validity is preserved by the rules, as can be verified without too much trouble. □

The real game, as usual, is in proving completeness. Following, for example, Baaz *et al.* (1993) (for *n*-sided sequents) and Takeuti (1987) (for two-sided sequents), I'll use the method of *reduction trees*. This method works by defining a construction that yields, for any given sequent, either a proof of that sequent or a model that does not satisfy the sequent. Clearly, if there is such a construction, our proof system is complete.

Some useful vocabulary is provided by the notions of *subsequent* and *sequent union*, which are defined from their set relatives pointwise, as follows:

DEFINITION 5.2. *A sequent* $S = \wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int$ *is a subsequent of a sequent* $S' = \wr\ \Xi' \mid\mid \Theta' \mid\mid \Sigma'\ \int$ *(written* $S \sqsubseteq S'$*) iff* $\Xi \subseteq \Xi'$ *and* $\Theta \subseteq \Theta'$ *and* $\Sigma \subseteq \Sigma'$.

A sequent $S = \wr\ \Xi \mid\mid \Theta \mid\mid \Sigma\ \int$ *is the sequent union of a set of sequents* $\{\wr\ \Xi_i \mid\mid \Theta_i \mid\mid \Sigma_i\ \int\}_{i \in I}$ *(written* $S = \bigsqcup_{i \in I}\{\wr\ \Xi_i \mid\mid \Theta_i \mid\mid \Sigma_i\ \int\}$*) iff* $\Xi = \bigcup_{i \in I}\{\Xi_i\}$ *and* $\Theta = \bigcup_{i \in I}\{\Theta_i\}$ *and* $\Sigma = \bigcup_{i \in I}\{\Sigma_i\}$.

The construction proceeds by building a tree in stages, starting from a root sequent $S_0 = \wr\ \Xi_0 \mid\mid \Theta_0 \mid\mid \Sigma_0\ \int$. At each stage, we apply all applicable operational rules—plus Cut rules, since they are noneliminable—as it were in *reverse*, from conclusion-sequent (which we have) to generate premise sequents, branching upwards for multipremise rules and extending unbranched for single-premise rules; we ensure that we only add formulas at each stage, so that every branch of the tree is ordered by the subsequent relation.

When any branch has as its topmost sequent an axiom, we call the branch *closed*. When a branch is closed, we don't bother applying any more rules to it; a branch that's not closed is *open*. Then we repeat until either every branch is closed or there is an infinite open branch. If every branch has closed, there is a proof of our root sequent: the tree itself. If there is an infinite open branch B, we can use it to construct a countermodel to the root sequent.

To begin, we require an enumeration of the terms and an enumeration of the formulas, and a sequent $S_0 = \wr\ \Xi_0 \mid\mid \Theta_0 \mid\mid \Sigma_0\ \int$ to build the tree from.

Stage 0: Let $S_0 = \wr\ \Xi_0 \mid\mid \Theta_0 \mid\mid \Sigma_0\ \int$ be the tree's root.
Stage $n + 1$**:** We proceed in substages and subsubstages:

1. For all branches in the tree after stage n: if its tip is an axiom, close the branch.
2. Now to the open branches. For each formula A in a sequent position in each open branch, if A already occurred in that sequent-position in that branch at stage n (in other words, if A has not been generated during stage $n + 1$), and if A has not already been reduced during stage $n + 1$, then reduce A as follows:
   - If A is a negation $\neg B$, then
     - if A is in the (left/middle/right) position, extend the branch by copying its current tip and adding B to the (right/middle/left) position.
   - If A is a conjunction $B \wedge C$, then
     - if A is in the left position, extend the branch by copying its current tip and adding both B and C to the left position.
     - if A is in the middle position, split the branch in three: extend the first by copying the current tip and adding B to both the middle and right positions; extend the second by copying the current tip and adding C to both the middle and right positions; and extend the third by copying the current tip and adding both B and C to the middle position.
     - if A is in the right position, split the branch in two: extend the first by copying the current tip and adding B to the right position; and extend the second by copying the current tip and adding C to the right position.
   - If A is a universal quantification $\forall x B(x)$, then
     - if A is in the left position, extend the branch by copying its current tip and adding $B(t)$ to the left position, where t is the first term in the enumeration not already used in a reduction of A in the left position before stage $n + 1$.
     - if A is in the middle position, split the branch in two: extend the first by copying the current tip and adding $B(a)$ to both the middle and right positions, where a is the first term in the enumeration not to occur anywhere in the current tip; extend the second by copying the current tip and adding $B(t)$ to the middle position, where t is the first term in the enumeration not already used in a reduction of A in the middle position before stage $n + 1$.
     - if A is in the right position, extend the branch by copying its current tip and adding $B(a)$ to the right position, where a is the first term in the enumeration not to occur anywhere in the current tip.
   - If A is an identity $s = t$, then
     - if A is in the left position, then extend the branch by copying its current tip and adding $B(t)$ to the (left/middle/right) position for

every formula $B(s)$ that occurs in the (left/middle/right) position of the current tip, as well as adding $B(s)$ to the (left/middle/right) position for every formula $B(t)$ that occurs in the (left/middle/right) position of the current tip.
  - ◦ if A is in the middle or right position, do nothing.

- If A is an atomic predication of truth to a distinguished name $T\langle B\rangle$, then
  - if A is in the (left/middle/right) position, then extend the branch by copying its current tip and adding B to the (left/middle/right) position.
- If A is something else (an atomic predication other than one of the form $T\langle B\rangle$), then do nothing.

3. Finally, we need to allow for Cut, since it is not admissible. Take the nth formula in the enumeration of all formulas—call it A—and extend each branch via the rule of Derived-Cut. That is, for each open branch, if its tip is $\wr\ \Xi \,||\, \Theta \,||\, \Sigma\ \int$, split it three ways and extend the three new branches with $\wr\ \Xi, A \,||\, \Theta, A \,||\, \Sigma\ \int$, $\wr\ \Xi, A \,||\, \Theta \,||\, \Sigma, A\ \int$, and $\wr\ \Xi \,||\, \Theta, A \,||\, \Sigma, A\ \int$, respectively.

Repeat this procedure until every branch closes, or, if that doesn't ever happen, until there is an infinite open branch. If every branch has closed, the tree is a proof of the root sequent S_0 (as can be verified by comparing the reduction steps to the sequent rules), so it has a proof. On the other hand, if there is an infinite open branch, we can use the branch to construct a countermodel to S_0. We do so as follows. First, collect all of B into a single sequent S_ω: let $S_\omega = \wr\ \Xi_\omega \,||\, \Theta_\omega \,||\, \Sigma_\omega\ \int = \bigsqcup\{S : S \text{ is a sequent in } B\}$. Note that, because the Cut step has been applied to every formula somewhere on B, every formula occurs in (at least) two places in S_ω. Note as well that, since the branch is ordered by the subsequent relation, any finite set of formulas appearing in S_ω must have appeared in a sequent at some point on B. Now, to specify a countermodel, we need a domain D and an interpretation function I.

To construct the domain, we start from the set T of terms that appear anywhere in S_ω. Define an equivalence relation $\equiv$ on T as follows: $\equiv\ = \{\langle s, t\rangle : s = t \in \Xi_\omega\}$. This is indeed an equivalence relation: it is reflexive, since all formulas of the form $t = t$ must appear in two of Ξ_ω, Θ_ω, Σ_ω, but they cannot appear in Σ_ω or the branch would have contained an axiom and so closed. Therefore, all formulas of the form $t = t$ must appear in Ξ_ω. $\equiv$ is symmetric, since if $s = t$ appears in Ξ_ω, it must have at some stage resulted in a substitution of s for t in $t = t$, resulting in $t = s$ appearing in Ξ_ω as well. And finally, $\equiv$ is transitive, since if both $s = t$ and $t = u$ appear in Ξ_ω, then they must have at some stage resulted in a substitution of s for t in $t = u$, resulting in $s = u$ appearing in Ξ_ω as well.

Let T' be the set of equivalence classes so created. We write $[s]$ for the equivalence class of a term s. T' is almost our domain, but we have distinguished names to address—they must name formulas, not equivalence classes of terms. Our final domain D is $\{a : a \in T' \text{ and } a \neq [\langle A\rangle] \text{ for any formula } A\} \cup \mathcal{L}^+$: T' without the equivalence classes of distinguished names, combined with the language itself.

Now, to specify I. For terms s, let $I(s) = [s]$, unless $s \equiv \langle A\rangle$ for some formula A, in which case let $I(s) = A$. This has the result that every identity formula appearing in Ξ_ω has value 1 (and so not 0), and that every identity formula appearing in Σ_ω has value 0 (and so not 1). Since no identity formula can take value $\frac{1}{2}$, every identity formula appearing in

Θ_ω has value distinct from $\frac{1}{2}$. Thus, no identity formula receives the value corresponding to its location in S_ω.

For n-ary predicates P (including the truth predicate T), let $I(P)$ be determined by S_ω as follows: $I(P)(\langle I(s_1), I(s_2), \ldots, I(s_n)\rangle) = 0/\frac{1}{2}/1$, respectively, iff $P(s_1, s_2, \ldots, s_n)$ does not appear in $\Xi_\omega/\Theta_\omega/\Sigma_\omega$.

We should pause to make sure this is well-defined. First, we know that $P(s_1, s_2, \ldots, s_n)$ will appear in at least two of Ξ_ω, Θ_ω, Σ_ω. We know as well that it cannot appear in all three, since then some sequent on branch B would be an axiom, and the branch would have closed. So $P(s_1, s_2, \ldots, s_n)$ appears in exactly two of Ξ_ω, Θ_ω, Σ_ω. Next, we need to make sure that the choice of terms does not matter: that if $I(s_i) = I(t_i)$ for $1 \leq i \leq n$, then $P(s_1, s_2, \ldots, s_n)$ appears in the same two positions in S_ω as $P(t_1, t_2, \ldots, t_n)$. We know that $I(s_i) = I(t_i)$ iff $[s_i] = [t_i]$ iff $s_i \equiv t_i$ iff the formula $s_i = t_i$ appears in Ξ_ω, for $1 \leq i \leq n$. This must have resulted in (one-at-a-time) substitutions connecting $P(s_1, s_2, \ldots, s_n)$ to $P(t_1, t_2, \ldots, t_n)$: if one appears in Ξ_ω, Θ_ω, or Σ_ω, the other must appear there too. Thus, $I(P)$ is well-defined. What's more, it ensures that no atomic formula receives the value corresponding to any position in which it appears in S_ω.

For negations, conjunctions, and universal quantifications, we can use the rules by which we've extended B to show that if none of their components receives the value associated with any place in which it appears in S_ω, neither will the compound. For example, if a negation $\neg A$ appears in Ξ_ω, then we know that $A \in \Sigma_\omega$; if $I(A) \neq 1$, it follows that $I(\neg A) \neq 0$. For another example, if $\forall x A(x)$ appears in Θ_ω, then either $A(a)$ appears in both Θ_ω and Σ_ω for some term a, or else $A(t)$ appears in Θ_ω for every term t. If $A(a)$ appears both in Θ_ω and Σ_ω, then $I(A(a)) = 0$, and so $I(\forall x A(x)) = 0$. On the other hand, if $A(t)$ appears in Θ_ω for every t, then there is no term t such that $I(A(t)) = \frac{1}{2}$; this means that $I(\forall x A(x)) \neq \frac{1}{2}$. Either way, we have $I(\forall x A(x)) \neq \frac{1}{2}$; it does not take the value associated with Θ_ω. By completing the induction along these lines, we can show that we have a model on which no formula receives the value associated with any place in which it appears in S_ω.

Finally, we need to make sure that $I(T\langle A\rangle) = I(A)$ for every formula A. But this is straightforward: we know that both $T\langle A\rangle$ and A must appear in exactly two places in S_ω. But wherever $T\langle A\rangle$ appears, there too must A appear, by construction. Since neither $T\langle A\rangle$ nor A receives the value associated with any place in which it appears, and since both appear in all places but one, they are constrained to receive the same value. Thus, our model obeys the constraint on T that enables it to deal with transparent truth.

Finally, since our original sequent $S_0 \sqsubseteq S_\omega$, no formula in S_0 receives the value associated with any place in which it appears; the model countermodels S_0.

THEOREM 5.3. *For any disjunctive sequent S, either the sequent has a proof or it has a countermodel.*

5.2. Negated-conjunctive sequents. The proof of completeness for disjunctive sequents relied quite crucially on the rules of Cut (via the derived rule of Derived-Cut). First and foremost, this is because the proof system for disjunctive sequents is not complete without Cut. But the rules of Cut also served to make the proof more convenient: by ensuring that, in the sequent S_ω generated from an infinite open branch, every formula would appear twice, thus uniquely determining the (one remaining) value it needed to be assigned by the countermodel.

Here, we prove completeness for negated-conjunctive sequents without Cut. This is interesting in its own right, but also allows for a model-theoretic proof of Cut-admissibility along lines laid out in Kremer (1988). The completeness proof works along mainly the same lines, with some added complications owing to the absence of Cut, and so this section will move quickly over the parts of the proof that are more or less the same, focusing on the complications introduced by the absence of Cut.

Soundness can be proved in the usual way:

LEMMA 5.4 (Soundness). *If a sequent* $\int \Xi \mid\mid \Theta \mid\mid \Sigma \wr$ *is provable (with Cut), then it is valid.*

Proof. Again, the axioms are valid and validity is preserved by the rules, including the rule of Cut. □

On to completeness. We will extend the notions of *subsequent* and *sequent union*, defined above for disjunctive sequents, to negated-conjunctive sequents in the obvious way.

Again, we build a tree starting from a particular root sequent S_0, applying the rules in reverse. This time, however, we make no provision for Cut. We will use this proof of Cut-free completeness to demonstrate the admissibility of Cut for negated-conjunctive sequents. Reduction rules can again be straightforwardly predicted from the sequent rules, and I won't go through them in detail here. Construct the tree until every branch has closed (resulted in an axiom) or there is an infinite open branch.

If the full tree has closed, it is again a proof of S_0. This, time, however, it is a Cut-free proof. If there is an infinite open branch, we use it to construct a countermodel $M = \langle D, I \rangle$ to S_0. For a negated-conjunctive sequent, recall, a countermodel is a model such that: $I(\xi) = 0$ for all $\xi \in \Xi$, $I(\theta) = \frac{1}{2}$ for all $\theta \in \Theta$, and $I(\sigma) = 1$ for all $\sigma \in \Sigma$. That is, a countermodel is one in which each formula in the sequent receives the value associated with its place in the sequent. First, take the open infinite branch B and collect it into a single sequent $S_\omega = \int \Xi_\omega \mid\mid \Theta_\omega \mid\mid \Sigma_\omega \wr = \bigsqcup\{S : S \text{ is a sequent in } B\}$.

We construct the domain D as before, except that now we use the identity formulas in Σ_ω to define the equivalence relation that gets it all rolling. The value of I for terms is as before. The value of I for predicates is precisely reversed: $I(P)(\langle I(s_1), I(s_2), \ldots, I(s_n)\rangle) = 0, \frac{1}{2}, 1$, respectively, iff $P(s_1, s_2, \ldots, s_n)$ *does* appear in $\Xi_\omega, \Theta_\omega, \Sigma_\omega$. We know that it cannot appear in more than one, since then B would have contained an axiom, and would not be open. (We know that this is well-defined by the same considerations about identity, mutatis mutandis, as last time.) Now, we have the beginnings of a model on which every atomic formula appearing in S_ω receives the value associated with the place in which it appears. However, because of the absence of Cut, there may well be atomic formulas that don't appear in S_ω. For any such atomic A, let $I(A) = \frac{1}{2}$.

Now, we can show (by induction) that every formula in S_ω receives on M the value associated with the place in which it appears. A few examples will give the flavor. If a negation $\neg A$ appears in Θ_ω, then A must also appear in Θ_ω. Since $I(A) = \frac{1}{2}$, by the induction hypothesis, $I(\neg A) = \frac{1}{2}$ too. Another example is provided when a conjunction $A \wedge B$ appears in Θ_ω. If this happens, then either $A \in \Theta_\omega$ and $B \in \Sigma_\omega$, or $B \in \Theta_\omega$ and $A \in \Sigma_\omega$, or $A, B \in \Theta_\omega$. In all three cases, the minimum of $I(A), I(B)$ is $\frac{1}{2}$, and so $I(A \wedge B) = \frac{1}{2}$. Thus, M, if a model, is a countermodel to S_ω.

To make sure that M is a model, we need to show one more thing: that $I(T\langle A\rangle) = I(A)$, for every formula A. This time, though, there is no guarantee of this. If $T\langle A\rangle$ appears on B, then we have this guarantee, as A will appear in the same place. But if $T\langle A\rangle$ does not appear on B, then it has value $\frac{1}{2}$ (by the rule for atomics), yet there is no guarantee that A will have value $\frac{1}{2}$. Fortunately, this is enough for a technique from Kremer (1988) to apply. Crucially, we have that whenever $I(T\langle A\rangle) \neq \frac{1}{2}$, $I(T\langle A\rangle) = I(A)$.

This is enough for M to be *fixable*; we can apply the Kripke fixed-point construction detailed in Kripke (1975) to generate a new model, $M' = \langle D, I' \rangle$. M', by construction, really is a model: $I'(T\langle A \rangle) = I'(A)$ for all formulas A. It is M' that will serve as our countermodel. Since the Kripke construction only changes the values of formulas of the form $T\langle A \rangle$, and since it only changes those values when they don't match the value of A, it will only affect the values of formulas that *don't* occur in S_ω; for formulas that *do* occur in S_ω, M and M' will match. Thus, M' will assign each formula in S_ω to the value corresponding to the place in which it occurs: it is a countermodel.

As a result, we have:

THEOREM 5.5. *For any negated-conjunctive sequent S, either the sequent has a Cut-free proof or it has a countermodel.*

COROLLARY 5.6. *Cut is admissible for negated-conjunctive sequents.*

Proof. Suppose a sequent S is provable with Cut. By Lemma 5.4, S is valid, and by Theorem 5.5, this guarantees that it has a Cut-free proof. □

§6. Conclusion. This paper has shown how to conservatively extend a classical logic with a fully transparent truth predicate. In doing so, all classically valid inferences are extended to the new vocabulary, and metainferences such as the deduction theorem (with side-premises and side-conclusions) and proof by cases are preserved. However, not all metainferences are preserved; the resulting consequence relation is nontransitive in certain cases. Here, I've shown how to identify a large number of cases in which the resulting consequence relation is transitive, and presented two sound and complete proof systems, based on three-sided sequents. For one of these proof systems, but not the other, Cut is admissible. Despite this, because of the truth rules, neither system enjoys a subformula property. Future work, including Ripley (2011a), will explore the prospects of basing a philosophically-motivated theory of truth on a system like the one explored here.[12]

§7. Acknowledgements. This research was partially supported by the Agence Nationale de la Recherche, grant ANR-07-JCJC-0070, program "Cognitive Origins of Vagueness", and by the Government of Spain, program "Borderlineness and Tolerance" ref. FFI2010-16984, MICINN.

BIBLIOGRAPHY

Baaz, M., Fermüller, C. G., & Zach, R. (1992). Dual systems of sequents and tableaux for many-valued logics. Technical Report TUW-E185.2-BFZ.2-92.

Baaz, M., Fermüller, C. G., & Zach, R. (1993). Systematic construction of natural deduction systems for many-valued logics: Extended report. Technical Report TUW-E185.2-BFZ.1-93.

[12] This paper started as a result of a discussion with Ole Hjortland, and has benefited greatly from discussions with him, as well as with Sam Butchart, Pablo Cobreros, Paul Égré, Lloyd Humberstone, Hannes Leitgeb, Toby Meadows (who provided particularly invaluable assistance), Graham Priest, Greg Restall, Robert van Rooij, and Zach Weber. It benefited still further from the helpful comments of two anonymous referees.

Barwise, J., & Etchemendy, J. (1987). *The Liar: An Essay on Truth and Circularity*. Oxford, UK: Oxford University Press.

Beall, Jc. (2009). *Spandrels of Truth*. Oxford, UK: Oxford University Press.

Beall, Jc, & van Fraassen, B. C. (2003). *Possibilities and Paradox: An Introduction to Modal and Many-valued Logic*. Oxford, UK: Oxford University Press.

Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2011). Tolerant, classical, strict. *Journal of Philosophical Logic*. To appear.

Field, H. (2008). *Saving Truth from Paradox*. Oxford, UK: Oxford University Press.

Frankowski, S. (2004). Formalization of a plausible inference. *Bulletin of the Section of Logic*, **33**, 41–52.

Halbach, V. (2011). *Axiomatic Theories of Truth*. Cambridge, UK: Cambridge University Press.

Hyde, D. (1997). From heaps and gaps to heaps of gluts. *Mind*, **106**, 641–660.

Kremer, M. (1988). Kripke and the logic of truth. *Journal of Philosophical Logic*, **17**(3), 225–278.

Kripke, S. (1975). Outline of a theory of truth. *Journal of Philosophy*, **72**(19), 690–716.

Leitgeb, H. (1999). Truth and the liar in De Morgan-valued models. *Notre Dame Journal of Formal Logic*, **40**(4), 496–514.

Malinowski, G. (1990). Q-consequence operation. *Reports on Mathematical Logic*, **24**, 49–59.

Maudlin, T. (2004). *Truth and Paradox*. Oxford, UK: Oxford University Press.

Priest, G. (2006). *In Contradiction*. Oxford, UK: Oxford University Press.

Priest, G. (2008). *An Introduction to Non-Classical Logic: From If to Is* (second edition). Cambridge, UK: Cambridge University Press.

Restall, G. (2005). Multiple conclusions. In Hajek, P., Valdes-Villanueva, L., and Westerståhl, D., editors. *Logic, Methodology, and Philosophy of Science: Proceedings of the Twelfth International Congress*. London, UK: Kings' College Publications, pp. 189–205.

Ripley, D. (2011a). Paradoxes and failures of cut. *Australasian Journal of Philosophy*. To appear.

Ripley, D. (2011b). Sorting out the sorites. In Berto, F., Mares, E., and Tanaka, K., editors. *Paraconsistency: Logic and Applications*. Dordrecht, Netherlands: Springer. To appear.

Shoesmith, D. J., & Smiley, T. J. (1978). *Multiple-conclusion Logic*. Cambridge, UK: Cambridge University Press.

Takeuti, G. (1987). *Proof Theory*. (second edition). Amsterdam: Elsevier Science. Studies in Logic 81.

DEPARTMENT OF PHILOSOPHY, 101 MANCHESTER HALL
344 MANSFIELD RD., UNIVERSITY OF CONNECTICUT
STORRS, CT 06269-2054, USA AND
SCHOOL OF HISTORICAL AND PHILOSOPHICAL STUDIES,
OLD QUAD, UNIVERSITY OF MELBOURNE
PARKVILLE, VIC 3010
AUSTRALIA
E-mail: davewripley@gmail.com