Journal of Philosophical Logic (2022) 51:1367–1386
https://doi.org/10.1007/s10992-021-09624-6

Supervaluations and the Strict-Tolerant Hierarchy

Brian Porter[1]

Received: 14 July 2020 / Accepted: 12 July 2021 / Published online: 30 July 2021

Abstract
In a recent paper, Barrio, Pailos and Szmuc (BPS) show that there are logics that have exactly the validities of classical logic up to arbitrarily high levels of inference. They suggest that a logic therefore must be identified by its valid inferences at every inferential level. However, Scambler shows that there are logics with all the validities of classical logic at every inferential level, but with no antivalidities at any inferential level. Scambler concludes that in order to identify a logic, we at least need to look at the validities and the antivalidities of every inferential level. In this paper, I argue that this is still not enough to identify a logic. I apply BPS's techniques in a super/sub-valuationist setting to construct a logic that has exactly the validities and antivalidities of classical logic at every inferential level. I argue that the resulting logic is nevertheless distinct from classical logic.

Keywords Supervaluationism · Subvaluationism · Strict-Tolerant logic · Substructural Logic; metainferential hierarchy · Metainference · Classification problem

1 Introduction

There are several different logics available: classical logic, intuitionistic logic, the Strong Kleene logic K3, Priest's Logic of Paradox LP, supervaluationism, subvaluationism, etc. These logics are uncontroversially distinct: no one is under the impression that LP and intuitionistic logic are really two ways of presenting one and the same logic. However, the same logic can be formulated in different ways. Classical logic, for example, can be formulated in a natural deduction system, or a Hilbert-style calculus, or a multiple-conclusion sequent calculus. Classical logic can

✉ Brian Porter
bporter@gradcenter.cuny.edu

[1] Graduate Center, City University of New York, 365 5th Ave, Room 7113, New York, NY 11209, USA

also be given a variety of semantics, and can be given several different axiomatizations. Yet all of these, intuitively, are simply different ways of presenting the same logic.

How can we identify a logic? Given two systems, how can we determine whether they are distinct logics, or are merely different presentations of the same logic? This is harder than it first seems. As a first pass, we might be tempted to say that a logic should be identified by its axioms and rules of inference. But this is too strict: it would count different formulations of classical logic as distinct logics, when in fact they are just different presentations of the same logic.

We might say that a logic should be identified by its set of theorems: the sentences that are derivable in the proof theory or get designated values at all models in the semantics. But this is too lax; classical logic, supervaluationism and Priest's paraconsistent logic LP all have the same theorems, and yet are uncontroversially distinct logics. One reason they are distinct is that they validate different inferences. For example, $A, A \rightarrow B \Rightarrow B$ is valid in classical logic and supervaluationism, but not in LP, while $A \vee B \Rightarrow A, B$ is valid in classical logic and LP but not in supervaluationism.[1]

We might therefore try to identify a logic by looking at its set of valid inferences. However, this is also too lax. Two logics can agree on which inferences are valid, but disagree on which *meta*-inferences are valid. The logic ST, for example, introduced by Cobreros, Egré, Ripley and van Rooj (hereafter CERvR) in [7] and [8], has exactly the valid inferences of classical logic, but does not validate the meta-inferential Cut rule. It is therefore distinct from classical logic, which validates every instance of Cut.

We therefore must at least look at which metainferences are valid. But even the inferences and metainferences of a logic are not enough to identify a logic. In a recent paper, Barrio, Pailos and Szmuc (hereafter BPS) show that there are logics that have exactly the validities of classical logic up to arbitrarily high inference levels (in a sense to be made precise below), but then differ from classical logic after that [2]. They argue that this means that a logic must by identified by its valid inferences at every inferential level.

However, even this will not do the job. In a recent paper, Scambler builds on BPS's result, and shows that there are logics that have exactly the validities of classical logic at every level, and yet do not have the antivalidities of classical logic at any level (where an antivalidity is an inference such that every model is a counterexample to that inference) [24]. He argues that this means that for L and L' to be the same logic, they must at least have all of the same validities and antivalidities at every inferential level.

In this paper, I argue that this too is insufficient; there are logics that have exactly the same validities and exactly the same antivalidities at every inferential level, and yet are still intuitively distinct. In Section 2 I introduce the logical framework that I'll be using in the rest of the paper. In Section 3 I apply BPS's methods in a

[1] I specifically have in mind here supervaluationism with what Williamson [29] calls the *global* consequence relation.

super/sub-valuationist framework, to construct notions of validity for inferences of every inferential level that have exactly the validities and antivalidities of classical logic. In Section 4 I argue that the resulting logic is distinct from classical logic. In Section 5 I discuss what this means for the problem of identifying a logic. In Section 6 I consider whether this logic is paraconsistent, and what this means for ST and other logics with similar consequence relations. In Section 7 I close with some concluding remarks.

2 Background

In this section I will introduce the logical framework that I'll be using in the rest of the paper. Most of the notation comes from Scambler [24]. Most of the technical machinery for the slice hierarchy comes from [2] and [18]. However, I will be extending the hierarchy into the transfinite.[2] The machinery for a super/sub-valuationist mixed consequence relation comes from [9].

2.1 Languages

We define our set of languages inductively over the ordinals:

Base Case: Let $\mathcal{L}_0$ be a standard propositional language, with propositional constants p_i and the connectives $\neg, \wedge, \vee$.

Successor Case: Given a language L_α for some ordinal α, let $\mathcal{L}_{\alpha+1}$ be the set of all pairs of sets $\langle\Gamma, \Delta\rangle$ such that $\Gamma \cup \Delta \subseteq \mathcal{L}_\alpha$. I'll write $\langle\Gamma, \Delta\rangle$ in a sequent format, as $\Gamma \Rightarrow_{\alpha+1} \Delta$. We'll call $\Gamma \Rightarrow_\alpha \Delta$ an "inference of order α" or an "α-inference". So $p \wedge q \Rightarrow_1 p$ is a 1-inference, $\{p \Rightarrow_1 p \wedge q\} \Rightarrow_2 \{\Rightarrow_1 q\}$ is a 2-inference, and so on. I will omit set brackets when no confusion can result.

Limit Case: Given languages L_β for all $\beta < \lambda$, let $\mathcal{L}_\lambda$ be the set of all pairs of sets $\langle\Gamma, \Delta\rangle$ such that $\Gamma \cup \Delta \subseteq \bigcup_{\beta<\lambda} L_\beta$. We'll call $\Gamma \Rightarrow_\lambda \Delta$ an "inference of order λ" or a "λ-inference". I will again will omit set brackets when no confusion can result.

Note that the limit-ordinal languages will be cumulative, in the terminology of [25]: the sequents of limit-ordinal languages can contain inferences from *any* lower level. The successor-ordinal languages, however, will not be cumulative in this sense: $\alpha + 1$-sequents can only contain α-inferences.

2.2 Valuations and Models

Let a *Boolean valuation* be a function $v : \mathcal{L}_0 \to \{0, 1\}$ such that the connectives $\neg, \wedge, \vee$ obey the truth tables of classical logic. The set of Boolean valuations is

[2] Scambler [25] also extends the slice hierarchies into the transfinite, though his methods for doing so are slightly different than mine. I suspect that the results using my transfinite framework carry over to his, and vice-versa, but have not confirmed this.

the set of models for propositional classical logic. I'll call propositional classical logic CL.

Let a *supervaluation* (SV) model be a nonempty set of Boolean valuations. I'll use $\mathfrak{V}$ to denote the set of SV models. These models are based on the formal semantics for supervaluationist and subvaluationist logics.[3]

2.3 Notions of Validity

We say a *notion of α-validity* is a function $V : \mathfrak{V} \times \mathcal{L}_\alpha \to \{0, 1\}$. We say a model $m \in \mathfrak{V}$ *satisfies* an α-inference $\Gamma \Rightarrow_\alpha \Delta$ iff $V(m, \Gamma \Rightarrow_\alpha \Delta) = 1$. We can think of a notion of 0-validity as a notion of "truth-in-a-model" that determines which sentences get designated values at each model. A notion of $\alpha + 1$-validity tells us which inferences between sets of α-inferences are valid, and a notion of λ-inference for limit ordinal λ tells us which inferences between sets of lower-level inferences are valid.

I will write $V \vDash_m \Gamma \Rightarrow_\alpha \Delta$ in place of $V(m, \Gamma \Rightarrow_\alpha \Delta) = 1$, and $V \nvDash_m \Gamma \Rightarrow_\alpha \Delta$ in place of $V(m, \Gamma \Rightarrow_\alpha \Delta) = 0$.

When no confusion can result, I will sometimes use Φ in place of $\Gamma \Rightarrow_\alpha \Delta$ for convenience. When it is necessary to indicate that Φ is an α-inference for some ordinal α, I will write Φ as Φ_α.

We say that an inference Φ is *valid* on a notion of α-validity iff for all $m \in \mathfrak{V}$, $V \vDash_m \Phi$. We will write this as $V \vDash \Phi$.

We say that an inference Φ is *anti-valid* on a notion of α-validity iff for all $m \in \mathfrak{V}$, $V \nvDash_m \Phi$. We will write this as $V \vDash \Phi$.

In other words, an inference is valid iff no model is a counterexample; an inference is antivalid if every model is a counterexample.

For all α and all $\Phi \in \mathcal{L}_\alpha$, we say that $CL \vDash \Phi$ whenever Φ is a valid α-inference in classical logic CL.

2.4 Slice Hierarchies

BPS and Scambler use the following definition to produce logics of arbitrarily high (finite) inference levels with notions of α-validity:

Definition 2.1 Successor Slice Let V and U be notions of α-validity. Then the *slice* of V and U, which we write as V/U, is the notion of $\alpha + 1$ validity such that for all $m \in \mathfrak{V}$: $V/U \vDash_m \Gamma \Rightarrow_{\alpha+1} \Delta$ iff $(\exists \gamma \in \Gamma) V \nvDash_m \gamma$ or $(\exists \delta \in \Delta) U \vDash_m \delta$.

We can extend this to the transfinite by adding the limit case:

[3]For a general overview of supervaluationism, see [15]; for earlier presentations and defenses see [10, 16, 17], and [26]; see [27] for a discussion of different consequence relations compatible with supervaluationist semantics. For a general overview of subvaluationism, see [5]; for defenses and earlier presentations, see [6, 12–14], and [28].

Definition 2.2 Limit Slice: Let $\{V_\beta\}_{\beta<\lambda}$ and $\{U_\beta\}_{\beta<\lambda}$ be sets of notions of validity for all $\beta < \lambda$. Then the *limit slice* of $\{V_\beta\}_{\beta<\lambda}$ and $\{U_\beta\}_{\beta<\lambda}$, which we write as $V_{<\lambda}/U_{<\lambda}$ is the notion of λ-validity such that for all $m \in \mathfrak{V}$: $V_{<\lambda}/U_{<\lambda} \vDash_m \Gamma \Rightarrow_\lambda \Delta$ iff either:

$\exists\gamma \in \Gamma$ such that γ is a β-inference for some $\beta < \lambda$ and $V_\beta \nvDash_m \gamma$, or
$\exists\delta \in \Delta$ such that δ is a β-inference for some $\beta < \lambda$ and $U_\beta \vDash_m \delta$.

Given two notions of α-validity V and U, we can build the transfinite slice hierarchy over V and U:

$$V, U$$
$$V/V, V/U, U/V, U/U$$
$$\vdots$$

Using slices like the ones defined above, BPS show that one can use Strong Kleene 3-valued models to build a slice hierarchy based on CERvR's Strict-Tolerant logic ST. They prove that for every inferential level $n < \omega$, there is a logic that has exactly the validities of classical logic up to order n, but differs from classical logic at higher inferences levels. They suggest that classical logic therefore must be identified by its valid inferences at *every* (finite) inferential level. But in [24] and [25], Scambler shows that even a logic with exactly the validities of classical logic at every inferential level can differ from classical logic. He demonstrates that what he calls the "tolerant twist logic" has exactly the validities of classical logic at every inferential level, but unlike classical logic, has no antivalidities.

In what follows, I construct a transfinite slice hierarchy that gives us a logic with exactly the validities of classical logic at every inferential level, and exactly the antivalidities of classical logic at every inferential level. I will then argue that this logic still should not be identified with classical logic.

2.5 A Note about Local and Global Validity

It is important to note that the definitions of successor and limit slices given above will produce notions of *local* validity, rather than notions of *global* validity. When dealing with metainferences, there are at least two ways to define metainferential validity over a class of models: *local* validity and *global* validity.[4]

In traditional (unsliced) contexts, local validity can be thought of as preservation of satisfaction, while global validity can be thought of as preservation of validity. A metainference is locally valid iff at every model, either some conclusion inference is satisfied or some premise inference is not. A metainference is globally valid iff either some conclusion inference in valid or some premise inference is valid. For example, the metainference $\{p \Rightarrow_1 q\} \Rightarrow_2 \{p \Rightarrow_1 q \wedge r\}$ is globally valid in classical logic, because the premise inference $p \Rightarrow_1 q$ is not valid. But it is not locally valid in classical logic, because there are classical valuations at which $p \Rightarrow_1 q$ is satisfied

[4]For a detailed discussion of the distinction between local and global metainferential validity, see [11].

but $p \Rightarrow_1 q \wedge r$ is not. In general, local validity implies global validity, but global validity does not imply local validity.

We can generalize these notions to slice notions of validity: an inference is V/U locally valid iff at every model, either some premise is not V-satisfied or some conclusion is U-satisfied. Similarly, an inference is V/U globally valid iff either some premise is not V-valid or some conclusion is U-valid.

The notions of validity generated by the above definitions for successor and limit slices are local, rather than global, notions of validity. Following [2, 18], and [24], I will be restricting attention to *local* validity at each level, rather than global validity. I will therefore use $CL \vDash \Phi$ whenever Φ is a locally valid inference in classical logic. There is a sense in which local validity is more fine-grained than global validity: any two logics with the same locally valid inferences must have the same globally valid inferences, but not vice-versa. ST and classical logic have the same globally valid 2-inferences (in the empty signature), but have different locally valid inferences: Cut is locally valid in classical logic, but is not locally valid in ST (even in the empty signature). Therefore, when we see in the following sections that the resulting logic has exactly the local validities of classical logic, it immediately follows that it also has exactly the global validities of classical logic.

3 The *LM* Hierarchy

3.1 Two Notions of 0-validity

There are at least two interesting notions of 0-validity that can be defined over the set of SV models $\mathfrak{V}$. These correspond to the notions of 0-validity for supervaluationism and subvaluationism, so I'll call them P (for suPervaluationism) and B (for suBvaluationism):

$P \vDash_m \phi$ iff $v(\phi) = 1$ for all $v \in m$

$B \vDash_m \phi$ iff $v(\phi) = 1$ for some $v \in m$

It is easy to see that for all 0-inferences, $P \vDash \phi$ iff $B \vDash \phi$ iff $CL \vDash \phi$. If every Boolean valuation assigns 1 to ϕ, then every set of Boolean valuations contains only valuations that assign 1 to ϕ. It is also easy to see that $P \vDash \phi$ iff $B \vDash \phi$ iff $CL \vDash \phi$. If every Boolean valuation assigns 0 to ϕ, then no set of Boolean valuations has any valuations that assign 1 to ϕ.

3.2 Six Notions of 1-validity

We could slice together P and B to form a hierarchy of notions of validity. For example, there are four notions of 1-validity that we can get just by slicing P and B:

$P/P \vDash_m \Gamma \Rightarrow_1 \Delta$ iff either $(\exists\gamma \in \Gamma)(\exists v \in m)v(\gamma) = 0$ or $(\exists\delta \in \Delta)(\forall v \in m)v(\delta) = 1$

$B/B \vDash_m \Gamma \Rightarrow_1 \Delta$ iff either $(\exists\gamma \in \Gamma)(\forall v \in m)v(\gamma) = 0$ or $(\exists\delta \in \Delta)(\exists v \in m)v(\delta) = 1$

$B/P \vDash_m \Gamma \Rightarrow_1 \Delta$ iff either $(\exists\gamma \in \Gamma)(\forall v \in m)v(\gamma) = 0$ or $(\exists\delta \in \Delta)(\forall v \in m)v(\delta) = 1$

$P/B \vDash_m \Gamma \Rightarrow_1 \Delta$ iff either $(\exists\gamma \in \Gamma)(\exists v \in m)v(\gamma) = 0$ or $(\exists\delta \in \Delta)(\exists v \in m)v(\delta) = 1$

P/P and B/B correspond to the (global) 1-validity consequence relations for supervaluationism and subvaluationism, respectively. Neither has all of the 1-validities of classical logic. For example, $P/P \nvDash \{\} \Rightarrow_1 A, \neg A$, and $B/B \nvDash A, \neg A \Rightarrow_1 \{\}$.

The mixed-condition 1-validity consequence relations B/P and P/B are effectively the super/sub-valuation equivalents of the K3/LP consequence relations TS and ST.[5] Like TS, B/P is not reflexive, in the sense that $A \Rightarrow_1 A$ is not valid unless A is a classical tautology or classical contradiction. Like ST, P/B has exactly the same valid 1-inferences as classical logic. And like ST, there is a sense in which P/B is not transitive. In particular, $P/B \vDash_m \Gamma \Rightarrow_1 A$ and $P/B \vDash_m A \Rightarrow_1 \Delta$ do not imply $P/B \vDash_m \Gamma \Rightarrow_1 \Delta$.[6] Although ST and P/B both have exactly the 1-validities of classical logic, P/B has an additional feature that ST does not have: P/B has all of the 1-validities of classical logic *and* all of the 1-antivalidities of classical logic.

B/P and P/B could be sliced together to construct a notion of 2-validity, which will again have all of the 2-validities of classical logic. However, this notion of 2-validity will only have *some* of the antivalidities of classical logic. For example, the 2-inference from $A \Rightarrow_1 A$ to $\Rightarrow_1 A \wedge \neg A$ is classically antivalid, but it is not antivalid in BP/PB: there are SV models at which $A \Rightarrow_1 A$ is not B/P-satisfied. In order to construct a logic with all of validities and all of the antivalidities of classical logic, we need to build our hierarchy using different notions of 1-validity.

Instead of using any of these notions of 1-validity, we will define notions of 1-validity directly, and build the hierarchy out of those. Our first such notion of 1-validity, L, corresponds to what Williamson [29] calls the "local" consequence relation on supervaluation models:

$L \vDash_m \Gamma \Rightarrow_1 \Delta$ iff $(\forall v \in m)[(\exists\gamma \in \Gamma)v(\gamma) = 0 \vee (\exists\delta \in \Delta)v(\delta) = 1]$

A 1-inference is L-satisfied at a model iff it is classically satisfied at every valuation in the model. It's easy to see that this notion of 1-validity has exactly the 1-validities and 1-antivalidities of classical logic.

In addition to L, we can also define a more tolerant consequence relation M:

$M \vDash_m \Gamma \Rightarrow_1 \Delta$ iff $(\exists v \in m)[(\exists\gamma \in \Gamma)v(\gamma) = 0 \vee (\exists\delta \in \Delta)v(\delta) = 1]$

A 1-inference is M-satisfied at a model iff it is classically satisfied at *some* valuation in the model. Like L, M has exactly the 1-validities and 1-antivalidities of classical logic. But despite this, M is not transitive. To see this, let m be a model containing two valuations u and v, such that $v(p) = 1$, $v(q) = 0$, $u(p) = 0$, and $u(q) = 0$. In this case $M \vDash_m \Rightarrow_1 p$ and $M \vDash_m p \Rightarrow_1 q$, and yet $M \nvDash_m \Rightarrow_1 q$.

[5] For more information about the strict-tolerant approach in a super/sub-valuationist setting, see [9].

[6] There are many different ways to use the term "transitive" when discussing consequence relations. See [20] for a survey of several different notions. My use of the term is somewhat idiosyncratic, in that it is really transitivity of *satisfaction* that I have in mind, rather than transitivity of *validity*. However, I take this to be an important and distinctive feature of the logic: it means that the set of sentences that are true in a model are not closed under the valid inferences of the logic. This has significant consequences for how the logic handles non-logical axioms, which will be discussed more in Sections 4 and 6.

There is one fact about L and M that will be important in constructing the LM hierarchy: for any SV model $m \in \mathfrak{V}$, $L \vDash_m \{\} \Rightarrow_1 \phi$ iff $P \vDash_m \phi$, and $M \vDash \{\} \Rightarrow_1 \phi$ iff $B \vDash \phi$. Without this property, we couldn't use these notions of 0- and 1-validity together to construct anything that deserved to be called a "logic".

It is easy to see, but important to note before we continue, that all six of these notions of 1-validity have the following feature: given a singleton model $\{v\}$, each notion of validity gives us $\vDash_{\{v\}} \Gamma \Rightarrow_1 \Delta$ iff $CL \vDash_v \Gamma \Rightarrow_1 \Delta$. Notions of α-validity at any level can have the equivalent property for α-inferences. I will call this the *singleton property*:

Definition 3.1 Singleton Property
A notion of α-validity V has the *singleton property* if for all singleton models $\{v\}$ and all α-inferences $\Gamma \Rightarrow_\alpha \Delta$, $V \vDash_{\{v\}} \Gamma \Rightarrow_\alpha \Delta$ iff $CL \vDash_v \Gamma \Rightarrow_\alpha \Delta$.

We can now use P, B, L and M to construct a hierarchy of notions of validity that have all of the validities and antivalidities of classical logic at every level of inference.

3.3 The Hierarchy

We will define what I will call the LM hierarchy inductively:

Base cases: L_0 = P, M_0 = B
L_1 = L, M_1 = M

Successor Case: for $\alpha \geq 1$:
$L_{\alpha+1} = M_\alpha / L_\alpha$, $M_{\alpha+1} = L_\alpha / M_\alpha$

Limit Case: for limit ordinals λ:
$L_\lambda = M_{<\lambda} / L_{<\lambda}$, $M_\lambda = L_{<\lambda} / M_{<\lambda}$

Call M_∞ the logic that evaluates α-validity at every ordinal α in accordance with M_α. Similarly for L_∞.

My primary goal in this section will be to show that M_∞ has exactly the validities and antivalidities of classical logic at every inference level. First, we need to show that every notion of validity in the hierarchy has the singleton property:

Lemma 1 *If two notions of α-validity X_α and Y_α both have the singleton property, then their slice X_α / Y_α also has the singleton property.*

Proof $CL \vDash_v \Gamma \Rightarrow_{\alpha+1} \Delta$ iff either there is a $\gamma \in \Gamma$ such that $CL \nvDash_v \gamma$, or there is a $\delta \in \Delta$ such that $CL \vDash_v \delta$. Since X_α and Y_α both have the singleton property, $CL \nvDash_v \gamma$ iff $X_\alpha \nvDash_{\{v\}} \gamma$ and $CL \vDash_v \delta$ iff $Y_\alpha \vDash_{\{v\}} \delta$. $X_\alpha / Y_\alpha \vDash_{\{v\}} \Gamma \Rightarrow_{\alpha+1} \Delta$ iff there is a $\gamma \in \Gamma$ such that $X_\alpha \nvDash_{\{v\}} \gamma$ or there is a $\delta \in \Delta$ such that $Y_\alpha \vDash_{\{v\}} \delta$. Therefore $CL \vDash_v \Gamma \Rightarrow_{\alpha+1} \Delta$ iff $X_\alpha / Y_\alpha \vDash_{\{v\}} \Gamma \Rightarrow_{\alpha+1} \Delta$. □

This also holds for the limit slices:

Lemma 2 *If every notion of validity $X_i \in \{X_\beta\}_{\beta<\lambda}$ and $Y_i \in \{Y_\beta\}_{\beta<\lambda}$ has the singleton property, then $X_{<\lambda}/Y_{<\lambda}$ has the singleton property.*

Proof Suppose $CL \vDash_v \Gamma \Rightarrow_\lambda \Delta$. Then either $CL \nvDash_v \gamma$ for some $\gamma \in \Gamma$, or $CL \vDash_v \delta$ for some $\delta \in \Delta$. By IH, for all $\beta < \lambda$, $X_\beta \vDash_{\{v\}} \Phi_\beta$ iff $CL \vDash_v \Phi_\beta$, and $Y_\beta \vDash_{\{v\}} \Phi_\beta$ iff $CL \vDash_v \Phi_\beta$. If $CL \nvDash_v \gamma_\beta$ for some $\gamma \in \Gamma$ and $\beta < \lambda$, then $X_\beta \nvDash_{\{v\}} \gamma_\beta$, and so $X_{<\lambda}/Y_{<\lambda} \vDash_{\{v\}} \Gamma \Rightarrow_\lambda \Delta$. Otherwise, $CL \vDash_v \delta_\beta$ for some $\delta \in \Delta$ and some $\beta < \lambda$. Then $Y_\beta \vDash_{\{v\}} \delta_\beta$, and so $X_{<\lambda}/Y_{<\lambda} \vDash_{\{v\}} \Gamma \Rightarrow_\lambda \Delta$.

The reverse direction follows the same pattern. □

It follows that if two notions of validity V and U have the singleton property, then every notion of validity in the transfinite hierarchy over V and U has the singleton property. Since P, B, L, and M all have the singleton property, this means that every notion of validity in M_∞ and in L_∞ has the singleton property. We can use this fact to prove sufficient conditions for a slice notion of $\alpha + 1$-validity to have all of the validities and all of the antivalidities of classical logic.

Lemma 3 *Let X_α and Y_α be two notions of α-validity defined over the set of SV models $\mathfrak{V}$ that both have the singleton property. Suppose that for all SV models m, $X_\alpha \vDash_m \Phi$ implies $CL \vDash_v \Phi$ for all $v \in m$, and $Y_\alpha \nvDash_m \Phi$ implies $CL \nvDash_v \Phi$ for all $v \in m$. Then $X_\alpha/Y_\alpha \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$ iff $CL \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$.*

Proof The left-to-right direction follows from lemma 1. Since X_α and Y_α both have the singleton property, X_α/Y_α does too. So $CL \nvDash \Phi$ implies that $X_\alpha/Y_\alpha \nvDash \Phi$: if some valuation v is a CL counterexample to Φ, then $\{v\}$ will be a X_α/Y_α counterexample to Φ.

For the right-to-left direction, suppose $X_\alpha/Y_\alpha \nvDash \Gamma \Rightarrow_{\alpha+1} \Delta$. Then there is an SV model $m \in \mathfrak{V}$ such that $X_\alpha/Y_\alpha \nvDash_m \Gamma \Rightarrow_{\alpha+1} \Delta$. It follows that $X_\alpha \vDash_m \gamma$ for all $\gamma \in \Gamma$ and $Y_\alpha \nvDash_m \delta$ for all $\delta \in \Delta$. By assumption, $X_\alpha \vDash_m \Phi$ implies $CL \vDash_v \Phi$ for all $v \in m$, and $Y_\alpha \nvDash_m \Phi$ implies $CL \nvDash_v \Phi$ for all $v \in m$. Therefore every $v \in m$ is such that $CL \vDash_v \gamma$ for all γ and $CL \nvDash_v \delta$ for all $\delta \in \Delta$. Therefore $CL \nvDash \Gamma \Rightarrow_{\alpha+1} \Delta$. □

Lemma 4 *Let X_α and Y_α be two notions of α-validity defined over the set of SV models $\mathfrak{V}$ that both have the singleton property. Suppose that for all SV models $m \in \mathfrak{V}$, $CL \vDash_v \Phi$ for all $v \in m$ implies $X_\alpha \vDash_m \Phi$, and $CL \nvDash_v \Phi$ for all $v \in m$ implies $Y_\alpha \nvDash_m \Phi$. Then $X_\alpha/Y_\alpha \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$ iff $CL \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$.*

Proof The left-to-right direction follows from lemma 1. Since X_α and Y_α both have the singleton property, X_α/Y_α does too. So $CL \nvDash \Phi$ implies that $X_\alpha/Y_\alpha \nvDash \Phi$: if some valuation v is not a CL counterexample to Φ, then $\{v\}$ will not be a X_α/Y_α counterexample to Φ.

For the right-to-left direction, suppose $CL \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$. Then for every Boolean valuation v, $CL \vDash_v \gamma$ for all $\gamma \in \Gamma$ and $CL \nvDash \delta$ for all $\delta \in \Delta$. By assumption, $CL \vDash_v \Phi$ for all $v \in m$ implies $X_\alpha \vDash_m \Phi$, and $CL \nvDash_v \Phi$ for all $v \in m$ implies $Y_\alpha \nvDash_m \Phi$. It follows that for every SV model $m \in \mathfrak{V}$, $X_\alpha \vDash_m \Phi$ and $Y_\alpha \nvDash_m \Phi$. Therefore, $X_\alpha/Y_\alpha \vDash \Gamma \Rightarrow_{\alpha+1} \Delta$. □

These conditions also hold for the limit slices:

Lemma 5 *Let $\{X_\beta\}_{\beta<\lambda}$ and $\{Y_\beta\}_{\beta<\lambda}$ be two sets of notions of β-validity defined over the set of SV models $\mathfrak{V}$ for all $\beta < \lambda$ that all have the singleton property. Suppose that for all $\beta < \lambda$ and all SV models $m \in \mathfrak{V}$, $X_\beta \vDash_m \Phi$ implies $CL \vDash_v \Phi$ for all $v \in m$, and $Y_\beta \nvDash_m \Phi$ implies $CL \nvDash_v \Phi$ for all $v \in m$. Then $X_{<\lambda}/Y_{<\lambda} \vDash \Gamma \Rightarrow_\lambda \Delta$ iff $CL \vDash \Gamma \Rightarrow_\lambda \Delta$.*

Proof The left-to-right direction follows from the singleton property.

For the right-to-left direction, suppose that $X_{<\lambda}/Y_{<\lambda} \nvDash \Gamma \Rightarrow_\lambda \Delta$. Then there is an SV model $m \in \mathfrak{V}$ such that for all $\gamma \in \Gamma$ and all $\delta \in \Delta$ and all $\beta < \lambda$, $X_\beta \vDash_m \gamma_\beta$ and $Y_\beta \nvDash_m \delta_\beta$. It follows that for all $v \in m$, $CL \vDash_v \gamma$ and $CL \nvDash_v \delta$ for all $\gamma \in \Gamma$ and $\delta \in \Delta$. Therefore, $CL \nvDash \Gamma \Rightarrow_\lambda \Delta$. □

Lemma 6 *Let $\{X_\beta\}_{\beta<\lambda}$ and $\{Y_\beta\}_{\beta<\lambda}$ be two sets of notions of β-validity defined over the set of SV models $\mathfrak{V}$ for all $\beta < \lambda$ that all have the singleton property. Suppose that for all $\beta < \lambda$, all β-inferences Φ, and all SV models $m \in \mathfrak{V}$, $CL \vDash_v \Phi$ for all $v \in m$ implies $X_\beta \vDash_m \Phi$, and $CL \nvDash_v \Phi$ for all $v \in m$ implies $Y_\beta \nvDash_m \Phi$. Then $X_{<\lambda}/Y_{<\lambda} \vDash \Gamma \Rightarrow_\lambda \Delta$ iff $CL \vDash \Gamma \Rightarrow_\lambda \Delta$.*

Proof The left-to-right direction follows from the singleton property.

For the right-to-left direction, suppose that $CL \vDash \Gamma \Rightarrow_\lambda \Delta$. Then for every Boolean valuation v, $CL \vDash_v \gamma$ for all $\gamma \in \Gamma$, and $CL \nvDash_v \delta$ for all $\delta \in \Delta$. By assumption, $CL \vDash_v \Phi_\beta$ for all $v \in m$ implies $X_\beta \vDash_m \Phi_\beta$, and $CL \nvDash_v \Phi_\beta$ for all $v \in m$ implies $Y_\beta \nvDash_m \Phi_\beta$. It follows that for all SV models $m \in \mathfrak{V}$, $X_\beta \vDash_m \gamma_\beta$ and $Y_\beta \nvDash_m \delta_\beta$. Therefore $X_{<\lambda}/Y_{<\lambda} \vDash \Gamma \Rightarrow_\lambda \Delta$. □

So to show that the logic M_∞ has all of the validities and antivalidities of classical logic at every inferential level, it suffices to show that for all ordinals $\alpha \geq 1$ and all SV models $m \in \mathfrak{V}$, $L_\alpha \vDash_m \Phi$ iff $CL \vDash_v \Phi$ for all $v \in m$, and $M_\alpha \nvDash_m \Phi$ iff $CL \vDash_v \Phi$ for all $v \in m$.

Lemma 7 *For all $\alpha \geq 1$, $L_\alpha \vDash_m \Phi$ iff $CL \vDash_v \Phi$ for all $v \in m$, and $M_\alpha \nvDash_m \Phi$ iff $CL \nvDash_v \Phi$ for all $v \in m$.*

Proof By induction. The base case is L and M, and follows immediately from the definition of L and M.

For the successor case, we take the inductive hypothesis that for all SV models $m \in \mathfrak{V}$ and all $\Phi \in \mathcal{L}_\alpha$, $L_\alpha \vDash_m \Phi$ iff $CL \vDash_v \Phi$ for all $v \in m$ and $M_\alpha \nvDash_m \Phi$ iff $CL \nvDash_v \Phi$ for all $v \in m$.

For $L_{\alpha+1}$, we note that $L_{\alpha+1} \vDash_m \Gamma \Rightarrow_{\alpha+1} \Delta$ iff either $(\exists\gamma \in \Gamma)M_\alpha \nvDash_m \gamma$ or $(\exists\delta \in \Delta)L_\alpha \vDash_m \delta$. By IH, for all $\gamma \in \Gamma$ and all $\delta \in \Delta$, $M_\alpha \nvDash_m \gamma$ iff $CL \nvDash_v \gamma$ for all $v \in m$, and $L_\alpha \vDash_m \delta$ iff $CL \vDash_v \delta$ for all $v \in m$. By definition, $(\exists\gamma \in \Gamma)CL \nvDash_v \gamma$ for all $v \in m$ or $(\exists\delta \in \Delta)CL \vDash_v \delta$ for all $v \in m$ iff $CL \vDash_v \Gamma \Rightarrow_{\alpha+1} \Delta$ for all $v \in m$.

For $M_{\alpha+1}$, we note that $M_{\alpha+1} \nvDash_m \Gamma \Rightarrow_{\alpha+1} \Delta$ iff $(\forall\gamma \in \Gamma)L_\alpha \vDash_m \gamma$ and $(\forall\delta \in \Delta)M_\alpha \nvDash_m \delta$. By IH, for all $\gamma \in \Gamma$ and all $\delta \in \Delta$, $L_\alpha \vDash_m \gamma$ iff $CL \vDash_v \gamma$ for all $v \in m$, and $M_\alpha \nvDash_m \delta$ iff $CL \nvDash_v \delta$ for all $v \in m$. By definition, $CL \vDash_v \gamma$ for all $v \in m$ and $CL \nvDash_v \delta$ for all $v \in m$ iff $CL \nvDash \Gamma \Rightarrow_{\alpha+1} \Delta$ for all $v \in m$.

For the limit case, we take the inductive hypothesis that for all $\beta < \lambda$, for all SV models $m \in \mathfrak{V}$ and for all $\Phi \in \bigcup_{\beta<\lambda} \mathcal{L}_\beta$, $L_\beta \vDash_m \Phi$ iff $CL \vDash_v \Phi$ for all $v \in m$ and $M_\beta \nvDash_m \Phi$ iff $CL \nvDash_v \Phi$ for all $v \in m$.

For L_λ, we note that $L_\lambda \vDash_m \Gamma \Rightarrow_\lambda \Delta$ iff either $(\exists\gamma \in \Gamma)(\exists\beta < \lambda)M_\beta \nvDash_m \gamma$ or $(\exists\delta \in \Delta)(\exists\beta < \lambda)L_\beta \vDash_m \delta$. By IH, for all $\gamma \in \Gamma$ and all $\delta \in \Delta$ and all $\beta < \lambda$, $M_\beta \nvDash_m \gamma_\beta$ iff $CL \nvDash_v \gamma_\beta$ for all $v \in m$, and $L_\beta \vDash_m \delta_\beta$ iff $CL \vDash_v \delta_\beta$ for all $v \in m$. By definition, $(\exists\gamma \in \Gamma)CL \nvDash_v \gamma$ for all $v \in m$ or $(\exists\delta \in \Delta)CL \vDash_v \delta$ for all $v \in m$ iff $CL \vDash_v \Gamma \Rightarrow_\lambda \Delta$ for all $v \in m$.

For M_λ, we note that $M_\lambda \nvDash_m \Gamma \Rightarrow_\lambda \Delta$ iff $(\forall\gamma \in \Gamma)(\exists\beta < \lambda)L_\beta \vDash_m \gamma$ and $(\forall\delta \in \Delta)(\exists\beta < \lambda)M_\beta \nvDash_m \delta$. By IH, for all $\gamma \in \Gamma$ and all $\delta \in \Delta$ and all $\beta < \lambda$, $L_\beta \vDash_m \gamma_\beta$ iff $CL \vDash_v \gamma_\beta$ for all $v \in m$, and $M_\beta \nvDash_m \delta_\beta$ iff $CL \nvDash_v \delta_\beta$ for all $v \in m$. By definition, $CL \vDash_v \gamma$ for all $v \in m$ and $CL \nvDash_v \delta$ for all $v \in m$ iff $CL \nvDash_v \Gamma \Rightarrow_\lambda \Delta$ for all $v \in m$. □

With this, we can now prove the primary result of this section:

Theorem 8 *For every ordinal α, $M_\infty \vDash \Phi_\alpha$ iff $CL \vDash \Phi_\alpha$, and $M_\infty \vDash \Phi_\alpha$ iff $CL \vDash \Phi_\alpha$.*

Proof Follows immediately from lemmas 3, 4, 5, 6, 7, and the fact that B has exactly the 0-validities and 0-antivalidities of classical logic and M has exactly the 1-validities and 1-antivalidities of classical logic. □

M_∞ therefore has exactly the validities and antivalidities of classical logic at every inferential level.

4 Is M_∞ Classical Logic?

We have just shown that the logic M_∞ has exactly the validities and antivalidities of classical logic at every inferential level. One might therefore be tempted to identify M_∞ with classical logic. However, this would be a mistake.

To illustrate why, we need to look at how the two logics handle the addition of non-logical axioms. When presented with the same set of axioms in the same language, M_∞ and classical logic will generate different *theories*, at least in some cases. I take this to be a reason for thinking that the two logics are distinct.

We often want to use a logic to prove theorems from sets of axioms, as in the case of ZFC or Peano Arithmetic. This does not involve moving to a new logic; a single logic, like classical logic, can be used with a variety of theory-specific non-logical axioms and still be the same logic in each case. Any set of axioms will have certain consequences in a logic. Given a logic L and a set of axioms Γ, let the *theory* of Γ in

L be the set of formulae that are true (satisfied) at all L-models at which the members of Γ are true (satisfied).[7]

Suppose that we are presented with two logics, L and L$'$. Suppose that, given the same language $\mathcal{L}$ and the same set of non-logical axioms Γ, the theory generated by Γ in L is strictly greater than the theory generated by Γ in L$'$. In that case, I take it that L and L$'$ can safely be considered different logics. What follows from a given set of axioms does not and should not depend on how the logic is presented. It would be quite a shock to discover that whether or not the continuum hypothesis follows from ZFC depends on how we present first order classical logic: in one presentation, the continuum hypothesis is independent of the axioms of ZFC, but in another presentation it is a theorem. That simply cannot happen; there must be some determinate, presentation-independent fact of the matter as to what the consequences of such-and-such axioms are in a given logic. What follows from a set of axioms in a logic is not a matter of presentation; it is an essential feature of the logic.

With that in mind, let's consider how classical logic and M_∞ behave given the same set of axioms in the same language. In the standard propositional language used in the previous section, take p_1 and $\neg p_1$ as axioms in both classical logic and M_∞. In the sequent calculus presentation that we've been using, we can also use $\{\Rightarrow_1 p_1, \Rightarrow_1 \neg p_1\}$ as our set of axioms.[8] We can then examine the theory of $\{\Rightarrow_1 p_1, \Rightarrow_1 \neg p_1\}$ in each logic.

The theory generated by these axioms in classical logic is trivial. Every sequent whatsoever follows from $\{\Rightarrow_1 p_1, \Rightarrow_1 \neg p_1\}$ in classical logic, including $\Rightarrow_1 A$ for all formulae A. In classical logic, there is no way for $\Rightarrow_1 p_1$ and $\Rightarrow_1 \neg p_1$ to be satisfied at the same valuation; the two inferences are not *jointly satisfiable*. Therefore it is trivially true that A is true in all models in which p_1 and $\neg p_1$ are true, and so any sentence A is a member of the theory of p_1 and $\neg p_1$ in classical logic.

However, this is not the case in M_∞. To see why, note that both $\Rightarrow_1 p_1$ and $\Rightarrow_1 \neg p_1$ will be satisfied at any SV model in which at least one Boolean valuation assigns p_1 value 1, and at least one valuation assigns p_1 value 0. Such models need not be trivial in M_∞; any set of two Boolean valuations assigning different values to p_1 will do. $\Rightarrow_1 p_1$ and $\Rightarrow_1 \neg p_1$ are therefore *jointly satisfiable* in M_∞: there are models that satisfy both inferences. This means that the question of which formulae and inferences are satisfied at the models of $\{\Rightarrow_1 p_1, \Rightarrow_1 \neg p_1\}$ is not at all trivial. For some models m and some formulae A, there will be no valuation in m that assigns value 1 to A. For example, there is no Boolean valuation that assigns value 1 to $p_1 \wedge \neg p_1$. As a result, $\Rightarrow_1 p_1 \wedge \neg p_1$ is antivalid in M_∞: there is no model that satisfies $\Rightarrow_1 p_1 \wedge \neg p_1$. This is true even if we restrict our attention to models that satisfy both $\Rightarrow_1 p_1$ and $\Rightarrow_1 \neg p_1$. Therefore $\Rightarrow_1 p_1 \wedge \neg p_1$ is not in the theory of $\Rightarrow_1 p_1$ and $\Rightarrow_1 \neg p_1$ in M_∞ (even though $\{\Rightarrow_1 p_1, \Rightarrow_1 \neg p_1\} \Rightarrow_2 \{\Rightarrow_1 p_1 \wedge \neg p_1\}$ is valid). But it *is* in the theory of those axioms in classical logic; everything is. As

[7] Although "theory" is defined here semantically, this definition is equivalent to the proof-theoretic notion of a theory as the set of formulae provable from a set of axioms. I use the semantic definition here for the simple reason that I do not currently have a proof theory for the model-theoretically-defined logic M_∞.

[8] We can do this because in both M_∞ and classical logic, a model is a counterexample to a formula A iff it is a counterexample to $\Rightarrow_1 A$ iff it is a counterexample to $\Rightarrow_2 \Rightarrow_1 A$, and so on.

such, the same axioms in the same language can have different theories in M_∞ than in classical logic. M_∞ and classical logic are therefore distinct logics.

5 Identity Conditions for Logics

I've argued that classical logic and M_∞ are not the same logic, despite the fact that they have the same validities and antivalidities at every inferential level. We might then ask, under what conditions can logics L and L′ rightfully be said to be the *same* logic?

Together with the results of [2] and [24], I take the results here to show that having the same validities and antivalidities is not sufficient to identify two formal systems as the same logic, even if they have the same validities and antivalidities at every inferential level. M_∞ and classical logic have exactly the same validities and antivalidities at every inferential level, and yet they can behave quite differently when given the same set of axioms. But the consequences of a set of axioms are not presentation-dependent features of a logic; the same axioms should not generate different theories depending on how we present the logic. M_∞ and classical logic must therefore be distinct logics, and so it is possible for two distinct logics to have all the same validities and antivalidities at every level of inference.

We must look beyond validities and antivalidities in order to determine whether two formal systems L and L′ are distinct logics, or are simply two presentations of the same logic. In light of the discussion in the previous section, I propose that we at least need to consider the sets of inferences that are jointly satisfiable in a given logic. Even if two logics agree on which inferences have counterexamples and non-counterexamples, the logics can still disagree regarding which *sets* of inferences share a single counterexample and which do not. In classical logic, the 1-inferences $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$ are not jointly satisfiable. In M_∞, they are. This seems to be precisely the reason that classical logic and M_∞ behave differently given p_1 and $\neg p_1$ as axioms: the axioms are jointly satisfiable in M_∞, but are not jointly satisfiable in classical logic. Having the same sets of jointly satisfiable inferences therefore seems promising as an identity condition to distinguish logics.

In fact if we take sets of jointly satisfiable inferences as an identity condition, then the antivalidities condition can be dropped. It is subsumed under the jointly satisfiable sets condition: if two logics have exactly the same sets of jointly satisfiable inferences, then they must also have the same antivalidities. This is because the α-inference Φ is antivalid iff Φ is not satisfiable iff the singleton set of inferences $\{\Phi\}$ is not jointly satisfiable.[9] However, the validities and jointly satisfiable sets

[9] Sets of inferences that are not jointly satisfiable therefore generalize the notion of antivalid inference: they are, in effect, the sets of inferences that are together antivalid. We could apply an analogous generalization to validity, and look at the sets of inferences that do not share a single counterexample. L, the local supervaluationist consequence relation, for example, has exactly the same valid 1-inferences as classical logic, but has different sets of 1-inferences that do not share a counterexample: $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$ cannot share a counterexample in classical logic, but they can in the local supervaluationist logic. Whether or not this generalization has any philosophically interesting applications remains to be seen.

conditions are independent. To see why, consider the trivial logic in which every inference is valid. In this logic, every set of inferences is jointly satisfiable. In at least some languages, every set of 1-inferences is jointly satisfiable in the strict-tolerant logic ST. This is because (without logical constants and the like) ST has a model in which every formula gets value $\frac{1}{2}$, and every 1-inference is satisfied at that model. So ST and the trivial logic have exactly the same sets of jointly satisfiable inferences, yet they have different validities: some inferences are invalid in ST, but no inferences are invalid in the trivial logic. The same-validities condition therefore cannot be subsumed under the same-sets-of-jointly-satisfiable-inferences condition. Joint satisfiability is a matter of there being a non-counterexample, and the existence of non-counterexamples cannot by itself tell us whether there are *no* counterexamples.

It is worth noting that the new identity condition offered here is a semantic condition. Unlike validity (and possibly antivalidity), "sets of jointly satisfiable inferences" is an inherently semantic notion, defined in terms of models and satisfaction conditions. Some, especially those who take a purely instrumentalist approach to the models for a logic, may object to this as an identity condition on logics.[10]

However, the identity condition offered here is a semantic condition in part because, like the slice-hierarchy logics that came before it, M_∞ is constructed model-theoretically. As such, its consequence relation is defined semantically. There may be proof theoretic ways to formulate these logics without appealing to any model-theoretic definitions. Once this is done, it will hopefully be clear what the proof-theoretic equivalent of jointly satisfiable inferences might be; it may be some sort of closure operator on sets of non-logical axioms. By introducing this semantic identity condition, I do not mean to suggest that there is no equivalent proof-theoretic condition that could serve the same purpose. There may well be a way to distinguish M_∞ from classical logic without making any appeal to models or satisfaction conditions. The important point is that, whatever that equivalent proof-theoretic condition might be, it will have to be go beyond valid and antivalid inferences.

It is also worth noting that sets of jointly satisfiable inferences, as an identity condition on logics, is more fine-grained than the usual properties used to identify logics, like the set of valid inferences or a counterexample relation.[11] One lesson we can learn from comparing M_∞ and classical logic (or comparing ST_∞ and classical logic) is that the valid and antivalid inferences of a logic do not by themselves tell us what consequences a set of axioms will have in the logic. M_∞ and classical logic have exactly the same valid inferences, yet in some cases they will give us different consequences for the same set of axioms in the same language.

If we were to look only at validities, or only at validities and antivalidities, then we would not have enough information to determine what we could or could not prove in the logic from non-logical axioms. So in order to understand how the logic behaves and what is or isn't provable in the logic, we need to look to more fine-grained details of the logic beyond just which inferences are valid. Valid inferences alone are too coarse-grained to tell us what follows from a set of axioms in the logic,

[10]Thanks to an anonymous reviewer for raising this issue.

[11]Thanks to an anonymous reviewer for raising this issue.

and are therefore too coarse-grained to tell us whether or not two formal systems are in fact the same logic. As such, we need to look to more fine-grained distinctions in order to determine whether two formal systems will give us the same consequences for the same set of axioms.

6 Paraconsistency and Nontransitive Consequence Relations

6.1 Is M_∞ Paraconsistent?

In [24], Scambler argues that the hierarchy logic based on ST introduced in [2] and [18] as a fully classical logic, which I will call ST_∞, is not in fact classical logic. He argues that, unlike classical logic, ST_∞ is really a paraconsistent logic: "In the case of [ST_∞], we have not really gotten rid of paraconsistency: we have merely thoroughly repressed it, so that it does not affect validity at any orders. Nevertheless, it is still present: there are valuations on which $p \wedge \neg p$ comes out valid" [24].

I take Scambler's point here to be something like the following: when we discuss paraconsistency, we often discuss it in terms of the validity or invalidity of the various rules of Explosion, like $A, \neg A \Rightarrow_1 B$ and $A \wedge \neg A \Rightarrow_1 B$. However, in these discussions, we are not interested in the validity of these schema for their own sake. Part of our interest in the validity or invalidity in these schema is that we take them to give us information as to whether or not the logical system can tolerate inconsistency. But ST_∞ and M_∞ can tolerate inconsistencies in their models: both logics have models at which both $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$ are satisfied. This, Scambler suggests, means that they are really paraconsistent logics.

Normally, a logic is called "paraconsistent" only if some version of Explosion (usually $A, \neg A \Rightarrow_1 B$ or $A \wedge \neg A \Rightarrow_1 B$) is invalid in that logic. But ST_∞ and M_∞ validate every rule of Explosion that classical logic validates, including the metainferential Explosion rule discussed in [3]. As such, ST_∞ and M_∞ are neither strongly nor weakly paraconsistent, in Hyde's terminology [12, 13].[12] So by any of the usual definitions of "paraconsistent", ST_∞ and M_∞ simply are not paraconsistent.

However, per Scambler's point, ST_∞ and M_∞ certainly have some paraconsistent-ish features. Both logics have models at which both $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$ are satisfied.[13] As a result, the logics can tolerate inconsistent axioms in a way that classical logic cannot. In classical logic, theories are closed under Explosion, in the following sense: if $A, \neg A \in T$ for theory T, then $B \in T$. Although Explosion is valid in ST_∞ and M_∞, theories are not closed under Explosion in these logics: there are theories containing A and $\neg A$ but not B, for some sentences A and B. So although ST_∞ and M_∞ are not paraconsistent by the usual definitions of paraconsistency, they should be considered at least psuedo-paraconsistent logics.

[12] Hyde credits this distinction to Arruda [1]. Equivalently, we could say in Ripley's terminology that the two logics are neither *conjunctively* nor *collectively* paraconsistent [21].

[13] It is worth noting that, although M_∞ has no models at which both A and $\neg A$ get value 0, L_∞ does. L_∞ is therefore in a similar situation: it validates the Law of Excluded Middle, yet has models that satisfy neither $\Rightarrow_1 A$ nor $\Rightarrow_1 \neg A$.

6.2 Scambler's Tortoise Objection to ST_∞ and M_∞

Scambler argues that this paraconsistency (or pseudo-paraconsistency) shows that ST_∞ (and by analogy, M_∞) cannot really be considered a presentation of classical logic. On this point, Scambler and I are in agreement. But Scambler further argues that there is something *wrong* with these logics. He argues that hierarchy logics like ST_∞ are not "closed under their own laws" in an important sense, and that this raises potential problems not just for any proponents of these logics, but for proponents of ST and similar logics as well.

Scambler [24] compares ST_∞ to Lewis Carroll's Tortoise [4]. The Tortoise accepts A, and B, and accepts C := "if A and B are true, Z must be true" but still does not accept Z. The Tortoise continues to accept statements of the form "if A and B and C and... are true, then Z must be true", but the Tortoise still refuses to accept Z.

I take the primary lesson of Carroll's paper to be that accepting or endorsing the *statement* of a rule is very different from actually *obeying* a rule. Asserting that an inference is valid is not the same thing as actually making that inference. Scambler's objection to ST_∞, and by extension to M_∞, is that these logics in effect "accept" classical inferences as valid, without in fact allowing us to *make* those inferences.

Scambler illustrates this issue by introducing a liar constant to the language, but we can make the same point without moving to a new language by looking at inconsistent axioms. ST_∞ and M_∞ both validate the Explosion Rule $A, \neg A \Rightarrow_1 B$, as well as the Metainferential Explosion Rule $\{\Rightarrow_1 A, \Rightarrow_1 \neg A\} \Rightarrow_2 \{\Rightarrow_1 B\}$. However, if we take A and $\neg A$ to be axioms (or $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$, in our sequent calculus presentation), we see that the theory generated by these axioms is not closed under explosion. M_∞, for example, has models in which $\Rightarrow_1 A$ and $\Rightarrow_1 \neg A$ are satisfied, but $\Rightarrow_1 B$ is not. The theory generated by these axioms in M_∞ therefore does *not* contain every sentence B whatsoever.

Scambler says that ST_∞ is not "closed under its own laws". For our purposes, we might instead say that theories in ST_∞ and M_∞ are not closed under valid inferences. The end result, however, is the same: the valid inferences of these logics do not necessarily correspond to *rules of inference* that we can use when proving theorems from non-logical axioms. Like Carroll's Tortoise, these logics accept the Explosion Rule, but do not allow us to infer *according to* the Explosion Rule.

Scambler [24] suggests that this poses a potential problem for these logics. ST_∞ and M_∞ seem to make precisely the same move that the Tortoise makes: endorsing rules that one does not follow. Defenders of these logics would therefore seem to be endorsing inferences without actually *making* those inferences. But then it is not clear exactly why one would want a logic that validates rules of inference that one cannot use. This certainly does appear to be a problem for these logics. Any defenders of ST_∞ and M_∞ would have to explain what purpose validating rules of inference that we cannot use might have.

6.3 The case of ST

Scambler argues that this is not just a problem for hierarchy logics like ST_∞ and M_∞; he argues that it also poses a problem for ST. ST has a nontransitive consequence relation for 1-inferences, but transitive consequence relations for every higher level of inference.[14] As Scambler puts it, "If the problem [with logics like ST_∞ and M_∞] is (as I suggested) that the logic is not closed under its own laws, then why isn't the fact that logics like [ST] also aren't closed under their laws similarly problematic? Don't we have essentially the same structure in each case?" [24].

Scambler's objection, I take it, is this: if nontransitive consequence relations are simply endorsing rules that they don't obey, then this seems like it will be a problem at *any* level. The defenders of ST, who endorse a nontransitive consequence relation at the level of 1-inferences, therefore have to explain why they are in any better position than Carroll's Tortoise, or the hierarchy logics that have nontransitive consequence relations at *every* level.

I agree with Scambler that the full hierarchy logics like ST_∞ and M_∞ seem to face a serious problem that defenders of those logics would have to address. But I do not think that this is *necessarily* a problem for defenders of ST.

In the case of ST, this is potentially a serious problem *if* we want to use the 1-inferences of ST to reason normally. Because ST holds premises to a different standard than it holds conclusions, valid 1-inferences do not preserve any nice properties like truth in a model (i.e. 0-inference satisfaction). In particular, this means that the theory of a set of axioms is not necessarily closed under the valid 1-inferences in ST.[15] The 1-inference $A, \neg A \Rightarrow_1 B$ is valid in ST, yet there are models of ST in which both A and $\neg A$ are satisfied 0-inferences but B is not. So if we try to use the valid 1-inferences of ST as rules of inference applied to axioms, we will end up "proving" sentences that do not actually follow from those axioms in ST.

As I said, this is *potentially* a serious problem for ST. However, this is not a problem if we want to use ST in a different way. For example, Ripley [22] presents a bilateralist interpretation of ST, according to which validity is understood in terms of assertion and denial.[16] On Ripley's bilateralist interpretation of ST, a 1-inference $\Gamma \Rightarrow_1 \Delta$ is valid iff the "position" of asserting all of the γs and denying all of the δs is incoherent. So according to Ripley, "$\Gamma \Rightarrow_1 \Delta$ can now be read as *the claim that* the position [of asserting the γs and denying the δs] is out of bounds." (emphasis mine; notation slightly altered) [22].

Thus on the bilateralist interpretation of ST, valid 1-inferences need not be understood as representing rules of inference that are safe to use. Rather, they should be understood as claims about what is impermissible to assert and deny. ST is then not a tool for reasoning about *sentences*, but a tool for reasoning about *positions*. On this interpretation, the 1-inferences are the claims *about which* we are making inferences; they do not themselves correspond to rules that we use to make inference.

---

[14] Recall that by "transitive", I mean that $\vDash_m \Gamma \Rightarrow_1 A$ and $\vDash_m A \Rightarrow_1 \Delta$ imply $\vDash_m \Gamma \Rightarrow \Delta$.

[15] Valid 1-inferences in ST *do* have the disjunctive property of either preserving what we might call "tolerant truth" (having value 1 or $\frac{1}{2}$) from left to right *or* preserving "tolerant untruth" (having value 0) from right to left. Unfortunately, this does not suffice to close theories under the valid 1-inferences. This is in part because the validity of an inference does not by itself tell us which property is preserved by that inference. As a result, neither property is preserved in all cases. Thanks to an anonymous reviewer for raising this issue.

[16] For earlier defenses of bilateralism independent of ST, see [23] and [19].

In some sense, this means that the 1-inferences of ST cannot be used in the way that we usually use inferences. But ST is still a perfectly usable logic for reasoning about positions, because the valid 2-inferences preserve 1-inference validity. Furthermore, the theory generated by any set of 1-inferences that we take as axioms will be closed under the (locally) valid 2-inferences of ST. If we take a set of 1-inferences as axioms, we can therefore use (locally) valid 2-inferences as rules of inference.

This interpretation of ST therefore avoids Scambler's objection, because it does not endorse rules that it refuses to obey. It obeys the 2-inference rules that it endorses, and it does not consider valid 1-inferences to be rules at all.[17]

However, this same approach will not work for ST_{∞} or M_{∞}. It is crucial to the bilateralist account of ST that we can understand metainferences in the usual way: as formal representations of rules of inference that we can use. Without that, it's not clear how we could use the logic *as a logic*. But ST_{∞} and M_{∞} have nontransitive consequence relations at *every level of inference*. As a result, there is no level of inference at which the valid inferences can be understood as rules of inference: there is no ordinal α at which theories are closed under all valid α-inferences. The bilateralist interpretation of ST reinterprets valid inferences at one level as claims instead of rules, but we can still use the valid inferences of the next level as rules of inference. In ST_{∞} and M_{∞}, *every* level has to be reinterpreted. This leaves no level at which valid inferences can be understood as rules of inference that it is safe to use. It's therefore not clear how we are to use these logics, if we are to use them.

There are many ways to use a formal construction. For example, in this paper I have used M_{∞} as an example in an argument for certain claims about identity criteria for logics. But in doing so, I wasn't really using M_{∞} *as a logic*. To use a logic as a logic, "from the inside" so to speak, we need to be able to use the formal construction as a tool for making inferences. I take Scambler's objection to hierarchy logics like ST_{∞} and M_{∞} to be that, due to the non-transitive nature of the logics, they cannot really be used as logics in this sense. In that, I agree. What purpose these logics might serve depends in part on how these logics can be usefully interpreted. At the moment, it is not clear how this is to be done. And this does indeed present a problem for any defenders of these logics.

7 Conclusion

I've argued that a logic cannot be identified by its valid and antivalid inferences, even at every inferential level. At a minimum, we suggest that we must also look to which sets of inferences are jointly satisfiable. Ultimately, we need to look to the rules of inference that the logic allows us to use. Two logics having exactly the same validities and antivalidities does not suffice to guarantee that the two logics obey the same rules of inference, or that they will allow us to prove the same consequences from the same set of axioms.

[17] This is not to say that the bilateralist interpretation is free from objections; only that it is free from this particular objection.

I take it that when we attempt to characterize logics and logical properties by the validity or invalidity of inferences, we often do so because we make certain assumptions about what those inferences represent. The recent development of mixed-condition consequence relations has demonstrated that these assumptions can be broken. In particular, the valid inferences of a logic can come apart from the rules of inference that the logic allows us to use. This means that instead of looking only at the inferences of a logic, we should be looking directly at the properties and rules of inference that the logic allows. Although inferences *can* represent these properties and rules in some settings, I take the results here and in [2] and [24] to show that inferences do not always do so.[18]

References

1. Arruda, A. In R. Priest (Ed.) (1989). *Norma aspects of the historical development of paraconsistent logic*. Philosophia: Munchen.
2. Barrio, E., Pailos, F., & Szmuc, D. (2019). A hierarchy of classical and paraconsistent logics. *Journal of Philosophical Logic*, *49*, 93–120.
3. Barrio, E., Pailos, F., & Szmuc, D. (2018). What is a paraconsistent logic?. In J. Malinowski, & W. Carnielli (Eds.) *Contradictions, from Consistency to Inconsistency*. Verlag: Springer.
4. Carroll, L. (1895). What the tortoise said to achilles. *Mind*, *4*(14), 278–280.
5. Cobreros, P. (2013). Vagueness: Subvaluationism. *Philosophy Compass*, *8*(5), 472–485.
6. Cobreros, P. (2011). Paraconsistent vagueness: a positive argument. *Synthese*, *183*(2), 211–227.
7. Cobreros, P., Egré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, *41*, 347–385.
8. Cobreros, P., Egré, P., Ripley, D., & van Rooij, R. (2013). Reaching Transparent Truth. *Mind*, *122*(488), 841–866.
9. Cobreros, P., Egré, P., Ripley, D., & van Rooij, R. (2012). Tolerance and mixed consequence in the S'valuationist setting. *Studia Logica*, *100*(4), 855–877.
10. Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*, 265–300.
11. Humberstone, L. (1996). Valuational semantics of rule derivability. *Journal of Philosophical Logic*, *25*(5), 451–461.
12. Hyde, D. (2008). *Vagueness logic and ontology*. Ashgate: Aldershot.
13. Hyde, D. (1997). From heaps and gaps to heaps of gluts. *Mind*, *106*(424), 641–660.
14. Hyde, D., & Colyvan, M. (2008). Paraconsistent Vagueness: Why not?. *Australasian Journal of Logic*, *6*, 107–121.
15. Keefe, R. (2008). Vagueness: Supervaluationism. *Philosophy Compass*, *3*, 315–324.
16. Keefe, R. (2000). *Theories of vagueness*. Cambridge: Cambridge University Press.
17. McGee, V., & McLaughlin, B. (1995). Distinctions without a difference. *Southern Journal of Philosophy*, *33*, 203–251.
18. Pailos, F. (2019). A fully classical truth theory characterized by substructural means. *The Review of Symbolic Logic*.
19. Restall, G. (2005). Multiple conclusions. In P. Hajek, L. Valdes-Villanueva, & D. Westerståhl (Eds.) *Logic, Methodology, and Philosophy of Science: Proceedings of the Twelfth International Congress* (pp. 189–205). London: Kings College Publications.

[18]Earlier versions of this paper were presented at the Buenos Aires Logic Group WIP Seminar, and at a workshop on substructural logics and hierarchies thereof hosted by the CUNY Logic and Metaphysics Workshop and the Saul Kripke Center. I am grateful to the audiences of these talks for their feedback. I am indebted to Eduardo Barrio, Paul Egré, Federico Pailos, Graham Priest, David Ripley, and Chris Scambler for many helpful discussions of these ideas and for all of their feedback on earlier drafts of the paper. Many thanks also to the two anonymous referees, whose helpful comments improved the paper significantly.

20. Ripley, D. (2018). On the ‘Transitivity’ of consequence relations. *Journal of Logic and Computation*, *28*(2), 433–450.
21. Ripley, D. (2015). Paraconsistent logic. *Journal of Philosophical Logic*, *44*, 771–780.
22. Ripley, D. (2013). Paradoxes and failures of cut. *Australasian Journal of Philosophy*, *91*(1), 139–164.
23. Rumfitt, I. (2000). Yes” and No”. *Mind*, *104*(436), 781–823.
24. Scambler, C. (2020). Classical logic and the strict tolerant hierarchy. *Journal of Philosophical Logic*, *49*, 351–370.
25. Scambler, C. (forthcoming). Transfinite Meta-inferences. Journal of Philosophical Logic. https://doi.org/10.1007/s10992-020-09548-7.
26. Van Fraassen, B. C. Singular Terms, Truth-Value Gaps, and Free Logic, (Vol. 63.
27. Varzi, A. (2007). Supervaluationism and its Logics. *Mind*, *116*, 633–676.
28. Varzi, A. (1997). Inconsistency without contradiction. *Notre Dame Journal of Formal Logic*, *38*, 621–640.
29. Williamson, T. (1994). *Vagueness*. London: Routledge.