RESPONSES

Respuestas

Daᴠɪᴅ Rɪᴘʟᴇʏ *ᵃ*
davewripley@gmail.com

ᵃ Philosophy Department, Monash University, Australia.

Abstract

These are the responses!

Key words: Paradox; Truth; Vagueness; Validity; Meaning.

Resumen

¡Estas son las respuestas!

Palabras clave: Paradojas; Verdad; Vaguedad; Validez; Significado.

Barrio & Pailos, "Why a Logic is not only its Set of Valid Inferences"

In both the precis and draft of *Uncut*, I say as little as possible about logic. I certainly use a lot of logical tools and techniques to develop the picture, but at the end of the day *Uncut* is simply not about logic. Rather, it is about meaning in natural language. I do use logic in *Uncut*, but I'm not engaged with questions *about* logic at all.

Despite this, a number of readers of the draft, and a number of audience members at presentations, have come away with the impression that I've made some claims about logic itself. As I do not share this impression, I'm often at a bit of a loss how to respond to it. I'm thankful to Barrio & Pailos for having written this impression down, so that I can take a crack in print at dispelling it. In terms of the main issue they raise, I'm afraid my response here can only be disappointing: I'm just going to say explicitly that I do not endorse the claims that Barrio & Pailos say that I implicitly endorse. But I clearly need to work to express this more clearly than I have yet managed to! So let me try again here.

Barrio & Pailos attribute to me the claim that 'a logic can be identified with the set of its validities,' that 'a logic is characterized

through its set of valid inferences,' that 'two different logics are the same logic if and only if their systems identify the same set of validities.' Barrio & Pailos do not say that I have actually made these claims, which is good, because I have not. Rather, they say that this view is 'implicit' in what I do say. I disagree.

What, then, about the textual evidence that Barrio & Pailos compile from the draft of *Uncut* to support the attribution of these claims? I have to admit some confusion on this point: the quotations they offer don't bear on these claims at all. The draft Barrio & Pailos take these quotations from explicitly defines '**CL**' as a particular sequent-style proof system, '⊢$_{CL}$' as the set of sequents derivable in this proof system, 'CFOLE' as the set of sequents that are valid in ordinary classical first-order logic with equality, and '⊩' as the set of sequents that are out of bounds. With these interpretations held in mind, the chosen quotations simply claim that the system "CL" derives exactly the valid sequents of classical first-order logic with equality, and that all the rules of **CL** preserve out-of-bounds-ness. The first of these claims is flat-out provable; it's just a usual claim of soundness and completeness for a particular proof system. The second claim is more controversial; as it is about our conversational norms, it's not the kind of thing that admits of proof in the same way. Neither of these claims, though, commits anyone to any views at all about what a logic is or when two logics are identical. There's just nothing here about the topics Barrio & Pailos are engaged with.

Metainferences in Uncut

However, it's not like it's accidental that *Uncut* looks in some depth at sets of sequents. I aim to give a precise and workable framework for theorizing about natural-language meaning in terms of positions, and when positions are in or out of bounds. Since positions are basically just sequents, thinking about which positions are out of bounds amounts to thinking about a set of sequents. So sets of sequents are definitely important to my project in *Uncut*.

Moreover, the metainferences such sets are or are not closed under also play an important role. For example, a compositional semantics for negation in this framework gives the conditions under which a position involving an assertion or denial of ¬A is out of bounds, in terms of the conditions under which a position involving an assertion or denial of A is out of bounds. *Uncut* puts forward the semantics according to which an assertion of ¬A is out of bounds whenever a denial of A is, and a denial of ¬A is out of bounds whenever an assertion of A is. Using [Γ⊩Δ]

to indicate that the position [Γ⇒Δ] is out of bounds, and using the double horizontal line for 'if and only if,' this semantics can be expressed like so:

$$\frac{[\Gamma \Vdash \Delta, A]}{[\neg A, \Gamma \Vdash \Delta]} \qquad \frac{[A, \Gamma \Vdash \Delta]}{[\Gamma \Vdash \Delta, \neg A]}$$

So metainferences too are definitely important to the project; they're the very terms in which I give compositional semantics to various pieces of vocabulary.

I don't see that metametainferences, or higher metainferences than that, have any particular importance to the project of giving semantics for natural language. I'm not particularly invested in this claim, though; indeed, I'd probably prefer it if it turns out to be wrong, since that would be more interesting! But for now that's how things look to me. (Barrio & Pailos certainly don't offer any arguments to the contrary; they don't consider natural language semantics at all, as far as I see.)

The formal tools I draw on and develop in *Uncut* reflect this approach. These tools work with sets of sequents and level-1 metainferences, and ignore higher metainferences entirely. If Barrio & Pailos are right about when logics are distinct, that means that my work in *Uncut* fails to focus in on any single logic. That's fine with me; I don't claim otherwise, and as far as I can tell nothing hangs on it.

Dittrich, "A Nontransitive Theory of Truth over PA"

In this paper, Dittrich shows the way forward for combining nontransitive theories of truth with rich and full arithmetic theories. This question is broached very briefly in Cobreros et al. (2013, sec. 3.4), but there it is handled model-theoretically, taking for granted the standard model of arithmetic. Dittrich's approach, in contrast, is proof-theoretic, involving sequent calculi that extend familiar cut-free sequent calculi for classical first-order logic with three additional kinds of rules: truth rules, arithmetic rules, and a restricted cut rule.

The setup is carefully constructed to allow for very strong theories of arithmetic to operate more or less as they usually do, while at the same time supporting various proof-theoretic arguments involving the truth predicate T. For this latter purpose, the proof systems I explore in *Uncut* will not work; they are too anarchic to admit the kind of convenient proof-theoretic analysis Dittrich brings. This is one reason why the main technical aruguments of *Uncut* are model-theoretic. Dittrich's systems,

by contrast, are tightly manicured in their logical and truth-theoretic components, and so support careful proof-theoretic reasoning.

The arithmetic content of these systems, however, is far out. Dittrich shows how to extend any classical theory of truth S to a naive theory of truth STAT[S], resolving the resulting paradoxes via nontransitivity, but retaining all the power of the underlying S.

One of the things that I find most interesting about nontransitive theories of truth is their potential to combine aspects of classical and nonclassical theories of truth. The consequence relations of these systems extend the consequence relations of classical logic, and so proof-theoretic approaches to the systems can draw a great deal on classical proof theory. On the other hand, models for these systems are more closely related to (in some cases by being identical to!) models for nonclassical theories of truth. This means that nontransitive theories provide a potential meeting ground, a place where insights both from so-called classical and nonclassical theories of truth can be combined. Dittrich's work, both in this paper and in the larger Dittrich (2020), reveals some of the potential benefits of this combination.

Lemma 1.2

Above I referred to STAT[S] as 'naive' rather than 'transparent.' By this I mean what Dittrich notes in their penultimate sentence: that the sequent $[\Rightarrow T\langle\phi\rangle\leftrightarrow\phi]$ is derivable in STAT[S] for every ϕ. This is already a lot; adding a full rule of cut to STAT[S] would render it trivial. However, it is not yet full transparency; and indeed, I'm not sure that STAT[S] really is fully transparent.

STAT is the set-set sequent calculus for $\mathcal{L}_T$ consisting of reflexivity, the rules $\neq$ and $=$,[1] the connective rules for $\neg, \supset, \forall$, the substitution rules *Sub*, the Robinson rules Q1–Q7, the induction rule *Ind*,[2] Cut for arithmetic formulas only, and the T-insertion rules TL and TR.

Lemma 1.2 claims that T-outL and T-outR are admissible in STAT, where these are the converses of TL and TR, respectively. I don't know this claim to be false, but I don't believe the given proof succeeds. Consider the following application of T-outL:

$$
\begin{array}{ll}
& \text{Ref:} \quad \dfrac{}{[Tx \Rightarrow Tx]} \\[2mm]
\text{SubL:} & \dfrac{}{[x = \langle p\rangle, T\langle p\rangle \Rightarrow Tx]} \\[2mm]
\text{T-outL:} & \dfrac{}{[x = \langle p\rangle, p \Rightarrow Tx]}
\end{array}
$$

---

[1] In these rules, I take s and t to be arbitrary *closed* terms.

[2] I take this rule to be unrestricted, including nonarithmetic formulas ϕ.

In terms of understanding Dittrich's given proof, a question arises: does the T⟨p⟩ that results here from the rule SubL count as principal or not? However, whether or not this occurrence counts as principal, the given proof does not show how to eliminate this use of T-outL. If this occurrence counts as principal, we are told 'the conclusion is immediate'; while this is certainly so if the occurrence is principal in an application of TL or TR, it is not so here. On the other hand, if this occurrence does not count as principal, we are told 'simply apply the Induction Hypothesis' to the line [Tx⇒Tx]. However, this line is not of the right form for this; it applies the truth predicate T to the *variable x*, where the rule T-outL we're working with only concerns applications of T to certain numerals.

In general, the derivations I worry about are those ending in one of the following forms:

$$
\text{SubL:}\quad \cfrac{\cfrac{\vdots}{[Ts, \Gamma \Rightarrow \Delta]}}{\cfrac{[s = \langle\phi\rangle, T\langle\phi\rangle, \Gamma \Rightarrow \Delta]}{[s = \langle\phi\rangle, \phi, \Gamma \Rightarrow \Delta]}} \text{T-outL:} \qquad\qquad \text{SubR:}\quad \cfrac{\cfrac{\vdots}{[\Gamma \Rightarrow \Delta, Ts]}}{\cfrac{[s = \langle\phi\rangle, \Gamma \Rightarrow \Delta, T\langle\phi\rangle]}{[s = \langle\phi\rangle, \Gamma \Rightarrow \Delta, \phi]}} \text{T-outR:}
$$

Here, the inductive step only seems to work in case the term s is the numeral of the code of ϕ. In other cases, either we can't apply the T-out rule at all (in case s is not the numeral of the code of any formula), or else the wrong result sequent is produced (in case s is the numeral of the code of some other formula).

Some of the time this is repairable. In particular, Dittrich (2020) shows how to handle the case where s is a closed term not denoting the code of ϕ; in this case the rule ≠ makes the needed final sequent an initial sequent.

In addition, for cases like the above example, in which the key term s is a variable, there is also a repair. STAT has the convenient (and usual) property that where σ is any substitution of arbitrary terms for variables, then the following rule is height-preserving admissible:[3]

$$
\text{Subv:}\quad \cfrac{[\Gamma \Rightarrow \Delta]}{[\sigma(\Gamma) \Rightarrow \sigma(\Delta)]}
$$

[3] This kind of rule has long proved key for proof theory for first-order languages; see e.g. Gentzen (1969, 3.10, 3.113.33). In *Uncut*, I include this rule as a primitive (and therefore derivable) rule of the proof system I call **CL**, in order to keep variables variabley.

Wlog, focus on the left form of the trouble case. If s is a variable, then since [Ts,Γ⇒Δ] is derivable, so is [T⟨φ⟩,σ(Γ)⇒σ(Δ)], where σ is the substitution of ⟨φ⟩ for the variable s. Moreover, this latter can be derived in the same number of steps. So we can apply the inductive hypothesis to get a proof of [φ,σ(Γ)⇒σ(Δ)]. Now repeated applications of SubL and SubR can bring us to a derivation of [t=⟨ψ⟩,ψ,Γ⇒Δ].[4]

Applying this strategy to the above example, and handling the initial sequent exactly as Dittrich explains, yields the following derivation:

$$
\begin{array}{ll}
\text{Ref:} & \dfrac{\rule{2.2cm}{0.4pt}}{[p \Rightarrow p]} \\[1.2em]
\text{TR:} & \dfrac{}{[p \Rightarrow T\langle p\rangle]} \\[1.2em]
\text{SubR:} & \dfrac{}{[x = \langle p\rangle, p \Rightarrow Tx]}
\end{array}
$$

Taking stock, we have seen that the proof can be made to work when s is the numeral of the code of φ, when s is a closed term not denoting the code of φ, and when s is a variable. This leaves two cases open: when s is a closed term denoting the code of φ but not the numeral for this code, and when s is an open term other than a variable. I don't have quick fixes for either of these cases, but neither do I have a demonstration that the Lemma fails (and indeed I wouldn't be too surprised if it holds). For now, then, I think we still don't know whether Dittrich's systems are truly transparent—and myself, I wouldn't want to guess.

Let me, then, briefly consider each of two possibilities: first, that these systems are indeed transparent, or that adding transparency to them would not yield triviality; and second, that they are not transparent, and adding transparency would yield triviality.[5]

First, let's suppose that Dittrich's systems really do feature a transparent truth predicate, or that transparency can be added to them without triviality. If this is so, then it seems to me Dittrich has built, or at least pointed the way towards, something very special: a transparent truth predicate added to classical arithmetic in a way that remains nontrivial while allowing for the formation of paradoxes of all sorts,

---

[4] Note that this makes the admissibility of the T-out rules not height-preserving, since multiple substitutions might be needed here. If height-preserving admissibility were desired, the substitution rules could be liberalized to allow multiple simultaneous substitutions.

[5] It's also possible that Dittrich's systems are not uniform in this regard: that various base systems S turn out to fall in different categories here. This too would be very interesting, but I won't comment directly on the possibility here.

and also allows us to reach the bewildering heights of proof-theoretic strength achieved by nontransparent axiomatic theories of truth.

Now, I myself happen not to care a great deal about proof-theoretic strength. I'm in it for the natural-language semantics, and the strength can fall where it may. But it's worth noticing that the central argument given against transparent truth by advocates of nontransparent axiomatic theories is based entirely on proof-theoretic strength.[6] This kind of strength may not matter for my purposes, but it certainly does for theirs.

That is, advocates of nontransparent axiomatic theories of truth don't really seem to value, make use of, or even really care about the nontransparency directly. They're optimizing for something else—proof-theoretic strength—and their key argument against transparent truth is that it gets in the way of this optimization. However, that argument has always proceeded by way of examples. If Dittrich's systems are transparent, or can be made to be, then they show that this conflict between transparency and proof-theoretic strength was only an appearance created by a particular choice of examples. Truth theorists concerned about proof-theoretic strength, even those willing to let proof-theoretic strength outweigh all other considerations, would for the first time have transparent systems to study as well. This would be a major breakthrough, and if you ask me which possibility I hope is correct about Dittrich's systems, it's this one.

Second, let's suppose that Dittrich's systems are not in fact transparent, and that adding transparency to them would result in triviality. This would also be a fascinating outcome. Recall that Dittrich's systems by definition include TL and TR, are fully naive, and are certainly not trivial. So we would still have a significant advance on the state of the art: the only other theories of truth we have with anything like this proof-theoretic strength are not only nontransparent, but also non-naive, and not closed under both TL and TR. So even if Dittrich's systems cannot be brought all the way to full transparency, they still provide a breakthrough of roughly the same kind.

However, if things do turn out to go this way, additional questions (and hopefully answers) arise as well. Just what is it that gets in the way of transparency? It's been thought that even naivete would get in the way of this kind of strength, and Dittrich's systems already show

[6] See for example Halbach (2011, sec. 20.1), which explicitly disavows multiple other arguments before making only this one. See also Halbach and Nicolai (2018); Picollo (2020); Fischer, Nicolai, and Dopico (2021).

us that this was a mistake. So if theorists turn out to have been right about transparency conflicting with this kind of strength, they must have been right for the wrong reasons. If we found out that Dittrich's systems could not be made transparent, this would hopefully set us on the path to appreciating the right reasons.

Whatever the situation turns out to be regarding the transparency of Dittrich's systems, then, I think this is exciting work, pointing the way towards a promising new family of systems for blending truth and arithmetic.

Ferguson, "Secrecy, Content, and Quantification"

In this paper, Ferguson raises a range of issues around the interpretation of the bounds in bounds consequence, and develops a fascinating treatment of non-weak-Kleene quantification in a weak-Kleene propositional base, all in the style of ST. Here, I divide my response into two halves: first, thinking a bit about how to interpret the bounds; and second, questioning the application of weak Kleene disjunction in looking at secrecy.

Which bounds?

By 'out of bounds,' Ferguson means something much broader than the understanding I put forward in *Uncut*. They understand a position to be out of bounds when 'to adopt this position is to flout conversational norms in some way.' That's definitely not the understanding of the bounds I develop in *Uncut*. So at least some of what Ferguson's article casts as disagreement seems to me much more like changing the subject.

The subject Ferguson raises—the totality of conversational norms—is certainly an interesting one! In a moment, I'll respond a bit to what they say about this notion. First, though, I think it's a good idea to take a crack at explaining how *Uncut* develops the idea of the bounds, since it certainly isn't this.[7]

The bounds I am concerned with in *Uncut* are certainly set by conversational norms, but not just any conversational norm is relevant. The idea is rather this: one important aspect of our conversational practices involves attributing positions to each other. We do this as a way to understand each other's point of view. Moreover, at least some conversational moves require this for their appropriateness; it's not

---

[7] The draft itself, as well as Ripley (2017a), develops these ideas as well.

felicitious to assert 'You're wrong; I didn't eat it' in a conversational context without attributing some commitment to 'You' to the effect that the speaker did eat it.

Attributing these positions to each other, though, is not simply a matter of keeping track of what's actually happened. We also hold each other to norms of coherence or fit. Sometimes a position simply doesn't fit together; it clashes with itself. The norms I'm concerned with are these norms of fit in particular.

I certainly agree with Ferguson that these norms don't exhaust all conversational norms. One way to see this is to look at how these norms are enforced. When someone seems to have adopted a position that doesn't fit together in the course of a conversation, they can be corrected. This correction can take many forms; I mention some of the most usual in the precis: reinterpretation (often by way of hypothesizing an ambiguity), asking for clarification, and outright dismissal. These corrections are all forms of norm enforcement.

These kinds of corrections, though, are not appropriate responses to violations of politeness or secrecy norms like the violations Ferguson focuses on. Supposing a context in which uttering 'fuck' violates a politeness norm, if A asserts 'Dave is a fuckhead' (or for that matter 'Dave is not a fuckhead'), they're clearly in violation of a conversational norm. Now, consider five possible responses someone might make to such an utterance:

1. I don't appreciate that kind of language.
2. I know what you mean, but do you have to say it like that?
3. Maybe you mean that Dave is lovely?
4. Do you mean always, or only at conferences?
5. That doesn't make any sense; you're just spouting garbage.

Of these five, only 1 and 2 would normally serve as appropriate attempts to enforce the politeness norm.

An attempt to reinterpret A's assertion, as in 3, is not really a sensible way to try to enforce the norm against saying 'fuck.' I note that there is a reading we can give to the sentence in 3 where it's a suggestion to A that they should communicate the same message as 'Dave is a fuckhead' through irony, by uttering 'Dave is lovely' ironically. On this reading, it can be an appropriate attempt to enforce the norm against saying 'fuck.' But on this reading, it's not an attempt at reinterpretation at all.

An attempt to clarify A's meaning, as in 4, doesn't do anything at all to enforce the norm against saying 'fuck'; it just takes A's utterance in stride. And outright dismissal, as in 5, is more or less the same; it also isn't an appropriate way to enforce the norm.

So politeness norms, like the norms Ferguson considers, are enforced through totally different mechanisms from the norms of coherence or fit I focus on in *Uncut*. Secrecy norms are enforced through different mechanisms yet again. Ferguson's conception of the bounds, then, is just a different, and much broader, conception than the one I'm focused on. I think it's a mistake to see this as a disagreement at all; there are just different topics here.[8]

Secrecy and weak Kleene

As I've pointed out above, Ferguson's topic—the totality of conversational norms—is not my topic in *Uncut*. But it's a fascinating topic, and it's exciting to see the logical work they develop to explore the idea.

I think Ferguson is right to say that usual analogies between quantifiers and connectives should fail, if we are focused on the totality of conversational norms. This is because the analogy relies on seeing *instances* of quantified sentences as analogous to *components* of compound sentences. But instances of quantified sentences do not *occur in* those quantified sentences, while the components of compound sentences do occur in the compounds. Since at least some conversational norms—for example politeness, as Ferguson points out—are sensitive to occurrence, they do not treat these cases alike.[9] Thus, any logical approach we develop should also not treat these cases alike.

However, I want to look in a bit more detail at Ferguson's password examples, and in particular how these examples relate to disjunction. I think they're too hasty to hold to a weak Kleene theory of disjunction here. The reason is that secrecy is not quite like politeness

[8] There is a separate misunderstanding also worth clearing up, but this one's quick. Ferguson attributes to me the view that 'the truth or falsity of the constituents of a position is sufficient to determine whether or not a position is in-bounds.' I have never said anything like this, and I certainly don't hold this view. Indeed, in both *Uncut* and the precis I put forward an example of context-dependence of the bounds that is straightforwardly incompatible with this view. Since I don't know why Ferguson attributes this to me, I don't think I have much more to say about it than that, though.

[9] This kind of case has a lot in common with those discussed in Freud (1925/1961). Here's Freud: "[T]he content of a repressed image or idea can make its way into consciousness, on condition that it is *negated*. Negation is a way of taking cognizance of what is repressed; indeed it is already a lifting of the repression, though not, of course, an acceptance of what is repressed…With the help of negation only one consequence of the process of repression is undone—the fact, namely, of the ideational content of what is repressed not reaching consciousness" (235-236). See also Ginn (2013).

in how it handles occurrences. Let me start by setting out some points of agreement: if the password is 'marriedcowboy,' then I agree with Ferguson that it violates secrecy norms to assert 'The password is "marriedcowboy",' to assert 'The password is not "marriedcowboy",' and to assert 'Either the password is "marriedcowboy" or it's "uglyduckling".' Moreover, it doesn't violate any secrecy norm to assert 'There exists a password.'

However, Ferguson maintains that it would violate secrecy norms to assert any finite disjunction of the form 'The password is X or the password is Y or the password is Z or …,' as long as 'marriedcowboy' is somewhere among X, Y, Z, …. This is not an incidental commitment; it follows directly from the weak Kleene treatment of disjunction. If secrecy *did* work like politeness, moreover, this would be true; if there's a norm against saying 'fuck,' then saying it once in the middle of a long disjunction still violates that norm.

However, real password systems feature only finitely many possible passwords. Often, a maximum password size will be fixed by the developer of the password system, but the mere fact that computer systems have to work with finite resources is sufficient to guarantee that there is always some maximum size for working passwords, no matter the system. Moreover, it's easy to arrive at upper bounds on the maximum size. For example, Hilbert and López (2011) estimates that in 2007, the total digital data storage capacity on Earth was around 300 exabytes, or 3×10^{20} bytes. Even today, no doubt, no password system supports passwords of this length.

Because of this, there are finite disjunctions of the form 'The password is X or the password is Y or the password is Z or …,' where X, Y, Z, … exhaust all possible passwords. Such disjunctions are not practically possible to assert, of course; for any realistic password system they come out way too long. But not all assertions of such disjunctions violate secrecy norms.[10] Indeed, asserting such a disjunction doesn't amount to much more than asserting 'There exists a password.' If that's right, then the weak Kleene treatment of disjunction Ferguson adopts doesn't correctly track the norms in play. Disjunctions over known finite domains really are a lot like particular quantifications.

---

[10] There are certainly some such disjunctions whose assertion might violate secrecy norms. For example, if 'marriedcowboy' is the password in the very first disjunct, or the very last, or some other salient location, or in every second disjunct, or something like that, then even asserting such an exhaustive disjunction might still have the effect of raising 'marriedcowboy' itself to salience, violating the secrecy norm. But most such disjunctions are not like this.

What logical tools could we use, then, to explore this issue? They won't be value-functional, I think, at least if there is only one value in use to represent 'norm-violating.' This is because at least some disjunctions of false claims with norm-violating claims are themselves norm-violating, as Ferguson points out, while at least some disjunctions of false claims with norm-violating claims are not, as I've pointed out. I won't enter into the project of trying to build a better model here; I just want to call attention to the fact that the weak Kleene approach seems to oversimplify the issue.

French, "Tolerance and the Bounds"

French introduces the notion of an *umbral position*: a position made of tolerant assertions and denials rather than strict. They define a relation of *shadowing* between umbral positions and positions: for an umbral position to shadow a position means that the position does not implicitly assert anything the umbral position tolerantly denies, and that the position does not implicitly deny anything the umbral position tolerantly asserts. That is, if we think of these umbral positions simply as ruling out certain strict acts, then for an umbral position to shadow a position is just for the position to obey the rulings-out made by the umbral position.

French goes on to show that, so defined, the umbral positions that shadow a given position obey a number of interesting constraints involving negation and conjunction that stem from the way the bounds on positions relate to these connectives. For example, the umbral positions $[\Gamma{:}\Delta,A]$ and $[\neg A,\Gamma{:}\Delta]$ shadow the very same positions as each other. This reflects one way in which the tolerant assertion of $\neg A$ is connected to the tolerant denial of A. This is analogous to, and depends on, but remains importantly distinct from, the connection imposed by the $\neg L$ rules between the strict assertion of $\neg A$ and the strict denial of A.

With this understanding in place, French defines a consequence relation $\Vdash^B$ as follows: $\Sigma \Vdash^B \Theta$ iff there is an in bounds position P not shadowed by $[\Sigma{:}\Theta]$. The idea, I take it, is that $\Vdash^B$ reflects a kind of 'being out of bounds' for umbral positions, with an umbral position counting as 'out of bounds' iff it rules out some in-bounds position.

This is an interesting partner relation to *Uncut*'s bounds consequence $\Vdash$. Where $\Gamma \Vdash \Delta$ means that the position $[\Gamma \Rightarrow \Delta]$ is so far gone as to be entirely out of bounds all on its own, $\Sigma \Vdash^B \Theta$ means that the umbral position $[\Sigma{:}\Theta]$ is just in a little bit of trouble: at the very least, it's not the kind of thing that can be taken entirely for granted, since

it rules out some in-bounds position. As French shows, $\Vdash^B$ exhibits a bunch of the kinds of properties exhibited by more familiar consequence relations. However, $\Vdash^B$ also exhibits a number of properties that might be unexpected for a consequence relation in this area, as French points out in Propositions 3.4 through 3.6.

A *different approach*

There's something a bit odd, I think, about counting an umbral position as out of bounds if there's just *some* in-bounds position it doesn't shadow. For example, take some claim p that can be either asserted or denied on its own without violating the bounds, like 'I like hamburgers' or something. That is, take some p where both [p⇒] and [⇒p] are in bounds. Note that [:p] does not shadow [p⇒], and that [p:] does not shadow [⇒p]. So on French's understanding, both [:p] and [p:] are out of bounds. This seems at least a bit weird: tolerant acts are meant to be weaker than their corresponding strict acts, and yet here we're counting strict acts as in bounds where their corresponding tolerant acts are not.

A different choice would be to count an umbral position as out of bounds when every in-bounds position isn't shadowed by it. Going this way would mean that we'd count an umbral position as out of bounds only when there's no in-bounds position that obeys its constraints. Supposing, as both French and I do, that the bounds on positions are closed under dilution, this is equivalent to counting an umbral position as out of bounds iff it does not shadow the empty position [⇒]—and that's the understanding I'll work with, since it's simpler.

On this understanding, tolerant acts indeed do come out as weaker than their corresponding strict acts. Suppose, for example, that [p:] is out of bounds in this new sense, that it does not shadow [⇒]. Then [⇒] implicitly denies p. Thus, [p⇒] must be equivalent to [p⇒p]; since the latter is out of bounds, so is the former. That is, when a tolerant assertion is out of bounds, its corresponding strict assertion is as well; the same goes for denials.

This also gives us a new consequence relation based on bounds on umbral positions. Say that $\Sigma \Vdash^z \Theta$ iff the umbral position [Σ:Θ] does not shadow the empty position. There are a lot of similarities between $\Vdash^B$ and $\Vdash^z$. For example, $\Vdash^z$ also obeys all the rules French displays for $\Vdash^B$. But there are also differences.

Both are not reflexive, but $\Vdash^z$ more so. We only have A$\Vdash^z$ A when the empty position implicitly asserts or denies A. It takes a lot for the empty position to implicitly assert A: that means that asserting A could

never bring any in-bounds position out of bounds, that assertions of A don't clash with anything at all. (Similarly if the empty position denies A.) Without getting too committal about the bounds, I think it's safe to say that most sentences are not like this.

Moreover, none of Propositions 3.4 through 3.6 hold of $\Vdash^z$. For example, suppose that the position [p⇒q] is in bounds. That gives no reason to think that either the empty position implicitly asserts p or it implicitly denies q. The bounds themselves might have nothing at all to say here. If that's so, however, then q⊮^{z} p, so Proposition 3.4 does not hold of $\Vdash^z$.

Proposition 3.5 shows that $\Vdash^B$ is not closed under cut. But $\Vdash^z$, on the other hand, is closed under cut, as long as we suppose that the empty position itself is in bounds. To see this, suppose that $\Sigma\Vdash^z$ Θ,A and A,$\Sigma\Vdash^z$ Θ. That is, neither [Σ:Θ,A] nor [A,Σ:Θ] shadows the empty position. So either: the empty position implicitly asserts some θ∈Θ, or it implicitly denies some σ∈Σ, or it both implicitly asserts and implicitly denies A. In either of the first two cases, [Σ:Θ] does not shadow the empty position for the same reason. In the third case, the empty position must be out of bounds already—which we've supposed is not the case. So $\Sigma\Vdash^z$ Θ.

Finally, Proposition 3.6 shows that ∧R↓ is not correct for $\Vdash^B$. But it is correct for $\Vdash^z$. To see this, suppose that $\Sigma\Vdash^z$ Θ,A and $\Sigma\Vdash^z$ Θ,B. That is, neither [Σ:Θ,A] nor [Σ:Θ,B] shadows the empty position. So either: the empty position implicitly asserts some θ∈Θ, or it implicitly denies some σ∈Σ, or it implicitly asserts both A and B. In either of the first two cases, [Σ:Θ,A∧B] does not shadow the empty position for the same reason. In the third case, the empty position must also implicitly assert A∧B,[11] and so again [Σ:Θ,A∧B] does not shadow it. Either way, then, $\Sigma\Vdash^z$ Θ,A∧B.

So I definitely agree with French's claim that '[if $\Vdash^B$] has a story to tell about when, for example, tolerant denial of a conjunction is out of bounds it cannot be the standard story.' But $\Vdash^z$ does not seem to be so limited. I think, then, that there is nothing in French's idea of umbral positions or shadowing that results in nonstandard approaches to the connectives. It's just that the standard approaches appear in the relation $\Vdash^z$, not (or not fully) in the relation $\Vdash^B$.

---

[11] Showing that implicit assertion itself allows for conjunction like this requires ∧L↑ to hold of bounds consequence.

Szmuc, "Inferentialism and relevance: the case of connexivity"

The system ST^c gives a fascinating lens on one way of thinking about the bounds, and the ways they might relate to consequence relations. I take a simple-minded approach in *Uncut*, looking at a consequence relation that matches the bounds exactly, in the sense that a position is in the relation iff it is out of bounds. Here, on the other hand, Szmuc develops ST^c as a consequence relation that contains only out-of-bounds positions, but does not contain all of them.

However, Szmuc offers a number of different understandings of ST^c, and I don't think these line up with each other. One of them is this:

$$\Gamma \underset{ST^c}{\vdash} \varDelta \text{ iff } \begin{cases} \text{there is no SK-valuation } v \text{ such that} \\ \quad v(A) = 1 \text{ for all } A \in \Gamma \text{ and } v(B) = 0 \text{ for all } B \in \varDelta, \text{ and} \\ \text{there is an SK-valuation } v \text{ such that } v(A) = 1 \text{ for all } A \in \Gamma, \text{ and} \\ \text{there is an SK-valuation } v \text{ such that } v(B) = 0 \text{ for all } B \in \varDelta \end{cases}$$

Another understanding Szmuc offers, though, is this:

$$\Gamma \underset{ST^c}{\vdash} \varDelta \text{ iff } \begin{cases} \text{it is out of bounds} \\ \quad \text{to both assert all the sentences in } \Gamma \text{ and deny all the sentences in } \varDelta, \\ \text{but it is not out of bounds to assert all the sentences in } \Gamma, \\ \text{nor is it out of bounds to deny all the sentences in } \varDelta \end{cases}$$

This is very different. These two understandings would align if the bounds on collections of assertions and denials were precisely given by the strong Kleene valuations. But that's not a plausible view about these bounds at all. For example, consider the example I use in the precis to illustrate the bounds: asserting 'Melbourne is bigger than Canberra' and 'Canberra is bigger than Darwin' while denying 'Melbourne is bigger than Darwin' is out of bounds. However, moving over into strong Kleene models, it's easy to find a model where the first-order sentences Bmc and *Bcd* both get the value 1 while the sentence *Bmd* gets the value 0. The example relies on bounds-related features of 'bigger' that are not reflected when we turn to the full space of strong Kleene models. So while Szmuc seems to treat these two understandings as interchangeable, they are not. Here, I want to look at each understanding in turn; I think they're both interesting and worth thinking about!

Strong Kleene models

First, the understanding in terms of strong Kleene models. Szmuc calls this 'a non-transitive subsystem of' the system P_S, where

P_S is defined just like ST^c except for using two-valued classical models instead of strong Kleene models. I think it's worth highlighting here Szmuc's footnote 15, and expanding on it a bit.

The main topic of that footnote has to do with what it takes to be 'non-transitive.' Unfortunately, that issue is seriously clouded by wild variations in terminology from author to author. To try to sort it out would be a bigger project than I can do here.[12] So I'll leave that issue to one side. The footnote does, however, highlight the importance of being precise about what vocabulary is present in the language before making claims about systems like ST^c, and that's what I want to focus on. This matters in particular here for getting clear on the relationship between ST^c and P_S.

If we work in an ordinary propositional or first-order language, ST^c and P_S validate exactly the same sequents as each other. The reasoning is just as for the systems Szmuc calls ST and CL; see for example Ripley (2012, Lemma 2.8) for that. So while it's true in this case that ST^c is a subsystem of P_S, it's perhaps misleading to say it that way without further comment; it's just as true that P_S is a subsystem of ST^c.

On the other hand, if we add a transparent truth predicate to the language, P_S becomes the empty logic: as there are no classical models at all with a transparent truth predicate, owing to the paradoxes, the conditions that must obtain for P_S validity are never met. ST^c does not collapse in this way. So in this case, it's not true that ST^c is a subsystem of P_S at all; it's rather the other way around.

The bounds

Whether we understand ST^c in terms of strong Kleene models or in terms of the bounds, it treats premises and conclusions very differently from each other. For example, working with the strong-Kleene-model understanding, we have $p \vdash_{ST^c} p,q$ but $\neg p, p \nvdash_{ST^c} q$. Both cases are cases where we cannot assign all the premises value 1 and all the conclusions value 0. But in the former case, the reason why has to do with a premise and a conclusion, while in the latter case, the reason has to do with two premises. ST^c is sensitive to this difference: basically, when the reason has to do with two premises (or two conclusions), it gets special treatment.

We can see a similar phenomenon working with the bounds understanding. Let me use a first-order language to represent

---

[12] I try at least a little bit in Ripley (2017b), but there I only consider a few of the available understandings.

English, with the predicate B for 'is bigger than,' and the names m,c,d for 'Melbourne,' 'Canberra,' 'Darwin,' respectively. Then we have Bmc,Bcd⊢$_{ST^c}$ Bmd. But we do not have Bmc,Bcd,¬Bmd⊢$_{ST^c}$. In either case, the positions represented are out of bounds: it is out of bounds to assert Bmc and Bcd while denying Bmd; and it is also out of bounds to assert all of Bmc,Bcd, and ¬Bmd. But because the latter case involves assertion only, ST^c treats it differently from the former, which involves both assertion and denial.

That is, ST^c on either understanding cares a great deal about the difference between having a claim as a conclusion and having the negation of that claim as a premise. The same kinds of differences obtain between having a claim as a premise and having its negation as a conclusion.

If we're working with strong Kleene models, I see no problem here; we're free to set up whatever kind of formal system we like, and this is a perfectly good one! But once we're interpreting these systems as telling us something about our conversational practices, I start to worry a bit about this phenomenon. In terms of the bounds, ST^c is sensitive to the difference between a claim's being denied and its negation being asserted, and likewise between a claim's being asserted and its negation being denied.

When it comes to actually recognizing speech acts, however, it can sometimes be difficult to tell the difference between a denial of A and an assertion of ¬A, or between an assertion of A and a denial of ¬A. Indeed, some have claimed that these are the very same acts: that every denial of A is at the same time an assertion of ¬A, and that every assertion of A is at the same time a denial of ¬A.[13]

Let's suppose for a moment that that's right. Then ST^c's treating assertions so separately from denials would seem unmotivated. But there is a relation in the area that might be more natural, suggested by Szmuc's remarks about 'relevant assertion' and 'relevant denial.' Say that position is *tight* iff it's out of bounds and all its proper subpositions are in bounds. Then a tight position is one in which every assertion and every denial is relevant, in Szmuc's sense. For example, the position containing assertions of 'Melbourne is bigger than Canberra' an 'Canberra is bigger than Darwin' and a denial of 'Melbourne is bigger than Darwin' is a tight position; it's out of bounds, but removing any one of these acts leaves a position that's in bounds.

Now, say that Γ⊢$_{ST^t}$ Δ iff the position [Γ⇒Δ] is tight. ST^t seems to fit better with the motivations Szmuc gives than ST^c does. In particular,

[13] See for example, (Frege 1897/1979; Price 1983; Ripley 2020).

as long as the bounds do not distinguish between an assertion of ¬A and a denial of A, then ST^t will not either. In this regard, it's unlike ST^c. However, it seems to sit better with Szmuc's motivation in terms of relevant assertions and denials than ST^c does.

Teijeiro, "Strength and Stability"

Teijeiro's paper considers in depth a number of conceptions of validity for metainferences. I fully agree with the paper's closing remark: 'it is important not to get caught in false controversies.' There is no need to get caught up (as for example Dicher & Paoli (2019, sec. 18.3.2) seems to) in questions about what the best notion of metainferential validity is, whatever that might mean. Instead, it's important to understand the different notions available and how they connect to each other, as well as how they connect to other notions of interest. Teijeiro's paper helps us to do just this.

In particular, I want to call attention to Tejeiro's Fact 4.1. This fact strikes me as important and well worth knowing about, and it's helped me to make more precise something I've been worried about since I started working on nontransitive extensions of classical logic. In the rest of this response, I want to say a bit about why I think Fact 4.1 is important and explain how it's helped me to make precise and I think nonarbitrary sense of a distinction I had previously struggled to see.

A question

The issue is this: there's clearly *something* nonclassical about logical approaches to vagueness and truth like those I explore in *Uncut*. At the very least, that's an initial response that a number of people have had to these approaches, and one that I think is worth understanding. But how can we understand it precisely?

Again, it's important not to get caught in false controversies: the question here isn't about what *really* counts as a classical theory of vagueness or truth. That's a pure terminological question, which everyone is free to settle by stipulation as they choose. Rather, the question here is about diagnosing a certain kind of informed but inarticulate intuition: is there an interesting distinction that these intuitions are tracking?

There is of course an obvious sense of 'nonclassical' available, call it 'nonclassical$_1$': a logical system is nonclassical$_1$ iff it fails to validate

some classically-valid argument. This can't be a way of understanding the feeling of nonclassicality, though, since the nontransitive systems in question are all classical$_1$.

In Cobreros et al. (2013) and Ripley (2013), my collaborators and I began to articulate a different way to see these systems as nonclassical, based on metainferences. While the validities of these systems are closed under many of the same metainferences as validity in classical logic, they do not match perfectly. Most notably, of course, validity in classical logic is closed under the metainference cut, unlike validity in nontransitive systems. So this gives us a second sense, call it 'nonclassical$_2$': a logical system is nonclassical$_2$ iff it fails to be closed under some classical metainference, where a classical metainference is a metainference obeyed by pure classical logic. Owing to the failure of cut (among many others), the nontransitive systems in question are certainly nonclassical$_2$.

The trouble with understanding nonclassicality in this way, as we point out in Cobreros et al. (2013), is that *every* nontrivial extension of classical logic fails to obey *some* metainference obeyed by classical logic itself. So sure, nontransitive theories of truth are nonclassical$_2$, but so is classical Peano arithmetic.[14] This too, then, fails to articulate the feeling.

Still, it seemed that metainferences matter somehow. So given a set Σ of metainferences obeyed by classical logic, let a logical system be classical$_Σ$ iff it obeys every metainference in Σ. Now we've got uncountably many distinct notions of classicality. This might help us precisely articulate all kinds of things, but it still doesn't answer the initial question: what's driving the robust intuition of nonclassicality here? Any choice of a particular Σ seems arbitrary.

The answer, I now think, is in Teijeiro's paper, and in particular in Fact 4.1. (The answer is probably also implicit in Barrio, Rosenblatt, & Tajer (2015), but I see that only with hindsight.) It is the schematic metainferences that provide a particularly relevant choice of Σ.

Schematic classicality and Boolean valuations

Here, I'll work with a propositional language and SET-SET sequents. Let a metainference be *schematically classical* if classical

---

[14] After all, classical logic is closed under the metainference from [⇒S0+S0=SS0] to [⇒0=S0], since neither is classically valid; but PA is certainly not closed under this metainference.

logic obeys every substitution instance of the metainference, and let a set of sequents be be *schematically classical* iff it is closed under all schematically classical metainferences.

There turns out to be an interesting relationship between the schematically classical sets of sequents and ordinary two-valued Boolean valuations. Half of this relationship is provided by Teijeiro's Fact 4.1, and the other half by Humberstone (2011, sec. 1.16). I'll state the relationship here and then briefly explain how this helped me to understand the intuitions indicated above.

Let a *valuation* be any function from the language to the set {1,0}. Say that a valuation is Boolean iff it obeys the usual constraints connecting values of compound sentences to values of their components.[15] Say that a valuation v is a counterexample to the sequent [Γ⇒Δ] iff v(γ)=1 for every γ∈Γ and $v(\delta)=0$ for every δ∈Δ. Given any set A of sequents, let V(A) be the set of all valuations that are not counterexamples to any sequent in A; and let B(A) be the set of all Boolean valuations in V(A). Finally, given any set V of valuations, let L(V) be the set of sequents that do not have any counterexamples in V.

With all that in mind, the following holds:

Fact 1 (Teijeiro). *For any set B of Boolean valuations, L(B) is schematically classical.*

Proof. By Teijeiro's Fact 4.1, every metainference that is schematically classical is also locally valid for Boolean valuations. Thus, each such metainference is globally valid for B, and so L(B) is closed under them all. □

Fact 2 (Humberstone). *For any schematically classical set A of sequents, A=L(B(A)).*

Proof. While Humberstone (2011, sec. 1.16) does not directly make this claim, all the materials needed to prove it are there. Let A be schematically classical. Then, in particular, A is a 'generalised consequence relation' in the sense of Humberstone (2011, p. 73), since this requires only obeying certain schematically classical metainferences. Moreover, A is '#-classical' for every # ∈{∧,∨,¬,⊤,⊥} (see definitions on

---

[15] For example, if the valuation v is Boolean we must have $v(φ∧ψ)=$min $(v(φ),v(ψ))$ and so forth, while if u is not Boolean we can't conclude anything in general about relationships between $u(φ∧ψ)$, $u(φ)$, and $u(ψ)$.

Humberstone (2011, pp. 62, 76)), since this too requires only obeying certain schematically classical metainferences.

By a somewhat well-known fact, we have A=L(V(A)), simply because A is a generalised consequence relation. By a less well-known fact, we can conclude that V(A)⊆B(A), by Humberstone (2011, Theorem 1.16.6), drawing on the fact that A is #-classical for every # ∈{∧,∨,¬,⊤,⊥}. Since B(A) is defined to be a subset of V(A), we can conclude from this that B(A)=V(A), and thus that A=L(B(A)). □

Putting Facts 1 and 2 together, we see that the schematically classical sets of sequents are exactly those sets of sequents determined by sets of Boolean valuations. Every set of Boolean valuations determines a schematically classical set of sequents; and every schematically classical set of sequents is determined by some set of Boolean valuations.

This gives us two ways to identify these sets of sequents. One is purely consequence-theoretic, looking only at which sequents are contained in a set, and checking for closure under particular metainferences. The other is valuation-theoretic, working instead just with sets of Boolean valuations, and letting validity fall where it may.

Via the connections above, we can see a purely consequence-theoretic reflection of the use of Boolean valuations only. It's also certainly true that many people associate classicality particularly with the use of Boolean valuations. So it would make sense for such people to use 'classical,' as applied to consequence relations, in a way restricted to *schematically* classical sets of sequents. This, I think, is what's driving at least some of the intuitions of nonclassicality I've run into over the years. I'm thankful to Teijeiro (and Humberstone) for having developed and demonstrated the facts needed to help me see this precisely.

References

Barrio, E., Rosenblatt, L., & Tajer; D. (2015). The logics of strict-tolerant logic. *Journal of Philosophical Logic, 44*(5), 551-571.

Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2013). Reaching transparent truth. *Mind, 122*(488), 841-866.

Dicher, B., & Paoli, F. (2019). ST, LP, and tolerant metainferences. In C. Başkent & T. M. Ferguson (Eds.), *Graham Priest on dialetheism and paraconsistency* (pp. 383-408). Springer.

Dittrich, J. (2020). Paradox, arithmetic, and nontransitive logic. PhD thesis, Ludwig Maximilians Universität München.

Dunn, J. M., & Hardegree, G. M. (2001). *Algebraic methods in*

philosophical logic. Oxford University Press.

Fischer, M., Nicolai, C., & Dopico, P. (2021). Nonclassical truth with classical strength: A proof-theoretic analysis of compositional truth over HYPE. *The Review of Symbolic Logic,* 1-24. https://doi.org/10.1017/S1755020321000137

Frege, G. (1897/1979. Logic. In H. Hermes, F. Kambartel & F. Kaulbach (Eds.), *Posthumous writings* (pp. 127-151). Basil Blackwell.

French, R., & Ripley, D. (2019). Valuations: bi, tri, and tetra. *Studia Logica, 107*(6), 1313-1346.

Freud, S. (1925/1961). Negation (1925). In J. Strachey, A. Freud, A. Strachey & A. Tyson (Eds.), *The Ego and the Id and other works* (vol. 19, pp. 235-239). The standard edition of the complete psychological works of Sigmund Freud. The Hogarth Press.

Gentzen, G. (1969). Investigations into logical deduction (1935). In M. E. Szabo (Ed.), *The Collected papers of Gerhard Gentzen* (pp. 68-131). North-Holland.

Ginn, M. [@shutupmikeginn]. (2013, November 21). *My "Not involved in human trafficking" T-shirt has people asking a lot of questions already answered by my shirt.* [Tweet]. Twitter. https://web.archive.org/web/20210605083513/https://twitter.com/shutupmikeginn/status/403359911481839617

Halbach, V. (2011). *Axiomatic theories of truth*. Cambridge University Press.

Halbach, V., & Nicolai, C. (2018). On the costs of nonclassical logic. *Journal of Philosophical Logic, 47*(2), 227-257.

Hilbert, M., & López, P. (2011). The world's technological capacity to store, communicate, and compute information. *Science, 332*(6025), 60-65. https://doi.org/10.1126/science.1200970.

Humberstone, L. (2011). *The connectives*. The MIT Press.

Picollo, L. (2020). Truth in a logic of formal inconsistency: How classical can it get? *Logic Journal of the IGPL, 28*(5), 771-806.

Price, H. (1983). Sense, assertion, Dummett, and denial. *Mind, 92*(366), 161-173.

Ripley, D. (2012). Conservatively extending classical logic with transparent truth. *Review of Symbolic Logic, 5*(2), 354-378.

Ripley, D. (2013). Revising up: Strengthening classical logic in the face of paradox. *Philosophers' Imprint, 13*(5), 1-13.

Ripley, D. (2017a). Bilateralism, coherence, warrant. In F. Moltmann & M. Textor (Eds.), *Act-based conceptions of propositional content: Contemporary and historical perspectives* (pp. 307-324). Oxford University Press.

Ripley, D. (2017b). On the 'transitivity' of consequence relations. *Journal of Logic and Computation, 28*(2), 433–450.

Ripley, D. (2020). Denial. In V. Déprez & M. T. Espinal (Eds.), *The Oxford handbook of negation* (pp. 47-57). Oxford University Press.

Scott, D. (1974). Completeness and axiomatizability in many-valued logic. In L. Henkin (Ed.), *Proceedings of the Tarski symposium* (pp. 411–436). American Mathematical Society.

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