STEVEN GROSS

TRIVALENT SEMANTICS AND THE VAGUELY VAGUE

ABSTRACT. Michael Tye responds to the problem of higher-order vagueness for his trivalent semantics by maintaining that truth-value predicates are “vaguely vague”: it’s indeterminate, on his view, whether they have borderline cases and therefore indeterminate whether every sentence is true, false, or indefinite. Rosanna Keefe objects (1) that Tye’s argument for this claim tacitly assumes that every sentence *is* true, false, or indefinite, and (2) that the conclusion is any case not viable. I argue – *contra* (1) – that Tye’s argument needn’t make that assumption. A version of her objection is in fact better directed against *other* arguments Tye advances, though Tye can absorb this criticism without abandoning his position’s core. On the other hand, Keefe’s *second* objection does hit the mark: embracing ‘vaguely vague’ truth-value predicates undermines Tye’s ability to support validity claims needed to defend his position. To see this, however, we must develop Keefe’s remarks further than she does.

1. INTRODUCTION

Any trivalent semantics for languages with vague predicates must deal with the *prima facie* problem of higher-order vagueness. One way of pressing the problem is as follows. What motivates the third truth-value, according to the proponent of trivalent semantics, is the existence of borderline cases for vague predicates and the corresponding absence of a sharp transition from true to false applications of the predicate in a sorites series. But it would seem that there are borderline cases of the truth-value predicates themselves and a corresponding absence of sharp transitions both from true to indefinite and from indefinite to false applications of the vague predicate in a sorites series. This, however, would seem to provide equal motivation for the positing of even further truth-values, and so trivalence should be denied.

Michael Tye’s response distinguishes vague predicates and *vaguely* vague predicates.[1] The former have borderline cases, cases (according to his gloss) such that it’s indeterminate whether the predicate applies. With the latter, it’s indeterminate whether they could have borderline cases – thus, indeterminate whether there could be cases

Synthese (2007) 156:97–117
DOI 10.1007/s11229-005-4542-9

such that it's indeterminate whether the predicate applies. Tye claims that truth-value predicates are vaguely vague. On his account, it's not true that they have borderline cases – though also not true that they don't. Tye thus asserts neither the presence nor the absence of sharp transitions from true to indefinite and from indefinite to false applications of a vague predicate in a sorites series. More generally, he holds that it's indeterminate whether all sentences are true, false, or indefinite – and thus not true, though not false, that his semantics is in this sense *exhaustive*.

The problem of higher-order vagueness has been a major stumbling block for a variety of approaches to vagueness. It's thus important to determine whether a line like Tye's could succeed. In this paper, I discuss two objections raised by Rosanna Keefe (2000, 121–122), rejecting the first but defending a development of the second.

Keefe's first objection is that Tye tacitly assumes that every sentence is either true, false, or indefinite in arguing that it's *indeterminate* whether every sentence is true, false, or indefinite: in defending his position, Tye deploys an assumption that is inconsistent with it. But I argue that Tye can defend his claim without making this assumption. Keefe's first objection is in fact better directed against *other* arguments Tye makes. Even here, however, Tye can absorb the criticism without abandoning his position's core.

Keefe's second objection is that it's in any event not viable to maintain that it's indeterminate whether every sentence is true, false, or indefinite, since the use of a three-valued logic requires that this classification be exhaustive. I argue that this objection indeed ultimately presents a fundamental challenge to Tye: if it's not true that Tye's semantics is exhaustive, it has difficulty supporting assumptions about validity needed to develop and defend the position itself. To bring this out, I develop Keefe's remarks further than she does and reply to some possible responses.[2]

2. TYE'S SEMANTICS

It will be useful first to rehearse the main elements of Tye's semantics for a language containing vague predicates – a variant of a truth-functional account using strong-Kleene tables.

On Tye's view, a monadic singular sentence Fc is true iff i_c belongs to S; false iff i_c belongs to S'; and indefinite iff it's indeterminate whether i_c belongs to S (or to S') – where i_c is the object in domain D assigned to individual constant c, and S and S' are the extension and counter-extension, respectively, of F. Note that S and S' are vague sets if F is vague – where a set is vague iff, first, it has borderline cases (that is, there are objects such that it's indeterminate whether they are members), and, second, it's indeterminate whether there are objects that are neither members, borderline members, nor non-members.

Quantifiers are introduced as follows: $(Ex)Fx$ is true if Fx is true for some assignment of an object of D to x, false if Fx is false for all assignments, and indefinite otherwise; $(x)Fx$ is true if Fx is true for all assignments of objects of D to x, false if Fx is false for some assignments, and indefinite otherwise.

Tye presents the semantics for logical connectives in the form of strong-Kleene tables, motivated by five principles. (1) The negation of a statement of given truth-value is its opposite in truth-value. (Indefiniteness is its own opposite.) (2) A conjunction is true if both its conjuncts are true and false if either conjunct is false; otherwise it's indefinite. (3) A disjunction is true if either disjunct is true and false if both disjuncts are false; otherwise it's indefinite. (4) The truth-value of $P \rightarrow Q$ is to be the same as that of $\sim P v Q$. (5) The truth-value of $P \leftrightarrow Q$ is to be the same as that of $(P \rightarrow Q)$ & $(Q \rightarrow P)$.

About the sentence-operator 'it's indeterminate whether,' we are told that 'It's indeterminate whether P' is equivalent to 'It's not determinate that P, and it's not determinate that not-P;' and that 'It's determinate that P' is true if P is true, and is false if P is false or indefinite.

Finally, an argument is valid so long as it can't be true that, when the premises are true, the conclusion is anything other than true.[3]

Tye's semantics is underdeveloped in various ways. One difficulty presents itself immediately. What his semantic claims *say* remains unclear to the extent that it's unclear how to understand the conditionals and biconditionals used therein. In his first presentation, Tye (1990, 545, n. 21) states that they express entailment relations. In a later recapitulation of his views, however, he only precludes a classical interpretation, leaving open how they *are* to be understood. (Tye 1994a, 203–204, n. 18/285, n. 8) One *possibility* he mentions there is understanding them in terms of the semantics for conditionals and biconditionals given by (4) and (5) above. So understood, however, the semantic claims are not true. For example, if either sentence flanking such a biconditional is indefinite, then the biconditional itself is indefinite and thus not true. And Tye does allow – indeed, it's central to his view – that 'i_c belongs to S'' (a constituent of the biconditional truth-clause for monadic singular sentences) can be indefinite.[4] But if the semantic claims are not true, then it's hard to save the appearance that Tye presents them *assertorically*.

I want to allow the purveyor of vaguely vague truth-value predicates to assert these semantic claims – for instance, in attempting to close argumentative gaps.[5] I will therefore assume that the semantic claims express entailment relations, in the sense of relations of logical implication. That is, I will understand claims of the form 'if A, then B' in Tye's semantics to express that B follows from A, that an inference from A to B is valid. My intention is to give the position its best shot and yet to show that, even conceding just this much, it is open to objection. If a defender of vaguely vague truth-value predicates would want to finger this minimal assumption as the source of the troubles discussed below, he would have to shoulder the burden of developing an alternative construal that does not itself lead to trouble.

3. KEEFE'S FIRST OBJECTION

3.1. *The First Objection*

Tye maintains that the truth-value predicates are vaguely vague. If the truth-value predicates are vaguely vague, then it's indeterminate whether they have borderline cases. So, if they are vaguely vague, then it's indeterminate whether there's a sentence that's neither true, false, nor indefinite; and thus:

(*) Every sentence is true, false, or indefinite

is indefinite.

If (*) is indefinite, then Tye certainly can't appeal to (*) in defending and developing his position. But Keefe objects that Tye does just this – in particular, in arguing for (*)'s indefiniteness:

> He attempts to show that (*) must be indefinite by the following argument. If true, (*) would commit us to sharp boundaries between the true sentences and the indefinite ones; but there are no sharp boundaries (as the series of predications along a sorites series illustrates), so (*) cannot be true. On the other hand, he continues, its falsity would require the introduction of a new truth-value. Therefore (*) is not true or false, so must be indefinite. Notice, however, that in

this very argument he assumes (*). For to argue that a sentence must be indefinite because it is neither true nor false is to make the very assumption in question. But then he is absurdly relying on assuming that (*) is true to argue that it is indefinite.[6]

3.2. *Reply to Keefe's First Objection*

Keefe claims that arguing for (*)'s indefiniteness from its non-truth and non-falsity assumes (*). But then Keefe must be assuming that nothing else can plausibly bridge the gap between premises and conclusion. An obvious strategy in response is to supply an alternative that's not itself objectionable – in particular, that doesn't entail (*). So, what alternative to invoking (*) might one offer in arguing as above for (*)'s indefiniteness?

Recall Tye's semantics for universal generalizations: $(x)Fx$ is true if Fx is true for all assignments of objects of D to x, false if Fx is false for some assignments, and indefinite otherwise. Tye makes clear that his "otherwise" condition is to be so understood that, if it's *not true* that Fx is true for all assignments of objects of D to x, and *not true* that Fx is false for some assignments, then the "otherwise" condition obtains, and so $(x)Fx$ is indefinite.[7]

Now, (*) is a universal generalization. So, one might argue that (*)'s indefiniteness can be inferred from its non-truth and non-falsity as follows. Suppose (*) is not true. Then it follows that it's not true that '*x* is a sentence and is true, false, or indefinite' is true for all assignments of objects of D to x. (Recall that I'm construing Tye's meta-linguistic (bi)conditionals as expressing inferential relations.) Suppose (*)'s not false. Then it follows that it's not true that '*x* is a sentence and is true, false, or indefinite' is false for some assignments. Thus, on the supposition that (*) is neither true nor false, the "otherwise" condition obtains for '*x* is a sentence and is true, false, or indefinite – and it follows that (*) is indefinite.

This line of reasoning makes two assumptions: first, that Tye's semantics for universal generalizations – or something parallel – holds for meta-linguistic claims like (*); second, that, if A entails B, then not-B entails the non-truth of A. Pending argument to the contrary, the first assumption seems reasonable, as Tye certainly allows indefinite meta-level universal generalizations, (*) being a case in point. The second assumption is needed to infer, from the non-truth of (*), that it's not true that its condition for truth obtains—similarly for the inference from its non-falsity. The assumed inference is indeed valid if *Ersatz Reductio ad Absurdam* is. *Ersatz Reductio* licenses an inference from a set of premises to the *non-truth* of A if A and those premises entail a contradiction. Unlike classical *Reductio*, *Ersatz Reductio* is valid in an exhaustive strong-Kleene system with truth as designated value. We'll ask later whether it's valid in Tye's system. What's relevant here, however, is that Keefe doesn't challenge Tye on this point – and indeed perhaps tacitly concedes *Ersatz Reductio* to him. For Tye employs *Ersatz Reductio* in his arguments for the non-truth and non-falsity of (*), which Keefe does not question. She questions only the step from these claims to (*)'s indefiniteness.[8] Her objection thus does not disallow that one may conclude that (*) is indefinite by appeal, not to (*) itself, but to the semantics for universal generalizations.[9]

3.3. *Keefe's First Objection Elsewhere*

Keefe (2000, 122) notes that Tye's argumentation elsewhere also seems to assume (*). The further example she herself supplies – that he concludes, from its non-truth and non-falsity, that the sentence 'every conditional in a sequence of premises of a conditionally-formulated sorites paradox is either true, false, or indefinite' is indefinite – can be handled as above. However, this is not the case with all the arguments Tye presents. It's worthwhile to bring this out, not only to show that Keefe's first objection applied elsewhere does indeed require Tye to adjust his defense of his position, but also because it will enable us to put on the table a crucial contention of Tye's that will play a role in our discussion in the next section.

Tye claims that it is indeterminate whether there are sentences of the form 'Mi is true' that are indefinite (where what replaces 'Mi' refers to an atomic sentence ascribing the relevant vague predicate to the ith object in some object-level sorites series) – similarly for other singular monadic truth-value imputations. Let's label the claim that there's an indefinite truth-imputation '(IT)' – so, Tye's claim is that (IT) is indefinite. (IT)'s indefiniteness plays a central role in Tye's development of his position: he invokes it in defending his response to meta-linguistic versions of the sorites, and it's also deployed in arguing for the non-falsity of the claim that 'is true' is vague.[10] But it's unclear how to save Tye's *argument* for (IT)'s indefiniteness from the need to assume (*).

Tye argues for (IT)'s indefiniteness as follows:

> ... (1) suppose that there is a statement "*Mj* is true" that is indefinite. (2) Then it cannot be true that *Mj* itself is either true or false or indefinite. (3) So it is not true that there is a statement "*Mj* is true" that is indefinite. But neither is it false that there is such a statement. For [(4) suppose it were false ...] then (5) every statement of the type "*Mn* is true" would be either true or false, with the result that (6) there would be a sharp transition from the true statements of the type "*Mn* is true" to the false ones. Intuitively, (7) it is not true that there are such transitions. [So, repeating the remark just above, (8) it's not false that there is such a statement.] So (9) it is, I maintain, indeterminate whether there are statements of the type "*Mn* is true" that are indefinite.[11]

The argument clearly falls into three parts: (1)–(3) is an argument for the non-truth of (IT); (4)–(8) is an argument for the non-falsity of (IT); and then it's concluded at (9) that (IT) is indefinite. The step to (9), on the basis of (3) and (8), is the first of three places that the argument can seem to assume (*). But this inference can again be defended as above.

The argument for (IT)'s non-falsity can seem problematic in two places. First, how does (8) follow from the inconsistency of (6) and (7)? Tye's semantics allows counter-instances to *Reductio*, and *Ersatz Reductio* only allows one to conclude that it's *not true* that it's false that there's an indefinite truth-imputation. It seems reasonable, however, that, *if* one accepts *Ersatz Reductio*, one should also allow Negative *Ersatz Reductio*: if not-A and some premises entail a contradiction, then, from those premises one may infer the non-falsity of A (which is truth-preserving according to exhaustive strong-Kleene tables). But then 'It's not true that it's false that A' entails 'It's not false that A'.[12] In other words, what we may call Restricted *Reductio* F – that is, *Reductio* where the contradiction is blamed on an imputation of falsity – *is* valid given Negative *Ersatz Reductio*.[13]

The second place the argument for (IT)'s non-falsity can seem problematic is in the step from (4) to (5). (5) would follow from (4) on the assumption that every statement of the type 'Mn is true' is true, false, or indefinite. But whence this assumption? Here is the second place it can seem that (*) is being assumed, since, if we had (*), we could infer the assumption as an instance. But the move from (4) to (5) must be made on other grounds – and here the appeal to the semantics for universal generalizations doesn't help us. There is, however, a route to the non-falsity of (IT) – albeit one that requires assumptions beyond those explicitly justified by Tye's semantics. Given that 'It's false that A' entails not-A, we may, from (4), infer that every j is such that 'Mj is true' is not indefinite. If we have Negation Elimination (truth-preserving according to exhaustive strong-Kleene tables) and can assume (reasonably enough) that 'It's determinate that A' entails 'It's true that A', then, using All-Introduction, we can indeed reach (5). I won't try to defend these assumptions, however, since there's a further, more difficult problem with the argument for (IT)'s non-truth.

Given *Ersatz Reductio*, (2) does indeed follow from (1).[14] But how is (3) then supposed to follow? We have yet to discharge the supposition of (1), and (3) just says that that supposition is not true. So, presumably (3) is meant to follow by *Ersatz Reductio*. It only does, however, if (2) – on the supposition of (1) – leads to a contradiction. We need either:

(2a) it's true that Mj itself is either true, false, or indefinite

or some other claim which contradicts something that follows from (1) and (2). Here is the third place that the argument can seem to require (*), since (2a) would follow from it. But, of course, we can't appeal to (*) – nor are the semantic clauses of use, since the Mi's are monadic singular sentences. This last apparent assumption of (*) proves the most difficult to avoid.[15] What's more, there's a parallel gap in Tye's argument for the non-truth of the claim that 'is true' is vague.[16]

Tye's argument for (IT)'s indefiniteness – a central plank of his position – thus seems a failure. But how damaging is this? Perhaps Tye could accept the correction and simply *maintain* (IT)'s indefiniteness (justification must come to an end somewhere). Or perhaps his position is potentially stronger than that: insofar as maintaining (IT)'s indefiniteness contributes to an overall attractive theory of vagueness in comparison to its competitors, the claim arguably receives indirect support. In fact, he can do even better than that. Given *Ersatz Reductio*, (IT)'s indefiniteness follows from the truth-value predicate's being vaguely vague.[17] So, even if his own argument fails, there is after all an alternative argument for (IT)'s indefiniteness.

It may be objected that we have introduced a circularity: it has already been noted that Tye appeals to (IT)'s indefiniteness in arguing for the non-falsity of the claim that 'is true' is vague (part of his argument that it's vaguely vague), and the circularity is compounded if one closes the gap in the argument for the non-truth of the claim that 'is true' is vague (as one can) by appealing again to (IT)'s indefiniteness. Of course, Tye could decide to declare one of these claims – (IT)'s indefiniteness or the vaguely vagueness of the truth-value predicates – the more fundamental and adjust his argumentation accordingly. But another, perhaps more attractive option would be to suggest that this argumentative circularity is not vicious: by showing how his central claims mutually support one another, it rather displays his position's coherence. Either way, our development of Keefe's first objection, though effective, arguably forces only a retrenchment. If we are to argue that Tye's position is truly problematic, we need to push further.

4. KEEFE'S SECOND OBJECTION

4.1. *The Second Objection*

Keefe's second objection is that it's not viable to maintain (*)'s indefiniteness. If (*)'s indefiniteness is not viable, then neither is it viable that the truth-value predicates are vaguely vague, since the latter entails the former. What's more, any *argument* for either must be unsound, even if the argument needn't inconsistently advert to (*) itself. Why is it not viable? She writes (2000, 122):

> The use of a three-valued logic for a vague language requires the assumption that the three values provide an exclusive and exhaustive classification of declarative sentences; if not, it suffers from the same defects as the rejected two-valued system. The claim that the three values *do* provide an exhaustive classification would commit us to the truth of (*). Since Tye maintains that (*) is indefinite, we can assume that he would likewise claim that it is indefinite whether the classification is exhaustive. But this alone is unsatisfactory. A three-valued logic for a language is inadequate if it is not true that all its sentences take one of the three values.

Taken by itself, this remark might seem to beg the question against Tye. To maintain that (*) is indefinite just is to maintain that it's indeterminate whether – and thus not true that – the semantics is exhaustive. So, surely Tye would simply reject exhaustiveness as an adequacy condition. To demand exhaustiveness is to demand precisely what Tye urges vagueness requires us to reject. But Keefe does not simply beg the question. Her point is that Tye's use of a three-valued logic *requires* exhaustiveness. She does not here spell out why, instead referring to "similar reasons" (2000, 122) she deploys against vague degree-theoretic truth-value predicates. Her worry there, in a nutshell, is that, if sentences are assigned degrees of truth from the range [0, 1] in a *vague* meta-language, then there will be sentences that can*not* be correctly assigned a value from that interval – and this is in tension with the degree-theorists' conception of logical truth and validity, for their claims concerning logical truth and validity follow only if they cover *all* cases (2000, 120–121).[18] Keefe's second objection is that one can enter an analogous charge against Tye.

We've been granting Tye *Ersatz Reductio*. Indeed, Tye's discussion suggests that he assumes more generally that inferences that are truth-preserving according to exhaustive strong-Kleene tables are valid in his system as well.[19] But, if Keefe is right, this is a significant – and unwarranted – assumption. This threatens both the viability of Tye's system – including the viability of maintaining (*)'s indefiniteness – and the arguments Tye advances in support of his claims. (Since the latter includes arguments in support of (*)'s indefiniteness, developing this objection vindicates a variant of Keefe's first objection, albeit on completely different grounds.)

In the next section, I spell out the basic objection – adjusted to apply to Tye – and then, in the sections that follow, consider two possible replies. Adjustments are necessary because the degree-theoretic case Keefe discusses is not precisely parallel. In particular, *vague* degree-theoretic truth-value predicates are *not* exhaustive; but it's *indeterminate* whether Tye's *vaguely* vague truth-value predicates are exhaustive. This difference in fact opens up room for maneuver not available in the other case. Keefe's objection, as applied to Tye, thus requires not only adaptation, but also development.

4.2. *The Objection Spelled Out*

When a logically complex claim's truth-value is a function of those of its constituents, an exhaustive truth-table enables positive claims concerning validity – at least as validity is standardly understood – precisely because the truth-table covers *all* cases. A judgment of *in*validity requires only one counter-instance. But a judgment of validity requires that *no* distribution of truth-values violates the requirements of validity. If validity is truth-preservation, then *any* distribution of truth-values across the constituents of the premises that renders the premises true must also render the conclusion true. This is so whether one is dealing with classical tables in a bivalent semantics or with three-valued tables.

But when it's not true that a semantics is exhaustive, it can't be claimed that the tables cover *all* cases. And this limits our ability to draw conclusions concerning validity from them. The tables can enable judgments concerning *non*-validity. For example, if we employ classical or strong-Kleene tables and take validity to be truth-preservation, we can conclude that A does not entail A&B. They can also enable us to draw *some* positive conclusions concerning validity. For example, again taking validity to be truth-preservation, the tables support positive judgments of validity when the conclusion is truth-functionally constructed from the premises. For instance, they support the inference from A and B to A&B. The relevant portion of the table for conjunction tells us that the conclusion must be true when the premises are true; so, in *this* case, it doesn't matter what the rest of the table says or whether it exhausts all remaining cases. However, when the conclusion is *not* truth-functionally constructed from the premises, it *does* matter that it's not true that the tables are exhaustive. Consider the classically valid inference from A to AvB. Both the classical table for disjunction and the strong-Kleene table are such that, when A is true, so is AvB for all cases covered by the tables. So neither table supplies a counter-instance. But if it's not true that the tables are exhaustive, we can't exclude there being a further truth-value that renders AvB false when A is true and B has the further truth-value. According to Tye, it's *not true* that there's a further truth-value beyond true, false, and indefinite. But it's also not true that there's not. So, it seems we can neither affirm nor deny that Disjunction Introduction is valid. Similarly for various other positive claims concerning validity.

Holding that (*) is indefinite and that therefore it's not true that the semantics is exhaustive thus threatens both (1) the viability of Tye's system and (2) the cogency, by his own lights, of his arguments – in particular, his argument for (*)'s indefiniteness. For (1) if underwriting judgments concerning validity is among the goals of a semantics, then exhaustiveness is after all a condition of adequacy, one that Tye's fails to satisfy. Moreover, (2) this failure undermines his ability to defend the reasoning he employs to support his claims. Without reason to accept *Ersatz Reductio*, Tye can no longer, for example, cogently advance his argument for (*)'s indefiniteness. Indeed, assuming (as it seems he must) that the metalanguage is similarly non-classical, (2) strengthens the case for (1). For Tye can't reply to (1) by arguing that his semantics *does* underwrite judgments concerning validity – albeit, in many cases, the judgment that it's indeterminate whether the schema is valid. Even putting aside whether such a conclusion would be in any event palatable, Tye's position precludes his providing an *argument* for this very claim concerning arguments.

4.3. *First Reply: Principles, Not Tables*

How might the defender of vaguely vague truth-predicates reply to this objection? A first move is to de-emphasize the strong-Kleene truth-tables in favor of the principles that motivate them; for these principles carry more information than the truth-tables. Recall that Tye claims, for instance, that a disjunction is true if either disjunct is true and false if either disjunct is false; otherwise it's indefinite. This does indeed yield the strong-Kleene disjunction table for disjuncts that are true, false, or indefinite. But it also goes beyond that. The principle in effect provides a recipe for constructing larger tables were there more truth-values. Note in particular that the clause stating that a disjunction is true if either disjunct is true tells us all we need to know to validate Disjunction Introduction. If A is true, then so is AvB no matter what B's truth-value is – whether it's true, false, indefinite, or something else.

The problem with this move, however, is that the semantic claims only yield so much. Consider, for example, Disjunction Elimination. It's tempting to reason as follows. If both not-A and AvB are true, then B must be true. For consider the problematic case in which B is neither true, false, nor indefinite – that is, in which a sentence is not covered by the three-valued table. In that case, it's not true that AvB has a true disjunct, and it's not true that both its disjuncts are false. But then, given Tye's construal of it, the 'otherwise' clause of the semantics for disjunction kicks in – and so AvB is I. Such a case thus would not be a counter-instance to the truth-preservation of Disjunction Elimination. Again, one can construe Tye's semantic claims as providing a recipe for constructing an expanded truth-table, were there further truth-values: in this case, telling us to fill in I's for such disjunctions.

We can see that something is amiss in this argument by noting that, if it's correct, then it follows that all disjunctions are T, F, or I. But Tye can't affirm this restriction of (*). For it is easy to construct sorites series from disjunctions, and Tye would want to claim that such series lack sharp boundaries. The problem above is that we can't claim to have covered all cases. In particular, Tye can't affirm that all disjunctions are such that either one disjunct is true, both are false, or it's not true that one disjunct is true and not true that both are false.[20] But if he cannot affirm this, then he cannot claim to have *ruled out* the possibility of AvB's being true though A is false and it's not true that B is true. Parallel remarks hold for various other positive validity claims.

Tye's semantics would thus still leave unsettled many questions concerning validity. Moreover, it would leave unsettled, in particular, inferences crucial to Tye's own argumentation, such as *Ersatz Reductio*. If a semantics should not leave such questions concerning validity unsettled, then Tye's semantics would be unviable. And his arguments (in particular, for (*)'s indefiniteness) would contain crucial gaps – albeit in places other than the one Keefe highlights in her first objection.

4.4. *Second Reply: Non-Truth of Non-Truth-Preservation, Not Truth-Preservation*

There is an alternative reply that naturally suggests itself. Instead of searching for beefed up semantic claims that overcome the limitations of indeterminately exhaustive truth-tables, we might attempt to overcome these limitations by altering the characterization of validity. The idea is this. It's not true, according to Tye that there are truth-values beyond truth, falsity, and indefiniteness. So, it's not true that there's a distribution of truth-values that provides a counterinstance to the claim that the inferences deemed truth-preserving by exhaustive strong-Kleene tables are so. Thus, even though indeterminately exhaustive truth-tables underdetermine what preserves truth, they yield information concerning what inferences are such that it's *not true* that they *don't* preserve truth. Why not then characterize validity so that it's *this* property that a good inference ought to have? Let's label this notion of validity—the non-truth of non-truth-preservation – *weak* validity. Truth-preservation according to exhaustive strong-Kleene tables would then entail weak validity according to indeterminately exhaustive strong-Kleene tables. In particular, *Ersatz Reductio* is weakly valid.

A *prima facie* problem with this move, however, is that it doesn't license the assertion of a conclusion of a weakly-sound argument (a weakly-valid argument with true premises). If C is a weakly-sound conclusion, weak-validity only guarantees that it's not true that C's not true. Since it's not true that the semantics is exhaustive, one can't then conclude that C is true. This is simply what results when one departs from truth-preservation as validity. Such departures are not formally – and need not be in other ways – objectionable. That depends on one's aims. But if one argues in accordance with the inferences thus validated in order to support one's position, as Tye does, then there does indeed seem to be a problem.[21]

It may be replied that this objection assumes that assertability requires truth: perhaps the non-truth of the relevant claim's non-truth suffices. This reply, however, founders on a second problem. As we've seen, Tye maintains – and, given *Ersatz Reductio*, is committed to the claim – that it's indeterminate whether there are truth-value imputations that are indefinite. But then it's not true that there are truth-value imputations that are neither true nor false. So, inferences that preserve truth according to exhaustive *classical two-valued* tables are weakly valid when restricted to truth-value imputations. In particular, every instance of the schema 'A or not-A' is weakly valid when A is a truth-value imputation. So, it's weakly valid that every truth-value imputation is true or not. It may not follow that it's *true* that every truth-value imputation is true or not, but *asserting* a sharp boundary is problematic in its own right.[22]

Furthermore, it's problematic in its own right to maintain that it's not true that it's not true that every truth-value imputation is true or not. Tye accepts that, from object-level sentences describing a typical sorites series, one can construct meta-linguistic sorites series ('Sentence 1 is true,' Sentence 2 is true,' etc.). And, for reasons parallel to his reasons for maintaining the indefiniteness of (*), he must maintain that it's indeterminate whether every truth-value imputation is true or not. But then it's not true that every truth-value imputation is true or not. And so it's true that it's not true – which contradicts the claim that it's not true that it's not true.

Finally, parallel considerations display the dire consequences of responding to the first *prima facie* problem, not by adjusting what assertion requires, but by trying somehow to show that weakly-sound conclusions in fact *are* true. For if weakly-sound conclusions are true, then the purveyor of vaguely vague truth-value predicates both commits herself to the *truth* of the sharp boundaries claim and contradicts her claim that it's indeterminate whether there are sharp boundaries in a meta-linguistic sorites.

5. CONCLUSION

Keefe's first objection to Tye leaves him room for reply. And its application to other arguments of Tye's does only limited damage. But her second objection, adjusted and developed to respond to possible replies, appears to deal the purveyor of vaguely vague truth-value predicates a deadly blow. If the position is to be resurrected, it must develop some alternative semantic rationale for the validity claims it needs or, alternatively, motivate them in some other, non-semantic manner.[23] Either way, there would be much work to be done.[24]

ACKNOWLEDGEMENTS

I thank Rosanna Keefe for correspondence and encouragement and Richard Heck for helping me clarify my thoughts on weak validity. I'm very grateful also to anonymous referees, some of whom provided feedback that shifted the very direction of parts of the paper.

NOTES

[1] Tye (1990) and (1994a). See also his (1994b). Tye (1994a) is reprinted with omissions in Keefe and Smith (1997). When providing pages references to the former, I will also provide page references to the latter following a slash.

[2] Keefe (2000, 110–112) also raises other objections to Tye's position: (i) that it does not respect penumbral connections (since they can be indefinite), (ii) that it allows certain classically valid inferences to fail, and (iii) that his brief discussion doesn't suffice to explain away why one might find these connections and inferences compelling. I do not discuss these objections, since they don't pertain in particular to Tye's response to the problem of higher-order vagueness. In addition, as Keefe notes, they are costs some might be willing to pay depending on the other candidate theories of vagueness available. Cf. Keefe (2000, 32–47) on reflective equilibrium.

NB: The classically valid inferences of (ii) are those not validated by *exhaustive* strong-Kleene tables. As will become clear, this objection is thus distinct from the objection that a *non*-exhaustive semantics fails to validate inferences needed to defend Tye's position.

There are no doubt other objections one might raise. Tye's commitment to vague objects, for example, is controversial. For Tye's development and defense of his position on vague objects, see Tye (1990) and (2000). Again, one might object that, if Tye's approach can defuse putative higher-order vagueness, it should be applicable to putative first-order vagueness as well – so, some motivation is needed for the differential treatment. It might be replied that it's simply the case that there are clear first-order borderline cases but not clear higher-order borderline cases. Cf. Sainsbury and Williamson (1997, 478–479). The present paper, however, does not claim to cover exhaustively all possible objections.

[3] The preceding paragraphs consist of near quotes, *modulo* some compression, from Tye (1990, 540–546) and (1994a, 194–196/282–285). (One difference between Tye (1990) and Tye (1994a), however, is that, in the former, Tye states the clauses for monadic singular sentences using 'if' instead of 'iff.')

As Tye moves freely between two ways of imputing truth-value, via the use of sentence-operators such as 'it's true that' and via the use of predicates such as 'is true,' so will I.

[4] Again, Tye holds that the truth-value predicates are vaguely vague (so it's *indeterminate* whether they can have borderline cases); so he wouldn't hold that sentences predicating a truth-value *can't* be indefinite. But if a semantic clause containing a constituent predicating a truth-value is true, then, if the clause is a biconditional, the constituent predicating a truth-value can't be indefinite (likewise if the clause is a conditional and either, one, the constituent predicating a truth-value is the antecedent and the consequent is not true or, two, the constituent predicating a truth-value is the consequent and the antecedent is true). So, it's indeterminate whether all instances of such a clause are true. Thus, given Tye's semantics for universal generalizations, the semantic claim itself is indefinite.

[5] Cf. Tye (1990, 547/286) on the argumentative limits imposed by an indefinite Axiom of Extensionality.

[6] Keefe (2000, 121–122). She cites Tye (1994a, 200/290) (cf. Tye 1990, 550–551). But the remarks there do not seem to support her reconstruction. Tye is emphasizing that his claim that the meta-linguistic sorites premise is not true does not commit him to holding that 'is true' is vague:

> ... if "is true" is extensionally vague then it follows that the set of true sentences has borderline members. This requires that there be sentences which are such that it is neither true nor false that they are true. And this, in turn, requires that there be sentences that are neither true nor false nor indefinite. I maintain that it is not true that there are such sentences. So I do not accept that "is true" is extensionally vague. ... in taking this view I am not committing myself to the position that these predicates are precise. Indeed, it is crucial to my account that they *not* be classified as precise. For if they were then every sentence would be either true or false or indefinite, and that would not only generate sorites difficulties of its own ... but also run counter to my claim that it is indefinite whether no statement of the form "Mn is true" is indefinite. Rather my view on the truth-value predicates is that they are vaguely vague: there simply is no determinate fact of the matter about [i.e., it is indeterminate] whether the properties they express have or could have borderline instances. So, it is indefinite whether there are any sentences that are neither true nor false nor indefinite.

Let (T-VAGUE) be the claim that 'is true' is vague. Tye argues that (T-VAGUE) is not true, partly on the basis of (T-VAGUE)'s entailing not-(*). He argues that not-(T-VAGUE) is not true – so, (T-VAGUE) is not false – partly on the basis of not-(T-VAGUE)'s entailing (*). He *maintains* (he doesn't explicit say that he *concludes*) that it is indeterminate whether (T-VAGUE) – i.e., he maintains that 'is true' is vaguely vague. And, *from the indefiniteness of (T-VAGUE)*, he concludes that it is indeterminate whether (*). It's not unreasonable to import an implicit 'because' into Tye's statement of what he maintains, but even so he would then be arguing rather for *(T-VAGUE)'s* indefiniteness from its non-truth and non-falsity.

Tye does, however, offer *elsewhere* an argument like the one Keefe discusses. Tye (1996, 220) writes:

Obviously, we cannot allow it to be true that every object-language sentence is either true or false or indefinite. For this would create sharp dividing lines in sorites sequences .... But neither can we allow it to be false .... For this would require further alternative truth-values.... The conclusion we should draw, then, in my view, is that the claim that every object language sentence is either true or false or indefinite is itself indefinite.

Tye here restricts himself to object-language sentences. One might thus be tempted to distinguish (*O) and (*M) – versions of (*) restricted to object-language and meta-language sentences, respectively – and then reply to Keefe that Tye can invoke (*M) to argue for (*O)'s indefiniteness. But if (*M) were true, there would be sharp dividing lines in meta-linguistic sorites sequences. In one form or another, as Tye (1994a, 206, n. 27) notes, "vagueness intrudes at all levels."

Incidentally, the passage that opens this note is the only place where Tye uses the construction 'it is indefinite whether.' In general, he reserves 'indeterminate' for the sentence-operator and 'indefinite' for the sentence-predicate, which he indeed takes pains to distinguish – see Tye (1990, 545–6) and (1994a, 196/285). I assume this is just a slip.

[7] Tye (1994a, 204, n. 24 and the text to which it's attached/290, n. 14). The otherwise-clause, so construed, allows for indefinite universal generalizations where it's indeterminate whether the universal generalization has indefinite instances. Such is the case, argues Tye, with typical meta-linguistic sorites premises of the form 'For all x, if Sx is true, then S(x+1) is true.' Regarding *object*-level sorites that use the sentences S1...Sn themselves, Tye argues on the basis of the strong-Kleene tables that the corresponding sorites premises are indefinite because they *do* have indefinite instances. The differing treatments of object- and meta-level sorites correspond to the difference between vague and vaguely vague predicates.

[8] Cf. also n. 19 below.

[9] It might be objected that, if one can thus infer that a universal generalization is indefinite from its being not true and not false, then one is after all committed to sharp boundaries: doesn't endorsing this inference commit one to (*U) – i.e., (*) restricted to universal generalizations – and thus to sharp boundaries in sorites series consisting of universal generalizations ('Every one-haired man is bald,' 'Every two-haired man is bald,' etc.)? But it's not clear that it does. To assert that a universal generalization's indefiniteness *follows from* its being not true and not false is not to assert that, if a universal generalization is not true and not false, then it's indefinite – at least not in a sense that allows us to conclude (*U). (Recall, again, that we're construing Tye's semantic clauses as asserting inferential relations. We see here again why this is charitable). Nor can we, given this inference, appeal to Conditional Introduction to reach (*U), as we could in a classical setting: Tye's semantics uncontroversially allows counter-instances to Conditional Introduction. One may, however, ask whether Conditional Introduction *restricted to truth-value imputations* is valid: with this, one could indeed infer (*U). As we'll see, Tye's position doesn't provide

one with the resources either to declare this schema valid *or* to declare it invalid. But this will be an instance of a broader problem.

Incidentally, note that, even allowing *Ersatz Reductio*, Tye's semantic clauses do not support the inference from non-truth and non-falsity to indefiniteness *generally*. They do not, for example, in the case of monadic singular sentences.

[10] Cf. notes 7 and 6 above.

[11] Tye (1994a, 200/289–290). Cf. Tye (1990, 550). (The numbering and other interpolations above are mine.) Tye is here in the midst of defending his treatment of a meta-linguistic sorites. Consider the following objection: if the sorites premise is indefinite, then some sentence of the form 'Mi is true' must be indefinite; but Tye is committed to holding that it's false that some sentence of the form 'Mi is true' is indefinite. Tye responds, first, that nothing in his view commits him to the falsity of 'some sentence of the form 'Mi is true' is indefinite' – indeed, as he argues in the quoted passage, he is committed to its indefiniteness – and, second, that in any event it doesn't follow from his semantics that an existential generalization is indefinite only if an instance is.

[12] Suppose it's not true that it's false that A. And suppose not-A. Then it's false that A. So it's true that it's false that A, which contradicts our first supposition. So, by Negative *Ersatz Reductio*, it's not false that A. (The argument assumes that not-A entails A's falsity and that A entails A's truth.)

[13] A parallel argument, using *Ersatz Reductio*, enables us to establish Restricted *Reductio* T – that is, *Reductio* where the contradiction is blamed on an imputation of truth. What of Restricted *Reductio* I – *Reductio* where the contradiction is blamed on an imputation of indefiniteness? We'll note in a moment that allowing this would seem to lead to trouble.

[14] If 'Mj is true' is indefinite, then it's not true that Mj is true and not false that Mj is true. Now, suppose in turn that Mj is true, that it is false, and that it is indefinite. Each supposition will allow the derivation of a contradiction. If Mj is true, then it's true that Mj is true. If Mj is false, then it's not true that Mj; so it's false that Mj is true. If Mj is indefinite, then it's not true that Mj; so it's false that Mj is true. This reasoning assumes that the truth-value predicates are exclusive (having one entails not having the others), that A entails 'it's true that A', and that not-A entails 'it's false that A.'

[15] If Restricted *Reductio* I were valid, then, from (2) and Restricted *Reductio* T and F, we could infer that it's not the case that Mj is either true, false, or indefinite. This would contradict the claim that (*) is indefinite. For, if (*) is indefinite, then it's not false; but, if Mj is neither true, false, nor indefinite, then (*) *is* false. So, assuming (*)'s indefiniteness, we could reach (3). The problem, however, is that it seems we could then run a variant of the argument from (1) to (3), now using an arbitrary M instead of the existential generalization. The Restricted *Reductios* would then allow us to strengthen our conclusion from 'It's *not true* that it's indeterminate whether M is true' to 'It's *not the case*.' But since M was arbitrary, we have that for *no* such sentence is it indeterminate whether it's true – which both contradicts (IT)'s indefiniteness and imposes sharp boundaries.

Similar problems arise if, attempting to find something to contradict (2), one affirms the Exhaustiveness Inferences – the inferences that license passing from a sentence's not having one truth-value to its having one or the other of the other two, and the inferences that license passing from a sentence's not having either of two truth-values to its having the truth-value that remains. (One might have otherwise hoped that affirming these *inferences* would mitigate the indefiniteness of exhaustiveness itself.) In general, as we'll see further below, the purveyor of vaguely vague truth-predicates has difficulty strengthening his position so as to defend his claims while not so strengthening them as to impose sharp boundaries after all.

[16] Tye (1994a, p.200/290 – cf. Tye 1990, 550–551]) writes:

... if "is true" is extensionally vague then it follows that the set of true sentences has borderline members. This requires that there be sentences which are such that it is neither true nor false that they are true. And this, in turn, requires that there be sentences that are neither true nor false nor indefinite. I maintain that it is not true that there are such sentences. So I do not accept that [i.e., it's not true that] "is true" is extensionally vague.

[The continuation of this passage is quoted above in n. 6.]

[17] From (1), we have that there's a truth-imputing sentence that's neither true nor false. So, 'is true' has a borderline case. Now, recall that, according to Tye, a predicate is vague iff it has borderline cases and it's indeterminate whether there are objects such that they are neither members, borderline members, or non-members of the predicate's extension. So, given (1), 'is true' is vague if it's indeterminate whether there are truth-imputing sentences that are neither true, false, nor indefinite. But the purveyor of vaguely vague truth-value predicates surely will maintain that this is indeed indeterminate. (Cf. Tye 1994a, 201/290–291, and 1990, 551.) For the claim that there's a truth-imputing sentence that's neither true, false, or indefinite is equivalent to (*) restricted to truth-imputing sentences; and we can argue for its indefiniteness in a manner parallel to the argument Tye presents for the indefiniteness of (*) itself. But the vagueness of 'is true' is inconsistent with its being vaguely vague. This gets us (IT)'s non-truth. We can then get its non-falsity as above – or also by an argument appealing to 'is true's being vaguely vague.

[18] Cf. Williamson (1994, 99 and 112).

[19] This is suggested by the fact that his discussion of validity focuses on the classically valid inferences that are not truth-preserving according to exhaustive strong-Kleene tables. See Tye (1990, 545, n. 20, and 1994a, 203, n. 13/283, n. 3). Keefe's discussion of Tye at one point might suggest that she also assumes that validity in his system coincides with truth-preservation according to exhaustive strong-Kleene semantics. For when, in assessing Tye's system, Keefe (2000, 111) objects that certain classically valid inferences fail, she presents only schemata that fail in *exhaustive* strong-Kleene systems. Moreover, in noting Tye's reply, she puts forward without objection his claim that, in his system, although there are no sentential validities, classical tautologies are never false and classical contradictions are never true. But we'll see – developing Keefe's own remarks elsewhere – that *prima facie* Tye's at best entitled to the claim that it's *not true* that the former are ever false or the latter ever true. Perhaps then at (2000, 111) Keefe does not intend to be read as speaking in her own voice and is holding off her more serious objection for later.

[20] Note that we do have that, from one of these conditions being met, it follows that the disjunction is T, F, or I. But we cannot therefore affirm the corresponding conditional.

[21] It's no good objecting that this very argumentation, if weakly-sound, doesn't establish that weakly-sound arguments don't establish the truth of their conclusions, but rather only establishes that it's not true that it's not true that they do. This objection itself presupposes that weakly-sound arguments don't establish that their conclusions are true.

[22] A variation on this move allows us to argue after all that all universal generalizations are true, false, or indefinite. According to Tye's semantics, if a universal generalization is neither true nor false, it follows that it's indefinite. But Conditional Introduction, though not generally weakly-valid, is weakly-valid when restricted to truth-value imputations. Suppose, then, that A is a universal generalization. If A is neither true nor false, then one can infer its indefiniteness. Since these are truth-value imputations, it follows that it's weakly-valid that A is either true, false, or indefinite. But A was an arbitrary universal generalization. So, all universal generalizations are true, false, or indefinite. But this imposes sharp boundaries on a sorites series consisting of universal generalizations ('All men with one hair are bald,' All men with 2 hairs are bald,' etc.). It's also inconsistent with Tye's claim that, in the corresponding meta-linguistic series ('Sentence 1 is true,' 'Sentence 2 is true,' etc.), it's indeterminate whether there's a sentence that's indefinite.

Both this argument and the one in the text employ All-Introduction. This is essential because, given Tye's claim that it's indeterminate whether there's a sentence that's neither true, false, nor indefinite, he will claim that no *particular* claim of the form 'A is true, false, or indefinite' is problematic. It's the generalization that spells trouble. Note that Tye has given us no reason to abandon All-Introduction and in fact, as noted above, employs it in his attempt to argue for (IT's) indefiniteness.

[23] In particular, as we've seen, it must motivate *Erstaz Reductio*, but not Conditional Introduction or the inference from 'It's not true that it's indeterminate whether A' to 'It's not indeterminate whether A' – while being careful not to otherwise reimpose sharp boundaries.

[24] It's natural to wonder whether parallel problems arise for non-epistemic accounts of vagueness that employ a *vague* meta-language. Not necessarily. Consider, for example, Keefe's own supervaluationist position (2000, chapters 7 and 8). She employs a vague meta-language – in particular, 'admissable specification' is vague as are thus some truth-value imputations. But this is consistent with maintaining, as she does, that all sentences are T, F, or I and that the semantics is thus exhaustive. On a supervaluationist approach, (*) can be true even if, for some instances, no disjunct is true. If Keefe's use of a vague meta-language is problematic, it's for other reasons.

REFERENCES

Keefe, R.: 2000, *Theories of Vagueness*, Cambridge University Press, Cambridge.

Keefe, R., and P. Smith, eds.: 1997, *Vagueness: A Reader*, MIT Press, Cambridge.

Sainsbury, M. and T. Williamson: 1997, Sorites, in B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*, Blackwell, Oxford, pp. 458–484.

Tye, M.: 1990, 'Vague Objects', *Mind* **99**, 535–557.

Tye, M.: 1994a, Sorites Paradoxes and the Semantics of Vagueness, *Philosophical Perspectives* 8, 189–206. Reprinted with omissions in Keefe and Smith (1997, 281–293).
Tye, M.: 1994b, 'Why the Vague Need *Not* be Higher-Order Vague', *Mind* **103**, 42–44.
Tye, M.: 1996, 'Fuzzy Realism and the Problem of the Many', *Philosophical Studies* **81**, 215–225.
Tye M.: 2000, 'Vagueness and Reality', *Philosophical Topics* **28**, 195–209.
Williamson, T.: 1994, *Vagueness*, Routledge, London.

Department of Philosophy
Georgetown University
Washington DC 20057
USA
E-mail: steven.gross@post.harvard.edu