

Tolerant, classical, strict

P. Cobreros, P. Égré, D. Ripley, R. van Rooij

Abstract

In this paper we investigate a semantics for first-order logic originally proposed by R. van Rooij to account for the idea that vague predicates are tolerant, that is, for the principle that if x is P , then y should be P whenever y is similar enough to x . The semantics, which makes use of indifference relations to model similarity, rests on the interaction of three notions of truth: the *classical* notion, and two dual notions simultaneously defined in terms of it, which we call *tolerant* truth and *strict* truth. We characterize the space of consequence relations definable in terms of those and discuss the kind of solution this gives to the sorites paradox. We discuss some applications of the framework to the pragmatics and psycholinguistics of vague predicates, in particular regarding judgments about borderline cases.

Our aim in this paper is to explore a semantic framework originally proposed by [?] in order to deal with the sorites paradox, and intended to formalize the idea that vague predicates are tolerant. Standardly, the idea of tolerance is expressed by means of the following principle: if some individual x is P , and x and y are only imperceptibly different in respects relevant for the application of the predicate P , then y is P as well. In classical logic, the principle of tolerance gives rise to the sorites paradox. Because of that, one influential strand of solutions to the sorites paradox consists in rejecting the principle and substituting weaker principles in its stead. A different approach consists in preserving the tolerance principle itself but appealing to a non-classical logic

The semantics originally proposed by van Rooij belongs to that second family: it allows us to validate the tolerance principle in its plain form, and it is non-classical. The framework rests on the interaction of three notions of truth for sentences involving vague predicates: the classical notion of truth, a notion of tolerant truth, and a dual notion of strict truth. Because of this, the framework leaves room for many different notions of logical consequence. In his earlier work, van Rooij suggested that the appropriate notion should be neither preservation of classical truth nor preservation of tolerant truth, but instead a mixed notion, on which we reason from classically-true premises to tolerantly-true conclusions.

As it turns out, however, the standard notions of logical consequence for tolerant truth and strict truth are also interesting *per se*. In particular, they bear an unexpected connection to more familiar many-valued logics: the Logic of Paradox (LP) proposed by Priest in [?], and its dual, so-called “Strong Kleene” logic (K3). Because of this, they cast a

new light on these many valued approaches, as applied to vagueness. Furthermore, the distinction between tolerant and strict truth also bears a connection to the frameworks of subvaluationism and supervaluationism that have been proposed to deal with vagueness. Because tolerant truth and strict truth are interdefined, however, the semantics makes distinct predictions, in particular regarding borderline cases, for which classical contradictions are predicted to hold tolerantly, and classical validities to fail strictly.

In section 1, we start out by rehearsing the main motivations behind van Rooij’s semantics for the notion of tolerance, and present some basic features of the semantics, in particular regarding the characterization of borderline cases for vague predicates. In section 2, we characterize logical truths for the notions of tolerant and strict truth, and establish a natural correspondence between tolerant/strict semantics and two well-known many-valued logics, LP and K3. In section 3, we enlarge the space of consequence relations and discuss various notions of mixed consequence, in particular van Rooij’s consequence from classical to tolerant truth and its kin, and discuss the application of this framework to the sorites paradox. In section 4, finally, we close this paper with the discussion of some applications of the semantics to the pragmatics psycholinguistics of vague predicates, in relation to recent experiments by [?], [?], [?]. The focus of that section concerns the predictions of tolerant and strict semantics for borderline cases, and in particular the choice between tolerant and strict interpretations for vague predicates.

1 Tolerant and strict semantics

1.1 Tolerance and indifference

Let us consider a vague predicate such as “tall”. The principle of tolerance corresponds to the following intuitive constraint, that if one individual is tall, and this individual is not visibly or relevantly taller or smaller than another individual, then the other is tall as well. Formally, the principle can be stated as follows:

- (1) $\forall x \forall y (P(x) \wedge x \sim_P y \rightarrow P(y))$, where P stands for “tall”, and \sim_P is the relevant indifference relation (namely not looking to have distinct heights).

Central to the discussion of this principle is the specification of the properties of the indifference relation. Arguably, a relation such as “not looking to have distinct heights” is reflexive and symmetric, but not transitive: a can look to have nearly the same height as b , b can look to have nearly the same height as c , but a and c may look to have distinct heights. In our approach, the non-transitivity of the indifference relation is a central feature of all vague predicates (see [?], [?]). One can think of indifference relations of this kind as defined from what [?] calls *semi-order* relations (see eg [?]). In the case of a predicate like “tall”, the semi-order associated with it would be the relation \succ_P such that $x \succ_P y$ expresses that x looks visibly or relevantly taller than y . From a semi-order relation, $x \sim_P y$ is definable as $\neg(x \succ_P y \vee y \succ_P x)$, namely neither of x and y looks significantly taller than the other.

We need not specify the properties of semi-orders in this paper (we refer to [?] for details). All we need is to assume is that any vague predicate P comes associated with an appropriate indifference relation \sim_P that is reflexive and symmetric, but possibly non-transitive. Van Rooij’s proposal is indeed that the tolerance principle can be validated if the semantics of vague predicates is made sensitive to such indifference relations. On van Rooij’s approach, we can say that x is tall *tolerantly* if there exists an individual y such that x is similar to y by way of how tall x looks, and y is tall *classically*.

One way to motivate this semantic conception could be the following: suppose that a subject, John, is given a first task, which is to draw as best as he can a sharp line between the tall and the non-tall individuals in a given set. The two sets thus delineated fix a perfectly classical extension for the predicate “tall” relative to John’s inner model of the situation (namely the sets are disjoint and exhaust the whole domain). Now suppose John can still remember where he drew the line, but is assigned a second task, namely is asked of some arbitrary individual x in the series whether x counts as tall or not. What we are assuming is that if x looks sufficiently similar by way of height to an individual y that was put in the set of tall people, then x can be called “tall” as well by John. This would happen even if x is in fact a member of the set of people that are non-tall, but is such that a slight shift of the line would have counted x as tall originally. Or to put it differently, this would be a situation in which x and y are on either side of the line drawn by John, but nevertheless look very similar in how tall they look. Importantly, however, John can still declare x not-tall in that case, since x looks also sufficiently similar to an individual that was put in the set of non-tall people (namely to x itself, or to non-tall individuals further off the line). On the other hand, however, if x happens to be sufficiently far off the line, namely if x does not look similar in how tall x looks to *any* of the individuals that have been counted as tall, then x will not count as tolerantly tall.

In brief, the intuition behind van Rooij’s understanding of tolerance is that whichever way one were to draw the line between the tall and the not-tall, there should remain room to count as tall individuals that are on the other side of the line, provided they are sufficiently close to the tall ones in respects of how tall they look. Understood this way, tolerance corresponds to the possibility of *coarsening* the extension initially assigned to a predicate. In what follows, we shall spell out this idea more formally. We shall proceed in two main steps: we present a first way of articulating the semantics, and explain why it falls short of capturing the idea of tolerance. We then state the official understanding of the notion of tolerance, and explain why it leads us to introduce a dual concept of strict satisfaction for a predicate.

1.2 Preliminaries: language and models

Language. The language we are interested in in this paper is the language of first-order logic (for now, without identity). To make the exposition simpler, we furthermore restrict the language to monadic predicate logic, as the extension to n -ary predicates does not pose special problems.

Definition 1. Let \mathcal{P} be a denumerable set of unary predicate symbols, \mathcal{C} be a denumerable set of individual constants, and \mathcal{V} be a denumerable set of individual variables. An atomic formula is of the form $P(\underline{a})$ or $P(x)$, where $P \in \mathcal{P}$ and $\underline{a} \in \mathcal{C}$, $x \in \mathcal{V}$.

Definition 2. Well formed formulae (wff): if ϕ is an atomic formula, it is a wff. If ϕ is a wff, so is $\neg\phi$; if ϕ and ψ are wff, so is $(\phi \wedge \psi)$; If ϕ is a wff, so is $\forall x\phi$.

Everywhere, we assume that disjunction \vee and the conditional \rightarrow are defined classically in terms of negation and conjunction. Likewise, $\exists x\phi$ stands for $\neg\forall x\neg\phi$. Brackets are omitted where no ambiguity would result.

Models. Classically, satisfaction of first-order formulae is defined over structures of interpretation consisting of a domain of individuals and an interpretation function. We will be interested in expansions of such structures in which every predicate comes with a relation of indifference or similarity. We thus distinguish two kinds of models:

Definition 3. A C-model M is a tuple $\langle D, I \rangle$ such that:

- D is a non-empty domain of individuals
- I is an interpretation function (of the usual classical sort) for the non-logical vocabulary: for a constant \underline{a} , $I(\underline{a}) \in D$; for a predicate P , $I(P) \in \{0, 1\}^D$.

When no ambiguity results, we write a for $I(\underline{a})$.

Definition 4. A T-model M is a tuple $\langle D, I, \sim \rangle$ such that $\langle D, I \rangle$ is a C-model and \sim is a function that takes any predicate P to a binary relation \sim_P on D . For any P , \sim_P is reflexive and symmetric (but possibly non-transitive).

We define satisfaction for wff in a substitutional manner, assuming that given a C-model or T-model, every individual d of the domain has a name \underline{d} .¹ If ϕ is a formula, $\phi[\underline{d}/x]$ is the result of substituting \underline{d} for every free occurrence of x in ϕ . We first define classical truth in this way:

Definition 5. *c-truth in a model.* Let M be either a C-model such that $M = \langle D, I \rangle$, or a T-model such that $M = \langle D, I, \sim \rangle$.

- $M \models^c P(\underline{a})$ iff $I(P)(a) = 1$.
- $M \models^c \neg\phi$ iff $M \not\models^c \phi$.
- $M \models^c \phi \wedge \psi$ iff $M \models^c \phi$ and $M \models^c \psi$.
- $M \models^c \forall x\phi$ iff for every d in D , $M \models^c \phi[\underline{d}/x]$.

Definition 6. A formula ϕ is classically valid iff every C-model makes it c-true. A formula ϕ is c-valid iff every T-model makes it c-true.

¹This is unimportant, but it simplifies exposition. It can be replaced with objectual quantification without any trouble.

Fact 1. *Classical validity and c -validity coincide.*

The proof is immediate, since every C-model can be seen as a reduct structure of the corresponding T-model, every T-model as an expansion of the corresponding C-model, and c -truth does not rest on the properties of \sim . A consequence is that in what follows, we will be able to work everywhere with T-models; and we will do so, except when we explicitly specify otherwise.

1.3 Tolerance: first approximation

1.3.1 The semantics

Let us define a first approximation of the notion of tolerant satisfaction, which we shall write $\models^{t'}$:

Definition 7. *t' -truth. Let M be a T-model:*

$$\begin{aligned} M \models^{t'} P(\underline{a}) &\text{ iff } \exists d \sim_P a : M \models^c P(\underline{d}) \\ M \models^{t'} \neg\phi &\text{ iff } M \not\models^{t'} \phi \\ M \models^{t'} \phi \wedge \psi &\text{ iff } M \models^{t'} \phi \text{ and } M \models^{t'} \psi \\ M \models^{t'} \forall x\phi &\text{ iff for all } d \in D : M \models^{t'} \phi[\underline{d}/x] \end{aligned}$$

Definition 8. *Similarity predicates. For each intended relation of indifference \sim_P over the model, we assume that there is a binary predicate of the language I_P such that by definition $M \models \underline{a}I_P\underline{b}$ iff $M \models^{t'} \underline{a}I_P\underline{b}$ iff $a \sim_P b$. That is, similarity predicates are classically interpreted, even when the relevant notion of satisfaction is tolerant satisfaction.*

1.3.2 Evaluation

The semantics we just defined implements the basic idea we described above, but it has two related shortcomings. Let t' -validity for sentences be defined in the expected way, namely as t' -truth in every T-model. Firstly, the principle of tolerance does not come out as a t' -validity in the semantics. Secondly, negation is defined in a very strong way: to say that the negation of a formula is tolerantly true means that the formula is not tolerantly true. A consequence is that tolerance fails to capture the idea of a uniform coarsening of the semantic value of a formula.

Tolerance. It is not the case that $\models^{t'} \forall x\forall y(P(x) \wedge xI_Py \rightarrow P(y))$. Consider a structure M with three elements a, b, c such that $a \sim_P b \sim_P c$ but $a \not\sim_P c$, and $I(P) = \{a\}$. Clearly, $M \models^{t'} P(\underline{b})$, but $M \not\models^{t'} P(\underline{c})$. Hence the principle of tolerance is not tolerantly valid on this understanding of tolerance.

Negation. Let us write $\llbracket P \rrbracket^{c,M} = \{d \in M; M \models^c P(\underline{d})\}$, and $\llbracket P \rrbracket^{t',M} = \{d \in M; M \models^{t'} P(\underline{d})\}$. Let us call $\llbracket P \rrbracket^{c,M}$ the classical extension of P in M , and $\llbracket P \rrbracket^{t',M}$ its tolerant extension. Clearly, for every atomic predicate P of the language, $\llbracket P \rrbracket^M \subseteq \llbracket P \rrbracket^{t',M}$, namely

the tolerant extension increases the classical extension of the predicate. Let us write $\llbracket \neg P \rrbracket^{c,M} = \{d \in M; M \models^c \neg P(\underline{d})\}$, and similarly $\llbracket \neg P \rrbracket^{t',M} = \{d \in M; M \models^{t'} \neg P(\underline{d})\}$. This time, $\llbracket \neg P \rrbracket^{t',M} \subseteq \llbracket \neg P \rrbracket^M$, but the converse is not true. This means that it is not true of arbitrary formulae that their tolerant extension in a model is more inclusive than their classical extension. In order to get a uniform notion of coarsening for formulae, it is necessary to weaken the semantics we have here for negation.

1.4 Tolerant and strict semantics

To circumvent both these limitations, [?] introduced a second notion of tolerant satisfaction, in terms of a dual notion of strict satisfaction, and of classical satisfaction. We write $M \models^t \phi$ and $M \models^s \phi$ for tolerant and strict satisfaction respectively. The two notions of satisfaction are defined by simultaneous induction.²

1.4.1 The semantics

Definition 9. *t-truth and s-truth. Let M be a T -model:*

$$\begin{aligned} M \models^t P(\underline{a}) &\text{ iff } \exists d \sim_P a : M \models^c P(\underline{d}) \\ M \models^t \neg\phi &\text{ iff } M \not\models^s \phi \\ M \models^t \phi \wedge \psi &\text{ iff } M \models^t \phi \text{ and } M \models^t \psi \\ M \models^t \forall x\phi &\text{ iff for all } d \in D : M \models^t \phi[\underline{d}/x] \end{aligned}$$

$$\begin{aligned} M \models^s P(\underline{a}) &\text{ iff } \forall d \sim_P a : M \models^c P(\underline{d}) \\ M \models^s \neg\phi &\text{ iff } M \not\models^t \phi \\ M \models^s \phi \wedge \psi &\text{ iff } M \models^s \phi \text{ and } M \models^s \psi \\ M \models^s \forall x\phi &\text{ iff for all } d \in D : M \models^s \phi[\underline{d}/x] \end{aligned}$$

Remark 1. *By definition, for every formula ϕ : $M \models^t \phi$ iff $M \not\models^s \neg\phi$, and $M \models^s \phi$ iff $M \not\models^t \neg\phi$, so \models^s and \models^t are duals.*

Remark 2. *As above, we make the assumption that similarity predicates are crisply interpreted relative to each of the notions of truth we have introduced, that is we have: $M \models^c \underline{a}I_P \underline{b}$ iff $M \models^t \underline{a}I_P \underline{b}$ iff $M \models^s \underline{a}I_P \underline{b}$ iff $a \sim_P b$. We make a similar assumption for identity predicates.*

This assumption rules out introducing further indifference relations (of the form \sim_{I_P} or $\sim_{=}$) for the tolerant and strict interpretation of similarity and identity predicates themselves. We are interested in such an approach, but we will not pursue it in this paper.

²We realized after developing the present account that the clauses given here for atomic satisfaction and negation are quite similar to those explored in [?, p. 259]. However, Kamp's treatment of conditionals and quantification, as well as his overall framework, is considerably more complex than what we consider here.

1.4.2 Evaluation

Three specific features of the present semantics can be distinguished.

Tolerance. First of all, the semantics makes the principle of tolerance t -valid.

Fact 2. For every atomic predicate P , $\models^t \forall x \forall y (P(x) \wedge x I_P y \rightarrow P(y))$.

Proof. By the clauses, this is equivalent to checking that for every M and for all a, b in M : $M \not\models^s P(\underline{a})$ or $M \not\models^s \underline{a} I_P \underline{b}$ or $M \models^t P(\underline{b})$. Suppose: $a \sim_P b$. Then either $a \in I(P)$, and $M \models^t P(\underline{b})$; or $a \notin I(P)$, and then $M \not\models^s P(\underline{a})$. \square

Consider, for instance, a T-model M with four elements a, b, c, d such that $I(P) = \{a, b\}$, $a \sim_P b \sim_P c \sim_P d$, and nothing else is related by \sim_P , except as required by symmetry and reflexivity. In this model, $M \models^t P(\underline{a})$, and $M \models^t \neg P(\underline{d})$ (indeed, $M \not\models^t P(\underline{d})$). However, $M \models^t P(\underline{a}) \rightarrow P(\underline{b})$, $M \models^t P(\underline{b}) \rightarrow P(\underline{c})$, and $M \models^t P(\underline{c}) \rightarrow P(\underline{d})$. This means that all premises of a standard sorites can be tolerantly true, without forcing the conclusion to be tolerantly true. Importantly, this implies that *modus ponens* is not a valid inference, if validity is understood as preservation of t -truth. As we shall discuss in section 3, however, less radical departures from classical logic are still compatible with the t -validity of the tolerance principle.

Negation. That the semantics weakens the meaning of negation can be seen from the new clauses. The previous semantics was such that M tolerantly satisfied $\neg\phi$ provided M did not tolerantly satisfy ϕ . Here, M tolerantly satisfies $\neg\phi$ provided M does not *strictly* satisfy ϕ , which is a weaker requirement. A consequence of this is that: $\llbracket P \rrbracket^{c,M} \subseteq \llbracket P \rrbracket^{t,M}$, and similarly, $\llbracket \neg P \rrbracket^{c,M} \subseteq \llbracket \neg P \rrbracket^{t,M}$. This property of coarsening is preserved by conjunction, and therefore transfers to all formulae, as we shall prove in the next section. Conversely, it is easy to see that $\llbracket P \rrbracket^{s,M} \subseteq \llbracket P \rrbracket^{c,M}$, and similarly, $\llbracket \neg P \rrbracket^{s,M} \subseteq \llbracket \neg P \rrbracket^{c,M}$. This means that in the same way in which the tolerant extension *coarsens* the classical extension of a predicate, the strict extension *sharpens* it.

Borderlines. A third and central feature of the semantics is that it allows us to define what it is to be a borderline case of application of a vague predicate in a natural way. Given a T-model M , borderline cases of the application of a predicate P may be defined as those that fall between the tolerant extension and the strict extension of a predicate:³

³See also [?] where a very similar definition of borderlineness is proposed, but in a metric setting, in terms of Voronoi diagrams. More generally, our present definition of borderline cases bears a direct analogy to the definition of the *boundary* of a set in topology. Given a topology, the boundary of a set is defined as the difference between the closure of the set and the interior of that set (see e.g. [?]). Tolerant and strict extensions play the same role relative to the classical extension of a predicate in a T-model as do closure and interior relative to a set given a suitable topology. In our setting, in which the relation \sim_P is possibly non-transitive, we cannot straightforwardly equate the operators $\llbracket \cdot \rrbracket^s$ and $\llbracket \cdot \rrbracket^t$ with interior and closure operators I and C , so as to satisfy for every P : $I(\llbracket P \rrbracket^{c,M}) = \llbracket P \rrbracket^{s,M}$ and $C(\llbracket P \rrbracket^{c,M}) = \llbracket P \rrbracket^{t,M}$. However, we could get this correspondence rigorously by transforming non-transitive T-models into transitive models

Definition 10. Let $b(P)^M$, the borderline region of a predicate P , be defined as follows $b(P)^M := \llbracket P \rrbracket^{t,M} \setminus \llbracket P \rrbracket^{s,M}$.

Equivalently, the borderline area of a predicate P can be defined as the set of cases that are neither strictly P , nor strictly not P . This definition is reminiscent of one standard definition of a borderline case of P : a case that is neither definitely P nor definitely not P . Due to the duality of tolerant and strict truth, borderline cases can also be described on the present account as cases that are both tolerantly P and tolerantly $\neg P$.

An important consequence of this is that some contradictions can be tolerantly true. Consider, for instance, the same model M . (Recall that M consists of four elements a, b, c, d such that $a \sim_P b \sim_P c \sim_P d$, nothing else is \sim_P related except as required by symmetry and reflexivity, and $\llbracket P \rrbracket^{c,M} = \{a, b\}$.) In this model, the two individuals b and c around the cutoff between $\llbracket P \rrbracket^{c,M}$ and $\llbracket \neg P \rrbracket^{c,M}$ are both tolerantly P and tolerantly $\neg P$. The idea that borderline cases support contradictory responses for a predicate is not new. We find it, for instance, in the subvaluationist treatment of vagueness advocated by [?]. In the present case, however, unlike in subvaluationism, borderline cases furthermore satisfy the conjunction of these two properties, namely $M \models^t P(\underline{b}) \wedge \neg P(\underline{b})$, and similarly for c .

At this point, we should note that the semantics does not commit us to linking assertion exclusively to t -truth rather than s -truth or even c -truth. Because of that, further work needs to be done before we can evaluate whether the present predictions are welcome or unwelcome. Thus, we defer until section 4 a discussion of the connection between tolerance, strictness, and facts concerning assertion. In the next section, we first investigate the logical properties of our framework in greater detail.

2 Validities and Entailment: tolerant and strict

In this section, we characterize both tolerant and strict validities, and the corresponding notions of logical consequence for each notion, namely preservation of tolerant truth and preservation of strict truth. The first part of the section states some basic lemmas concerning validities. The second part gives us a generalization of those results by means of a correspondence between t -validity and LP-validity, and s -validity and K3-validity. An important caveat: in much of this section we restrict the characterization of validity and entailment to formulae that are free of I_P and identity predicates. We'll be explicit about when these special predicates are included (we'll call this the *full vocabulary*) and when they are not (the *restricted vocabulary*). We close the section with a brief comparison between the present framework and the frameworks of subvaluationism and supervaluationism on the one hand, and with more familiar semantics for LP and K3 on the other.

(see [?], where the operation is called layering).

The same analogy holds with the notions of inner and outer approximation to a set in the theory of rough sets (see [?]). Usual rough sets require an underlying set with an equivalence relation; if we allow the relation to be only reflexive and symmetric, the approach becomes very close to the present one.

2.1 t -validities and s -validities

Definition 11. A formula ϕ is t -valid ($\models^t \phi$) iff for every T -model M , $M \models^t \phi$; it is s -valid ($\models^s \phi$) iff for every T -model M , $M \models^s \phi$; and it is c -valid ($\models^c \phi$) iff for every T -model M , $M \models^c \phi$.

Definition 12. A formula ϕ is t -unsatisfiable ($\not\models^t \phi$) iff no T -model M is such that $M \models^t \phi$; it is s -unsatisfiable ($\not\models^s \phi$) iff no T -model M is such that $M \models^s \phi$; and it is c -unsatisfiable ($\not\models^c \phi$) iff no T -model M is such that $M \models^c \phi$.

Lemma 1. for any formula ϕ in the full vocabulary, and any T -model M , $M \models^c \phi \Rightarrow M \models^t \phi$, and $M \models^s \phi \Rightarrow M \models^c \phi$.

Proof. by simultaneous induction on \models^t and \models^s .

- Atomic predication: if $M \models^c P(\underline{a})$, then clearly $M \models^t P(\underline{a})$, since $a \sim_P a$. And if $M \models^s P(a)$, then clearly $M \models^c P(a)$ for the same reason.

- I_P and $=$: We've required already that $M \models^c aI_P b$ iff $M \models^t aI_P b$ iff $M \models^s aI_P b$, and similarly for $=$.

- Negation: if $M \models^c \neg\phi$, then $M \not\models^c \phi$, so by induction hypothesis, $M \not\models^s \phi$, which implies $M \models^t \neg\phi$. If $M \models^s \neg\phi$, then $M \not\models^t \phi$, and by induction hypothesis, $M \not\models^c \phi$, ie $M \models^c \neg\phi$.

- Conjunction: if $M \models^c \phi \wedge \psi$, then $M \models^c \phi$ and $M \models^c \psi$, and by induction hypothesis, $M \models^t \phi$ and $M \models^t \psi$, ie $M \models^t \phi \wedge \psi$. The case for \models^s is analogous.

- Universal quantification: if $M \models^c \forall x\phi$, then for all d in M , $M \models^c \phi[d/x]$, and by induction hypothesis, for all d , $M \models^t \phi[d/x]$, ie $M \models^t \forall x\phi$. The case for \models^s is analogous. \square

Corollary 1. If $\models^c \phi$, then $\models^t \phi$.

Proof. if $\models^c \phi$, then for every T -model M , $M \models^c \phi$. By Lemma 1, for every M , $M \models^t \phi$, hence $\models^t \phi$. \square

Corollary 2. If $\phi \models^c$, then $\phi \models^s$.

Proof. if $\phi \models^c$, then for every T -model M , $M \not\models^c \phi$, hence by Lemma 1, every M is such that $M \not\models^s \phi$; hence, $\phi \models^s$. \square

Lemma 2. Let M be a C -model of the form $\langle D, I \rangle$, and $M' = \langle D, I, \sim \rangle$ be the T -model obtained from M by letting $a \sim_P b$ iff $a = b$, for every P . Then for every formula ϕ in the restricted vocabulary, $M \models^c \phi$ iff $M' \models^c \phi$ iff $M' \models^t \phi$ iff $M' \models^s \phi$.

Proof. Obviously \sim_P is an equivalence relation in this case, hence M' is well-defined. The proof is by induction on ϕ .

- Atomic case: Suppose $M \models^c P(\underline{a})$, then there is a d in M , namely a itself, such that $d \sim_P a$, and $M' \models^c P(\underline{d})$, so $M' \models^t P(\underline{a})$. Likewise, since only a is P -similar to itself, for every $d \sim_P a$, $M' \models^c P(\underline{d})$, hence $M' \models^s P(\underline{a})$.

- Negation: suppose $M \models^c \neg\phi$. Then $M \not\models^t \phi$, and by induction hypothesis, $M' \not\models^t \phi$ and $M' \not\models^s \phi$. By definition, this means that $M' \models^s \neg\phi$, and that $M' \models^t \neg\phi$.

- Conjunction: immediate for both cases.

- Universal quantification: $M \models^c \forall x\phi$ iff for all d in D , $M \models^c \phi[\underline{d}/x]$ iff by induction hypothesis, $M' \models^t \phi[\underline{d}/x]$ iff $M' \models^s \phi[\underline{d}/x]$, iff $M' \models^t \forall x\phi$, iff $M' \models^s \forall x\phi$.

- Finally, we know from Lemma 1 that if strict and tolerant satisfaction coincide on M' , then they also coincide with classical satisfaction on M' . \square

We can now strengthen Corollaries 1 and 2 to biconditionals:

Theorem 1. *For every formula ϕ in the restricted vocabulary, $\models^c \phi$ iff $\models^t \phi$, and $\phi \models^c$ iff $\phi \models^s$.*

Proof. From Corollary 1, we know that $\models^c \phi$ entails $\models^t \phi$. Conversely, if $\not\models^c \phi$, then it means that there is a C-model M such that $M \not\models \phi$. From Lemma 2, it follows that the T-model M' obtained from M by taking \sim_P to be identity for every P is such that $M' \not\models^t \phi$. Hence $\not\models^c \phi$ entails $\not\models^t \phi$.

Similarly, we know from Corollary 2 that $\phi \models^c$ entails $\phi \models^s$. To show the converse, suppose that $\phi \not\models^c$. Then there is a C-model M such that $M \models \phi$. From Lemma 2, we know that the T-model M' derived from M as the Lemma specifies is such that $M' \models^s \phi$. Thus, $\phi \not\models^c$. \square

Despite the affinities between t and s on the one hand and classical logic on the other, there are some striking differences. For example, the set of s -validities in the restricted vocabulary is empty (and dually, every sentence in the restricted vocabulary is t -satisfiable). To establish this, the following lemma more than suffices:

Lemma 3. *For every formula ϕ in the restricted vocabulary, there is a T-model M such that $M \not\models^s \phi$ and $M \models^t \phi$.*

Proof. Let M be a T-model in which every atomic predicate P has a classical extension that is neither empty nor equal to the whole domain. For every pair of elements a and b in the domain of M , and for every predicate P , let $a \sim_P b$. By induction, one can show that for every ϕ , $M \not\models^s \phi$ and $M \models^t \phi$:

- Atomic case: clearly, for every formula of the form $P(\underline{a})$, $M \models^t P(\underline{a})$, since one can find a d P -similar to a that is classically P . Consequently, $M \models^t P(\underline{a})$. Dually, for every a , $M \not\models^s P(\underline{a})$, since a must be P -similar to some element that is not classically P .

- Negation: if $\phi = \neg\psi$. By induction hypothesis, $M \models^t \psi$, and $M \not\models^s \psi$. If $M \models^s \phi$, then by definition $M \not\models^t \psi$, which is a contradiction, so $M \not\models^s \phi$. Since $M \not\models^s \psi$, then $M \models^t \neg\psi$; that is, $M \models^t \phi$.

- Conjunction and Universal quantification: both cases are straightforward. □

From the previous lemma, it follows immediately that no formula ϕ in the restricted vocabulary is s -true in every T-model M , and every ϕ in the restricted vocabulary is t -true in at least one model; hence:

Theorem 2. *No formula ϕ in the restricted vocabulary is s -valid. Every formula ϕ in the restricted vocabulary is t -satisfiable.*

Over the restricted vocabulary, we see that tolerant validities coincide with classical validities, and that strict unsatisfiability coincides with classical unsatisfiability. On the other hand, we can see that no formula is tolerantly unsatisfiable, and that no formula is strictly valid.

Because of that, we can already observe that the logics induced by s -truth and t -truth do not coincide with supervaluationism ($[?]$, $[?]$) nor with subvaluationism ($[?]$). These frameworks associate the language with a set of admissible (classical) precisifications. Then, a sentence is supervaluationistically true if and only if it is classically true in *every* admissible precisification, and it is subvaluationistically true if and only if it is true in *at least one* admissible precisification. Based on these quantification patterns, one may have expected t -truth to coincide with sub-truth, and s -truth with super-truth. In both sub- and supervaluationism, however, validity for formulas coincides with classical validity; this implies in particular that s -validity is distinct from supervaluationist validity.⁴ Dually, classical contradictions are not subvaluationistically satisfiable; this implies that t -satisfaction does not coincide with subvaluationist satisfaction.

However, t -validities and s -validities appear to coincide exactly with logical validities in two well-known extensions of the logic FDE of first-degree entailment, namely with Priest's Logic of Paradox on the one hand (LP), and the strong Kleene logic on the other (K3). This coincidence is not fortuitous, as we proceed to show in the next subsection.

2.2 Correspondence with LP and K3

LP and its dual K3 are often presented as three-valued logics, and we present them here in this way. In what follows we use the values 1, 1/2, and 0. 1 and 0 can be read as truth and falsity, respectively; 1/2 indicates an intermediate status. Advocates of LP often understand the value 1/2 as applying in the overlap of truth and falsity, and advocates of K3 often understand it as applying in the gap between truth and falsity. For our immediate formal purposes, of course, it doesn't matter how we interpret this value. Entailment in K3 correspond to preservation of the value 1 from premises to conclusions, and entailment

⁴For detailed discussions of supervaluationist systems of consequence, see eg $[?]$, $[?]$, and $[?]$.

in LP on the other hand corresponds to preservation of nonzero values from premises to conclusions.

If we define s -entailment as preservation of strict truth, and t -entailment as preservation of tolerant truth, a natural correspondence immediately arises between the two frameworks: we let value 1 represent strict truth, 0 represent strict falsity, and $1/2$ represent borderline truth (in the sense we defined in the previous section). Tolerant truth then corresponds to assigning value 1 or $1/2$ to a formula (that is, nonzero value), since a formula is tolerantly true either if it is strictly true, or if it is borderline true.

2.2.1 MV-models and entailment

To establish the correspondence more formally, we first introduce the notion of an MV-model (for many-valued model). We use MV-models only over the restricted vocabulary, and do not at any point extend them to include I_P or $=$ predicates. We let $\min(A)$ denote the minimum value in the set A , and use $\min(A, B)$ to abbreviate $\min(\{A, B\})$.

Definition 13. An MV-model M is a tuple $\langle D, I \rangle$ such that:

- D is a non-empty domain of individuals; and
- I is a three-valued interpretation that works as follows:
 - For any name \underline{a} , $I(\underline{a}) \in D$;
 - For any predicate P , $I(P) \in \{1, 1/2, 0\}^D$;
 - $I(P(\underline{a})) = I(P)(I(\underline{a}))$;
 - $I(\neg\phi) = 1 - I(\phi)$;
 - $I(\phi \wedge \psi) = \min(I(\phi), I(\psi))$;
 - $I(\forall x\phi) = \min(\{I(\phi[\underline{d}/x]) : d \in D\})$

Definition 14. An MV-model $M = \langle D, I \rangle$ LP-satisfies a wff ϕ ($M \models^{LP} \phi$) iff $I(\phi) > 0$. An MV-model $M = \langle D, I \rangle$ K3-satisfies a wff ϕ ($M \models^{K3} \phi$) iff $I(\phi) = 1$.

We subsume all notions of c -consequence, s -consequence, t -consequence, LP-consequence, and K3-consequence under the following definition (for $X = c, s, \text{ or } t$, an X -model is a T-model; for $X = \text{LP or K3}$, an X -model is an MV-model):

Definition 15. For any logic X , let X -consequence be defined as follows: $\Gamma \models^X \Delta$ iff for every X -model M such that $M \models^X \gamma$ for every $\gamma \in \Gamma$, $M \models^X \delta$ for some $\delta \in \Delta$.

For our purposes here, we rely on available axiomatizations of logical consequence in LP and K3. In particular, we refer to [?, ?] or [?] for details. We rehearse some prominent features of these logics: both logics validate the classical rules of conjunction introduction, conjunction elimination, de Morgan laws for conjunction and negation, double negation introduction as well as elimination, universal generalization, and universal instantiation. In K3, moreover, every formula is entailed by a classical contradiction; in LP, dually, every classical validity is entailed by any formula.

2.2.2 Model correspondence

To transfer these results to s -consequence and t -consequence, we show that for every MV-model, we can construct an equivalent T-model, and vice versa.

Definition 16. A T-model M is equivalent to an MV-model M' over a set \mathcal{L} of wffs iff for every wff $\phi \in \mathcal{L}$:

- $M \models^t \phi$ iff $M' \models^{LP} \phi$, and
- $M \models^s \phi$ iff $M' \models^{K3} \phi$

Lemma 4. Let \mathcal{L} be our language, using only the restricted vocabulary. For every T-model $M = \langle D, I, \sim \rangle$, there is an MV-model equivalent over the language \mathcal{L} .

Proof. We define the equivalent MV-model $M' = \langle D', I' \rangle$ as follows:

- $D' = D$
- For any name \underline{a} , $I'(\underline{a}) = I(\underline{a})$
- For any predicate P and any $d \in D$:
 - If $M \models^s P(\underline{d})$, then $I'(P)(d) = 1$
 - If $M \not\models^t P(\underline{d})$, then $I'(P)(d) = 0$
 - Otherwise, $I'(P)(d) = 1/2$

By Lemma 1, we know that $\{\phi : M \models^s \phi\} \subseteq \{\phi : M \models^t \phi\}$, so the cases here are exclusive and exhaustive.

Now we show that M is equivalent to M' , by an induction on formula construction. The base case is immediate. Inductive case:

- \neg : Suppose $\phi = \neg\psi$, and the inductive hypothesis holds for ψ . Then $M \models^t \phi$ iff $M \not\models^s \psi$ iff $M' \not\models^{K3} \psi$ iff $I'(\psi) < 1$ iff $I'(\phi) > 0$ iff $M' \models^{LP} \phi$. Similarly, $M \models^s \phi$ iff $M \not\models^t \psi$ iff $M' \not\models^{LP} \psi$ iff $I'(\psi) = 0$ iff $I'(\phi) = 1$ iff $M' \models^{K3} \phi$.
- \wedge : Suppose $\phi = \psi \wedge \chi$, and the inductive hypothesis holds for ψ and χ . Then $M \models^t \phi$ iff $(M \models^t \psi \text{ and } M \models^t \chi)$ iff $(M' \models^{LP} \psi \text{ and } M' \models^{LP} \chi)$ iff $\min(I'(\psi), I'(\chi)) > 0$ iff $M' \models^{LP} \phi$. Similarly, $M \models^s \phi$ iff $(M \models^s \psi \text{ and } M \models^s \chi)$ iff $(M' \models^{K3} \psi \text{ and } M' \models^{K3} \chi)$ iff $\min(I'(\psi), I'(\chi)) = 1$ iff $M' \models^{K3} \phi$.
- \forall : Suppose $\phi = \forall x\psi$, and the inductive hypothesis holds for $\psi[d/x]$ for every d . Then $M \models^t \phi$ iff $(M \models^t \psi[d/x] \text{ for all } d \in D)$ iff $(M' \models^{LP} \psi[d/x] \text{ for all } d \in D)$ iff $\min(\{I'(\psi[d/x]) : d \in D\}) > 0$ iff $M' \models^{LP} \phi$. Similarly, $M \models^s \phi$ iff $(M \models^s \psi[d/x] \text{ for all } d \in D)$ iff $(M' \models^{K3} \psi[d/x] \text{ for all } d \in D)$ iff $\min(\{I'(\psi[d/x]) : d \in D\}) = 1$ iff $M' \models^{K3} \phi$.

□

Lemma 5. *For every MV-model $M = \langle D, I \rangle$, there is a T-model equivalent over the set of all wffs.*

Proof. The equivalent T-model $M' = \langle D', I', \sim \rangle$ will operate with an expanded domain. We define the T-model as follows:

- $D' = \{ \langle d, i \rangle : d \in D, i \in \{0, 1\} \}$
- For any name \underline{a} in the old language, $I'(\underline{a}) = \langle I(\underline{a}), 0 \rangle$. Add a new name \underline{a}' to the language for every old name \underline{a} , and let $I'(\underline{a}') = \langle I(\underline{a}), 1 \rangle$.
- For any predicate P and old name \underline{a} :
 - $I'(P)I(\underline{a}) = 1$ iff $I(P)I(\underline{a}) = 1$, and $I'(P)I(\underline{a}) = 0$ otherwise;
 - $I'(P)I(\underline{a}') = 1$ iff $I(P)I(\underline{a}) > 0$, and $I'(P)I(\underline{a}') = 0$ otherwise;
 - \sim_P contains $\langle I'(\underline{a}), I'(\underline{a}) \rangle$, $\langle I'(\underline{a}), I'(\underline{a}') \rangle$, $\langle I'(\underline{a}'), I'(\underline{a}) \rangle$, and $\langle I'(\underline{a}'), I'(\underline{a}') \rangle$
- For any predicate P , \sim_P contains nothing more than is given for each old name \underline{a} in the last clause

We now show that the models are equivalent over the old language. Proof is by induction on formula construction. Base step is where all the action is.

$M \models^{K3} P(\underline{a})$ iff $I(P)I(\underline{a}) = 1$ iff $I'(P)I(\underline{a}) = I'(P)I(\underline{a}') = 1$ iff $(M' \models^c P(\underline{a}))$ for all $d \in D$ such that $d \sim_P a$ (since $a \sim_P a$, $a' \sim_P a$, and nothing else $\sim_P a$) iff $M' \models^s Pa$. Similarly, $M \models^{LP} P(\underline{a})$ iff $I(P)I(\underline{a}) > 0$ iff $I'(P)I(\underline{a}') = 1$ iff $(M' \models^c P(\underline{a}))$ for some $d \in D$ such that $d \sim_P a$ iff $M' \models^t P(\underline{a})$.

Inductive step is quick:

\neg : $M \models^{LP} \neg\phi$ iff $I(\phi) < 1$ iff $M \not\models^{K3} \phi$ iff $M' \not\models^s \phi$ iff $M' \models^t \neg\phi$. Similarly, $M \models^{K3} \neg\phi$ iff $I(\phi) = 0$ iff $M \not\models^{LP} \phi$ iff $M' \not\models^t \phi$ iff $M' \models^{K3} \neg\phi$.

\wedge : $M \models^{LP} \phi \wedge \psi$ iff $\min(I(\phi), I(\psi)) > 0$ iff $(M \models^{LP} \phi \text{ and } M \models^{LP} \psi)$ iff $(M' \models^t \phi \text{ and } M' \models^t \psi)$ iff $M' \models^t \phi \wedge \psi$. Similarly, $M \models^{K3} \phi \wedge \psi$ iff $\min(I(\phi), I(\psi)) = 1$ iff $(M \models^{K3} \phi \text{ and } M \models^{K3} \psi)$ iff $(M' \models^s \phi \text{ and } M' \models^s \psi)$ iff $M' \models^s \phi \wedge \psi$.

\forall : $M \models^{LP} \forall x\phi$ iff $\min(\{\phi[d/x] : d \in D\}) > 0$ iff $(M \models^{LP} \phi[d/x] \text{ for all } d \in D)$ iff $(M' \models^t \phi[d/x] \text{ for all } d \in D')$ iff $M' \models^t \forall x\phi$. Similarly, $M \models^{K3} \forall x\phi$ iff $\min(\{\phi[d/x] : d \in D\}) = 1$ iff $(M \models^{K3} \phi[d/x] \text{ for all } d \in D)$ iff $(M' \models^s \phi[d/x] \text{ for all } d \in D')$ iff $M' \models^s \forall x\phi$.

□

Theorem 3. *For all sets of wff Γ and Δ , $\Gamma \models^t \Delta$ iff $\Gamma \models^{LP} \Delta$, and $\Gamma \models^s \Delta$ iff $\Gamma \models^{K3} \Delta$.*

Proof. Immediate. □

2.3 The full vocabulary

So far, our discussion has focused mainly on our restricted vocabulary—in which neither identity nor our family of similarity relations is expressible. Here, we consider the effects created by the expansion to our full vocabulary, in which both identity and similarity predicates occur.

Language As before, we use the language of the quantified monadic predicate calculus, including an identity relation $=$, and, for every unary predicate P , a binary relation I_P , which will express the similarity relation \sim_P .

These are interpreted as described above, in §1. Recall that both identity and similarity are always interpreted classically; there is no difference between a model's strictly satisfying, classically satisfying, or tolerantly satisfying any sentence built entirely from identity or similarity relations.⁵

2.3.1 Identity

In this section, we consider the effect that introducing identity has on our consequence relations. The first thing to note is that introducing identity breaks the proofs that $\models^t = \models^{LP}$ and $\models^s = \models^{K3}$. For consider sentences like $(\forall x \forall y (x = y)) \rightarrow (Pa \vee \neg Pa)$. Although this sentence is not valid in K3 (there might be only one thing, and still that thing might satisfy neither P nor its negation), it is strictly satisfied (and therefore both classically and tolerantly satisfied) by every T-model. After all, although $Pa \vee \neg Pa$ can fail to be strictly satisfied in a T-model M , it can only do so when there are two things a and b in M 's domain such that $a \sim_P b$, $M \models^c Pa$ and $M \models^c \neg Pb$. This means a and b must be distinct.

Similarly, although $(\forall x \forall y (x = y)) \wedge Pa \wedge \neg Pa$ is satisfiable in LP (there might be only one thing, and still that thing might satisfy both P and its negation), it cannot be tolerantly satisfied (and therefore cannot be classically or strictly satisfied) by any T-model. After all, although $Pa \wedge \neg Pa$ can be tolerantly satisfied by a T-model M , it can only do so in precisely the same circumstances as are required for $M \not\models^s Pa \vee \neg Pa$. And again, that requires two distinct objects in the domain.

So although strict and tolerant consequence are very similar to K3 and LP consequence, and indeed are the same in the restricted vocabulary, once we expand our vocabulary to include identity we see that they are distinct. This is because, while MV-models allow us to assign nonclassical values directly to atomic predications, T-models allow us to do so only via covert quantification over the domain. If the domain includes only one thing, T-models can no longer provide non-classical values for any atomic predications. Thinking along these lines yields the following:

Fact 3. *If $\Gamma \models^c \Delta$, then, where \models^m is either \models^s or \models^t :*

⁵If we were to relax these constraints, much of what we are about to claim would fail; we do not know precisely what the resulting systems would look like, although we are interested in pursuing the issue in future work.

- $\Gamma \cup \{\forall x \forall y (x = y)\} \models^m \Delta$, and
- $\Gamma \models^m \Delta \cup \{\neg \forall x \forall y (x = y)\}$

Proof. Suppose $\Gamma \models^c \Delta$, and suppose M is a countermodel for either of the four inferences in the consequent of Fact 3. If M is a countermodel to the first inference, it must (strictly or tolerantly, whichever matters) satisfy $\forall x \forall y (x = y)$, and thus there must be only one member of M 's domain. If M is a countermodel to the second inference, it must fail to (strictly or tolerantly, again) satisfy $\neg \forall x \forall y (x = y)$; again, there must be only one member of M 's domain. So no matter which inference M is supposed to be a counterexample to, there is only one member of M 's domain.

By examination of the clauses for atomic predication, we can see that this requires that $M \models^s Pa$ iff $M \models^c Pa$ iff $M \models^t Pa$, for any atomic sentence Pa . What's more, we know that identity and similarity predications are satisfied in all three ways if at all. Induction on formula complexity shows that, for any sentence ϕ , $M \models^s \phi$ iff $M \models^c \phi$ iff $M \models^t \phi$. We know that it's not the case both that $M \models^c \gamma$ for every $\gamma \in \Gamma$ and that $M \not\models^c \delta$ for every $\delta \in \Delta$ (since $\Gamma \models^c \Delta$). But that means it can't be both that $M \models^s \gamma$ for every $\gamma \in \Gamma$ and that $M \not\models^s \delta$ for every $\delta \in \Delta$; nor can it be both that $M \models^t \gamma$ for every $\gamma \in \Gamma$ and that $M \not\models^t \delta$ for every $\delta \in \Delta$. Thus, M is not a counterexample after all to either of the inferences in question. Contradiction. \square

Similar effects can be created by restricting the domain in other ways. For example, $\forall x \forall y \forall z (x = y \vee x = z \vee y = z)$, $Pa \wedge \neg Pa \models^t Pb \wedge \neg Pb$. Given the first premise (that there are at most two things), the only way for the second premise to be tolerantly satisfied by a model M is if there are two objects in M 's domain that bear \sim_P to each other, exactly one of which classically satisfies P . Whichever one of these objects b picks out, $M \models^t Pb \wedge \neg Pb$. This argument is not valid in LP, however. For similar reasons, $\forall x \forall y \forall z (x = y \vee x = z \vee y = z)$, $Pa \vee \neg Pa \models^s Pb \vee \neg Pb$, but the argument is not K3-valid. Still more validities can be found along these lines: so long as one object in a model's domain tolerantly satisfies $Px \wedge \neg Px$, another object in the domain must as well; and so long as every object but one in a model's domain strictly satisfies $Px \vee \neg Px$, the last one must as well.

2.3.2 Similarity

The biggest difference introduced by similarity has to do with the principle of tolerance for a predicate P : $\forall x \forall y ((Px \wedge x I_P y) \rightarrow Py)$. This principle, strictly speaking, cannot be stated in LP, since the language of LP does not include our special I_P predicates. We might state an analogue of the principle by adding to LP I_P predicates required to be reflexive and symmetric, but even then the principle would not be LP-valid.⁶ However, the principle is tolerantly valid on our models, as shown in §1. Similarly, although the negation of this tolerance principle is satisfiable in K3-augmented-with-reflexive-symmetric- I_P -predicates,

⁶For a countermodel, consider an LP-model with two members of the domain, a and b . Let I_P be the universal relation on the domain, and let $I(Pa) = 1$ and $I(Pb) = 0$. This model is an LP-counterexample to tolerance. This is possible because LP has no way to recognize the connection between I_P and P .

it is not strictly satisfiable on our models (since to be strictly unsatisfiable just is to have a tolerantly-valid negation).

There will be more differences created by these similarity relations when we examine more articulated notions of consequence in §3; we'll discuss those differences there.

2.4 Comparisons

The correspondence we established between tolerant semantics and LP on the one hand, and between strict semantics and K3 on the other, is worth commenting on for several reasons.

First of all, as briefly emphasized at the end of section 2.1, one might have expected *s*-truth and *t*-truth to behave like subvaluationist truth and supervaluationist truth, based on the *prima facie* analogy between the quantification patterns involved in each case. However, what we see is that the semantics make quite distinct predictions, in particular regarding borderline cases. In sub- and super-valuationism, borderline cases are predicted to satisfy classical validities. In particular, every individual is predicted to be tall or not tall, including an individual who is a borderline case of tallness. Conversely, no individual can be both tall and not tall, not even borderline cases. In the present case, by contrast, every individual is *tolerantly* tall or not tall, but some individuals, namely borderline cases, are tolerantly both. By contrast, not all individuals are *strictly* tall or not tall in a model, since borderline cases are predicted to be neither strictly tall, nor strictly not tall. In our view, and as argued by [?] in relation to the application of LP and K3 to vagueness, these specific predictions for borderline cases are not unwelcome. Rather, unlike for sub- and super-valuationism, they imply that a special semantic status is acknowledged of borderline cases.

A second feature of our target semantics is that, while it coincides with the predictions of the many-valued logics LP and K3, it answers to a distinct motivation. Rather than seeing truth as a unified notion to which sentences might answer in three (or more) different ways, our approach posits distinct notions of truth, each of which a sentence may have or fail to have, but none of which is many-valued.

A third feature of the present semantics is that it gives us a psychologically plausible characterization of borderline cases as cases equisimilar with cases that would support opposed categorizations if subjects were forced to be bivalent. This characterization of borderline cases agrees with other accounts based on the notion of similarity.⁷ Furthermore, the characterization of borderline cases within T-models allows us to make sense of the idea that borderline cases are shifty or ambivalent cases, namely cases that can be conceptualized under opposed points of view. We shall say more about this in §?? below.

A fourth feature of the present approach concerns the characterization we gave of logical consequence. Above in §1.4, we pointed out that *t*-semantics can make the main premise of a sorites true without paradox, but only because modus ponens is no longer a *t*-valid (or

⁷See [?], where borderline cases of color predicates, in particular, are basically characterized as cases equidistant between prototypes in the relevant conceptual space.

LP-valid) rule of inference. Just as in LP, one objection to the present treatment might be that we fail to adequately capture the real meaning of the conditional when we model it in terms of negation and conjunction in the present framework.⁸ More generally, the definition we adopted for t -entailment may appear to depart too much from classical logic to provide a decent solution to the sorites.

In the next section, however, we examine various alternatives to the definition of tolerant entailment we examined here. Given that we have three notions of truth we can work with, namely tolerant, classical, and strict, there is indeed room for more consequence relations that just preservation of classical truth, preservation of tolerant truth, or preservation of strict truth. At the end of the section, we explain how the present framework allows us to defuse the sorites paradox.

3 Mixed consequence

Although we have discussed three distinct notions of consequence thus far—strict, classical, and tolerant—our models in fact give us the materials to define a number of additional notions of consequence. Some of these additional notions, we believe, are philosophically interesting in their own right. In fact, we believe that the most natural notion of consequence flowing from these models is not any of the three we have discussed so far.

First, a structured way to talk about many different consequence relations:

Definition 17. *Where m and n are s , c , or t , and Γ and Δ are sets of formulas: $\Gamma \vDash^{mn} \Delta$ iff every T -model M such that $M \vDash^m \gamma$ for every $\gamma \in \Gamma$ is also such that $M \vDash^n \delta$ for some $\delta \in \Delta$.*

That is, for an argument to be mn -valid is for every model that m -satisfies all of the premises to n -satisfy at least one of the conclusions. From our three notions of satisfaction, this immediately generates nine notions of consequence. The first thing to notice is that three of our nine consequence relations have already been characterized. That's because $\vDash^{tt} = \vDash^t$, $\vDash^{cc} = \vDash^c$, and $\vDash^{ss} = \vDash^s$. These unmixed relations have been dealt with earlier in the paper. We also generalize the notions of formula validity and unsatisfiability:

Definition 18. *A formula ψ is mn -valid ($\vDash^{mn} \psi$) iff $\emptyset \vDash^{mn} \psi$. A formula ψ is mn -unsatisfiable ($\psi \vDash^{mn} \emptyset$) iff $\psi \vDash^{mn} \emptyset$.*

Note that mn -validity amounts to n -validity (since whatever m is, every model m -satisfies the empty set), and that mn -unsatisfiability amounts to m -unsatisfiability (since whatever n is, every model n -satisfies the empty set).

The purpose of this section is to explore these notions of consequence, characterize the relations between all nine notions, and offer some philosophical reasons for being interested in ‘mixed’ consequence relations (that is, consequence relations where $m \neq n$). We first define the notion of *duality* relating some of our nine consequence relations, then we compare these relations attending to their respective strength.

⁸See [?] for discussion of this point on LP.

3.0.1 Duality

Definition 19 (Dual consequence relation). Let \models^x be a notion of logical consequence. Its *dual* is the notion of logical consequence \models^y such that: $\Gamma \models^x \Delta$ iff $\neg(\Delta) \models^y \neg(\Gamma)$ (where $\neg(\Delta) = \{\neg\delta \mid \delta \in \Delta\}$)

The duality between notions of logical consequence is based on the duality of the notions of satisfaction and the duality of \forall and \exists in the definition of logical consequence. Though this fact is perhaps a bit too obvious to require a proof, we might express the relations of duality in a synthetic way as follows.

Definition 20. For a notion of satisfaction indexed as x , $d(x) = y$ just in case, for any M , $M \models^x \neg\varphi$ iff $M \not\models^y \varphi$

Lemma 6. \models^{mn} is the dual of $\models^{d(n)d(m)}$

Proof. Assume: $\Gamma \models^{mn} \Delta$, then:

For every M : if $\forall\gamma \in \Gamma, M \models^m \gamma$ then $\exists\delta \in \Delta, M \models^n \delta$
iff

For every M : if $\forall\delta \in \Delta, M \not\models^n \delta$ then $\exists\gamma \in \Gamma, M \not\models^m \gamma$
iff

For every M : if $\forall\delta \in \Delta, M \models^{d(n)} \neg\delta$ then $\exists\gamma \in \Gamma, M \models^{d(m)} \neg\gamma$
iff

$\neg(\Delta) \models^{d(n)d(m)} \neg(\Gamma)$ □

This yields immediately the duality relations over the possible combinations of consequence relations since we already know that $d(c) = c$, $d(s) = t$ and $d(t) = s$. More particularly:

1. \models^{cc} , \models^{st} and \models^{ts} are self-dual.
2. \models^{ss} is the dual of \models^{tt} .
3. \models^{sc} is the dual of \models^{ct} .
4. \models^{cs} is the dual of \models^{tc} .

3.0.2 Relations between consequence relations

Lemma 7. For any notion of satisfaction m , $\models^{tm} \subseteq \models^{cm} \subseteq \models^{sm}$. For any notion of satisfaction m , $\models^{ms} \subseteq \models^{mc} \subseteq \models^{mt}$.

Proof. Since we know that, for any model M , $\{\phi : M \models^s \phi\} \subseteq \{\phi : M \models^c \phi\} \subseteq \{\phi : M \models^t \phi\}$ (Lemma 1), it follows that if a model M is a sm -counterexample to an argument, it is also a cm -counterexample, and if it's a cm -counterexample, it's also a tm -counterexample. Similarly, if a model is a mt -counterexample to an argument, it must also be a mc -counterexample, and if it's an mc -counterexample, it must also be an ms -counterexample. □

This lemma answers some questions about the relations between our nine notions, but not all. We go on to complete the picture, first in the restricted vocabulary we used earlier (restricted to exclude both identity and similarity relations), and then in the full vocabulary.

3.1 Restricted vocabulary

It is useful to keep in mind the following facts:

1. For any T-model M , $M \models^s \varphi \Rightarrow M \models^c \varphi \Rightarrow M \models^t \varphi$ (Lemma 1).
2. For any C-model M , there is a T-model M' s.t. $M' \models^c \phi$ iff $M' \models^t \phi$ iff $M' \models^s \phi$ iff $M \models \phi$ (Lemma 2).
3. Every formula is t -satisfiable and no formula is s -valid (Lemma 3).

\models^{st} , \models^{cc} , \models^{sc} , and \models^{ct} coincide

Lemma 8. $\Gamma \models^{st} \Delta \Rightarrow \Gamma \models^{cc} \Delta$

Proof. Assume $\Gamma \not\models^{cc} \Delta$, then:

$$\begin{aligned} \exists M : \forall \gamma \in \Gamma \ M \models^c \gamma \text{ and } \forall \delta \in \Delta \ M \not\models^c \delta \\ \Downarrow \\ \exists M' : \forall \gamma \in \Gamma \ M' \models^s \gamma \text{ and } \forall \delta \in \Delta \ M' \not\models^t \delta \end{aligned}$$

□

Since we know from Lemma 7 that $\models^{cc} \subseteq \models^{st}$, this shows that $\models^{cc} = \models^{st}$. Lemma 7 also gives us $\models^{cc} \subseteq \models^{sc} \subseteq \models^{st}$ and $\models^{cc} \subseteq \models^{ct} \subseteq \models^{st}$, so we can conclude that $\models^{cc} = \models^{sc} = \models^{ct} = \models^{st}$.

\models^{tc} is strictly weaker than \models^{tt} and \models^{cs} is strictly weaker than \models^{ss}

We know from Lemma 7 that $\models^{tc} \subseteq \models^{tt}$ and that $\models^{cs} \subseteq \models^{ss}$. Now we show that $\Gamma \models^{tt} \Delta \not\Rightarrow \Gamma \models^{tc} \Delta$ and $\Gamma \models^{ss} \Delta \not\Rightarrow \Gamma \models^{cs} \Delta$. In general, $\varphi \models^{tt} \varphi$ and $\varphi \models^{ss} \varphi$ but, for example, $Px \not\models^{tc} Px$ and $Px \not\models^{cs} Px$. After all, we can have a model M such that $M \models^c Px$ but $M \not\models^s Px$, or such that $M \models^t Px$ but $M \not\models^c Px$. Thus, the relations \models^{tc} and \models^{cs} are not reflexive.

On the other hand, some instances of reflexivity do hold. For example, $\forall x Px \models^{cs} \forall x Px$, and $\exists x Px \models^{tc} \exists x Px$. One must be wary here: $\forall x (Px \vee \neg Px) \not\models^{cs} \forall x (Px \vee \neg Px)$, and $\exists x (Px \wedge \neg Px) \not\models^{tc} \exists x (Px \wedge \neg Px)$. Thus, these logics do not obey uniform substitution, either.

These two logics are, however, transitive.⁹ Let's consider \models^{tc} . Suppose $\Gamma, A \models^{tc} \Delta$ and $\Gamma \models^{tc} A, \Delta$. There must be no model M such that $M \models^t \gamma$ for every $\gamma \in \Gamma \cup \{A\}$ and

⁹The version of transitivity we consider here—if both $\Gamma, A \models \Delta$ and $\Gamma \models A, \Delta$ are valid, then $\Gamma \models \Delta$ is valid—is one of many possible variations. In fact, tc and cs satisfy any version of transitivity we're aware of.

$M \not\models^c \delta$ for every $\delta \in \Delta$. Similarly, there must be no model M such that $M \models^t \gamma$ for every $\gamma \in \Gamma$ and $M \not\models^c \delta$ for every $\delta \in \Delta \cup \{A\}$. Now, suppose for reductio there is a model M' such that $M' \models^t \gamma$ for every $\gamma \in \Gamma$ and $M' \not\models^c \delta$ for every $\delta \in \Delta$. It must be that $M' \not\models^t A$, and it must be that $M' \models^c A$. But this is impossible (by Lemma 1). So there is no such M' , and $\Gamma \models^{tc} \Delta$. The proof is the same for \models^{cs} , *mutatis mutandis*.

Exactly what inferences hold in these logics? Later, in §3.3, we will provide a tableau-based proof theory that is sound and complete for all nine notions of consequence. For now, we hope these remarks are enough to convey the flavor; \models^{tc} and \models^{cs} are odd logics indeed.

\models^{ts} is strictly weaker than both \models^{tc} and \models^{cs}


Lemma 9. $\Gamma \models^{ts} \Delta$ is empty.

Proof. $\Gamma \not\models^{ts} \Delta$ just in case there is a model M such that: $\forall \gamma \in \Gamma M \models^t \gamma$ and $\forall \delta \in \Delta M \not\models^s \delta$. Now, Lemma 3 shows that there is a model M' in which for every formula φ , $M' \models^t \varphi$ and $M' \not\models^s \varphi$. Thus, $\Gamma \not\models^{ts} \Delta$ for any Γ and Δ . \square

Since both \models^{tc} and \models^{cs} are non-empty, \models^{ts} is strictly weaker than both. (Indeed, \models^{ts} —the empty relation—is the weakest possible consequence relation.)

3.1.1 Summing up


Regarding the nine notions of logical consequence that we might define in the present framework, there are six distinct consequence relations on the restricted vocabulary. \models^{cc} , \models^{ct} , \models^{sc} , and \models^{st} coincide. \models^{tt} and \models^{ss} are distinct and both strictly weaker than \models^{cc} .¹⁰ \models^{tc} and \models^{cs} are also distinct, and are strictly weaker than \models^{tt} and \models^{ss} respectively. Finally, \models^{ts} is strictly weaker than both \models^{tc} and \models^{cs} .



¹⁰This can be proved parallel to the above inclusions, or it follows directly from the facts that *tt*-consequence is LP-consequence, *ss*-consequence is K3-consequence, and *cc*-consequence is classical consequence.

3.2 Full vocabulary

Unlike the restricted vocabulary, in the full vocabulary all nine notions of consequence result in distinct consequence relations, arranged by strength as follows:



We can demonstrate the distinctness of the four strongest relations as follows: the inference from $\{Pa, aI_Pb\}$ to $\{Pb\}$ is sc , ct , and st -valid, but not cc -valid. Tolerance—the sentence $\forall x\forall y((Px \wedge xI_Py) \rightarrow Py)$ —is ct and st -valid, but not cc or sc -valid. Tolerance’s negation is sc and st -unsatisfiable, but cc and ct -satisfiable. What’s more, st -validity is not simply the union of sc - and ct -validity, since the inference from $\{Pa, aI_Pb, bI_Pc\}$ to $\{Pc\}$ is st -valid, but is neither sc - nor ct -valid.

Lemma 7, recall, holds for the full vocabulary as well. Further, all the examples that showed distinctness between systems in §3.1 still hold. The only remaining question, then, is whether \models^{tt} and \models^{ss} are both still weaker than \models^{cc} , or whether the new vocabulary disrupts that relationship. In fact, the new vocabulary does disrupt that relationship: recall that tolerance is tt -valid, but not cc -valid, and its negation is ss -unsatisfiable, but not cc -unsatisfiable. So no inclusions hold between \models^{tt} , \models^{ss} , and \models^{cc} in the full vocabulary.

The deduction theorem (in the form $\Gamma \models^{mn} \Delta$ iff $\models^{mn} \bigwedge \Gamma \rightarrow \bigvee \Delta$)¹¹ holds for some of our relations, but not all:

Lemma 10. *The deduction theorem ($\Gamma \models^{mn} \Delta$ iff $\models^{mn} \bigwedge \Gamma \rightarrow \bigvee \Delta$) holds for a consequence relation \models^{mn} iff $m = d(n)$ (that is, iff \models^{mn} is self-dual).*

Proof. For the right-to-left direction of Lemma 10, assume that $m = d(n)$. Then:

- $\Gamma \models^{mn} \Delta$ iff
- every T-model M is such that either $M \not\models^m \bigwedge \Gamma$ or else $M \models^n \bigvee \Delta$ iff

¹¹Since we do not include infinitary conjunction or disjunction in our language, here Γ and Δ must be assumed to be finite sets.

- every T-model M is such that either $M \models^{d(m)} \neg \bigwedge \Gamma$ or else $M \models^n \bigvee \Delta$ iff
- every T-model M is such that either $M \models^n \neg \bigwedge \Gamma$ or else $M \models^n \bigvee \Delta$ iff
- every T-model M is such that $M \models^n \bigwedge \Gamma \rightarrow \bigvee \Delta$ iff
- $\models^{mn} \bigwedge \Gamma \rightarrow \bigvee \Delta$

For the left-to-right direction: Assume the deduction theorem holds for \models^{mn} . Recall that $\models^{mn} \phi$ iff $\models^{xn} \phi$, for any x, ϕ ; that is, formula validity depends only on the second notion of satisfaction involved. We know that the deduction theorem holds for $\models^{d(n)n}$, so we can conclude that $\Gamma \models^{mn} \Delta$ iff $\models^{mn} \bigwedge \Gamma \rightarrow \bigvee \Delta$ iff $\models^{d(n)n} \bigwedge \Gamma \rightarrow \bigvee \Delta$ iff $\Gamma \models^{d(n)n} \Delta$. But as we've seen, there are no notions x, m, n of satisfaction such that $\models^{mn} = \models^{xn}$ and $m \neq x$. So $m = d(n)$. □

Thus, the deduction theorem holds for *st*, *cc*, and *ts* only, in the full vocabulary. (In the restricted vocabulary, it holds in addition for *sc* and *ct*, since those are the same as *cc* and *st* in the absence of = and I_P .)

3.2.1 Choosing a consequence relation: modus ponens, tolerance, and the deduction theorem

Let's see how the stronger notions of consequence here manage to validate tolerance while avoiding soritical reasoning. Consider a soritical sequence of people, arranged by height. Let P be 'tall'; thus I_P will be the usual similarity relation in respect of height. Thus, we have a sequence $\langle a_1, a_2, \dots, a_n \rangle$, where the first i members are classically P and the remainder are not, and such that $a_i \sim_P a_{i+1}$ for $1 \leq i \leq n$. Let's focus on *ct*-consequence, for concreteness. We know that Pa_1 , $a_1 I_P a_2$, and $a_2 I_P a_3$ all hold, and we know that $Pa_1 \wedge a_1 I_P a_2 \wedge a_2 I_P a_3 \models^{ct} Pa_2 \wedge a_2 I_P a_3$. What's more, we know that $Pa_2 \wedge a_2 I_P a_3 \models^{ct} Pa_3$. It seems we are being led down the soritical series, forced to conclude first Pa_2 , then Pa_3 , and so on. (We might need to conjoin in more facts about the similarity relation, but nothing should stop us from doing that.) However, this is not the case. Although the above-mentioned inferences do hold in *ct*, $Pa_1 \wedge a_1 I_P a_2 \wedge a_2 I_P a_3 \not\models^{ct} Pa_3$. That is, \models^{ct} is not transitive.¹² So although it validates each step of the soritical reasoning, it does not validate chaining those steps together. *sc* and *st* behave similarly in this regard.¹³

¹²That is, it fails even the simple transitivity principle: if $\phi \models \psi$ and $\psi \models \chi$, then $\phi \models \chi$. There are more general transitivity principles one might consider (with side premises and/or side conclusions), but violation of this simple transitivity principle suffices for violation of these as well.

¹³The only difference here occurs with *st*, because $Pa \wedge a I_P b \wedge b I_P c \models^{st} Pc$. It stops there, though, as $Pa \wedge a I_P b \wedge b I_P c \wedge c I_P d \not\models^{st} Pd$. Informally, *sc* and *ct* allow us to take only one step, and no more, along the similarity relation, while *st* allows us two steps, and no more. In all cases, though, it's the 'and no more' that blocks the threatening soritical reasoning.

Why would we be interested in nontransitive consequence relations?¹⁴ There are a few reasons. For one thing, we seem to reason nontransitively in the presence of soritical sequences. If we want to capture that reasoning, we must use a nontransitive relation to do the capturing. Additionally, one might think, as we do, that simply using a consequence relation on which tolerance is valid (as it is on \models^{tt}) isn't quite getting at what we want, if modus ponens isn't valid on that consequence relation (as it is not on \models^{tt}). In fact, modus ponens is valid on all four of our strongest relations, since it's *cc*-valid, and anything *cc*-valid is *sc*-, *ct*-, and *st*- valid too. Thus, both *ct*- and *st*-validity are relations on which both tolerance and modus ponens are valid. Such a relation can't reasonably be transitive; it would force us to reason soritically. But a non-transitive relation fits the bill nicely.

Indeed, when it comes to choosing a consequence relation to focus on, we prefer *st* to *sc* on the grounds that tolerance is *st*-valid, and it seems to us that tolerance ought to be valid. We prefer *st* to *ct* because the deduction theorem holds for *st*, and it seems to us that the deduction theorem ought to hold. *st* is the only notion of consequence that satisfies these two desiderata. Thus, we think *st* is the best-motivated of our consequence relations; it validates tolerance, satisfies the deduction theorem, and supports modus ponens. It is quite a reasonable relation for reasoning with vague predicates.

It's worth noting that the semantics of *st* bear similarities to other semantics that have been proposed for consequence relations for vague language. For example, [?] presents an orthodox fuzzy semantics (with sentences taking values from the closed real interval $[0,1]$) for atomic sentences and the connectives, but defines consequence as follows (notation changed):

$\Gamma \models \delta$ iff every model on which every $\gamma \in \Gamma$ takes a value strictly greater than .5 is also such that δ takes a value greater than or equal to .5.

Note that this consequence relation sets a stricter standard for its premises than for its conclusion, just as *st* does (and *ct* and *sc* do). However, Smith's consequence relation, unlike ours, is transitive; in fact, it is the consequence relation of classical logic. (Recall that, in the restricted vocabulary, the same is true of *st*, *ct*, and *sc*.)

Another relative of the present semantics is to be found in [?]. Zardini considers a different sort of model, but again defines consequence as depending upon two distinct standards for satisfaction, with the premises held to a higher standard than the conclusions, just as Smith and we do. Like our semantics, and unlike Smith's, Zardini's semantics results in a nontransitive consequence relation. Unlike our semantics, however, Zardini's consequence relation is a weakening of the classical one. For example, the inference from ϕ , $\phi \rightarrow \psi$, and $\psi \rightarrow \chi$ to χ is invalid for Zardini, but it is *cc*-valid, and so *st*, *ct*, and *sc* valid as well.

When it comes to our weaker relations (*cs*, *tc*, and *ts*), they are no less weird in the full vocabulary than they are in the restricted vocabulary. *cs* and *tc* still don't obey uniform substitution, for one thing. *ts*, however, is now no longer the empty relation, since we've

¹⁴Some might argue that if it's nontransitive it's not a consequence relation. This seems to us merely a terminological issue.

introduced vocabulary (identity and similarity) that don't differ from tolerant to strict satisfaction. Thus, some inferences are now *ts*-valid. For example, $a = b \vDash^{ts} b = a$; $aI_P b, b = c \vDash^{ts} cI_P a$, $\forall x \forall y (x = y) \vDash^{ts} Pa \vee \neg Pa$, and so on. Of course, these inferences are then valid in all nine of our relations, as all are extensions of *ts*. In fact, Fact 3 holds, not just for \vDash^{ss} and \vDash^{tt} , as it was stated, but for all nine notions of consequence, including *ts*.

3.3 Tableaux

Although it is widely recognized that using a non-transitive entailment relation might solve the sorites paradox (see eg [?, p. 20]), this is sometimes taken to be problematic. For example, [?] claims that transitivity is essential to any notion of proof. One way to meet that objection is to provide a proof system. Notice that it is standardly taken to be hard to give a proof theory for such a logic (see [?], for instance).

In fact, there is a single tableau system that is sound and complete for all nine notions of consequence discussed above, in the full vocabulary.

Definition 21. A tagged sentence is something of the form ϕ, m where ϕ is a sentence and m is either *s*, *c*, or *t*.

Definition 22. A *T*-model M satisfies a tagged sentence ϕ, m iff $M \vDash^m \phi$.

The nodes on our tableaux are tagged sentences rather than just sentences.¹⁵ Depending on the main (and sometimes secondary) connective in a tagged sentence, we apply the appropriate rule to it to generate more tagged sentences. For familiar connectives ($\neg, \wedge, \vee, \forall, \exists, =$), we simply use ordinary tableau rules, and carry the tag along. A pair of examples:

$\phi \rightarrow \psi, s$	$\neg(\phi \rightarrow \psi), s$
$\neg\phi, s \quad \psi, s$	ϕ, s $\neg\psi, s$

There are a few novel rules, however. The first batch covers the interaction between predication and similarity relations:

Pu, s $uI_P v, c$	Pu, t $uI_P v, c$
$P[v/u], c$ (for every such v)	$P[v/u], c$ (for a new v)
$\neg Pu, s$ $uI_P v, c$	$\neg Pu, t$ $uI_P v, c$
$\neg P[v/u], c$ (for every such v)	$\neg P[v/u], c$ (for a new v)

¹⁵For details on how tableaux standardly work in first-order logic with identity, see eg [?] or [?]. Here, we assume familiarity with the general idea.

The second batch ensures that similarity and identity do not differ from tolerant to classical to strict:

$uI_Pv, s/t$	$\neg(uI_Pv), s/t$
uI_Pv, c	$\neg(uI_Pv), c$

$u = v, s/t$	$\neg(u = v), s/t$
$u = v, c$	$\neg(u = v), c$

(Here, s/t can be either s or t .) The third and final batch encodes the reflexivity and symmetry of the I_P relations:

.	uI_Pv, c
uI_Pu, c	vI_Pu, c

A branch *closes* iff it includes two tagged sentences of the form ϕ, c and $\neg\phi, c$. A tableau *closes* iff every branch on it closes.

Fact 4 (Soundness). *If a tableau built on a set Γ of tagged sentences closes, then there is no T -model M that satisfies every tagged sentence in Γ .*

Fact 5 (Completeness). *If a tableau built on a set Γ of tagged sentences does not close, then there is a T -model M that satisfies every tagged sentence in Γ .*

These facts hold both for our restricted vocabulary and for the full vocabulary. We omit the proofs here; they are simple modifications of usual soundness and completeness proofs for tableaux. (For an example of the usual proofs, see [?].) Note that these facts allow us to use our tableau for any of our nine notions of consequence:

Fact 6. $\Gamma \models^{mn} \Delta$ iff a tableau built on Λ closes, where $\Lambda = \{\gamma, m : \gamma \in \Gamma\} \cup \{\neg\delta, d(n) : \delta \in \Delta\}$.¹⁶

We can also use these tableaux to explore still more articulated questions that do not reduce to questions about mn -consequence for any m, n . For any sets Γ, Λ, Σ of sentences, these tableaux will answer: is there any model where every sentence in Γ holds classically, every sentence in Λ holds tolerantly, and every sentence in Σ holds strictly?

3.4 The sorites

How does this framework address the sorites paradox? In order to be as clear as possible, we consider two different versions of the sorites paradox. As above, we favor the relation st , and so we respond to the paradox from this perspective. Throughout, we assume a sorites series of objects a_1, \dots, a_n for the predicate P .

¹⁶As before, $d(c) = c$, $d(s) = t$, and $d(t) = s$.

Version 1 One version of the sorites proceeds directly from similarity relations:

$$\frac{\begin{array}{l} Pa_1 \\ \forall i \in [1, n](a_i I_P a_{i+1}) \end{array}}{Pa_k}$$

This version of the sorites is *st*-invalid. (There’s no funny business with the quantification here; everything stays the same if we replace the quantified premise with n many atomic premises.) Of course, it’s classically invalid as well; the interest in the *st* response comes not just from its invalidating this argument, but from invalidating this argument while *validating* each step. That is, the following argument is *st*-valid:

$$\frac{\begin{array}{l} Pb \\ b I_P c \end{array}}{Pc}$$

The system *st* can accomplish this balancing act because of its nontransitivity; transitive logics must validate both or neither of these inferences. We think, though, that the second inference is a good one, and that the first is not.

Version 2 A different version of the paradox has tolerance as a premise:

$$\frac{\begin{array}{l} Pa_1 \\ \forall i \in [1, n](a_i I_P a_{i+1}) \\ \forall x \forall y((Px \wedge x I_P y) \rightarrow Py) \end{array}}{Pa_k}$$

(Again, we can replace the quantified premises with their instances without changing anything.) This version of the paradox is *st*-valid. After all, it’s classically valid, and, as we’ve seen, *st* is stronger than classical logic. Here, the reason we do not conclude that Pa_k holds, even tolerantly, is because there is an untrue premise. The third premise, tolerance, is not strictly true, and it is strict truth we require of our premises in *st*.

It is quite difficult for a conditional to be strictly true. Remember, we understand \rightarrow as a material conditional: $\phi \rightarrow \psi$ is to be read as $\neg\phi \vee \psi$. This allows even $\phi \rightarrow \phi$ to fail when ϕ is a borderline sentence; we should not be surprised that tolerance does not meet this high a standard.

However, as we’ve mentioned before, tolerance does meet the lower standard of tolerant validity. This version of the sorites paradox reminds us that one must be careful using even valid sentences as premises, if the standard for validity differs from the standard that premises must meet.

It may at first seem to be in tension with our tolerance-preserving approach to call this argument valid, and refuse to strictly assent to tolerance, but we think that once an appropriate pragmatic framework is in place, the appearance of tension dissipates. We turn to this issue in §??.

4 The pragmatics of vague predicates

In this section we discuss some applications of our framework to the semantics and pragmatics of vague predicates. Our framework allows us to define two dual notions of interpretation for a predicate on top of the classical one, namely strict and tolerant. This raises the question of which of these interpretations is likely to be preferred in interpreting and using vague predicates. We think that both strict and tolerant interpretations are needed in an empirically adequate treatment of vague predicates. This obligates us to give some account of when each sort of interpretation plays a role. When someone asserts a vague sentence, are we to interpret it strictly, classically, or tolerantly? We appeal to an independently motivated pragmatic mechanism (the “strongest meaning hypothesis”) that delivers either strict or tolerant interpretations, depending on the context. In this our treatment of assertion agrees with the pragmatic account of vague predicates recently proposed in [?]. In the second part of this section, we relate our framework to the experimental data they obtained, as well as to the earlier findings of [?] and [?].

5 Conclusion

In this paper we proposed a new semantic framework for the treatment of vague predicates. As we discussed along the way, our framework shares a number of features with extant semantics for vagueness, though we believe it differs from each in some important respects.

First of all, our similarity-based semantics for first-order logic rests on the idea that vagueness is tied in an essential way to non-transitivity, whether of indifference or indiscriminability. In this, the framework agrees with one of the central hypotheses of epistemicism in particular (see [?]), but it does not commit us to the view that the classical extensions we stipulated in T-models necessarily reflect the “objective” meaning of vague predicates. Rather, on our approach these models can be used to describe the internal representations of subjects confronted with the task of categorizing objects based on how similar they look.

Secondly, we have argued that the duality between the notions of tolerant and strict truth gives a natural characterization of borderline cases. It also allows us to validate the tolerance principle; finally, the definition of a mixed relation of logical consequence, from strict to tolerant, allows us to preserve most of the classical rules of inference, in particular modus ponens, only at the expense of having a non-transitive notion of logical consequence. In our opinion, the loss of transitivity for logical consequence is less dramatic a cost than the loss of modus ponens in response to the sorites paradox.

Finally, we have seen that the very duality of strict and tolerant interpretations can be

seen to match the experimental data recently established regarding how people evaluate complex sentences containing borderline cases.

Several issues remain to be investigated. As mentioned in the last section, one aspect we did not go into concerns the possibility of presenting a tolerant/strict semantics based on partial rather than total interpretations. Secondly, in this paper we mostly focused on the sorites paradox and on the status of borderline cases. We deliberately set aside the issue of higher-order vagueness, in particular because it concerns the behavior of an operator like “definitely” or “clearly” in a richer language. As briefly pointed out, it would be very natural to introduce a “definitely” operator to mirror the notion of strict truth syntactically (see [?]).¹⁷

¹⁷Acknowledgements...