Keeping Vague Score

Sam Carter

The Journal of Philosophy

Abstract

This paper introduces a novel theory of vagueness. Its main aim is to show how naïve judgments about tolerance and indeterminacy can be preserved while departing from classical logic only in ways which are independently motivated.

The theory makes use of a bilateral approach to acceptance and rejection. Combined with a standard account of validity, this approach gives rise to an entailment relation which is non-transitive. I argue that this is desirable: it is both pre-theoretically plausible and provides a compelling solution to a number of longstanding puzzles. The theory aims to preserve the naïve picture of vagueness. It combines tools from expressivist and dynamic treatments of information in conversation to show how principles which are generally assumed to be in tension are compatible. The resulting logic departs from classical logic in precisely those places we should expect it to.

1 Introduction

Imagine a series of 50 monochrome cards, arranged from darkest to lightest, like this:

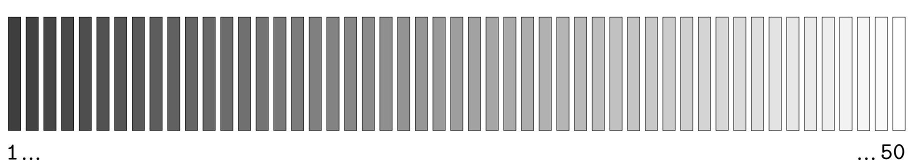

Some of the cards are dark and some of them are not. Which cards are dark, however, is a vague matter. This vagueness shows itself in a variety of ways. One of these is in how people talking about the cards are permitted/required to classify them.

Some cards are permitted to be classified one way and that way only. Cards at the extremes of the series, like #1 and #50, are required to be classified as dark and required to be classified as not dark, respectively Not all the cards are like this, however. Sometimes, the classification of a card is left at an individual's discretion.

Take, for example, #25. There are contexts in which #25 is neither required to be classified as dark nor required to be classified as not dark. Nothing about the card's shade alone settles the question of whether it is dark in such contexts. In fact, we can go further. When dealing with a card in the middle of the series, not being prohibited from denying it is dark is not the same as being required to affirm it is not. There are contexts in which #25 is permitted to be classified as dark and is also permitted to be classified as not dark. The lack of a requirement to classify either way need not be due to a requirement to classify neither way.

To permit choice like this in some contexts is, I'll take it, just part of what it is to be vague. However, the choices we face about the use of a vague expression are not fixed once and for all. How individuals are permitted and/or required to classify different cases can change over the course of a conversation.

Take some n such that card $\#n$ is pretty close to the middle of the sequence. We can expect to find a context in which the cards $\#n{-}1$ through $\#n{+}1$ could each be permissibly classified either way. Once we start making choices about how to classify some cards, however, the way we are permitted to classify the others quickly becomes more constrained. Here, it is worth distinguishing two different ways in which the constraints on how $\#n$ is to be classified can change:

(i) Suppose that we classify $\#n{+}1$ as dark. Then we will be required to classify $\#n$ as dark too. Despite the classification of the cards originally being at the speaker's discretion, classifying $\#n + 1$ as dark (or $\#n{-}1$ as not dark) produces a requirement to classify $\#n$ the same way.

(ii) Suppose that we classify $\#n{-}1$ as dark. Then we will not be permitted to classify $\#n$ as not dark. Despite the classification of the cards originally being at the speaker's discretion, classifying $\#n{-}1$ as dark (or $\#n{+}1$ as not dark) produces a requirement not to classify $\#n$ differently.

Nothing particularly mysterious is going on here. In both types of case, the changes in requirements on how $\#n$ is classified are the product of two factors: first, permissibly classifying a card as dark (not dark) changes the context, so that that card is settled to be dark (not dark) in the context that results. Second, how a card is permitted/required to be classified in a context is determined (at least in part) by which cards have are settled as dark (and as not dark).

Sometimes, as in the first case, this is attributable to penumbral connections between cases (Fine (1975)). In a context in which it is settled that a card is dark (not dark), any card at least as dark (light) is required to be classified as being dark (not dark), too. But this is not the only way that our permissions and requirements can be constrained. In the second case, the change in how $\#n$ is permitted to be classified is attributable to the tolerance of the vague expression (Wright (1975)). In a context in which it is settled that some card is dark (not dark), any card only marginally darker (lighter) is prohibited from being classified as being not dark (being dark).

There is an important asymmetry between the two types of change. Although, after classifying $\#n{-}1$ as dark, $\#n$ is not permitted to be classified as not dark, plausibly, it is not required to be classified as dark either. Refusing to classify it either way remains permissible. Put another way, classifying a card as dark (not dark) merely restricts the range of classifications permitted for marginally darker (lighter) cards in the series, it does not expand the range of classifications required for them.[1]

We can think about a conversation's context on the model of a scoreboard (Lewis (1979)). Licit assertions (moves) in the conversation (game) determine the state of the context (scoreboard). Correspondingly, at any given stage of the conversation (game), the context (scoreboard) determines what assertions (moves) are licit. In what follows, I'll argue that accounting for the interaction between moves and scoreboard is essential to resolving traditional puzzles associated with vagueness.

These ideas are not novel. That vagueness permits choice has been widely observed (see, e.g., Wright (1987, 1995); Tappenden (1993); Raffman (1994, 1996, 2014a); Sainsbury (1996); Soames (1998); Shapiro (2006); Gaifman (2010); MacFarlane (2016)). And so has the fact that what choices are permitted varies according to what is said (see, e.g. Kamp (1981); Pinkal (1983); Ballweg (1983); Eikmeyer & Rieser (1983); Raffman (1994, 1996); Soames (1998); Kyburg & Morreau (2000); Barker (2002, 2003); Shapiro (2006); Gaifman (2010); Ludlow (2014); MacFarlane (2016)). The limited aim of the present paper is to show how these ideas can be combined to explain our naïve picture of vagueness.

Our naïve picture of vagueness comprises the collection of pre-theoretic judgments which together characterize the target phenomenon. It is sometimes claimed that our naïve picture is incoherent—that the judgments which comprise it cannot be jointly satisfied (Dummett (1975); Horgan (1994, 1998); Eklund (2002, 2005, 2019)). The theory below, if successful, demonstrates that this charge is mistaken. Our pre-theoretic judgments about vague matters can all be accommodated within a single framework.

[1] Note that, given our mnemonic and perceptual limitations, we may often be unable to tell with perfect reliability whether a given object has already been classified as positive (or negative) case. For example, we may forget which cards have been classified as dark or be unable to distinguish cards which are similar in shade. As a result, we will not always be in a position to tell what semantic requirements we are subject to in classifying an object. This is not distinctive of vague language, however. Even restricting attention to precise expressions, there is independent reason to think that agents are frequently subject to semantic requirements which, due to their cognitive limitations, they are unable to follow perfectly reliably. For example, given the fallibility of our perceptual capacities, we will often be unable to tell whether an object whose weight is around a quarter of a kilogram should be classified as a positive case or negative case of the predicate 'weighs 247 grams'. But if the rules governing our application of precise expressions often outstrip our cognitive capacities, it is unclear why the rules governing the application of vague expressions should be different. I am grateful to a referee at *The Journal of Philosophy* for encouraging me to discuss this point.

Here's the plan. §3 introduces a new approach to theorizing about vagueness, one which incorporates, in a more formal setting, the informal ideas in this section. On this approach, entailment is treated in a familiar fashion. An argument is valid just in case no-one who accepts the the premises could go on to permissibly reject the conclusion. As we've just seen, however, a vague claim may be incapable of being rejected despite not yet being accepted. As a result, our entailment relation will be non-transitive—concatenating arguments is not guaranteed to preserve validity. §§4-6 show how the resulting framework deals with many of the traditional problems of vagueness and discuss its relationship to existing theories (in particular, dynamic ones, such as Kamp (1981), and non-transitive ones, such as Zardini (2008); Cobreros *et al.* (2012)). Before starting on the positive proposal, however, we need to see what the traditional problems are that it aims to solve.

2 The Problems of Vagueness

Puzzles to do with vagueness come in two broad kinds: (i) those associated with the (apparent) tolerance of vague expressions and (ii) those associated with the (apparent) indeterminacy of vague expressions. In this section, we'll look at each in turn.

In order to frame these puzzles, we need a language which is sufficiently expressive. Within this constraint, we will be austere, limiting our language to a single unary predicate (F), a single binary predicate ($\sim$) and finitely many variables ($x_1, ..., x_n$, for some n).[2]

Definition (Language).

L_0 is the smallest set containing $\{\top, \bot\}$, $\{F(x_i)|0 \leq \ i \leq n\}$ and $\{x_i \sim x_j|0 \leq \ i, j \leq n\}$ which is closed under the boolean connectives of negation ($\neg$) and conjunction ($\wedge$). L_1 is the smallest set containing L_0 and which is closed under the boolean connectives, a conditional connective ($\rightarrow$) and non-vacuous universal quantification ($\forall x_i$).[3] L_2 is the smallest set containing L_1 and which is closed under boolean connectives, the conditional, non-vacuous universal quantification and truth/falsity operators (T, F).

Under our intended interpretation, $x_i \sim x_j$ says that x_i and x_j differ marginally along the dimension(s) relevant for F. $\mathsf{T}(\phi)$ and $\mathsf{F}(\phi)$ say that ϕ is true/false, respectively. We define disjunction ($\vee$), material implication ($\supset$) and existential quantification ($\exists x_i$) in the usual way. We define $\phi \leftrightarrow \psi$ as the conjunction of $\phi \rightarrow \psi$ and $\psi \rightarrow \phi$. We'll assume that the conditional is at least as strong as the material conditional. We won't, however, assume it is to be identified with it (though (for now) we won't rule it out, either).

[2] Given this austerity, we'll need to limit ourselves to sorites series with finitely many members. That'll easily be enough for the puzzles we're interested in, though.

[3] That is, $\forall x_i \phi \in \mathsf{L}_1$ if and only if $\phi \in \mathsf{L}_1$ and x_i occurs free in ϕ.

2.1 Tolerance

Vague expressions (often) display tolerant behavior. One way this shows up is in the pattern of requirements on what claims can be accepted together. As we saw, an individual is prohibited from classifying card $\#n$ as dark and classifying its marginally lighter successor, $\#n+1$, as not dark.

Corresponding to this prohibition, however, there appears to be a requirement to accept certain conditional claims. At least naïvely, it seems we should accept that for any n, if $\#n$ is dark, then $\#n$+1 is dark, too. Call this, simply, **Tolerance**.

Tolerance $\forall x_i \forall x_j ((F(x_i) \wedge x_i \sim x_j) \rightarrow F(x_j))$

Despite its appeal, **Tolerance** is generally taken to be untenable. That's because it is classically in tension with the conditions characterizing a sorites series. We'll take a sorites series to be a situation in which the following pair of conditions hold:

Limits $\exists x_i \exists x_j (F(x_i) \wedge \neg F(x_j))$

Continuity $\forall x_1 \forall x_n \exists x_2 ... \exists x_{n-1} (x_1 \sim x_2 \wedge ... \wedge x_{n-1} \sim x_n)$

Limits says that F has a positive case and a negative case. **Continuity** says that for any pair of objects in the domain, we can find a sequence of objects starting with the first, ending with the second, and in which each object is a marginal variant of its successor.[4]

Fact 1. **Tolerance**, **Limits** and **Continuity** are classically inconsistent.

It is important to distinguish **Tolerance** (the schema), and tolerance (the property of vague languages which makes **Tolerance** appear valid). Even those who deny **Tolerance** is valid will hold that vague predicates are tolerant (in this sense).

Tolerance is only one of the ways tolerance manifests. Closely related and seemingly emerging from the same basic aspect of vagueness is the feeling that vague expressions can't give rise to sharp cutoffs (Campbell (1974); Sanford (1975, 1976); Wright (1987); Sainsbury (1989, 1996)). At least naïvely, it seems we should reject that there is any $\#n$ such that $\#n$ is dark and $\#n+1$ is not. Call this **Sharp Cutoffs**:

Sharp Cutoffs $\exists x_i \exists x_j (F(x_i) \wedge \neg F(x_j) \wedge x_i \sim x_j)$

Despite its lack of appeal, **Sharp Cutoffs** is often taken to be unavoidable. That's because it is a classical consequence of the conditions characterizing a sorites series.

[4] More carefully, **Limits** and **Tolerance** are schemata whereas **Continuity** is (an abbreviation for) a sentence of L_1. Playing loose in order to play fast, I will talk of logical relations holding between sentences and schemata where what is meant, strictly, is that they hold between the sentences and every combination of instances of the schemata.

Fact 2. Limits and **Continuity** classically imply **Sharp Cutoffs**.

In the setting of the sorites series, adherence to classicality commits us to accepting a sharp boundary between the cards which are dark and the cards which are not dark.

2.2 Indeterminacy

Vague expressions (often) also display indeterminacy. One way this shows up is in the apparent absence of truth and falsity. That is, it is part of the naïve picture that vagueness gives rise to failures of **Bivalence**:

Bivalence $\mathsf{T}(\phi) \vee \mathsf{F}(\phi)$

That **Bivalence** should be denied has been widely entertained (see, e.g., Fine (1975); Sanford (1975); Tye (1989, 1994); Tappenden (1993); Keefe (2000); Shapiro (2006)). And, again at least naïvely, it seems very natural to resist attributing either truth or falsity to vague claims. For it to be vague whether $\#n$ is dark is, in part, for it to be neither true nor false that it is (and neither true nor false that it isn't).

Whether this is viable, however, is less clear. Consider the following pair of principles:

Transparency $\phi \leftrightarrow \mathsf{T}(\phi)$

Polarity $\neg\phi \leftrightarrow \mathsf{F}(\phi)$

Transparency says that any claim is equivalent to the claim that it is true. **Polarity** says that the negation of any claim is equivalent to the claim that it is false. Together, these principles form the core of a standard, disquotational picture of truth and falsity.[5] And yet, at least in a classical setting, they rule out the possibility of rejecting **Bivalence** (Williamson (1992, 1994); Horwich (1998b); Field (2000)).

Fact 3. Transparency and **Polarity** classically imply **Bivalence**.

As long as we remain committed to the disquotational picture, classical logic commits us to accepting that, for every card, it is either true that it is dark or it is false.

2.3 Summary

Our naïve picture of vagueness takes vague expressions to be tolerant and indeterminate. Our naïve picture is in tension with classical logic, however.

[5] This is, importantly, not the same as a disquotational theory of truth. I take it that many theorists who would reject a disquotational theory would nevertheless accept both **Transparency** and **Polarity**. For discussion of disquotation in relation to vagueness, see, e.g., Peacocke (1981); Field (2000); Williamson (1994); Horwich (1998b,a, 2000).

One response is to deny the relevant parts of the naïve picture (while explaining why they were naïvely appealing in the first place). In this way, **Facts 1** and **2** are often treated as reasons to deny **Tolerance** and accept **Sharp Cutoffs**. Different theorists offer different explanations of why the naïve picture appeals to us. Some propose that no instance of the former is false (or known to be false) and no instance of the latter true (or known to be true) (Fine (1975); Williamson (1997); Keefe (2000); Williams (2008). Some propose that knowing the antecedent of an instance of the former guarantees the truth of the consequent and knowing one conjunct of the latter precludes the truth of the other (Williamson (1994, §8.4)). And others propose that the cutoff (though sharp) moves according to what we pay attention to, so that it never lies just where we are looking (Raffman (1994, 1996); Fara (2000); Kennedy (2010)). Equally, **Fact 3** is often treated as a reason to accept **Bivalence**, even for vague claims. Apparent failures are attributed to our alleged tendency to conflate truth and falsity with determinate (or known) truth and determinate (or known) falsity (Williamson (1994, 1997); McGee & McLaughlin (1995)).

The alternative response is to deny some part(s) of classical logic. Various targets have been suggested. Some propose giving up modus ponens (Jaśkowski (1969); Machina (1976); Hyde (1997)). Others propose denying non-contradiction, excluded middle or both (Machina (1972, 1976); Tye (1989, 1994); Burgess (1990); Burgess & Humberstone (1987)). Some have even suggested that there may be no logic of vague languages (Wright (1975)) or that vague expressions are incoherent (Dummett (1975)).

Neither response is completely satisfactory. Regarding the latter, if the logic of vague language is non-classical, it is crucial that its deviations from classical logic are motivated. Nothing in our naïve understanding of vagueness suggests that modus ponens, non-contradiction or excluded middle should fail in cases of vagueness, however. From the claims that $\#n$ is dark and that if $\#n$ is, then $\#n{+}1$ is too, the inference that $\#n{+}1$ is dark is impeccable. Equally, we have no inclination, for any n, to reject that $\#n$ is either dark or it isn't or to accept that it is both dark and isn't.

Regarding the former, part of the job of a theory of vagueness is to account for our naïve judgements. All things equal, we would prefer a theory which preserved the naïve picture to one which dismisses it (cf. Zardini (2008)). Tolerance and indeterminacy are (the) central phenomena of vagueness. A theory which rejects them is liable to leave us at least somewhat disappointed. Of course, such a theory may be defensible if no alternative is available (or no available alternative is tenable). However, its defense will depend on pessimism about the prospects of finding such an alternative.

As stated above, my main goal in this paper is to show that this pessimism is unwarranted. We can develop a consistent theory which vindicates the naïve picture, while deviating from classical logic only in ways which are independently motivated.

3 A Model of Vagueness

Our theory will combine two core ideas, both introduced in §1:

(i) judgments about validity reflect facts about what moves speakers are permitted/required to make;

(ii) facts about what moves speakers are permitted/required to make change in response to permissible moves made by speakers.

In representing the two-way relationship between contexts and utterances, we will take sentences' semantic content to determine acceptance and rejection conditions. Each sentence will be associated with a pair of sets of contexts: the contexts at which it is accepted and the contexts at which it is rejected. These correspond, intuitively, to the contexts at which endorsing the sentence is required and the contexts at which it is impermissible, respectively. To reflect changes in what is required we also need to encode dynamic facts about the way conversations develop over time. We'll do this by equipping our model with an update operation. Updating a context with a sentence returns the minimal change to the former required to ensure that the latter is accepted (cf. Kamp (1981)).

We'll understand entailment in a standard way. An argument is valid (in the sense we're interested in) just in case no-one who accepts the premises could reject the conclusion. As indicated above, it is important to distinguish this from the property an argument has just in case anyone who accepts the premises must accept the conclusion. There are some contexts in which a sentence cannot be denied, even though it is not required to be endorsed.

3.1 Contexts

A model comprises a domain of objects, $\mathcal{D}$, a relation of marginal variance, $\approx$, and an interpretation function, $[\![\cdot]\!]$. $\mathcal{D}$ is an finite linearly ordered set. Intuitively, the ordering, $\geq$, over the set corresponds to a sorites series whose greatest member is a clear positive case and whose least member is a clear negative case.

$\approx$ is a reflexive, symmetric (but potentially non-transitive) relation over $\mathcal{D}$, subject to two constraints. First, if $d \approx d'$ and $d \geq d'' \geq d'$, then $d \approx d''$ and $d'' \approx d'$. Second, any pair of objects in the domain are related by the transitive closure of $\approx$. Intuitively, $\approx$ can be thought of as the relation which holds between d and any objects sufficiently 'near' to d in the sorites series encoded by the ordering.

A context, c, is a pair $\langle c^+, c^- \rangle$ of subsets of $\mathcal{D}$. Contexts represent possible states of a conversation. Classifying an object is a public act, aimed at changing the conversation's state. Where successful, it settles the object classified either as a positive case or as a negative case. Contexts track how objects are settled.

Among these we identify some contexts as admissible. Each admissible context represents a state that a non-defective conversation could be in.

A context is admissible iff it satsifies the following three conditions:[6]

Boundaries $c^+ \neq \varnothing$ and $c^- \neq \varnothing$.

Convexity If $d \in c^+$ and $d' \geq d$, then $d' \in c^+$; and
If $d \in c^-$ and $d \geq d'$, then $d' \in c^-$.

Coherence If $d \in c^+$ and $d' \in c^-$, then $d \not\approx d'$.

Boundaries says that every admissible context must settle some object as a positive case and some object as a negative case. **Convexity** says that if an admissible context settles an object as a positive (negative) case, then it must settle every object higher (lower) in the order as a positive (negative) case, too. Together, they ensure that every admissible context corresponds to an initial segment and final segment of the ordering encoding the sorites series. **Coherence** says that no admissible context may settle two marginal variants as positive and negative cases, respectively. It is an immediate consequence that every admissible object must leave some object unsettled (neither a settled positive case nor a settled negative case).[7] In what follows, $c, c', ...$ will range over admissible contexts only, unless explicitly noted otherwise. Likewise, we will restrict our use of 'context' to refer only to admissible contexts.

We say that c' is an extension of c when c' settles every object settled by c the same way. Intuitively, we can think of the extensions of a context as permissible ways that a conversation in the state of c could develop by classifying additional cases.

Definition (Extension). $c \preccurlyeq c'$ iff $c^+ \subseteq c'^+$ and $c^- \subseteq c'^-$.

Given our relation of extension, we can introduce an operation of minimal change. Where C is a set of contexts:

Definition (Update). $c + \mathrm{C} = Min\{c' \in \mathrm{C} | c \preccurlyeq c'\}$.

That is, $c + \mathrm{C}$ is the set of minimal admissible extensions of c in C. Intuitively, we can think of update with C as a process whereby we obtain one or more admissible members of C by settling cases left previously unsettled, with the restriction that no unsettled cases are settled unnecessarily. Derivatively, we let $\mathrm{C} + \mathrm{C}' = \bigcup\{c + C' | c \in C\}$.

[6] Our admissible contexts are structurally equivalent to (partial) precisifications in supervaluational models (Fine (1975); cf., in particular, Shapiro (2006)). More loosely, they can be thought of as analogous to what Restall (2005) and Ripley (2013) refer to as states or positions, respectively (although restricted to atomic sentences).

[7] We can see the coherence constraint as a rejection of what Fine (1975) calls *completeability* (for precisifications) and what Restall (2005) calls *extensibility* (for states): the claim that for any admissible context and sentence ϕ, it must be possible to extend it into an admissible context at which ϕ is accepted or ϕ is rejected.

As we will see, **Coherence** will play an important role in our framework. It encodes the idea that there is something defective about a conversation in which a sharp boundary is imposed between marginal variants. A number of authors have proposed that contexts are subject to a constraint of this form (see e.g., Kamp (1981); Tappenden (1993); Soames (1998); Shapiro (2006); Gaifman (2010); Ripley (2013)). What it is attributable to, however, remains a point of dispute.

One alternative is that it arises from the semantics of vague expressions. It is a brute fact about the meaning of tolerant predicates, according to this proposal, that they cannot have marginally differing positive and negative cases. The requirement that contexts be coherent is simply an instance of the requirement that contexts not disregard the meanings of our expressions. To attempt to reach a context which violated this constraint would be to misuse language. Some version of this option appears to be endorsed by Dummett (1975) and Wright (1975).

A second alternative is that it arises, not from the semantics of vague expressions, but from the pragmatics of their use. One way to develop this idea appeals to interlocutors communicative goals. Proponents of this idea sometimes suggest that cooperativity requires speakers to leave (some) flexibility in the classification of further cases (Tappenden (1993); MacFarlane (2016)) or that it can never be in our practical interest to settle marginally different objects differently (Kennedy (2010), cf. Fara (2000)).

A different way to develop the idea appeals to facts about our psychology. Raffman (2011, 2012, 2014b) reports an experiment (conducted in collaboration with Delwin Lindsey and Angela Brown) which found that (a) participants match color-patches to different hues depending on what other patches under attention and (b) when under attention, adjacent color-patches are routinely judged to appear to be the same hue. The coherence constraint could be explained by a general pragmatic requirement not to classify objects differently if they appear the same (given the assumption that classifying an object requires it to be attended to).

3.2 Basic Semantics

Our semantics associates each sentence with acceptance conditions and rejection conditions. A bilateral interpretation, $[\![\cdot]\!]_g$, maps each ϕ to a pair of sets of admissible contexts, $\langle [\![\phi]\!]_g^+, [\![\phi]\!]_g^- \rangle$. Intuitively, c is a member of $[\![\phi]\!]_g^+$ ($[\![\phi]\!]_g^-$) iff ϕ is accepted (rejected) at c and g. Let $[\![\top]\!]_g^+ = [\![\bot]\!]_g^-$ be the set of admissible contexts, and let $[\![\bot]\!]_g^+ = [\![\top]\!]_g^-$ be $\varnothing$. We start by fixing the interpretation of the other atoms and the boolean connectives.

Definition (Basic Semantics).

i. $c \in [\![F(x_i)]\!]_g^+$ *iff* $g(x_i) \in c^+$
$c \in [\![F(x_i)]\!]_g^-$ *iff* $g(x_i) \in c^-$

ii. $c \in [\![x_i \sim x_j]\!]_g^+$ *iff* $g(x_i) \approx g(x_j)$
$c \in [\![x_i \sim x_j]\!]_g^-$ *iff* $g(x_i) \not\approx g(x_j)$

iii. $c \in [\![\neg\phi]\!]_g^+$ *iff* $c \in [\![\phi]\!]_g^-$
$c \in [\![\neg\phi]\!]_g^-$ *iff* $c \in [\![\phi]\!]_g^+$

iv. $c \in [\![\phi \wedge \psi]\!]_g^+$ *iff* $c \in [\![\phi]\!]_g^+ \cap [\![\psi]\!]_g^+$
$c \in [\![\phi \wedge \psi]\!]_g^-$ *iff* $c \in [\![\phi]\!]_g^- \cup [\![\psi]\!]_g^-$

$F(x_i)$ is accepted (rejected) at c and g iff $g(x_i)$ is settled as a positive (negative) case at c. The acceptance and rejection conditions of $F(x_i)$ need not be exhaustive: where $g(x_i) \notin c^+ \cup c^-$, $F(x_i)$ is neither accepted nor rejected at c and g. $x_j \sim x_j$ is accepted (rejected) at c and g iff $g(x_i)$ and $g(x_j)$ are (not) marginal variants. Since the relation of marginal variance is fixed by the model, $x_i \sim x_j$ is context insensitive. At an assignment, it is either accepted at every context or else rejected at every context.

$\neg\phi$ is accepted (rejected) at c and g iff ϕ is rejected (accepted) there. That is, we define negation in terms of rejection (and not *vice versa*). In this respect the framework follows other bilateralist approaches (Price (1990); Smiley (1996); Rumfitt (2000); Restall (2005); Cobreros *et al.* (2012), for overviews, see Ripley (2011b, 2020)). The core idea is that acceptance and rejection are to be treated symmetrically. Neither is prior to the other—each sentence is associated directly with conditions for both. Negation simply toggles between them.

$\phi \wedge \psi$ is accepted (rejected) at c and g iff ϕ and (or) ψ are accepted (is rejected) there. The acceptance conditions of a conjunction are the intersection of the acceptance conditions of its conjuncts; the rejection conditions are the union of their rejection conditions. Recall that we defined $\vee$ in terms of $\wedge$ and $\neg$ in the usual way. So the acceptance conditions of a disjunction are union of the acceptance conditions of its disjuncts and its rejection conditions are the intersection of their rejection conditions.

As we saw in §1, it is important to distinguish between a sentence (i) being accepted and (ii) being incapable of being rejected. We'll say a context *supports* ϕ at g iff there is no extension of the context at which ϕ is rejected (relative to g).

Definition (Support). $c \models_g \phi$ iff for all c': if $c \preccurlyeq c'$, then $c' \notin [\![\phi]\!]_g^-$.

Derivatively, where C is a set of contexts, we'll say that $\mathrm{C} \models_g \phi$ iff every member of C supports ϕ at g. Where Δ is a set of sentences, we'll say $C \models_g \Delta$ iff for all

$c \in C$, there is some $\phi_i \in \Delta$ such that $c \models_g \phi_i$. Sometimes, where no confusion is liable to arise, I will elide reference to an assignment when glossing support, acceptance and rejection.

To see the difference between acceptance and support in our framework, consider an atom, $F(x_i)$. c' extends c iff c' is a permissible way for a conversation in state c to develop. At a context at which $F(x_i)$ is accepted, every permissible way the conversation could develop is one in which $F(x_i)$ is accepted. At a context at which $F(x_i)$ is supported, there is no permissible way for the conversation to develop in which $F(x_i)$ is rejected. Crucially a sentence may be supported at a context despite not being accepted.

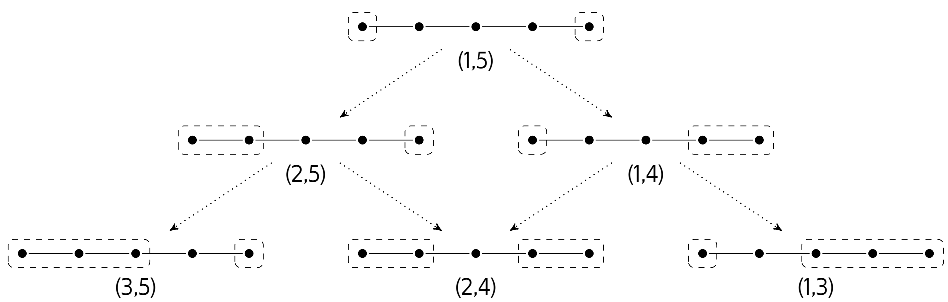

Figure 1: A five-object model.

Consider a model such that (i) $d_i \leq d_j$ iff $i \leq j$ and (ii) $d_i \approx d_j$ iff $|i-j| \leq 1$. Let $c_{(n,k)}$ be the context which settles d_i as a positive case iff $i \leq n$ and as a negative case iff $i \geq k$. For example, **Figure 1** depicts such a model with five objects. Let g^* be the privileged assignment for which $g^*(x_i) = d_i$ (for $1 \leq i \leq 5$). It is easy to see that, in the model depicted, d_3 is not settled as a positive case in $c_{(2,5)}$. Yet, given **Coherence**, there is also no admissible extension of $c_{(2,5)}$ at which d_3 is settled as a negative case. Accordingly, $F(x_3)$ is supported at $c_{(2,5)}$ and g^* despite not being accepted there.

Finally, we define the conditional in terms of update and support.

Definition (Conditional).

$$c \in [\![\phi \to \psi]\!]^+_g \textit{ iff } c + [\![\phi]\!]^+_g \models_g \psi.$$
$$c \in [\![\phi \to \psi]\!]^-_g \textit{ iff } c + [\![\phi]\!]^+_g \not\models_g \psi.$$

$\phi \to \psi$ is accepted (rejected) at c and g iff the contexts obtained by updating c with ϕ support (don't support) ψ. Our conditional implements a version of the same idea behind the dynamic strict conditional (Kamp (1981); Dekker (1993); Gillies (2004)). To check whether a conditional is accepted, we first make the minimal modification necessary to accept its antecedent and then check the status of the consequent at the result.

A context supports $\phi \to \psi$ (at g) iff it supports $\phi \supset \psi$ (at g).[8] However, the acceptance conditions for the two differ: the material conditional is accepted iff either the consequent is accepted or the antecedent is rejected. Yet a conditional can be accepted even if neither the consequent is accepted nor the antecedent rejected. Accordingly, updating with the material conditional can have a different effect to updating with the conditional. As a result, while $\to$ obeys the analogue of the R-rule for $\supset$, it does not obey the analogue of the L-rule (see **Figure 2**, below).[9]

3.3 Validity

Consider the property an argument has iff no-one who accepts the premises could reject the conclusion. We'll say that an argument is conversationally valid iff it has this property.[10]

Definition (Conversational Validity).

$$\phi_i, ..., \phi_j \models \psi \text{ iff for all } c \text{ and } g\text{: } c + ([\![\phi_i]\!]^+_g \cap ... \cap [\![\phi_j]\!]^+_g) \models_g \psi$$

To check whether an argument is conversationally valid, we ask whether updating with the premises always produces a context which supports the conclusion. In what follows, we'll say that $\phi_i, ..., \phi_j$ are conversationally consistent iff the inference from the premises to $\bot$ is conversationally invalid. Finally, we'll say that $\phi_i, .., \phi_j \models \psi_{i'}, ..., \psi_{j'}$ iff for every c and g, $c+([\![\phi_i]\!]^+_g \cap ... \cap [\![\phi_j]\!]^+_g) \models \psi_{i'}, ..., \psi_{j'}$.[11]

The goal of the discussion below is not to argue that conversational validity is the only property we ought to care about in evaluating whether an argument is good. Instead, more modestly, it is to establish two points. First, that our naïve judgements about vague matters are succesfully explained by the hypothesis that they track facts about conversational validity. And, second, that conversational validity is a well-behaved model of entailment for vague languages, one which retains all desired features of classicality. Together, these points leave us in a position to vindicate the coherence of our naïve picture of vagueness.

[8] Or, more carefully, where $\phi, \psi \in \mathsf{L}_1$. *Proof sketch:* Observe that $c \not\models_g \phi \to \psi$ iff there is some c' and c'' such that $c \preccurlyeq c'$, $c' + [\![\phi]\!]^+_g \preccurlyeq c''$ and $c'' \in [\![\psi]\!]^-_g$. The latter condition holds iff there is some c' such that $c \preccurlyeq c'$ and $c' \in [\![\phi]\!]^+_g \cap [\![\psi]\!]^-_g$ (NB: this step depends on ϕ being persistent (see page 21), hence the restriction to L_1). But this holds iff $c \not\models_g \phi \supset \psi$. Contraposition completes the proof.

[9] *Proof:* Observe that $F(x_i) \models F(x_i)$ and $\models \top, F(x_i)$. However, $F(x_i) \to \top \not\models F(x_i)$, since $c + [\![\phi \to \top]\!]^+_g = c$, for all c.

[10] The role of update (i.e., $+$) in the definition of conversational validity is inessential. The definition is equivalent to the requirement that for all c and g: if $c \in [\![\phi_i]\!]^+_g \cap ... \cap [\![\phi_j]\!]^+_g$, then $c \models_g \psi$. The above formulation is preferred in virtue of displaying the symmetry between $\models$ and $\to$.

[11] Intuitively, conversational validity requires that the conclusion is supported when the premises are accepted together. An alternative would be to define validity in terms of sequential acceptance instead, so that $\phi_i, ..., \phi_j \models \psi$ holds iff for all c and g: $(c+[\![\phi_i]\!]^+_g)+...+[\![\phi_j]\!]^+_g \models_g \psi$. The differences between these approaches do not matter for the present fragment: the two entailment relations are equivalent over L_1.

Even at this intermediate point, we can already see how the framework captures some of the tolerant behavior of vague predicates. Consider **Deductive Tolerance**:[12]

Deductive Tolerance $\quad F(x_i), x_i \sim x_j \mid\!\!- F(x_j)$

Deductive Tolerance is a one-shot, inferential version of **Tolerance**. Its apparent validity is brought out by the appeal of the so-called 'forced march' sorites series (Horgan (1994)). Imagine being presented with each card in the monochrome series, starting with #1 and proceeding sequentially through the lighter shades. For an individual who has accepted that card $\#i$ is dark, it appears impermissible to reject the claim that $\#i{+}1$ is dark too. Our entailment relation reflects this, by validating deductive tolerance.

Fact 4. **Deductive Tolerance** is conversationally valid.

To see why, observe that if $d \approx d'$ and c settles d as a positive case, then there is no extension of c which settles d' as a negative case. Consider an arbitrary context c and assignment g, where $g(x_i) = d$ and $g(x_j) = d'$. If $d \not\approx d'$, then updating with $x_i \sim x_j$ at g will return the empty set (at which everything is supported). So assume otherwise. Then updating with $F(x_i)$ and $x_i \sim x_j$ will return the minimal extension of c at which d is settled as a positive case. Call this c'. We know that no extension of c' settles d' as a negative case. So c' supports $F(x_j)$ at g.

Note that we can distinguish two variants of the forced march scenario. In the first variant, an individual is given the option to leave cards unsettled. In this variant, impermissible moves can in principle be avoided by an agent who refrains from classifying some cards around the middle of the series.In the second variant, an individual is forced to classify each card as either a positive case or a negative case. As a result, at some point they must switch from classifying cards as dark to classifying them as not dark (on pain of classifying the last card as dark). This switch in classification involves an (apparent) violation of the coherence constraint. This is unsurprising, however. It is a familiar observation that, where individuals' options are artificially restricted, they may be forced to act in ways that are impermissible.[13]

[12] It is important to distinguish the use of $\mid\!\!-$, $\models$ and $\models_g$ (for all g). $\mid\!\!-$ and $\models$ relate formulae. $\Delta \mid\!\!- \Gamma$ denotes the inference from Δ to Γ. $\Delta \models \Gamma$ says that $\Delta \mid\!\!- \Gamma$ is conversationally valid. In contrast, $\models_g$ relates contexts and formulae. It denotes support at g .

[13] It is not obligatory to interpret this switch as involving a violation of the coherence constraint, however. Alternatively, it could be interpreted as involving a tacit modification of how previously settled objects are classified. On this interpretation, to conform to the coherence constraint, the the classification of some cards earlier in the sequence must be revised, so that some cards previously classified as dark are now unsettled again. This revision would also allow us to explain a form of hysteresis effect, as discussed (Raffman (2014a, 2017)), in which some cards which have previously been classified as dark could now be permissibly classified as not dark, instead, when approached in a forced march from the opposite direction. Whether these revisions should be modeled semantically (as a special kind of update) or non-semantically (as a kind of pragmatic repair) is an interesting question which is worth further

In addition to validating **Deductive Tolerance**, our framework validates a range of other core features of classicality.

LEM $\mid\!\!- \phi \vee \neg\phi$

LNC $\mid\!\!- \neg(\phi \wedge \neg\phi)$

MP $\phi, \phi \rightarrow \psi \mid\!\!- \psi$

Above, I suggested that vagueness gives us no reason to think that any of these features of classicality fail. And, indeed, each is preserved in our framework.

Fact 5. **LEM**, **LNC** and **MP** are conversationally valid.[14]

The conversational validity of **LEM** and **LNC** guarantees that there is no admissible context at which they are rejected. It is important to distinguish this from the claim that they are accepted at every context. In fact, as we will see, there are some contexts which fail to accept instances of each. In the following section, I will argue that this is the right prediction. Although a conversation can never reach a state in which instances of excluded middle or non-contradiction can be permissibly denied, it can reach a state in which some instances cannot be endorsed.

Since the framework validates **Deductive Tolerance**, you might anticipate that it must depart from classicality somewhere or other. And it does. In our framework, entailment is non-transitive. From the fact that ϕ entails ψ and ψ entails χ, we cannot conclude that ϕ entails χ. More generally, we allow for failures of Cut, of which transitivity is a special instance.

$$\frac{\Gamma \mid\!\!- \phi, \Delta \qquad \Gamma', \phi \mid\!\!- \Delta'}{\Gamma, \Gamma' \mid\!\!- \Delta, \Delta'} \quad \text{Cut}$$

To see why, observe that $F(x_i)$ and $x_i \sim x_j$ entail $F(x_j)$. And $F(x_j)$ and $x_j \sim x_k$ entail $F(x_k)$. However, $F(x_i)$, $x_i \sim x_j$, and $x_j \sim x_k$ do not entail $F(x_k)$.[15]

This is exactly what we should expect. From the fact that, after update with ϕ, ψ cannot be rejected and after update with ψ, χ cannot be rejected, it does not follow that after update with ϕ, χ cannot be rejected. To reason this way would

consideration.

[14] *Proof:* **LNC** is conversationally valid iff there is no c and g such that $c \in [\![\phi]\!]^+_g$ and $c \in [\![\phi]\!]^-_g$. This is trivial for $x_i \sim x_j$. For $F(x_i)$, it is guaranteed by **Coherence** and the reflexivity of $\approx$. The proof is completed by induction on the complexity of ϕ. **LEM** is conversationally valid iff **LNC** is conversationally valid. Finally, for **MP**, observe that if $c \in [\![\phi]\!]^+_g$, then $c + [\![\phi]\!]^+_g = c$. Furthermore, if $c \in [\![\phi \rightarrow \psi]\!]^+_g$, then $c + [\![\phi]\!]^+_g \models_g \psi$. So if $c \in [\![\phi]\!]^+_g \cap [\![\phi \rightarrow \psi]\!]^+_g$, $c \models_g \psi$.

[15] We can see this by returning to **Figure 1**. At $c_{(2,5)}$ and g^*, updating with $F(x_2)$, $x_2 \sim x_3$, and $x_3 \sim x_4$ returns $c_{(2,5)}$. And yet $c_{(2,5)} \not\models_{g^*} F(x_4)$.

be to ignore the difference between accepting a sentence and being incapable of rejecting it. There may be contexts at which ϕ cannot be rejected but at which updating with ϕ nevertheless has a non-trivial effect.

Failures of transitivity are part of our naïve picture of vagueness, in the following sense (Ziff (1974); Machina (1976); Zardini (2008)). Let S_i name the sentence $\ulcorner \#i$ is dark $\urcorner$ (for $1 \leq i \leq 50$), and consider the sequence of sentences $S_1, ..., S_{50}$. The deductive sorites puzzle arises as a result of inferences between these sentences which appear to be valid (and inferences which appear invalid). Specifically: (i) each sentence in the sequence appears to follow from its predecessor; but (ii) the last sentence in the sequence appears not to follow from the first. In this way, there is a sequence of inferences, each of which seems valid, but whose concatenation seems invalid.[16]

While our framework departs from classicality, it does so minimally. That is, it validates as much of classical logic as possible while resisting transitivity.

$$\frac{\Gamma, \phi/\psi \vdash \Delta}{\Gamma, \phi \wedge \psi \vdash \Delta} \wedge_{\mathrm{L}} \qquad \frac{\Gamma, \phi \vdash \Delta \qquad \Gamma, \psi \vdash \Delta}{\Gamma, \phi \vee \psi \vdash \Delta} \vee_{\mathrm{L}}$$

$$\frac{\Gamma \vdash \phi, \Delta \qquad \Gamma \vdash \psi, \Delta}{\Gamma \vdash \phi \wedge \psi, \Delta} \wedge_{\mathrm{R}} \qquad \frac{\Gamma \vdash \phi/\psi, \Delta}{\Gamma \vdash \phi \vee \psi, \Delta} \vee_{\mathrm{R}}$$

$$\frac{\Gamma \vdash \phi, \Delta}{\Gamma, \neg\phi \vdash \Delta} \neg_{L} \qquad \frac{\Gamma, \phi \vdash \Delta \qquad \Gamma \vdash \psi, \Delta}{\Gamma, \phi \supset \psi \vdash \Delta} \supset_{\mathrm{L}}$$

$$\frac{\Gamma, \phi \vdash \Delta}{\Gamma \vdash \neg\phi, \Delta} \neg_{R} \qquad \frac{\Gamma, \phi \vdash \psi, \Delta}{\Gamma \vdash \phi \supset \psi, \Delta} \supset_{\mathrm{R}}$$

Figure 2: Classical operational rules for sequent calculus.

Specifically, our framework validates each of the classical operational rules of the sequent calculus (Gentzen (1935b,a), **Figure 2**); it differs only in giving up the structural rule of Cut. In fact, we can prove that it coincides, over the

[16] Importantly, I am not claiming that the only theories which can adequately explain these judgments are those which deny that entailment is transitive. Contextualist and epistemicist theories both offer to explain the apparent validity of **Deductive Tolerance** while retaining the claim that entailment is transitive. Contextualists attribute the appearance of validity to the fact that whenever the inference from S_i to S_{i+1} is under consideration, the sharp boundary between positive and negative cases will not fall between the i^{th} and $i+1^{th}$ card Raffman (1994, 1996); Fara (2000); Kennedy (2010). Epistemicist theories attribute the appearance of validity to the fact that no instance of the inference can be known to be a counter-example or that whenever the premise is known, the conclusion will be true (Williamson (1994, §8.4); Williamson (1997); Williams (2008).

boolean fragment, with the logic ST (Cobreros *et al.* (2012)).

Fact 6. The classical operational rules for $\neg, \wedge, \vee$ and $\supset$, are conversationally valid in L_0.

In giving up Cut, certain other classically valid rules will be invalidated. In particular, proof by cases is not conversationally valid.

$$\frac{\Gamma \mid\!\!\!- \phi \vee \psi \qquad \Delta, \phi \mid\!\!\!- \chi \qquad \Pi, \psi \mid\!\!\!- \chi}{\Gamma, \Delta, \Pi \mid\!\!\!- \chi} \quad \text{PBC}$$

This should not be surprising. By the classical operational rules, proof by cases has the transitivity of entailment built in.[17] If we want to give up the transitivity of entailment in order to accommodate judgments about tolerance, we will also have to give up proof by cases.

4 Tolerance

Before addressing the status of **Tolerance** and **Sharp Cutoffs**, we need to extend our semantics to quantification. Our guiding idea will be as follows: In evaluating the universal generalization of a sentence, we check, for each object, whether settling the status of that object would produce a counter-instance. Acceptance requires that no way of settling any object would produce a counter-instance. Rejection requires that any way of settling some object would produce a counter-instance (and, moreover, that the relevant object can be settled permissibly).

Let $c[d, ..., d']$ be the set of extensions of c which settle each of $d, ..., d'$ as a positive or negative case (that is, $c[d, ..., d'] = \{c' \succcurlyeq c | d, ..., d' \in c'^{+} \cup c'^{-}\}$). Let $g^{[x_i/d]}$ be the assignment variant which differs from g (if at all) in mapping x_i to d.

Definition (Quantification).

$$c \in [\![\forall x_i \phi]\!]^{+}_{g} \textit{ iff } \forall d : c[d] \subseteq [\![\phi]\!]^{+}_{g^{[x_i/d]}};$$

$$c \in [\![\forall x_i \phi]\!]^{-}_{g} \textit{ iff } \exists d : c[d] \subseteq [\![\phi]\!]^{-}_{g^{[x_i/d]}} \text{ and } c[d] \neq \varnothing.$$

$\forall x_i \phi$ is accepted (at a context and assignment) iff for each object in the domain and any extension of the context settling the status that object: ϕ is accepted at that extension and the assignment variant mapping the object to x_i. $\forall x_i \phi$ is rejected (at a context and assignment) iff for some object in the domain there is an extension of the context settling the status that object and for any such

[17] To see why, observe that it implies that if $\Gamma, \phi \mid\!\!\!- \psi \vee \psi$ and $\Delta, \psi \mid\!\!\!- \chi$, then $\Gamma, \Delta, \phi \mid\!\!\!- \chi$. But, by $\vee_{\mathrm{R}}$, we have that if $\Gamma, \phi \mid\!\!\!- \psi$, then $\Gamma, \phi \mid\!\!\!- \psi \vee \psi$.

extension: ϕ is rejected at that extension and the assignment variant mapping the object to x_i. By the duality of $\forall/\exists$: $\exists x_i \phi$ is accepted (rejected) iff $\forall x_i \neg\phi$ is rejected (accepted).

It is easy to check that **Limits** and **Continuity** are accepted at every context. By the boundaries constraint, there are guaranteed to be positive and negative cases of the predicate at every context. Similarly, our constraints on the relation of marginal variance ensure that, for any pair of objects, it is possible to find a finite series of marginal variants starting with the first and ending with the second. More notably, every context also accepts **Tolerance**. By the coherence constraint, there is no way of assigning marginal variants to x_i and x_j such that after updating with $F(x_i)$, there is a possible extension at which $F(x_j)$ is rejected. Since each of the principles is accepted at every context and assignment, they are (trivially) jointly accepted at some context and assignment. So they are jointly conversationally consistent. Since acceptance and rejection are mutually exclusive, each principle is rejected at no context. So they are also each supported at every context and assignment.[18]

Fact 7. Limits, **Continuity** and **Tolerance** are conversational validities.

As we have seen, acceptance and support can come apart—as a result, sentences which are supported at every context but accepted at no context can be conversational validities without being conversationally consistent. Not so for **Tolerance**, **Limits** and **Continuity** however. Since each is accepted at every context, they are also conversationally consistent.

Turn to **Sharp Cutoffs**. **Sharp Cutoffs** is rejected at a context iff there are no marginal variants d and d', such that at some extension of the context which settles both: d is a positive case and d' is a negative case. By the **Coherence** constraint, **Sharp Cutoffs** is rejected at every context.

Fact 8. Limits and **Continuity** do not conversationally imply **Sharp Cutoffs**.[19]

[18] *Proof*: To prove each quantified claim is accepted at every context and assignment, it suffices to show that the corresponding open sentence is accepted at every context paired with the appropriate assignment variant(s). <u>For **Continuity**</u>: D is finite and connected under $\approx$. So for any $d_1, d_n \in D$, there is some $d_2, ..., d_{n-1}$ such that, for any c and g: $c \in [\![x_1 \sim x_2 \wedge ... \wedge x_{n-1} \sim x_n]\!]^+_{g^{[x_1/d_1]...[x_n/d_n]}}$. <u>For **Limits**</u>: By **Boundaries**, there is some $d, d' \in D$ such that for any c and g, $c \in [\![F(x_i) \wedge \neg F(x_j)]\!]^+_{g^{[x_i/d][x_j/d']}}$. <u>For **Tolerance**</u>: For any $d, d' \in D$ and an arbitrary g, let $g^{[x_i/d][x_j/d']} = g'$. Then for any c: $c + [\![F(x_i) \wedge x_i \sim x_j]\!]^+_{g'} \models_{g'} F(x_j)$. So $c \in [\![(F(x_i) \wedge x_i \sim x_j) \rightarrow F(x_j)]\!]^+_{g'}$.

[19] *Proof:* Since **Limits** and **Continuity** are conversationally consistent, it suffices to show that **Sharp Cutoffs** is rejected at every context and assignment. Consider an arbitrary c and g. We'll show that for any $d, d' \in D$, and any $c' \in c[d, d']$, one of $x_i \sim x_i$, $F(x_i)$, and $\neg F(x_j)$ is rejected at c' and $g' = g^{[x_i/d][x_j/d']}$. Trivially, if $d \not\approx d'$, then $c' \in [\![x_i \sim x_j]\!]^-_{g'}$. So suppose $d \approx d'$. Since $c' \in c[d, d']$, it follows by **Coherence** that either $d, d' \in c'^+$ or $d, d' \in c'^-$. In the former case, $c' \in [\![\neg F(x_j)]\!]^-_{g'}$; in the latter, $c' \in [\![F(x_i)]\!]^-_{g'}$. So for all d, d' :

In fact, we can observe something further. Consider the negation of **Sharp Cutoffs**.

No Sharp Cutoffs $\neg \exists x_i \exists x_j (F(x_i) \wedge \neg F(x_j) \wedge x_i \sim x_j)$

Since **Sharp Cutoffs** is rejected at every context, its negation is accepted at every context.[20] In this way, our framework vindicates the naïve judgment that, not only are we unable to draw a precise line between the positive and negative cases of a tolerant predicate, we should deny that there is such a line to be drawn (*pace* Fine (1975); Williamson (1994); Keefe (2000)).

4.1 At The Margins

Vagueness, on the view we're considering, requires incompleteness. A conversation cannot (permissibly) reach a state in which each object in a sorites series is settled as either a positive or a negative case. It can, however, come close. Say that a context is marginal iff there is some d, d', and d'' such that (i) $d \approx d'$ and $d' \approx d''$ and (ii) $d \in c^+$ and $d'' \in c^-$. Some marginal contexts are admissible (cf. Kamp (1981, 244)). However, at marginal contexts, the requirements imposed on us are distinctive.

Fact 9. c is marginal iff there is some g such that: $c \models_g F(x_i)$ and $c \models_g \neg F(x_i)$.

Marginal contexts can support contraries. To see why, consider again the model in **Figure 1**. $c_{(k-1,k+1)}$ is both admissible and marginal for $2 \leq k \leq 4$. Yet no extension of $c_{(k-1,k+1)}$ settles the kth element of the series as either a positive or negative case. So, both $F(x_k)$ and $\neg F(x_k)$ will be supported at $c_{(k-1,k+1)}$ and g^*.

At first glance, this might seem surprising. As we observed above, it is never permissible both to classify $\#n$ as dark and to classify $\#n$ as not dark. Shouldn't our framework surely rule inadmissible any context at which both claims are supported?

To see why this worry is misguided, it suffices to recall the role that support plays

$c[d, d'] \subseteq [\![F(x_i) \wedge \neg F(x_j) \wedge x_i \sim x_j]\!]^-_{g[x_i/d][x_j/d']}$.

[20] For the same reason, the material variant of **Tolerance**, $\forall x_i \forall x_j (F(x_i) \wedge x_i \sim x_j) \supset F(x_j)$, is also both conversationally consistent and valid. Note, however, that while the conjunction of its instances is also conversationally valid, it is not conversationally consistent. That is, although $\bigwedge_{x_i, x_j \in \mathcal{V}} ((F(x_i) \wedge x_i \sim x_j) \supset F(x_j))$ is not rejected at any context, there is no context at which it is accepted.

This brings out an important feature of the framework. Universally quantified closed sentences and the conjunction of their instances are mutually entailing. However, their acceptance conditions can come apart. This reflects the idea is that conjunctive claims differ from universal claims in forcing us to consider specific cases. Accepting a conjunctive claim requires settling cases in a way that ensures each conjunct is accepted. In contrast, a universally quantified claim can be accepted without its instances being settled, as long as it is guaranteed they won't be rejected.

in the framework. A context supports ϕ iff ϕ cannot permissibly be rejected. And $\neg\phi$ can permissibly be rejected iff ϕ can permissibly be accepted. The admissibility of contexts which support contraries, then, simply amounts to the admissibility of contexts at which there is a sentence which can neither be accepted nor rejected.

Admitting contexts like this fits well with our opening observations about vagueness. We can coherently imagine, for example, that it is permissible to classify #24 as dark and, simultaneously, to classify #26 as not dark. After doing so, however, it would be impermissible to classify #25 as dark and, equally, impermissible to classify it as not dark. In marginal contexts, refraining from classifying unsettled cases is our only option. Trivially, where we are not permitted to speak, we are required to be silent.

Marginal contexts also support complex claims involving contraries. Where c is marginal, there is some g such that both $\neg(F(x_i) \wedge \neg F(x_i))$ and $F(x_i) \wedge \neg F(x_i)$ are supported at c and g (cf. Priest (1979); Priest & Routley (1989)).

Fact 10. Where c is marginal, there is some g such that: $c \mathrel{|\!\!\!\equiv_g} F(x_i) \wedge \neg F(x_i)$ and $c \mathrel{|\!\!\!\equiv_g} \neg(F(x_i) \vee \neg F(x_i))$

That is, while we are never permitted to reject an instance of **LNC**, we are not always required to reject every counter-instance. And, similarly, while we are never permitted to accept a counter-instance to **LEM**, we are not always required to accept every instance.[21]

While more contentious, I want to suggest that this is also the right prediction. After classifying #24 as dark and #26 as not dark, it is odd to insist that #25 either is dark or it isn't. Since we are not permitted to accepted either disjunct, we should not accept their disjunction (cf. Tye (1994); Cobreros *et al.* (2012); Ripley (2013))). This reflects the idea that accepting a disjunction requires settling cases in such a way as to accept some disjunct (even if there is no disjunct one is required to accept). These considerations equally tell against the permissibility of accepting the relevant instances of **LNC** in marginal contexts (since they have the same acceptance conditions as the relevant instances of **LEM**, by the definition of disjunction).[22]

[21] Since an existentially quantified sentence is accepted (at a context) iff supported, the classical introduction rules for existential quantification fail. Where $\Gamma \models F(x_i) \wedge \neg F(x_i)$, $\Gamma \not\models \exists x_i(F(x_i) \wedge \neg F(x_i))$. We are always permitted to accept that there is no counterinstance to **LNC**, even if we may be faced with instances which we are unable to permissibly reject. The crucial difference is that accepting the existentially quantified claim does not require us to settle any permissible cases, and thus cannot lead us into incoherence.

[22] This judgment is consistent with results reported in Ripley (2011a), which found participants expressed low levels of agreement for contradictions involving objects at the extremes of a sorites series, but intermediate levels for at least some objects in the intermediate region. This tracks predictions for marginal contexts, on the assumption that participants' levels of agreement for supported sentences are higher than for rejected sentences but lower than for accepted sentences.

Crucially, what matters for validity is not what we are permitted to accept but what we are required not to reject. And, here, there is a significant asymmetry between the claim that $\#n$ is and isn't dark and the negation of that claim. In any non-marginal admissible context, the former can permissibly be rejected. In contrast, there is no admissible context (marginal or non-marginal) in which it is permissible to reject the latter (cf. Tappenden (1993, 566)). Accordingly, we classify **LNC** and **LEM** as theorems; no matter how a conversation develops, they can never be denied.

Note that our framework is not dialethist. No contradiction is accepted (at any context). Combined with the account of truth and falsity in the next section, this guarantees that there is no context at which we are permitted to accept that a contradiction is true. Nor is it paraconsistent. Every contradiction is inconsistent (that is, $\phi \wedge \neg\phi \models \bot$). We retain the non-triviality of our logic by virtue of giving up transitivity. From the fact that some premises entail an inconsistent claim, it does not follow that the premises are themselves inconsistent. Thus, for instance, while $F(x_i), \neg F(x_j), x_i \sim x_k$, and $x_k \sim x_j$ imply $F(x_k) \wedge \neg F(x_k)$, they do not imply $\bot$. Just because accepting some premises makes it both impermissible to accept ϕ and impermissible to reject ϕ, this does not mean accepting the premises is itself impermissible.

5 Indeterminacy

Finally, to evaluate the status of **Bivalence**, we need to extend our semantics to truth and falsity attributions. We'll take the acceptance and rejection conditions of claims about truth and falsity to be mutually exhaustive.

Definition (Truth/Falsity).

i. $c \in [\![\mathsf{T}(\phi)]\!]^+_g$ *iff* $c \in [\![\phi]\!]^+_g$
$c \in [\![\mathsf{T}(\phi)]\!]^-_g$ *iff* $c \notin [\![\phi]\!]^+_g$

ii. $c \in [\![\mathsf{F}(\phi)]\!]^+_g$ *iff* $c \in [\![\phi]\!]^-_g$
$c \in [\![\mathsf{F}(\phi)]\!]^-_g$ *iff* $c \notin [\![\phi]\!]^-_g$

$\mathsf{T}(\phi)$ is accepted at c and g iff ϕ is accepted there; otherwise, it is rejected. $\mathsf{F}(\phi)$ is accepted at c and g iff ϕ is rejected there; otherwise, it is rejected. Put another way, a sentence is accepted as true (false) at exactly those contexts at which it is accepted (rejected). At contexts in which it is neither accepted nor rejected, it is rejected as true and rejected as false.

Adding truth and falsity to our semantics has some important effects. Consider the property of persistence:

Definition (Persistence).

ϕ is persistent iff for all c and g: if $c \in [\![\phi]\!]^+_g$ and $c \preccurlyeq c'$, then $c' \notin [\![\phi]\!]^-_g$.

For sentences in L_1, acceptance and rejection behave monotonically: once a sentence is accepted (rejected), the conversation cannot develop into a state in which it is rejected (accepted). *A fortiori*, if $p \in \mathsf{L}_1$, ϕ is persistent. However, once truth and falsity are introduced this no longer holds in full generality.

At a context at which ϕ is not accepted, $\mathsf{T}(\phi)$ will be rejected (and its negation accepted). Nevertheless, as long as $\neg\phi$ is not supported at the context, the context will have an extension at which $\mathsf{T}(\phi)$ is accepted. Accordingly, $\neg\mathsf{T}(\phi)$ is non-persistent: some contexts at which it is accepted can be extended into contexts at which it is rejected. The same goes, *mutatis mutandis*, for $\neg\mathsf{F}(\phi)$.

Failures of persistence have significant implications. Updating with a non-persistent sentence is not guaranteed to produce a context at which that sentence is supported. Accordingly, we will have failures of **Idempotence**.

Idempotence $\quad \Gamma, \phi \mid\!\!\!- \phi$

The lesson is that, if we want to capture what is distinctive about claims of truth and falsity, we need to change how we think about update. Our existing update operation only modeled the effect of coming to accept a sentence. However, where that sentence is non-persistent, coming to accept it is compatible with permissibly rejecting it at some later point. To preserve **Idempotence**, we need to capture the way that update changes not only the present state of a conversation, but also the way that conversation can develop in the future. In addition to ensuring that a sentence is accepted, update should rule out any later changes to the context which would lead it no longer to be accepted.

5.1 Updating Update

A context space, C, is a set of admissible contexts. A context space represents the states of a conversation which have not been antecedently ruled out. Where $c \in C$ and there is no $c' \in C$ such that $c' \prec c$, c is a candidate for the present state of the conversation represented by C. Where there is more than one such context, we say that the present state of the conversation is indeterminate.[23]

Update on context spaces is defined as intersection.

Definition (Update$^+$). $\quad C \dotplus C' = C \cap C'$

Intuitively, $C \dotplus [\![\phi]\!]^+_g$ is the result of (i) coming to accept ϕ at the present state of the conversation and (ii) ruling out any state at which ϕ is not accepted. We define revised relations of support and entailment for context sets analogously to our original definitions.[24] Where an argument is valid in the revised sense, we will say it is conversationally valid$^+$.

[23] A determinate context can become indeterminate in a number of ways, e.g., as a result of update with a disjunction.

[24] Derivatively, were $\mathcal{C}$ is a set of context spaces, $\mathcal{C} \dotplus C' = \{C \dotplus C' \mid C \in \mathcal{C}\}$ and $\mathcal{C} \mid\!\models_g \phi$ iff for all $C \in \mathbb{C} : C \mid\!\models_g \phi$.

Definition (Support$^+$ & Validity$^+$).

i. $C \Vert\!\models_g \phi$ iff for all $c \in C$: $c \notin [\![\phi]\!]_g^-$.
ii. $\phi_i, ...\phi_j \Vert\!\models \psi$ iff for all C and g: $C \dot{+} ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]_g^+) \Vert\!\models_g \psi$.

Where c is an admissible context, $\uparrow c = \{c' \mid c \preccurlyeq c'\}$ is the context space comprising every admissible extension of c. $\uparrow c$ characterizes a conversation which is (determinately) presently in the state characterized by c and at which no admissible way the conversation could develop is ruled out. There are some nice relationships between our old and new frameworks over such context spaces.

Observation $\uparrow(c + [\![\phi]\!]_g^+) = (\uparrow c) \dot{+} [\![\phi]\!]_g^+$ (for persistent ϕ)

Observation $c \models_g \phi$ *iff* $\uparrow c \Vert\!\models_g \phi$ (for persistent ϕ)

Over an important class of context spaces, our new notions of update and support behave in the same way for persistent sentences as our old ones did. Updating$^+$ the admissible extensions of a context with $[\![\phi]\!]_g^+$ returns all and only those contexts which are admissible extensions of the result of updating the context with $[\![\phi]\!]_g^+$. Similarly, the admissible extensions of a context support$^+$ ϕ iff the context supports ϕ. From these observations, it follows that the old and new relations of entailment coincide over the language free of T and F.

Fact 11. $\phi_i, .., \phi_j \models \psi$ *iff* $\phi_i, .., \phi_j \Vert\!\models \psi$ (for $\phi_i, ...\phi_j, \psi \in \mathsf{L}_1$)[25]

Fact 11 is reassuring. It tells us that our new framework preserves the results of the previous sections. Note, however, that unlike in our old framework, **Idempotence** holds even for the non-persistent fragment of the language. Since $\dot{+}$ is intersective, after update$^+$ with ϕ, a context space is guaranteed to include only contexts at which ϕ is accepted. Since no context both accepts and rejects ϕ, it follows that the result of updating$^+$ with ϕ is guaranteed to support$^+$ ϕ in the new framework.

Finally, observe that the framework resolves the puzzle to do with indeterminacy which we started with. The acceptance conditions for ϕ and $\mathsf{T}(\phi)$ coincide, and, likewise, for $\neg\phi$ and $\mathsf{F}(\phi)$. As a result, updating$^+$ a context space with one member of either pair is guaranteed to result in a set of contexts which accept

[25] *Proof:* for the R⇒L direction, suppose that $\phi_i, ..., \phi_j \not\models \psi$. So there is some c such that $c + ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!] \not\models_g \psi$. It follows, by our second observation, that $\uparrow(c + ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]) \Vert\!\not\models_g \psi$. But we know, from our first observation, that $\uparrow(c + ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]) = \uparrow(c) \dot{+} ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!])$ So, since we have a counter-instance, we can conclude that $\phi_i, ..., \phi_j \Vert\!\not\models \psi$. To complete the proof, we simply contrapose.

For the L⇒R direction, suppose that $\phi_i, ..., \phi_j \Vert\!\not\models \psi$. Let C be a context space such that $C \dot{+} ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]^+) \Vert\!\not\models_g \psi$. There must therefore be some $c \in C \dot{+} ([\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]^+)$ such that $c \in [\![\psi]\!]_g^-$. However, by the definition of $\dot{+}$, we know that $c + [\![\phi_i]\!]_g^+ \cap ... \cap [\![\phi_j]\!]^+ = c$. So, since we have a counterinstance, we can conclude that $\phi_i, ..., \phi_j \not\models \psi$. Again, to complete the proof, we simply contrapose.

(and, *a fortiori*, do not reject) the other. So, **Transparency** and **Polarity** are both supported$^+$ at every context space.

While no context rejects both ϕ and $\neg\phi$, there are contexts at which $\mathsf{T}(\phi)$ and $\mathsf{F}(\phi)$ are both rejected. Accordingly, **Bivalence** is not supported$^+$ at every context space.

> **Fact 12.** **Transparency** and **Polarity** are conversationally valid$^+$.
> **Bivalence** is conversationally invalid$^+$.

Since **Transparency** and **Polarity** are everywhere accepted, update$^+$ with either has no effect on the context space. So, it follows immediately from **Fact 12** that:

> **Fact 13.** **Transparency** and **Polarity** do not conversationally imply$^+$ **Bivalence**.

Within a classical setting, **Bivalence** follows from **LEM** and the disquotational picture of truth and falsity via proof by cases. Our framework invalidates proof by cases. This is because, as we saw, proof by cases has the transitivity of entailment built in. There is no conversation in which it is permissible to reject that either $\#n$ is dark or it isn't. And, in any conversation, after coming to accept that $\#n$ is (not) dark, it is impermissible to reject that it is true (false) that $\#n$ is dark. However, we cannot conclude that there is no conversation at which it is permissible to reject that it is either true that $\#n$ is dark or false that $\#n$ is dark. To do so would be to conflate accepting a claim with being incapable of rejecting it.[26]

Finally, observe that the addition of truth and falsity operators to the language augments the ways that objects can be classified. In addition to classifying an object as a positive/negative case (via update$^+$ with $F(x_i)$ or $\neg F(x_i)$, respectively), an object can also be classified as unsettled. Updating$^+$ a context space with $[\![\neg\mathsf{T}(F(x_i)) \wedge \neg\mathsf{F}(F(x_i))]\!]^+_g$ eliminates all contexts at which $g(x_i)$ is settled as either a positive or negative case. Intuitively, this reflects the idea that classifying an object as an unsettled case makes it impermissible go on to classify it as either a positive or negative case.

[26]Within the updated system, there is a mismatch between the semantic clauses for the conditional (which are stated in terms of support) and the definition of entailment (which is stated in terms of support$^+$). The natural fix is to update the original clauses accordingly.

> **Definition (Conditional$^+$).**
> $c \in [\![\phi \to \psi]\!]^+_g$ *iff* $(\uparrow c) \oplus [\![\phi]\!]^+_g \Vdash_g \psi$.
> $c \in [\![\phi \to \psi]\!]^-_g$ *iff* $(\uparrow c) \oplus [\![\phi]\!]^+_g \nVdash_g \psi$.

The pair of observations on page 23 suffice to ensure that acceptance conditions of sentences in the persistent fragment of the language are unaffected by this revision.

5.2 Borderline Cases

The framework developed above does not make use of a category of borderline cases. One reason for this is that it is not needed in stating the original puzzles or developing the framework which aims to resolve them. Nevertheless, it is natural to ask: is there a reasonable candidate for this category in the framework?

'Borderline case' is used by many authors as a term-of-art, to describe certain objects in an intermediate region of a sorites series. Within the present framework, we can identify multiple conditions which characterize different subsets of this region, each of which appear reasonable candidates for being the borderline cases.

Say that two contexts approximate each other iff for any object on which they differ, there is a marginal variant of that object on which they agree.[27] Intuitively, we can think of the approximations of a context as differing only in ways which are irrelevant. Then, consider the condition ϕ satisfies at c and g iff:

(∇_1) ϕ is neither accepted nor rejected at every extension of c (relative to g).

(∇_2) ϕ is neither accepted nor rejected at c (relative to g).

(∇_3) ϕ is neither accepted nor rejected at some approximation of c (relative to g).

For each condition, **Figure 3** depicts the objects for which $F(x_i)$ satisfies that condition relative to the depicted context and an assignment mapping x_i to that object.

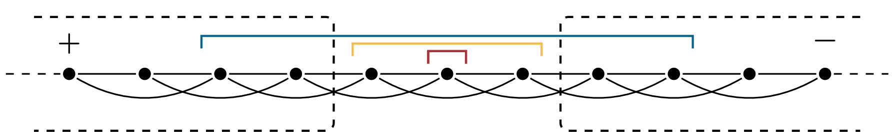

Figure 3: Intervals in the intermediate region.

The red region corresponds to the objects for which $F(x_i)$ satisfies (∇_1). These are the objects which are prohibited from being classified either way. It is these objects for which instances of **LEM** cannot be accepted. The yellow region corresponds to the objects for which $F(x_i)$ satisfies (∇_2). These are the objects which are not classified either way. It is these objects for which instances of **Bivalence** fails. The blue region corresponds to the objects for which $F(x_i)$ satisfies (∇_3). These are the objects such that leaving them unsettled would not make a relevant difference to the context.

[27] More carefully, a context c approximates c' iff:
(i) for all $d \in c^+ \cup c'^+$ there is some $d' \in c^+ \cap c'^+$ such that $d \approx d'$; and
(ii) for all $d \in c^- \cup c'^-$ there is some $d' \in c^- \cap c'^-$ such that $d \approx d'$.

None of these conditions has a uniquely privileged status. Instead, it seems better to conclude that there are a multiple overlapping but non-equivalent categories, each of which picks out different features of objects in the intermediate region, none of which have a special claim to the title 'borderline case'.

Of course, 'Borderline case' is not always used as a term-of-art. Raffman (2005, 2014a) argues that, in ordinary speech, 'borderline case' is used to describe objects which fall between the extensions of two proximate incompatible predicates—such as, e.g., 'light' and 'dark'. In Raffman's terminology, proximate incompatible predicates are contraries associated with the same order for which there exist objects which are permitted to be classified as either. A borderline case of 'dark' (relative to 'light'), for Raffman, is simply an object which is in the extension of neither.

Raffman's proposal is designed to capture the idea that, in ordinary speech, which objects are classified as, e.g., 'borderline dark' appears sensitive to what alternative classifications are operative (e.g., 'light', 'pale', 'nearly dark', etc.). The languages in this paper contain only a single unary predicate. However, it would be uncomplicated to extend them with additional predicates associated with the same ordering. Within such an extension, a definition in terms of incompatible predicates could be given, with the aim of capturing this feature of our everyday notion of borderline cases. Further exploration of the ordinary use of 'borderline case' within the present framework will need to be left to future work, however.[28]

6 Comparisons

The framework we have been considering has features in common with a number of previous accounts of vagueness. There is considerable overlap with both dynamic theories (Kamp (1981); Pinkal (1983); Ballweg (1983); Eikmeyer & Rieser (1983); Kyburg & Morreau (2000); Barker (2002, 2003)) and non-transitive theories (Zardini (2008); van Rooij (2010); Cobreros *et al.* (2012, 2015, forthcoming)).[29] In this section, I'll explore these comparisons in more detail, focusing in particular on the dynamic theory of Kamp (1981) and the non-transitive theories K_0 of Zardini (2008) and ST of Cobreros *et al.* (2012).

6.1 Dynamic Theories: Kamp (1981)

Dynamic approaches to vagueness assign a central role to the operation of updating on contexts. Among dynamic approaches, our framework has most in common with that of Kamp (1981). Informally, both appeal to a similar philosophical picture of vagueness. Like Kamp, our framework attempts to account for the way using vague expressions can settle previously unsettled cases. Both

[28] I am grateful to an anonymous referee at *The Journal of Philosophy* for encouraging me to discuss Raffman's account of borderline cases.

[29] Cf. Beall (2014) for a related non-transitive treatment of the conditional.

also draw on a similar methodological toolkit. Like Kamp, our framework employs a bilateral semantics, a restriction to coherent contexts and an entailment relation defined in terms of monotonic update. A signature feature of Kamp (1981) is validating **Deductive Tolerance**. And his framework achieves this while also retaining **MP**.[30]

However, there are also significant points of difference between the two frameworks. Kamp does not validate **Tolerance**. While each of its instances is valid, their universal generalization is not. In fact, not only is **Tolerance** invalid, it is inconsistent with **Limits** and **Continuity**. At any context at which the latter are accepted, the former with be rejected. Accordingly, Kamp's framework offers us no help in resolving our first puzzle to do with tolerance.

Second, both **LNC** and **LEM** are invalid for Kamp. As Kamp (2013) establishes, unless the class of models is unreasonably restricted, the logic of the framework can be no stronger than Strong Kleene (a.k.a. K_3). Since neither non-contradiction nor excluded middle is valid in the latter, neither will be valid in the former.

Most significantly, Kamp's framework fails to determine a unique consequence relation. Extensions are defined relative to a syntactic inference relation. Different choices of inference relation will give rise to different relations of semantic consequence. A key desideratum is that the consequence relation coincides with the inference relation that determines it. However, there are multiple such candidates. And, as Kamp observes, it is not obvious whether any of them provide a plausible logic of vagueness. Thus, the present framework also differs from Kamp in making concrete predictions about what inferences are valid.

6.2 Non-Transitive Theories: Zardini (2008) and Cobreros et al. (2012)

Non-transitive approaches to vagueness deny that valid arguments can be freely concatenated. An argument that results from stringing together a series of valid sub-arguments need not be valid itself. Among non-transitive theories, the two most developed are those of Zardini (2008) and Cobreros *et al.* (2012).

In both, as in our framework, transitivity fails for the same reason: entailment is defined in an asymmetric fashion. In evaluating whether an argument is valid, the standard to which the premises are held is strictly higher than the standard to which the conclusion is held. And clearly, from the fact that wherever one claim meets the higher standard, a second meets the lower, and wherever the second meets the higher, a third meets the lower, we cannot conclude that wherever the first meets the higher, the third meets the lower. This point is important. While update plays a role in the statement of conversational validity, this role is ultimately inessential (cf. fn.10). Failures of transitivity

[30] Note that the closely related rule, that if $\Gamma \vdash \phi$ and $\Gamma' \vdash \phi \to \psi$, then $\Gamma, \Gamma' \models \psi$ fails in Kamp's framework as it does in ours.

arise not from shifts in context but from this asymmetry between the evaluation of premises and conclusion.

Both Zardini and Cobreros et al. validate **Deductive Tolerance** and **MP**. The latter also validate the L/R rules of the sequent calculus for the propositional fragment of the language (Ripley (2013)). Differences emerge, however, in the treatment of quantification. Zardini focuses exclusively on a propositional language, and so doesn't consider the puzzles involving quantified sentences in §2. Cobreros et al. do consider a quantified language, and, as in our framework, classify **Tolerance** and **No Sharp Cutoffs** as valid. Crucially, however, their framework does not resolve either of the puzzles to do with tolerance with which we started. **Limits**, **Continuity** and **Tolerance** remain inconsistent. And, likewise, **Limits** and **Continuity** continue to imply **Sharp Cutoffs**.

This reflects a significant defect of the theory. The positive status which their theory accords **Tolerance** is also accorded to **Sharp Cutoffs**: the claim that there is a sharp boundary dividing the positive and negative cases. Yet, as we observed above, it is part of the naïve picture of vagueness that the latter should be rejected, while the former should be accepted. If it is to vindicate the naïve picture, a theory needs to capture the asymmetry in our judgments about these two sentence schema. Insofar as it affords **Tolerance** and **Sharp Cutoffs** the same status, the theory in Cobreros *et al.* (2012) fails to do so.

Finally, like Kamp, neither Zardini nor Cobreros et al. extends their non-transitive approach to the tension between the disquotational picture of truth and apparent failures of **Bivalence**.

7 Wrapping Up

The primary aim of this paper has been to show that our naïve picture of vagueness can be captured within a coherent and easily interpreted framework. Failures of transitivity are built into this naïve picture. Combining observations about acceptance and rejection with a familiar way of understanding entailment puts us in a position to explain how such failures arise.

Here is the main refrain, one more time: an argument is judged valid iff no-one who accepts its premises could reject its conclusion. If we restrict our attention to precise matters, it is plausible that a sentence must be accepted iff it may not be rejected. When it comes to vague matters, however, things are more complex; what we are not permitted to reject and what we are required to accept can come apart. By adopting a framework which reflects this, we can both capture the failures of transitivity which form part of our naïve picture, and, simultaneously, resolve classic puzzles to do with tolerance and indeterminacy.

References

Ballweg, Joachim. 1983. Vagueness or Context Dependence? *Pages 59–78 of:* Ballmer, Thomas T., & Pinkal, Manfred (eds), *Approaching Vagueness*. Elsevier.

Barker, Chris. 2002. The Dynamics of Vagueness. *Linguistics and Philosophy*, **25**(1), 1–36.

Barker, Chris. 2003. Clarity and the Grammar of Skepticism. *Mind & Language*, **24**(3), 253–273.

Beall, Jc. 2014. Finding Tolerance without Gluts. *Mind*, **123**(491), 791–811.

Burgess, J. A. 1990. The Sorites Paradox and Higher-Order Vagueness. *Synthese*, **85**(3), 417–474.

Burgess, J. A., & Humberstone, I. L. 1987. Natural Deduction Rules for a Logic of Vagueness. *Erkenntnis*, **27**(2), 197–229.

Campbell, Richmond. 1974. The Sorites Paradox. *Philosophical Studies*, **26**(3-4), 175–191.

Cobreros, Pablo, Egré, Paul, Ripley, David, & van Rooij, Robert. 2012. Tolerant, Classical, Strict. *Journal of Philosophical Logic*, **41**(2), 347–385.

Cobreros, Pablo, Egré, Paul, Ripley, David, & van Rooij, Robert. 2015. Vagueness, Truth and Permissive Consequence. *Pages 409–430 of:* Fujimoto, Kentaro, Martínez Fernández, José, Galinon, Henri, & Achourioti, Theodora (eds), *Unifying the Philosophy of Truth*. Springer Verlag.

Cobreros, Pablo, Egré, Paul, Ripley, Dave, & van Rooij, Robert. forthcoming. Tolerant Reasoning: Nontransitive or Nonmonotonic? *Synthese*, 1–25.

Dekker, Paul. 1993. *Transsentential Meditations: Ups and Downs in Dynamic Logic*. Ph.D. thesis, University of Amsterdam: ILLC.

Dummett, Michael. 1975. Wang's Paradox. *Synthese*, **30**(3-4), 201–32.

Eikmeyer, Hans-Jurgen, & Rieser, Hannes. 1983. A Formal Theory of Context Dependence and Context Change. *Pages 131–188 of:* Ballmer, Thomas T., & Pinkal, Manfred (eds), *Approaching Vagueness*. Elsevier.

Eklund, Matti. 2002. Inconsistent Languages. *Philosophy and Phenomenological Research*, **64**(2), 251–275.

Eklund, Matti. 2005. What Vagueness Consists In. *Philosophical Studies*, **125**(1), 27–60.

Eklund, Matti. 2019. Incoherentism and the Sorites Paradox. *In:* Zardini, Elia, & Oms, Sergi (eds), *The Sorites Paradox*. Cambridge University Press.

Fara, Delia Graff. 2000. Shifting Sands: An Interest Relative Theory of Vagueness. *Philosophical Topics*, **28**(1), 45–81.

Field, Hartry. 2000. Indeterminacy, Degree of Belief, and Excluded Middle. *Noûs*, **34**(1), 1–30.

Fine, Kit. 1975. Vagueness, Truth and Logic. *Synthese*, **30**(3-4), 265–300.

Gaifman, Haim. 2010. Vagueness, Tolerance and Contextual Logic. *Synthese*, **174**(1), 5–46.

Gentzen, Gerhard. 1935a. Untersuchungen Über Das Logische Schließen. I. *Mathematische Zeitschrift*, **35**, 176–210.

Gentzen, Gerhard. 1935b. Untersuchungen Über Das Logische Schließen. II. *Mathematische Zeitschrift*, **39**, 405–431.

Gillies, Anthony S. 2004. Epistemic Conditionals and Conditional Epistemics. *Noûs*, **38**(4), 585–616.

Horgan, Terence. 1994. Robust Vagueness and the Forced-March Sorites Paradox. *Philosophical Perspectives*, **8**(Logic and Language), 159–188.

Horgan, Terence. 1998. The Transvaluationist Conception of Vagueness. *The Monist*, **81**(2), 313–330.

Horwich, Paul. 1998a. *Meaning*. Vol. 50. Oxford University Press.

Horwich, Paul. 1998b. *Truth*. Clarendon Press.

Horwich, Paul. 2000. The Sharpness of Vague Terms. *Philosophical Topics*, **28**(1), 83–92.

Hyde, Dominic. 1997. From Heaps and Gaps to Heaps of Gluts. *Mind*, **106**(424), 641–660.

Jaśkowski, Stanisław. 1969. Propositional Calculus for Contradictory Deductive Systems. *Studia Logica*, **24**(1), 143–160.

Kamp, Hans. 1981. The Paradox of the Heap. *Journal of Symbolic Logic*, 225–277.

Kamp, Hans. 2013. *Meaning and the Dynamics of Interpretation: Selected Papers of Hans Kamp*. Brill.

Keefe, Rosanna. 2000. Supervaluationism and Validity. *Philosophical Topics*, **28**(1), 93–105.

Kennedy, Christopher. 2010. Vagueness and Comparison. *In:* Egre, Paul, & Klinedinst, Nathan (eds), *Vagueness and Language Use*. Palgrave-Macmillan.

Kyburg, Alice, & Morreau, Michael. 2000. Fitting Words: Vague Language in Context. *Linguistics and Philosophy*, **23**(6), 577–597.

Lewis, David. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, **8**(1), 339–359.

Ludlow, Peter. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press.

MacFarlane, John. 2016. Vagueness as Indecision. *Aristotelian Society Supplementary Volume*, **90**(1), 255–283.

Machina, Kenton F. 1972. Vague Predicates. *American Philosophical Quarterly*, **9**(3), 225–233.

Machina, Kenton F. 1976. Truth, Belief, and Vagueness. *Journal of Philosophical Logic*, **5**(1), 47–78.

McGee, Vann, & McLaughlin, Brian. 1995. Distinctions Without a Difference. *Southern Journal of Philosophy*, **33**(S1), 203–251.

Peacocke, Christopher. 1981. Are Vague Predicates Incoherent? *Synthese*, **46**(1), 121–141.

Pinkal, Manfred. 1983. Towards a Semantics of Precization. *Pages 13–57 of:* Ballmer, Thomas T., & Pinkal, Manfred (eds), *Approaching Vagueness*. Elsevier.

Price, Huw. 1990. Why 'Not'? *Mind*, **99**(394), 221–238.

Priest, Graham. 1979. The Logic of Paradox. *Journal of Philosophical Logic*, **8**(1), 219–241.

Priest, Graham, & Routley, Richard. 1989. Systems of Paraconsistent Logic. *Pages 142–155 of:* Priest, G., Routley, R., & Norman, J. (eds), *Paraconsistent Logic: Essays on the Inconsistent*. Philosophia Verlag.

Raffman, Diana. 1994. Vagueness Without Paradox. *Philosophical Review*, **103**(1), 41–74.

Raffman, Diana. 1996. Vagueness and Context-Relativity. *Philosophical Studies*, **81**(2-3), 175–192.

Raffman, Diana. 2005. Borderline Cases and Bivalence. *Philosophical Review*, **114**(1), 1–31.

Raffman, Diana. 2011. Vagueness and Observationality. *Pages 107–121 of:* Ronzitti, Giuseppina (ed), *Vagueness: A Guide*. Springer Verlag.

Raffman, Diana. 2012. Indiscriminability and Phenomenal Continua. *Philosophical*

Perspectives, **26**(1), 309–322.
Raffman, Diana. 2014a. Disjunctivism, Discrimination, and Categorization. *In:* Brogaard, Berit (ed), *Does Perception Have Content?* Oxford University Press.
Raffman, Diana. 2014b. *Unruly Words: A Study of Vague Language*. Oxford University Press.
Raffman, Diana. 2017. Vagueness, Hysteresis, and the Instability of Color. *In:* Silva, Marcos (ed), *How Colours Matter to Philosophy*. Springer.
Restall, Greg. 2005. Multiple Conclusions. *In:* Hájek, Petr, Valdés-Villanueva, Luis, & Westerståhl, Dag (eds), *Logic, Methodology and Philosophy of Science*. College Publications.
Ripley, David. 2011a. Contradictions at the Borders. *Pages 169–188 of:* Nouwen, Rick, van Rooij, Robert, Sauerland, Uli, & Schmitz, Hans-Christian (eds), *Vagueness in Communication*. Springer.
Ripley, David. 2011b. Negation, Denial, and Rejection. *Philosophy Compass*, **6**(9), 622–629.
Ripley, David. 2013. Revising Up: Strengthening Classical Logic in the Face of Paradox. *Philosophers' Imprint*, **13**.
Ripley, David. 2020. Denial. *Pages 47–57 of:* Deprez, Vivianne, & Espinal, M. Teresa (eds), *The Oxford Handbook of Negation*. OUP.
Rumfitt, I. 2000. Yes and No. *Mind*, **109**(436), 781–823.
Sainsbury, R. M. 1989. Tolerating Vagueness. *Proceedings of the Aristotelian Society*, **89**, 33–48.
Sainsbury, R. M. 1996. Concepts Without Boundaries. *Pages 186–205 of:* Keefe, Rosanna, & Smith, Peter (eds), *Vagueness: A Reader*. MIT Press.
Sanford, David H. 1975. Borderline Logic. *American Philosophical Quarterly*, **12**(1), 29–39.
Sanford, David H. 1976. Competing Semantics of Vagueness: Many Values versus Super-Truth. *Synthese*, **33**(2-4), 195–210.
Shapiro, Stewart. 2006. *Vagueness in Context*. Oxford University Press.
Smiley, Timothy. 1996. Rejection. *Analysis*, **56**(1), 1–9.
Soames, Scott. 1998. *Understanding Truth*. Oxford University Press USA.
Tappenden, Jamie. 1993. The Liar and Sorites Paradoxes: Toward a Unified Treatment. *Journal of Philosophy*, **90**(11), 551–577.
Tye, Michael. 1989. Supervaluationism and the Law of Excluded Middle. *Analysis*, **49**(3), 141–143.
Tye, Michael. 1994. Sorites Paradoxes and the Semantics of Vagueness. *Philosophical Perspectives*, **8**, 189–206.
van Rooij, Robert. 2010. *Vagueness, Tolerance and Non-Transitive Entailment*.
Williams, J. R. G. 2008. Supervaluationism and Logical Revisionism. *Journal of Philosophy*, **105**(4), 192–212.
Williamson, Timothy. 1992. Inexact Knowledge. *Mind*, **101**(402), 217–242.
Williamson, Timothy. 1994. *Vagueness*. Routledge.
Williamson, Timothy. 1997. Imagination, Stipulation and Vagueness. *Philosophical Issues*, **8**, 215–228.
Wright, Crispin. 1975. On the Coherence of Vague Predicates. *Synthese*, **30**(3-4), 325–65.
Wright, Crispin. 1987. Further Reflections on the Sorites Paradox. *Philosophical Topics*, **15**(1), 227–290.
Wright, Crispin. 1995. The Epistemic Conception of Vagueness. *Southern Journal of Philosophy*, **33**(S1), 133–160.

Zardini, Elia. 2008. A Model of Tolerance. *Studia Logica*, **90**(3), 337–368.
Ziff, Paul. 1974. The Number of English Sentences. *Foundations of Language*, **11**(1), 519–32.