

Inconstancy and inconsistency*

David Ripley
davewripley@gmail.com

In everyday language, we can call someone ‘consistent’ to say that they’re reliable, that they don’t change over time. Someone who’s consistently on time is *always* on time. Similarly, we can call someone ‘inconsistent’ to say the opposite: that they’re changeable, mercurial. A student who receives inconsistent grades on her tests throughout a semester has performed better on some than on others.

With our philosophy hats on, though, we mean something quite different by ‘consistent’ and ‘inconsistent’. Something consistent is simply something that’s not contradictory. There’s nothing contradictory about being on time, so anyone who’s on time at all is consistently on time, in this sense of ‘consistent’. And only a student with an unusual teacher can receive inconsistent grades on her tests throughout a semester, in this sense of ‘inconsistent’.

In this paper, I’ll use ‘consistent’ and ‘inconsistent’ in their usual philosophical sense: to mark the second distinction. By contrast, I’ll use ‘constant’ and ‘inconstant’ to mark the first distinction. And although we can, should, and do sharply distinguish the two distinctions, they are related. In particular, they have both been used to account for some otherwise puzzling phenomena surrounding vague language. According to some theorists, vague language is inconstant. According to others, it is inconsistent.

I do not propose here to settle these differences; only to get a bit clearer about what the differences amount to, and to show what it would take to settle them. In §1 I’ll give a brief overview of theories of vagueness that crucially invoke inconstancy, and theories that crucially invoke inconsistency. (I’ll also briefly mention inconsistency’s twin, incompleteness.) In §2, I present a formal framework (along the lines of that in [9]) for inconstancy and inconsistency. This will clarify just how the target theories of vagueness differ. §3 summarizes one strain of experimental research on speakers’ use of vague predicates: research into claims like ‘Man x is both tall and not tall’, when man x is a borderline case of ‘tall’. Such phenomena invite explanation in terms of inconstancy or inconsistency, and I’ll explore explanations of both sorts. These explanations will lead me to revisit the formal framework in §4, and prove that, in a certain sense to be clarified there, inconstancy and inconsistency are deeply related; each can do everything the other can do. Finally, §5 offers some lessons to be drawn from the preceding discussion, and points out ways around the

*Published as [17]. Research supported by the Agence Nationale de la Recherche, Project “Cognitive Origins of Vagueness”, grant ANR-07-JCJC-0070. This paper has benefited greatly from discussions with Emmanuel Chemla, Paul Egré, Patrick Greenough, Joshua Knobe, Graham Priest, Diana Raffman, François Recanati, Greg Restall, Robert van Rooij, Laura Schroeter, and Zach Weber, as well as the comments of two anonymous referees.

equivalence result proved in §4. The overall lesson is this: only with a theory of context in hand can we distinguish inconstancy from inconsistency.

1 Setting the scene

This discussion will operate broadly within the framework of [9]. Although it is familiar, I'll recount it briefly here. It's a story about how expressions come to have extensions, and so how predicates come to be associated with things that do or don't satisfy them, how names come to be associated with their bearers, how sentences come to be associated with truth-values, &c. Each (disambiguated) expression is first of all associated with a *character*. From character to extension, there is a two-step process. First, the character, together with the context of use, determines a *content*. For simplicity, I'll identify characters with functions from contexts of use to contents. Second, the content, together with circumstances of evaluation, determines an *extension*. Again for simplicity, I'll take contents to be functions from circumstances of evaluation to extensions. Note that this is so far terribly silent about what extensions, contexts, and circumstances of evaluation *are*. That's good; we'll be considering a few different hypotheses as we go along, but much of the discussion will stay at this abstract level.

1.1 Inconstant characters

I'll say that a character C is *inconstant* when there are contexts of use c_1, c_2 such that $C(c_1) \neq C(c_2)$. The classic examples of expressions with inconstant characters are indexicals. Consider the character of the indexical 'I'; call it I . If c_1 is a context of use in which I utter 'I', and c_2 is a context of use in which Napoleon utters 'I', $I(c_1)$ is a function that takes any circumstance of evaluation to me, and $I(c_2)$ is a function that takes any circumstance of evaluation to Napoleon.¹ Since I am not Napoleon, I is an inconstant character. Similar examples show that other indexicals, like 'here', 'now', 'three months ago', &c, have inconstant characters as well.

Indexicals, though, are not the only expressions with inconstant character. Demonstratives as well have inconstant character: the content determined by 'that' in one context of use (where one thing is demonstrated) can differ from the content determined by 'that' in another context of use (where something else is demonstrated).

Some (eg [20]) have alleged that vague words have inconstant character as well. (In fact, Soames takes vague words to *be* indexicals.) On these views, uses of the same vague predicate in distinct contexts of use can determine distinct contents. Thus, if context of use varies in the appropriate ways, two different occurrences of the same vague sentence might be used to claim different things. Here's an example. Let's consider a flea named Maximilian, and suppose that Maximilian is particularly large for a flea. In a context where we're categorizing fleas by size, Carrie might say 'Maximilian is huge', and thereby express that Maximilian is very large for a flea. By contrast, in a case where

¹This makes some irrelevant assumptions for concreteness; maybe 'I' doesn't determine a constant function at all! But we don't need to worry about that here; it's an example of how the framework might operate.

we're categorizing animals by size, Carrie might say 'Maximilian isn't huge', and thereby express that Maximilian isn't very large for an animal. In these two occurrences, 'huge' is expressing two distinct contents. In the first case, it expresses the property of being very large for a flea, and in the second case, it expresses the property of being very large for an animal.

This is an example of a particular sort of context-dependence: dependence on contextually-determined comparison class. It is relatively uncontroversial that vague words exhibit this sort of context-dependence. But this is not *why* they are vague. For one thing, some precise predicates depend on comparison class as well. Maximilian is larger than average, if the comparison class is the class of fleas, and he's not larger than average, if the comparison class is the class of animals. But 'larger than average' is perfectly precise. For another, dependence on comparison class does not exhaust vagueness; even once we fix a particular comparison class, say the class of fleas, it can still be a vague matter whether some particular flea is huge (relative to that comparison class). Contextualists like Soames acknowledge this, and propose that inconstant character can be used to account for the remaining vagueness as well.

On their view, there are no vague contents; the vagueness is in the character only. The inconstant character determines a range of possible contents, depending on the context of use. (Some aspects of the context of use might determine a comparison class, but let's suppose that the comparison class is fixed, so as to focus on vagueness in particular.) Take the example 'huge', and suppose our comparison class is fixed to the class of animals. In some contexts of use, this might determine a content *bigger than 1000kg*; in other contexts of use, a content *longer than 20m*. But on every content determined by 'huge' with respect to this comparison class, a full-grown blue whale counts as huge, and a wasp does not. Clear cases and countercases of 'huge' are things that have or lack the content in question *no matter which content it is*. Borderline cases, on the other hand, have some of the contents and lack others.

This is entirely neutral about just which aspects of a context of use matter for the content determined by vague characters. Of course, individual theorists need not remain neutral in this way, and may well disagree with each other. I won't get into those differences here; an inconstant-character theorist simply takes *some* aspect of context of use to matter.²

1.2 Inconstant contents

I'll say that a content C is *inconstant* when there are circumstances of evaluation c_1, c_2 such that $C(c_1) \neq C(c_2)$. Most contents are typically taken to be inconstant. Any content that has different extensions at different possible worlds, for example, is inconstant. If we take at least some time-relativity to occur in the step from content to extension (as eg [9] does), then there will be contents that have different extensions at different times. These, too, will be inconstant. Even precise contents, then, are often taken to be inconstant.

But worlds and times are not the only factors alleged to matter for inconstant contents. For example, [10] takes instances of so-called 'faultless disagreement' to involve inconstant content. On their view, when Jack and Jill disagree about

²In the formal framework of [9], contexts of use are taken to determine at least an agent, time, position, and world, but factors besides these might be relevant as well.

whether avocado is tasty (say Jack thinks it is and Jill thinks it isn't), they are not talking past each other because there is a single content involved. That content (call it AT) maps world/time/judge triples to truth values. Jack and Jill are faultless because from Jack's point of view, the relevant triple is $w/t/Jack$ (for some w, t), where $AT(w, t, Jack)$ is true, and from Jill's point of view the relevant triple is $w/t/Jill$ (for the same w, t), where $AT(w, t, Jill)$ is false. On this view, the content AT is inconstant, and indeed its inconstancy is sensitive to the judge parameter contributed by circumstances of evaluation. (Other inconstant contents, like 'is over a meter in height' might have different extensions at different worlds or times, but will not depend on the judge parameter.)

A theory of vagueness based on inconstant contents (as offered in eg [5]) allows for vague contents; once character and context have determined a content, there's still more to the vagueness story. These theories have a similar structure to inconstant-character theories, but deployed between content and extension, rather than between character and content. Thus, for an inconstant-content theorist, clear cases of a vague content C are in all of its extensions, clear counterexamples of C are in none of its extensions, and borderline cases of C are in some extensions and out of others.

Again, this is strictly neutral as to which aspects of a circumstance of evaluation matter. It's also neutral as to what a circumstance of evaluation is, and as to how (or whether) a sentence's context of use affects which circumstances of evaluation are relevant for judging it. Again, individual theorists are not neutral about these issues, but I won't enter into these debates here, except to point out that world- or time-dependence alone don't seem to suffice for vagueness. Because of this, below I will sometimes take inconstant content-views to suppose that vague contents determine distinct extensions *given* a world and time.³ I think this is fair to every extant inconstant-content theorist of vagueness.

1.3 Inconsistent extensions

Let's call a pair of a context of use and a circumstance of evaluation a *total parameter*. Then, given a total parameter, a character determines a single extension. On the Kaplanian picture, there is no room for an extension to be inconstant. An extension does not determine anything else that might vary. However, the picture so far is neutral on just what sort of thing an extension ought to be. Let's stick to predicates for the moment; although their extensions are typically taken to be sets, all that's vital is that they be the sorts of things that objects can be *in* or *out* of. Pa is true in a total parameter iff a is in P 's extension relative to that total parameter; and $\neg Pa$ is true in a total parameter iff a is out of P 's extension relative to that total parameter.

I'll say that an extension is *inconsistent* when some things are both in it and out of it. If we are using a classical metalanguage, we are barred from simultaneously doing all of: 1) taking extensions to be sets, 2) taking 'in' to be set membership, and 3) taking 'out' to be set non-membership. In a paraconsistent set theory (like those described in [12] or [24]), we could do all of 1)–3). I think the most natural exposition of an inconsistent-extension view would adopt a paraconsistent metalanguage, and do just that. But it's not required, and here I'll stick to a classical metalanguage. There's a familiar way

³The formal framework of [9] does not allow for this.

to do this:⁴ an extension is a pair $\langle I, O \rangle$ of (classical) sets. The members of I are in the extension, and the members of O are out of it. An inconsistent extension, on this understanding, is one in which I and O overlap.⁵

Inconsistent extensions, like inconstant characters and contents, have found application outside of vagueness. For example, [12] takes the predicate ‘true’, w/r/t some (perhaps all) total parameters, to have an inconsistent extension. Some sentences are both in and out; they are both true and not true. For example, the sentence ‘this sentence is not true’ is both true and not true, according to Priest. This is, of course, inconsistent.

Vague predicates have also been taken to have inconsistent extensions, in eg [25], [14], and [18]. (None of these papers explicitly relates inconsistency to Kaplan’s framework, but I take the idea of an inconsistent extension to be faithful to what all three authors propose.) For these authors, a vague predicate P , relative to a particular total parameter, determines an inconsistent extension. On this picture, vagueness persists at the level of extensions. Clear cases of a vague extension E are in its extension, clear counterexamples of E are out of its extension, and borderline cases of E are both in and out of its extension.

1.4 Incomplete extensions

This also gives us a way to understand *incomplete extensions*. An extension is incomplete iff there is something neither in it nor out of it. Given our pair-of-sets modeling of extensions, this is easy to capture. Simply take a pair $\langle I, O \rangle$ such that $I \cup O$ does not exhaust the domain.

Incomplete extensions have been taken to play an important role in understanding vagueness by a number of different authors, including [6], [20], [23], &c. I won’t have much to say about them here. As we’ll see, the core phenomenon I’m interested in exploring is one much more easily explained in terms of inconstancy or inconsistency than incompleteness. This doesn’t at all rule out incompleteness’s playing a role in a full theory of vague language.⁶ It just won’t occupy center stage in this paper.

2 A model theory

Definition 1. A model is a tuple $\langle D, C_1, C_2, I \rangle$ such that:

- D is a domain: a set of objects,
- C_1 is a set of contexts of use,
- C_2 is a set of circumstances of evaluation, and
- I is an interpretation function:

⁴See [2] or [13] for more details.

⁵Sometimes I is called an ‘extension’ and O an ‘antiextension’. This is clearly a different use of ‘extension’ than the one at play in the Kaplanian tradition. My ‘extension’ follows the Kaplanian tradition, and so picks out ‘extension’/‘antiextension’ pairs, in the other sense.

⁶Indeed, I think incompleteness should play such a role, and have argued in [3] for a theory of vagueness that integrates inconsistency with incompleteness.

- For every singular term a in the language, $I(a) \in D$,⁷ and
- for every predicate P in the language $I(P)$ is a function from members of C_1 to (functions from members of C_2 to (pairs $\langle E_I, E_O \rangle$ of subsets of D)).⁸

Abbreviations: It will be handy to talk about total parameters in what follows; as above, a total parameter is a pair of a context of use with a circumstance of evaluation. I'll write T_M for the set of total parameters included in a model M . That is, where $M = \langle D, C_1, C_2, I \rangle$, $T_M = C_1 \times C_2$. Also, it'll be handy to think of a predicate P 's extension in a total parameter t : let's use $I(P)(t)$ for the whole extension (a pair), $I_I(P)(t)$ for the set of things *in* the extension (the pair's first member), and $I_O(P)(t)$ for the set of things *out* of it (the pair's second member). If we need to be more explicit, we can separate out the context of use from the circumstance of evaluation: where c_1 is a context of use and c_2 is a circumstance of evaluation, $I(P)(c_1)(c_2)$ is a whole extension, $I_I(P)(c_1)(c_2)$ is what's in it, and $I_O(P)(c_1)(c_2)$ is what's out of it.

We can now recursively define a notion of *satisfaction*. Note that a model does not satisfy a sentence simpliciter. It does so only with respect to a total parameter.

Definition 2. For a model $M = \langle D, C_1, C_2, I \rangle$ and a total parameter $t \in T_M$:

- $M, t \models Pa$ iff $I(a) \in I_I(P)(t)$
- $M, t \models \neg Pa$ iff $I(a) \in I_O(P)(t)$
- $M, t \models A \wedge B$ iff $M, t \models A$ and $M, t \models B$
- $M, t \models \neg(A \wedge B)$ iff $M, t \models \neg A$ or $M, t \models \neg B$
- $M, t \models \forall xA$ iff every x -variant M' of M is such that $M', t \models A$
- $M, t \models \neg\forall xA$ iff some x -variant M' of M is such that $M', t \models \neg A$

An x -variant of a model M is defined in the obvious way; \forall and \exists can be defined similarly to the above or taken to be abbreviations. Definition 2 has the effect of taking compound sentences, in the presence of inconsistent or incomplete extensions, to be governed by the system FDE described in eg [2] or [13]. For arguments that this truth-functional way of treating inconsistency and incompleteness is superior to its subvaluationist and supervaluationist relatives, see eg [7] or [18]; I won't argue the point here.

We now have a formal framework within which to understand the various hypotheses considered above about vague language. The hypotheses considered here can all take our language to have a unique, 'intended' model. (That is, nothing in their treatment of vagueness prohibits such a hypothesis. This is unlike, say, [19]'s 'plurivaluationism', which crucially rejects this hypothesis.) Call this intended model $M_{@} = \langle D_{@}, C_{1@}, C_{2@}, I_{@} \rangle$.

A predicate P in our language has inconstant character iff there are $c, d \in C_{1@}$ such that $I_{@}(P)(c) \neq I_{@}(P)(d)$. A predicate P has inconstant content in a

⁷These models thus don't allow for inconstancy in which thing a singular term picks out. It would be easy to allow for, but since it will play no role here, I keep it simple.

⁸The pairs here are extensions; they tell us what is in (their first member) and what is out (their second). Thus, the functions from C_2 to the pairs are contents, and the functions from C_1 to contents are characters. This clause tells us: I must assign a character to every predicate.

context of use $c \in C_{1@}$ iff there are $e, f \in C_{2@}$ such that $I_{@}(P)(c)(e) \neq I_{@}(P)(c)(f)$. A predicate P has an inconsistent extension in a total parameter $t \in T_{M@}$ iff $I_{@_i}(P)(t) \cap I_{@_o}(P)(t) \neq \emptyset$. And a predicate P has an incomplete extension in a total parameter $t \in T_{M@}$ iff $I_{@_i}(P)(t) \cup I_{@_o}(P)(t) \neq D_{@}$.

Since features of our language are reflected in features of the intended model, we can view hypotheses about our language as hypotheses about the intended model. For example, if a theorist takes our language to be fully consistent, she will think that the intended model never assigns an inconsistent extension to any predicate, in any total parameter.

Call a model *consistent* iff it never assigns an inconsistent extension to any predicate in any total parameter, and *inconsistent* otherwise. Similarly, call a model *complete* iff it never assigns an incomplete extension to any predicate in any total parameter, and *incomplete* otherwise. (That is, consistency or completeness of a model requires absolute overall consistency or completeness of its extensions; a single inconsistent or incomplete extension results in an inconsistent or incomplete model.) Finally, call a model *classical* iff it is both consistent and complete.

So far, this model theory serves merely to systematize and clarify the issues at hand. But in §4, it will be used to show that inconsistency and inconstancy are more intimately related than one might at first suspect. First, though, let's get clearer on the contrast between them, by looking at their distinct explanations of some recent empirical work.

3 Experimental results

In this section, I briefly present data from two studies of speakers' responses to vague sentences. The studies are reported more fully in [16] and [1] respectively.

In [16], I report and discuss the results of an experiment conducted at the University of North Carolina at Chapel Hill. In this experiment, participants were shown several pairs of shapes. Each consisted of a circle and a square at some distance from each other. The distances ranged from quite far (as far as could be displayed on the projectors used) to quite close (touching). Each participant was randomly assigned to one of four groups, and a different sentence was assigned to each group. Participants were asked to indicate their level of agreement or disagreement with their assigned sentence as applied to each pair in turn, by choosing a number from 1 to 7, where 1 was labeled 'Disagree' and 7 labeled 'Agree'. The sentences were:

- The circle is near the square and it isn't near the square.
- The circle both is and isn't near the square.
- The circle neither is near the square nor isn't near the square.
- The circle neither is nor isn't near the square.

A majority of participants (76/149) gave responses that fit a hump pattern: the maximum level of agreement occurred for some pair(s) at intermediate distance, with agreement decreasing towards clear cases of 'near' or clear counterexamples of 'near'. Indeed, the maximum responses were significantly higher than the

responses to the extreme cases.⁹ This result held across sentences; there was no significant difference among sentence types in the frequency of responses fitting the hump pattern,¹⁰ or in level of maximum response.¹¹

Moreover, participants did not just tend to agree *more* with these sentences in the middle cases; they tended to agree *simpliciter*. Although different participants gave their maximum responses to different pairs, the mean of these maxima was significantly above the middle response of 4.¹²

In [1], the authors report and discuss a similar experiment. In this experiment, participants were shown a picture of five men of varying heights, from quite tall to quite short, and asked to judge four sentences about each man, for a total of twenty judgments; questions were randomly ordered. For man number n , participants judged the sentences:

- # n is tall
- # n is not tall
- # n is tall and not tall
- # n is neither tall nor not tall

Unlike the previous study, in this study participants were asked not for a degree of agreement, but for a simple categorical judgment. The possible responses for each sentence were: ‘True’, ‘False’, and ‘Can’t tell’.

Again, participants tended to agree more with the ‘both’ and ‘neither’ sentences for men of middling height.¹³ Because of the way Alxatib & Pelletier report their data, it is impossible to tell how many participants agree with ‘both’ or ‘neither’ sentences for some man or other in the middle ground. The best we can do is to tell how many participants agreed for the man in the exact center. This is a less informative measure, since some participants may not have taken the exact center man to be the best borderline case. We see that participants exhibit no clear preference about either the ‘both’ or the ‘neither’ sentence, as applied to the exact center man.¹⁴ However, even this lack of clear preference is a phenomenon to be explained. On many theories of language, there is simply no room for these sentences *ever* to be true.

How can we explain these patterns of judgments? There are a large number of available options, each of which makes its own predictions about what

⁹Maximum responses: mean 5.3, standard deviation 1.9. Extreme cases: mean 2.4, standard deviation 1.7. $t = 19.4$, $\text{dof} = 148$, $p < .001$. Restricted to responses that fit the hump pattern, the results are even more extreme: Maximum responses: mean 5.3, standard deviation 1.4; extreme cases: mean 1.2, standard deviation .5; $t = 25.4$, $\text{dof} = 75$, $p < .001$. The crucial number here is p ; it gives the probability of seeing a difference this extreme in the data if there were really no underlying difference. A probability less than .001 is very small—if there were no underlying difference, the observed data would be extremely unlikely.

¹⁰ $\chi^2(3, N = 149) = 4.3, p = .23$

¹¹ $F(1, 148) = .62, p = .6$; restricted to hump responses, $F(1, 75) = 1.41, p = .25$.

¹² $t = 8.15$, $\text{dof} = 148$, $p < .001$; restricted to hump responses, $t = 8.15$, $\text{dof} = 75$, $p < .001$.

¹³For ‘both’ sentences: 44.7% of participants agreed for the man of median height, as opposed to 14.5% for the shortest and 5.3% for the tallest. For ‘neither’: 53.9% agreed for the man of median height, as opposed to 27.6% for the shortest and 6.6% for the tallest.

¹⁴From 76 participants, Alxatib & Pelletier report 34 ‘True’ responses and 31 ‘False’ responses to the ‘both’ sentence. Measured by a sign test, $p = .80$; the data show no significant preference. For the ‘neither’ sentence, they report 41 ‘True’ responses and 32 ‘False’. Here, $p = .35$; still no significant preference.

would happen in further experiments. The way in which I'll proceed is to consider the options we began with—inconstant character, inconstant content, and inconsistent extensions—explore how they would explain the existing results, and draw out such predictions. I'll argue that inconstant-character explanations have trouble explaining the results in [16], but that inconstant-content and inconsistent-extension explanations can do better. Of course, there are many more possible explanations of the above data than just these three. Here, I take a narrow focus. (For discussion of a number of other options, see [16].)

3.1 Inconstancy explanations

The two varieties of inconstancy explanation have much in common; I'll use our notion of total parameter to ignore the differences between them until §3.2. Consider familiar cases of inconstancy: indexicals like 'I' and 'today'. In different total parameters, these take on different extensions. This can happen even within a single sentence. For example, the following sentence does not have to be contradictory: 'Jim's birthday is today, but it's not today'. After all, midnight may have passed while the sentence was being uttered. The two occurrences of 'today' can pick out different days because they *occur* at different times. The total parameter changes in between the occurrences, and it changes in a way that matters for the extension assigned to 'today'. The same can happen for 'I'. Here, what matters about the total parameter is who is speaking. Consider a sentence begun by one person and finished by another: 'I'm sleepy today, but I'm not'. Again, the sentence need not be contradictory, so long as total parameter shifts during its utterance in an appropriate way. (It's important that the two tokens of 'I' be uttered by different people.)

A theorist who takes vagueness to consist in inconstancy thus has a *prima facie* explanation for the results of the above experiments. According to this explanation, participants who agreed with the claim that a certain man was both tall and not tall, or that a certain pair of shapes was both near and not near, were not actually agreeing to a contradiction, but simply saying that the man is tall relative to one total parameter, but not relative to another. Since we know (from examples like those above) that total-context changes within the utterance of a sentence can affect the extensions determined by the sentence's predicates, this would explain why participants might agree to the sentences in question. Similarly, when participants agreed with the claim that a certain man was neither tall nor not tall, or that a certain pair of shapes was neither near nor not near, total-context shifts within the utterance could explain their response.

Of course, this is not in itself a complete explanation. We need to know more. In particular, we need to know why it is that total parameter shifts in just the right way to secure such agreement, especially since it must shift in different ways for different sentences. That is, for participants to agree to 'both', they must first take the man to be tall and then not tall, while to agree to 'neither', they must first take the man to be not tall and then tall. We also need some story about what the relevant aspect of total parameter is. What is it that changes when total parameter changes? Finally, we would need to know why it is that total parameter only shifts so as to allow such agreement in borderline cases, but not in clear cases.

Here's part of a potential explanation: suppose there is a range of extensions

a vague predicate can determine, and some extensions that are ruled out. For example, perhaps ‘tall’ can determine an extension that draws the line (for adult males, say) at 180cm, or an extension that draws the line at 181cm, but cannot determine an extension that draws the line at 210cm. Someone 210cm tall, then, will count as tall no matter which extension the character determines, but someone 181cm tall will be in some extensions and out of others. This is plausible enough, and is broadly in line with much thinking, contextualist and otherwise, about vagueness. Then we could explain why agreement to ‘both’ and ‘neither’ sentences doesn’t occur as much in extreme cases; in clear cases and clear countercases, the various possible extensions are all agreed.

This would tell us why participants don’t agree to ‘both’ and ‘neither’ sentences in clear cases, but it wouldn’t yet tell us why they *do* agree in borderline cases, just why they *could*. Here as well, it’s worth some care. One plausible thought is that participants will agree with sentences proposed to them as much as possible. That is, we might suppose that they have a bias towards ‘agree’ and ‘true’ responses. This would allow us to predict agreement from the possibility of agreement. But it does not seem to be a thought that can apply in full generality to the above results. Alxatib & Pelletier, recall, also asked their subjects about atomic sentences: ‘Man *n* is tall’ and ‘Man *n* is not tall’. Here, though, they found that participants tended *not* to call the sentences true when man *n* was a borderline case. This casts doubt on the thought that participants simply want to agree as much as possible. If there is indeed an extension available that allows participants to call the sentence true (as there must be in borderline cases, on an inconstancy explanation), then a bias toward agreement would predict agreement with atomic sentences in borderline cases. But this is not what is observed. But for our purposes here, let’s suppose that an explanation for these results in terms of inconstancy can be found. We’ve got other fish to fry.

3.2 Character, content, and occurrences

So much for the similarities. The inconstant-character theorist faces a particular problem that her inconstant-content relatives don’t seem to. The problem is this: our other examples of inconstant character—indexicals and demonstratives—exhibit a particular feature that vague predicates lack. Although an inconstant character can determine two different contexts within a single sentence (as we’ve seen), it can only do this when the word at issue occurs *twice*.

Return to our earlier examples. Although ‘Jim’s birthday is today, but it’s not today’ can be about two different days if its utterance is timed right, ‘Jim’s birthday is today, but it’s not’ is about only *one* day, no matter how its utterance is timed. Although ‘I’m sleepy, but I’m not sleepy’ can be about two different people, if uttered by them in the appropriate way, ‘I’m sleepy and not sleepy’ is about only *one* person, no matter how it is uttered. Demonstratives work the same way. Although ‘Mary’s buying that, unless Joan buys that’ can be about two different things (picked out by the different occurrences of ‘that’, perhaps along with two overt demonstrations), ‘Mary’s buying that, unless Joan does’ can only be about *one* thing, no matter how many overt demonstrations accompany it. The pattern is quite general (see [22]).

An inconstant-character explanation of the vagueness data should thus predict a striking difference between ‘The circle is near the square and it isn’t near

the square’ and ‘The circle both is and isn’t near the square’, as well as between ‘The circle neither is near the square nor isn’t near the square’ and ‘The circle neither is nor isn’t near the square’. After all, the first sentence in each pair has two occurrences of ‘near’; this allows for context shifts to matter, resulting in two distinct contents within each sentence. But the second sentence in each pair should not allow for this, since there is only one occurrence of ‘near’, and as we’ve seen, we can only get one content out of one occurrence of an inconstant character, no matter how context does or does not shift. But this prediction is not borne out. In fact, there is no significant difference in maximum response for the sentences within either pair.¹⁵ Neither is there a significant difference in maximum response overall between the two-‘near’ sentences and the one-‘near’ sentences.¹⁶ An inconstant-character explanation should not predict this, and could account for it only at the cost of supposing that although vague predicates have inconstant characters, they do not behave in their inconstancy like other expressions with inconstant character. That would be ad hoc, to say the least.

Inconstant-content explanations do better when it comes to this fact. The reason is this: it may well be, for all we know, that a single occurrence can determine a single content that is then evaluated with respect to two distinct circumstances of evaluation. We do not, as we did for inconstant characters, have evidence to the contrary. Thus, all inconstant-content theorists need suppose is that this is indeed so, and I see no reason to begrudge them this assumption. Thus, inconstant-content theorists have an available explanation for the above data.

3.3 Inconsistent extensions

Of course, inconsistent-extension explanations also have an available explanation for the data, and a quite different one. On such an account, there is no need for variation at all. The vague predicates, relative to a particular total parameter, determine an inconsistent extension, and the borderline cases are simply both in that extension and out of it. When participants agree to ‘both’ sentences, they are simply reporting this fact. It might at first seem that participants’ agreement to ‘neither’ sentences tells against this hypothesis, but that’s not so. If something is both P and not P , then it is indeed neither P (since it’s not P) nor not P (since it’s not not P).¹⁷ Thus, an inconsistent-extension approach predicts the observed responses to the ‘both’ sentences and the ‘neither’ sentences.¹⁸

So inconsistency views can explain the observed data, as can inconstancy views. Where one sees a single inconsistent extension, the other sees distinct

¹⁵For conjunction: $t = .35$, $\text{dof} = 82$, $p = .73$; for disjunction, $t = 1.2$, $\text{dof} = 63$, $p = .24$.

¹⁶ $t = .55$, $\text{dof} = 147$, $p = .59$

¹⁷This relies on one De Morgan law and a double-negation inference; both are valid in the logic—LP—recommended by the inconsistent-extension authors cited above, as well as the logic—FDE—of the (possibly-incomplete) models in §2.

¹⁸And indeed to the atomics, if the pragmatic hypothesis of [12, p. 291] is taken as part of an inconsistent-extension view. Priest proposes that a Gricean implicature can be generated in some contexts by the assertion of p : if the speaker believed both p and $\neg p$, an assertion of either one would be misleadingly incomplete, so the hearer is entitled to conclude that the speaker does not believe both. This could explain, if participants believe that borderline cases are both P and not P , why they might be reluctant to agree to either atomic claim in isolation.

consistent extensions. As far as the data collected in [16] and [1] shows, either of these explanations can work. Is there any data that could tell between them?

4 An equivalence

This section will suggest that there is not, at least given our current level of abstraction. As we've seen, inconstancy explanations for the above data crucially invoke *multiple* total parameters to do their work. If we suppose that each sentence is judged in a single total parameter, inconstancy explanations lose all grip on the data. This suggests that our earlier model-theoretic notion of satisfaction is too restrictive to understand the inconstancy theorist's predictions; it needs to be liberalized. We need the idea of a sentence being satisfied by a model together with a *set* of total parameters.

Definition 3. For a model $M = \langle D, C_1, C_2, I \rangle$ and a set $\{t_i\} \subseteq T_M$ of total parameters:

- $M, \{t_i\} \models Pa$ iff $I(a) \in I_I(P)(t)$, for some $t \in \{t_i\}$
- $M, \{t_i\} \models \neg Pa$ iff $I(a) \in I_O(P)(t)$, for some $t \in \{t_i\}$
- $M, \{t_i\} \models A \wedge B$ iff $M, \{t_i\} \models A$ and $M, \{t_i\} \models B$
- $M, \{t_i\} \models \neg(A \wedge B)$ iff $M, \{t_i\} \models \neg A$ or $M, \{t_i\} \models \neg B$
- $M, \{t_i\} \models \forall xA$ iff every x -variant M' of M is such that $M', \{t_i\} \models A$
- $M, \{t_i\} \models \neg \forall xA$ iff some x -variant M' of M is such that $M', \{t_i\} \models \neg A$

What is it for a sentence to be satisfied by a *set* of total parameters? It is for there to be some way of deploying these total parameters in the course of the sentence's interpretation that results in the sentence being satisfied.

Fact 1. From Definition 3 we can prove the following by induction on formula construction:¹⁹

- This definition subsumes Definition 2 as a special case: $M, t \models \phi$ iff $M, \{t\} \models \phi$, for any M, t, ϕ , as can be shown by induction on ϕ 's construction.
- Where $\{t_j\} \subseteq \{t_i\}$, $\{\phi : M, \{t_j\} \models \phi\} \subseteq \{\phi : M, \{t_i\} \models \phi\}$.

When experimental participants agree to 'The circle both is and isn't near the square', the inconstancy theorist takes them to do so only relative to a *set* of total parameters. In (at least) one total parameter in the set, the circle is in the (consistent) extension of 'near the square', and in (at least) one other total parameter in the set, the circle is out of the (different, still consistent) extension of 'near the square'. Since a conjunction, relative to a set of total parameters, may have its conjuncts evaluated in different total parameters, the overall sentence can be true.

(An inconstancy theorist need not, of course, accept this claim about conjunction. She might think that conjunctions are only ever evaluated in a single total parameter. Then she would reject Definition 3; but she would at the same

¹⁹Appropriate restrictions should be made throughout: M is a model, $\{t_i\} \subseteq T_M$, where M is the model in question, &c.

time lose any explanation for the observed results. It is crucial to the inconsistency theorist's explanation that the conjunctions evaluated as true by the participants are evaluated as true only because their conjuncts are evaluated in different total parameters.)

Relative to a set of total parameters, then, some sentences that look like contradictions can be true, *even in a classical model*. They only 'look like contradictions'; since the model is fully classical, nothing inconsistent is in play here.

It's fair to ask at this point: just how much of the behavior of inconsistent models can be simulated by inconstant models? Or vice versa: how much of the behavior of inconstant models can be simulated by inconsistent models? The short answer to both questions: *all of it*.²⁰

Fact 2. For any model $M = \langle D, C_1, C_2, I \rangle$, there is a consistent model $M' = \langle D', C'_1, C'_2, I' \rangle$ such that: for any total parameter $t \in T_M$ there is a set $\{t_i\} \subseteq T_{M'}$ of total parameters such that: for all sentences ϕ , $M, t \models \phi$ iff $M', \{t_i\} \models \phi$. (Where M is complete, there is a classical M' that fits the bill.)

Proof. Here's one such M' .²¹

- $D' = D$
- $C'_1 = C_1$
- $C'_2 = \{\langle c, i \rangle : c \in C_2 \text{ and } i \in \{0, 1\}\}$
- $I'(a) = I(a)$, for singular terms a
- Where $t = \langle c_1, \langle c_2, 0 \rangle \rangle$, $I'(P)(t) = \{\{d \in D' : d \in (I_I(P)(c_1)(c_2) - I_O(P)(c_1)(c_2))\}, \{d \in D' : d \in I_O(P)(c_1)(c_2)\}\}$
- Where $t = \langle c_1, \langle c_2, 1 \rangle \rangle$, $I'(P)(t) = \{\{d \in D' : d \in I_I(P)(c_1)(c_2)\}, \{d \in D' : d \in (I_O(P)(c_1)(c_2) - I_I(P)(c_1)(c_2))\}\}$

The basic idea is simple, although it takes a few brackets to spell out: Any extension is split into two consistent 'shadow extensions'. One shadow (tagged with a 0 above) is such that something's out of it iff it's out of the original extension, and in it iff it's in the original extension *and not also* out of the original extension. The other shadow (tagged with a 1) is such that something's in it iff it's in the original extension, and out of it iff it's out of the original extension *and not also* in the original extension. If the original extension is consistent, its shadows are identical to each other, and to the original extension. The machinery only matters where the original extension is inconsistent.

Note that M' is consistent. Any total parameter t is of the form $\langle c_1, \langle c_2, i \rangle \rangle$, where $i \in \{0, 1\}$. Where $i = 0$, something is in $I'(P)(t)$ iff it's in $I_I(P)(c_1)(c_2)$ *and not also* in $I_O(P)(c_1)(c_2)$. Something is out of $I'(P)(t)$ iff it's in $I_O(P)(c_1)(c_2)$. Nothing can meet both these conditions; it'd have to be both in and out of $I_O(P)(c_1)(c_2)$, and I've stipulated that our metalanguage is fully classical; that can't happen. Similar reasoning shows that where $i = 1$, $I'(P)(t)$ must be consistent as well.

²⁰The following facts owe much to remarks in [11].

²¹Here, I use inconstant content to mimic inconsistent extension. Inconstant character, or some mix between the two, would work equally well formally, but inconstant content is more philosophically satisfying for the task at hand, for reasons given in §3.2.

What's more, where $t = \langle c_1, \langle c_2, i \rangle \rangle$, something is either in or out of $I'(P)(t)$ iff it's either in $I_I(P)(c_1)(c_2)$ or in $I_O(P)(c_1)(c_2)$. Since $D = D'$, it quickly follows that M' is complete iff M is. Since we know M' is consistent, we know M' is classical iff M is complete.

The last part of the proof speaks to the equivalence between the two models, and works by induction on the construction of ϕ . For any total parameter $t = \langle c_1, c_2 \rangle \in T_M$, let $t_i = \langle c_1, \langle c_2, i \rangle \rangle$, where $i \in \{0, 1\}$; so $\{t_i\} = \{\langle c_1, \langle c_2, i \rangle \rangle : i \in \{0, 1\}\}$. Clearly $\{t_i\} \subseteq T_{M'}$.

Base cases:

$M, t \models Pa$ iff $M', \{t_i\} \models Pa$: $M, t \models Pa$ iff $I(a) \in I_I(P)(t)$ iff $I'(a) \in I'_I(P)(t_1)$ iff $M', \{t_i\} \models Pa$.

(For the last step RTL direction: note that $M', \{t_i\} \models Pa$ iff $M', t_i \models Pa$ for some $t_i \in \{t_i\}$. But we know there are only two such t_i : t_1 and t_0 . Reflection on the definition of I' reveals that $I'_I(P)(t_0) \subseteq I'_I(P)(t_1)$; so if $M', t_0 \models Pa$, then $M', t_1 \models Pa$; so if $M', \{t_i\} \models Pa$, then $M', t_1 \models Pa$, and so $I'(a) \in I'_I(P)(t_1)$.)

$M, t \models \neg Pa$ iff $M', \{t_i\} \models \neg Pa$: $M, t \models \neg Pa$ iff $I(a) \in I_O(P)(t)$ iff $I'(a) \in I'_O(P)(t_0)$ iff $M', \{t_i\} \models \neg Pa$.

(The last step RTL direction is justified in the same way as the previous parenthetical indicates, mutatis mutandis.)

The inductive cases follow immediately from Definitions 2 and 3, given the inductive hypothesis. □

Fact 3. For any consistent model M , there is a model M' such that: for any set of total parameters $\{t_i\} \subseteq T_M$, there is a total parameter $t \in T_{M'}$ such that: for all sentences ϕ , $M, \{t_i\} \models \phi$ iff $M', t \models \phi$. (Where M is classical, there is a complete M' that fits the bill.)

Proof. Here's one such M' :

- $D' = D$
- It doesn't matter what C'_1 is (so long as it's nonempty)
- $C'_2 = \emptyset(T_M)$
- $I'(a) = I(a)$, for singular terms a
- $I'(P)(c_1)(c_2) = \langle \{d \in D : d \in I_I(P)(t) \text{ for some } t \in c_2\}, \{d \in D : d \in I_O(P)(t) \text{ for some } t \in c_2\} \rangle$

Here, we take each set of total parameters in M to a single circumstance of evaluation in M' . Contexts of utterance in M' simply don't matter; we let circumstances of evaluation do all the work. (Note that c_1 doesn't appear on the right hand of the '=' in the definition of $I'(P)(c_1)(c_2)$.)²² M assigns to P a set E of extensions in a set of total parameters. That set of total parameters determines a single circumstance of evaluation in M' , and in M' P gets as its extension, in any total parameter involving this circumstance of evaluation, a blurring of all the extensions in E . Something is in the blurred extension iff it's

²²As before, this is formally arbitrary. Here, it's philosophically arbitrary as well. But it does the job.

in any member of E , and out of the blurred extension iff it's out of any member of E . (In fact, this technique doesn't rely on M being consistent; it works just as well for arbitrary M . We can show that M' is complete iff M is, and that the stated equivalence holds.)

The equivalence between M and M' is shown as in the proof of Fact 2, by induction on formula construction. Let t be some total parameter in M' whose second member is $\{t_i\}$. Then $M, \{t_i\} \models \phi$ iff $M', t \models \phi$.

Base cases:

$M, \{t_i\} \models Pa$ **iff** $M', t \models Pa$: $M', t \models Pa$ iff $I'(a) \in I'_I(P)(t)$ iff $I(a) \in I_I(P)(t')$ for some $t' \in \{t_i\}$ iff $M, \{t_i\} \models Pa$.

$M, \{t_i\} \models \neg Pa$ **iff** $M', t \models \neg Pa$: $M', t \models \neg Pa$ iff $I'(a) \in I'_O(P)(t)$ iff $I(a) \in I_O(P)(t')$ for some $t' \in \{t_i\}$ iff $M, \{t_i\} \models \neg Pa$.

As before, the remainder of the induction follows from Definitions 2 and 3. \square

Facts 2 and 3 show us that inconstancy can stand in for inconsistency, and vice versa. Whenever an inconsistent-extension theorist supposes that an assertion of a seemingly-contradictory sentence happens in a single total parameter, an inconstancy theorist can claim that the assertion happens across multiple contexts. Fact 2 shows that this strategy will always work. Similarly, when an inconstancy theorist supposes that an assertion happens across multiple contexts, an inconsistency theorist can claim that the assertion happens within a single context. Fact 3 shows that *this* strategy will always work.

For all that, though, inconstancy and inconsistency are quite different hypotheses. In the next section, I consider some ways we might try to tell between them, despite the equivalence demonstrated here.

5 How to discriminate

We've been supposing throughout this paper that inconstancy and inconsistency are empirical hypotheses. Typically, empirical hypotheses are judged by their consequences, but we've just seen that differences in predictions will be hard to come by in this case. It will not be impossible, though; in §5.2 I'll discuss ways to find empirical differences between inconstancy theories and inconsistency theories.

But first, it might be thought that there's an easier way to settle the issue. After all, if we require consistency of our theories, then an inconsistent-extension approach is simply a non-starter; so even if it were predictively equivalent to an inconstancy theory, we ought to prefer the inconstancy theory, simply because it is not inconsistent.

5.1 Consistency as a theoretical requirement

To stop there would be to flatly beg the question against certain inconsistent-extension views, of course. Some might decide, with [11], that begging the question is simply the thing to do when faced with inconsistency. That would be a mistake in this case; it would offer us no reason to accept an inconstancy view over an inconsistency view. Rather, it would be to concede that no such reason is offerable.

Of course, rather than simply begging the question, one might attempt to offer such reason. Maybe there are compelling arguments against inconsistent theories in general, or inconsistent theories as they apply to vague language. These are more general issues than there is room to consider here. I do not believe, however, that any such arguments are convincing. For a thorough consideration of the issue, see [12] (for the general case) or [8] (for the application to vagueness).

But even if we suppose that consistency is absolutely required of any theory we offer, this still does not rule out inconsistent-extension theories *per se*. The inconsistent-extension theorist can avoid contradicting herself, if she so chooses. She cannot avoid supposing that language users contradict themselves when it comes to borderline cases of vague predicates. But so long as she takes them to be *mistaken*, she can stay totally consistent. There is room here for an error theory of a certain type. In its structure, it would be reminiscent of theories offered in [21] and [4]. They take speakers to be led into error with vague predicates by their linguistic competence. As pointed out in [16], Sorensen and Eklund do not predict that speakers believe contradictions about borderline cases; they predict very different errors. But someone could offer a theory that takes speakers to assign inconsistent extensions to their vague predicates *in error*.

As far as I know, nobody has offered such a theory. Certainly the inconsistent-extension theories cited above—those in [25], [14], and [18]—*are* inconsistent. *They* would be ruled out by an absolute consistency constraint on theories. But theories much like these can remain consistent while explaining the present data in the same way. So consistency as a theoretical virtue isn't going to help us decide between inconstancy and inconsistency theories.

5.2 Different predictions

The real key to the decision is going to have to come from empirical predictions. In the light of Facts 2 and 3, these are going to be tricky to come by, but it's not impossible. There is a way to get inconstancy theories to make different predictions from inconsistency theories: develop rich empirical theories of *total parameter*. Here, I'll give a toy example that shows how this could be done, then point out that my earlier argument against inconstant-character theories fits this mold, and finally turn to an empirical theory of context offered in [15].

First, a silly example. Suppose we all agree (for whatever reason) that the only way for total parameter to change is for the speaker to be in a different room. If the speaker is in the same room in two total parameters, then they are in fact the same total parameter. Now suppose that speakers tend to agree with 'Man x is both tall and not tall' when man x is a borderline case, and that they do so *without moving from one room to another*. Given our supposition, speakers cannot be using two different total parameters in evaluating the sentence; there is only one context involved. Thus, the inconstancy theorist would have no available explanation for the (supposed) data, while the inconsistency theorist would.

It's easy to see how this gets around Fact 2: the proof of Fact 2 relied on playing fast and loose with contexts. If we have some understanding of what a context is and (more importantly) when it must be constant, the construction

in the proof of Fact 2 can't get off the ground. This silly example is enough to break the proof.

In fact, we've already seen a less silly example of this same phenomenon: I argued in §3.2 that an inconstant-character theory could not explain the experimental data, since we should take a single occurrence of a vague predicate to be evaluated in at most one context of use, and that's where an inconstant-character theory has to allow for variation to explain the data. Again, this is an example of relating an empirical phenomenon (one occurrence of a vague predicate) to a constraint on total parameters (no shifting the context-of-use part). As we saw, this allows us to get some empirical bite on our theories.

Finally, I turn to an example in the literature of a possible empirical constraint on total-context-shifting. In [15], two types of context are distinguished: *internal* and *external*. External context involves the full setting in which a judgment takes place: its location and time, along with other factors relevant for the judgment in question. For example, if the judgment in question is a color-judgment, lighting and background conditions are vitally important. Internal context, on the other hand, is a matter of the judge's psychology. Raffman gives an example of a subject being marched from red to orange along a sorites sequence, and imagines homunculi battling in the judge's head: one for a judgment of red, one for a judgment of orange. (Nothing hangs on the homunculus-talk being taken literally.) At first, she says, the red homunculus is the clear winner, but as the colors being judged move towards orange, the competition gets closer and closer. Eventually, the orange homunculus wins the day. This shift—from the red homunculus dominating to the orange homunculus dominating—is a shift in internal context. This shift will affect the judge's judgments, and since (according to Raffman) vague predicates are judgment-dependent, the effect on judgments will amount to an effect on the extensions of 'red' and 'orange'.

We don't need to worry about which aspects of Raffman's contexts (internal or external) correspond to which aspects of our total parameter (context of use or circumstance of evaluation). We can simply suppose that internal and external context together are enough for total parameter.

If this were all we had, it would still not be enough to derive any experimental predictions. We still need to know more about the conditions under which this total parameter changes, and the effects of such a change. Raffman provides us with just what we're after. She predicts that internal context shifts are *sticky*. That is, she thinks it is easier for the currently dominant homunculus to stay on top than it is for another homunculus to depose it. Now empirical predictions are forthcoming:

Once [an automatic] car has shifted to a new gear, it will continue to use that gear as long as possible, even if it slows to a speed previously handled by a lower gear. For example, if the car has shifted from second to third gear at 30 mph, it will remain in third even if it slows to 25 mph, a speed previously handled by second. (Shifting gears is hard work.) Analogously, once the competent speaker has shifted from 'red' to 'orange', if asked to retrace his steps down the series he will now call some patches 'orange' that he formerly called 'red'. The "winning" homunculus always strives to maintain her control as long as she can [15, p. 179].

If total parameters are sticky in this way, it's hard to see how to use a shift in total parameter to explain the data in §3. The inconstancy explanation crucially turned on rapidly varying total parameters; supposing stickiness is in play blocks that rapid variation.

Of course, a different empirical theory of context could well yield a different result. This in no way impugns the prospects of an inconstancy theory that can account for the data. But it makes a point clearer: inconstancy theories differ from inconsistency theories only insofar as they differ from each other. There is always an inconstancy theory to match any inconsistency theory, and vice versa; but the matching theories will differ about how total parameter works. If we have good reason to suppose that total parameter is constrained in some way or other, we can get at real differences between inconstancy and inconsistency. If we do not, we cannot.

6 Conclusion

In this paper, I've argued that inconstancy and inconsistency are closer relatives than one might think at first glance. To show this, I've pointed out experimental data that they both seem well-suited to explain, and proved a formal result showing that either phenomenon can simulate the other. But the simulation, as we've just seen, can happen only if we have no constraints on the notion of context. So the lesson here is for inconstancy theorists and inconsistency theorists alike: *get clear on what context is*. The difference between inconstancy and inconsistency just doesn't matter otherwise.

References

- [1] Sam Alxatib and Jeffrey Pelletier. The psychology of vagueness: Borderline cases and contradictions. *Mind and Language*, 201x. Forthcoming.
- [2] Jc Beall and Bas C. van Fraassen. *Possibilities and Paradox: An Introduction to Modal and Many-valued Logic*. Oxford University Press, Oxford, 2003.
- [3] Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 201x. Forthcoming.
- [4] Matti Eklund. What vagueness consists in. *Philosophical Studies*, 125(1):27–60, 2005.
- [5] Delia Graff Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28(1):45–81, 2000. (Originally published under the name "Delia Graff").
- [6] Kit Fine. Vagueness, truth, and logic. *Synthese*, 30:265–300, 1975.
- [7] Dominic Hyde. The prospects of a paraconsistent approach to vagueness. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds: Contemporary Research on Vagueness*. Oxford University Press, Oxford, 2010.
- [8] Dominic Hyde and Mark Colyvan. Paraconsistent vagueness: Why not? *Australasian Journal of Logic*, 6:107–121, 2008.

- [9] David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard K. Wettstein, editors, *Themes From Kaplan*, pages 481–564. Oxford University Press, Oxford, 1989.
- [10] Peter Lasersohn. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28:643–686, 2005.
- [11] David Lewis. Logic for equivocators. *Noûs*, 16(3):431–441, 1982.
- [12] Graham Priest. *In Contradiction*. Oxford University Press, Oxford, 2006.
- [13] Graham Priest. *An Introduction to Non-Classical Logic: From If to Is*. Cambridge University Press, Cambridge, 2 edition, 2008.
- [14] Graham Priest. Vague inclosures. In Francesco Berto, Edwin Mares, and Koji Tanaka, editors, *Paraconsistent Logic (tentative title)*. 201x. Forthcoming.
- [15] Diana Raffman. Vagueness and context-relativity. *Philosophical Studies*, 81:175–192, 1996.
- [16] David Ripley. Contradictions at the borders. In Rick Nouwen, Robert van Rooij, Hans-Christian Schmitz, and Uli Sauerland, editors, *Vagueness and Communication*. Springer, 201x. Forthcoming.
- [17] David Ripley. Inconstancy and inconsistency. In Petr Cintula, Christian Fermüller, Lluís Godo, and Petr Hajek, editors, *Reasoning Under Vagueness: Logical, Philosophical, and Linguistic Perspectives*. College Publications, 201x. Forthcoming.
- [18] David Ripley. Sorting out the sorites. In Francesco Berto, Edwin Mares, and Koji Tanaka, editors, *Paraconsistent Logic (tentative title)*. 201x. Forthcoming.
- [19] Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, 2008.
- [20] Scott Soames. Replies. *Philosophy and Phenomenological Research*, 65(2):429–452, 2002.
- [21] Roy Sorensen. *Vagueness and Contradiction*. Oxford University Press, Oxford, 2001.
- [22] Jason Stanley. Context, interest-relativity, and the sorites. *Analysis*, 63(4):269–280, 2003.
- [23] Michael Tye. Sorites paradoxes and the semantics of vagueness. *Philosophical Perspectives*, 8:189–206, 1994.
- [24] Zach Weber. Transfinite numbers in paraconsistent set theory. *Review of Symbolic Logic*, 3(1):71–92, 2010.
- [25] Zach Weber. A paraconsistent model of vagueness. *Mind*, 201x. Forthcoming.

David Ripley
 Email: davewripley@gmail.com