



**UNIVERSITY OF LEEDS**

This is a repository copy of *Vagueness as Indecision*.

White Rose Research Online URL for this paper:  
<http://eprints.whiterose.ac.uk/99567/>

Version: Accepted Version

---

**Article:**

Williams, JRG [orcid.org/0000-0003-4831-2954](https://orcid.org/0000-0003-4831-2954) (2016) *Vagueness as Indecision*.  
Proceedings of the Aristotelian Society, 90 (1). pp. 287-309. ISSN 0066-7374

<https://doi.org/10.1093/arisup/akw001>

---

© 2016 The Aristotelian Society. This is a pre-copyedited, author-produced PDF of an article accepted for publication in Proceedings of the Aristotelian Society following peer review. The version of record: 'Williams, JRG (2016) *Vagueness as Indecision*. Proceedings of the Aristotelian Society, 90 (1). pp. 287-309' is available online at: <https://doi.org/10.1093/arisup/akw001>. Uploaded in accordance with the publisher's self-archiving policy.

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

**Vagueness as Indecision**  
**JRG Williams**  
**University of Leeds**

Paint being red is one thing; it is another for a person to *treat paint as red* for some practical purpose. The first is a matter of the paint and its properties; the second involves activity: placing the pot on a particular shelf, fetching it in response to verbal instructions, etc.

What is treated-as-F need not be F. If there is a convention to stop at red-painted road signs, I might treat a sun-faded (now pink) sign as red, by stopping at it. In games of make-believe, we may treat tree stumps as bears.<sup>1</sup> In the law courts, we treat-as-innocent the guilty and innocent alike. In suitable circumstances we are treating-as-F, while being full-well aware that the items in question are not F.

Sometimes the practical issue of whether to treat something as F is where our interest lies, and the question of whether it's *really* F seems secondary or silly. Here's a lovely illustration due to Daniel Elstein. A recipe for meringues contains the instruction to whip egg whites until they just form stiff peaks and then add sugar. I treat the whipped eggs as just-stiff, for the practical purposes of making meringues, when I start adding sugar. There is a salient *standard* for when treating-as-just-stiff is a good thing to do—it's when the resulting mixture will produce excellent meringues. Asking the question whether the eggs are *really* just stiff would be perverse—the right way to address the alethic normative question is to address the practical normative question—is it right to treat these egg whites as just stiff, for culinary purposes?

A bold conjecture: borderline cases of vague predicates are examples of this second type, where alethic interest is subsumed to practical interest. Uncertainty over whether a borderline instance *x* of red/large/tall/good is then to be understood as practical uncertainty over whether to *treat x* as red/large/tall/good. This is the thesis to be explored below, informed by the expressivist/quasi-realist theories of vagueness developed by Daniel Elstein (ms) and John MacFarlane (this volume).

Section 1 introduces a question about our attitudes to borderline cases of vague predicates F. Section 2 explores the actions of *treating* and/or *counting* a thing as F. Section 3 reviews how we might *share* our practical plans to count-as-F and *evaluate* those plans. Section 4 looks at the shapes that the *best* plans to count-as-F may take. Section 5 links these *practical* evaluations to the *cognitive* evaluations of doxastic attitudes to vague predications. Sections 6 and 7 concern puzzles for the approach suggested here. Section 6 explores its treatment of normatively defective or contested terms, and section 7 raises a puzzle about the mechanics of MacFarlane's detailed implementation of the approach in connection to gradable adjectives.

### **1. The cognitive role question.**

Vague properties have borderline cases. A borderline case of redness, for example, is one where it is indeterminate whether the item in question is red. The borderline cases contrast with clear cases, where the item is either determinately red or determinately not red. So much, at least, is common ground.

The question of the cognitive role of indeterminacy is *what is the correct cognitive attitude to have to the proposition that *n* is red, when it is indeterminate whether *n* is red?* If something is determinately red, it is true that it is red, and a belief that it is red is accurate/correct (dually, it'd be inaccurate to *disbelieve* that it is red). If something is determinately not red, it is false that it is red, and a belief that it is red is inaccurate/incorrect (dually, it'd be accurate to *disbelieve* that it is red). Again, this is common ground between competing theories of vagueness. It is however contentious how to extend the story about what attitudes are correct to cover borderline cases—hence our question.

One can approach the cognitive role question by first settling the alethic status of borderline cases. When *n* is smack-bang borderline red, there are those who say that it is neither true nor false, those who say it is both true and false, and those who say that *n is red* is half-true. I'll call these *gappers*, *glutters* and *scalers*

---

<sup>1</sup> The example is drawn from Walton (1990).

respectively. Gappers should say that the correct attitude to the untrue proposition that  $x$  is  $F$  is to utterly disbelieve it (have credence zero). Glutters should say that the correct attitude to have to the true proposition is to accept it fully (have credence one). Scalars should say that the correct attitude to adopt varies with the degree of truth of the proposition (the correct credence is equal to the degree of truth of the proposition).<sup>2</sup>

Just as gappers, glutters and scalars in different ways reject classical constraints on truth value distributions, they equally reject classically-based accounts of rational constraints on belief. For example, a gapper will say that a perfectly accurate rational belief state will have credence zero in  $n$  being red, and credence zero in  $n$  being not red, if  $n$  is a borderline case of redness. But the leading classically-based account of rational requirements would say that it is rationally required that one's credence in  $p$ , and in not- $p$ , total 1.<sup>3</sup>

An approach to the cognitive role question is *semantically driven* iff:

- (i) it endorses a particular answer to the cognitive role question (say: when  $p$  is borderline, it's correct to utterly reject  $p$  and utterly reject not- $p$ ); and
- (ii) it explains *why* that answer is correct, by appeal to truth values (say: that the proposition in question is gappy, being such that neither it nor its negation is true).

My previous work on the cognitive role question presupposed a semantically-driven approach. The expressivist, among other things, offers an alternative.

## 2. Treating and counting

John MacFarlane's indecision account of vagueness (as well as Daniel Elstein's quasi-realist account) tells us to concentrate not on the question of *whether  $n$  is red*, but on *whether to treat  $n$  as red*.

To understand what practical question is being asked here, we need to know what it is to treat something as  $F$  in the first place. The first way of understanding that phrase is parasitic on a prior understanding of  $F$  itself. Consider the case of game-playing—we know what it is to treat a tree stump as a bear, because of our antecedent understanding of what bears are, their salient characteristics, and what behaviours might be appropriate in the presence of one.

This parasitic model could be transferred to the case of vagueness. We would presuppose familiarity with *determinate* cases of redness, and how (in fixed context) we are to behave towards them. To treat any other item (including a borderline case of redness) *as red* is to behave *in those same ways* towards it (in the same context).

The alternative is a model on which we have an independent grip on what it is to treat something as  $F$ . Consider the term “fragile”. There may be a certain suite of behaviours—certain *scripts*—such that  $x$  is treated as fragile iff those scripts are applied to  $x$ . Perhaps this would be to handle  $x$  carefully, to duck to avoid shards if it is hurled at the wall above one's head, etc.

---

<sup>2</sup> Here is how I see the debate. An alethic gapper posits three alethic statuses,  $A$ ,  $B$  and  $C$ , which she goes on to label as truth, falsity and neither-true-nor-false. She has to explain what about these statuses makes them deserve such labels. And I say that their normative role with respect to belief is the key here (in particular, that a proposition should be believed iff it is true, and rejected iff it is untrue). That makes the connections above particularly natural. For discussion see Field (2000, 2003), Priest (2006), Smith (2008, 2010), and my own Williams (2011, 2012a, 2012b, 2014b, 2016). Other answers to the cognitive role question are available! See Caie (2012), Williams (2014a).

<sup>3</sup> The revisionary theory of rational attitudes can make a practical difference. In other work, I've explored this (Williams 2014a, 2014b) through asking about a model in which you must choose to benefit or sacrifice the interests of a future individual  $x$ , where you know that it is indeterminate whether you are  $x$ . To the extent that personal identity determines prudential or moral choiceworthiness (or simply that your intrinsic desires treat self and others differently), these different theories can give quite different recommendations. In order to keep the recommendations of personal-identity-based decision-making in sync with Parfittian care for psychological connectedness between person-stages, I've argued that certain cases (Methuselah-style instances of fading personal identity) require the scalar's treatment, while others (instances of personal fission) require we follow the glutter.

The scripted model could be transferred to the case of vagueness, but only if we spot ourselves the assumption that vague predicates are associated with behavioural scripts that are specifiable without re-using the term in question. To treat something as red would be just to follow the associated script in one's behaviour towards that thing. However, it is extremely difficult to imagine spelling out what these scripts are for any but the most artificial terms (have a go at filling out the "etc." in the above stab at fragility). *Linguistic* behaviour associated with redness is scripted—one applies the term "red" to the thing, fetches it in response to verbal request for something "red", etc. But it's really not clear that there's any *general* non-linguistic script associated with redness or other vague predicates. The lovers of red things will gather red things to them; the haters of red will avoid them. Those with a desire that red things be piled up in front of Buckingham Palace will take the means to that end. I despair of finding a common behavioural pattern herein.

The best chance for the scripted model is to deploy it in a highly restricted way. We would focus on one highly constrained context—something illustrated in the case of *largeness* by MacFarlane's apple-sorting plant. In virtue of the particular institutional setting, there *is* now a behavioural script *for that particular environment* in which treating-as-large consists. There is a particular kind of sorting behaviour, with a rationale in terms of the economic goals of the plant, and treating apples as large in that one peculiar context can be characterized in terms of that behaviour (in addition to the linguistic ones). Consider Elstein's eggs; if we restrict attention to the case of *meringue making*, then there's a particular behavioural script to follow that constitutes treating whipped eggs as just stiff—it's to start adding sugar while continuing beating. Change the context to one where the goal is to create soufflé, and the activity you implement once the egg whites are stiff is quite different. But even if we replace the unqualified "treating as F" with "treating as F in situation c" there'll still be plenty of situations in which there's no specific content (beyond the linguistic) to treating as F.<sup>4</sup>

Whichever way we go, "treating as F" will vary along several dimensions:

**Depth** One can treat something as F in some respects but not others. The more respects in which one treats something as F, the greater the *depth* of treating as F. For example, when sorting your paint onto shelves, you may treat a borderline shade as red, by placing it alongside on the shelf of reds. But you may not treat it as red in your verbal behaviour—describing it using "red" or fetching it in response to the instruction "fetch a red paint". At the limit, something is treated as red to the *fullest* depth when you treat it as you do red things with respect to every aspect of behaviour. (On the parasitic model, this is to carry over all aspects of your behaviour towards paradigm Fs; on the scripted model, this is to implement all the relevant scripts).

**Extent** One can treat something as F *to greater or lesser degrees*. Suppose for example you are prepared to pay a pound for a red apple. If you are prepared to pay 95p for an apple of a certain diameter, then you are not treating it *exactly* as you do red apples, but you're pretty close—this scaled-down behaviour counts as *partially* treating it as red. At the limit, you *wholly* treat something as red (in a given respect) if you treat it *exactly* as you do red things (in that respect). (This is framed in terms that lend themselves to the parasitic understanding of treatment. On the scripted understanding, the scripts themselves have to allow for full or partial implementation.)

**Range** One can treat something as F in some situations and not others. I treat a tree stump as a bear so long as the game continues (to a full extent, extending to linguistic as well as non-linguistic behaviour), but the moment it's over I stop hiding from the "line of sight" of the tree stump. At the limit, I treat something as red with the *fullest* range when I treat it as I do red things with respect to *every* context.

---

<sup>4</sup> Thomas Brouwer raised the following issue: suppose the scripts vary from context to context as just floated, so the characteristic behaviours associated with "treating as F" vary. Then in virtue of what do they all get the label "treating as F"? Why, for example, is sorting behaviour on the basis of size in different plants, which varies considerably *extensionally*, to be regarded as *treating as large*? Something further must be added: perhaps a commonality or family resemblance among the scripts; or the shadow of a parasitic account (the scripts, though they vary, are always applied to *paradigmatic* Fs, and never applied to *paradigmatic* non-Fs); or perhaps some universal script found in all the cases that ties them together (a candidate would be the *linguistic* behaviour of applying the term "F" to the objects that are the target of the script).

If treat something to the maximal depth in a given context *c*, and with extent of degree *d*, we will say we count it as red to degree *d* in *c*.<sup>5</sup> If we treat something as *F* to the maximal extent and depth in a given context *c*, we'll say that it is *counted as F in c*. If we treat it as *F* to the maximal extent and depth and range, we'll say that it is *counted as F*. The terminology of “counting” thus marks off those cases where the *depth* is full, and so one is treating the item as red in *all* of one's linguistic and non-linguistic behaviour.

### 3. Sharing and evaluating plans.

Now we have some grip on what it might be to *treat* (or *count*) something as red, we can understand what it would be to *plan* to count something as red, in this or that circumstances. We can also understand the notion of a *hyperplan* (a complete contingency plan giving resolutions to every conceivable choice-point) and what it is for such a plan to exhibit a pattern whereby a thing is, for example, counted as red iff it has a certain size.<sup>6</sup>

Before we turn to the question of *evaluating* such counting plans or hyperplans, let us recap the work that MacFarlane puts this notion to.

A rough approximation to MacFarlane's semantics is the following: “*x* is large” is true at world *w* and hyperplan *h* iff *x* is of a certain size at *w*, and *h* includes a plan to count anything of that size or bigger, as large.<sup>7</sup>

If I utter “apple *A* is large”, then (according to the semantics) this will convey to my audience in the first instance that I plan to count anything of *A*'s size or larger as large. Depending on what the common ground between us is, this may have various effects in practice. If we are presupposing that apple *A* is either a 2cm tiddler or a 20cm giant, and it's already taken for granted that I'd count the former as not large and the latter as large, then the net effect of my utterance is to convey factual information to do with the size of the apple—eliminating the situations in which apple *A* is the tiddler. If the dimensions of apple *A* are already common knowledge—that it is 12cm in diameter, say—then I do not assert *factual* information concerning the apple. My audience will, however, learn something of my plans to count things as large. In the first case, points are eliminated from the conversational scoreboard on the basis of their factual component. In the second case, points are eliminated on the basis of their planning component.<sup>8</sup>

MacFarlane's proposed semantics and pragmatics tells us how agents *share* their plans by using sentences that express those plans. It doesn't tell us anything about how to *evaluate* those plans—but prima facie the full range of normative evaluations is applicable to practical plans to treat or count things as *F*. Lots of the examples that feature in the literature on vague predicates involve low-stakes predications. It doesn't matter all that much in the usual course of events whether a borderline case is or isn't counted as bald, or red, or large. But vagueness is ubiquitous, and so there are plenty of cases about which we care deeply which admit of borderline cases. *Being a woman* surely admits of borderline cases. It's certainly possible for you to adopt a plan to count a certain contested class of persons as *women* or as *not women* and convey this to others in

---

<sup>5</sup> In principle the extent to which you treat something as *F* in one degree may differ from the extent to which you treat it as *F* in another degree (maybe when playing a game, you treat a tree-stump as a bear wholly for linguistic purposes, but only partially for purposes of avoidance---you are a bit more laissez-faire in your behaviour towards the stump than you would be towards a real bear.

<sup>6</sup> For the notion of a hyperplan, and the model for MacFarlane's use of them, see Gibbard (2003).

<sup>7</sup> MacFarlane's favoured clauses here will be rather more complex, reflecting Kennedy's semantics for gradable adjectives—for example, relativity to “comparison class” is handled via a tacit domain restriction on the underlying degree-assignment function. I'll come back to special difficulties that face him because of that at the end, but this rough approximation will do for now.

<sup>8</sup> Compare Plunkett and Sundell (2013) on metalinguistic negotiation. A *locus classicus* of the style of dynamics MacFarlane uses is Stalnaker (1984). MacFarlane credits Yalcin with the extension of the framework to the dynamics of hyperplans. Note that it's a factual (psychological) matter what my plans are. This is something I could, if I wanted, directly convey information about by psychological self-ascription. The expressivist treatment has it that in uttering sentences involving “red”, “large” etc. I am expressing, though not asserting, certain of my plans. My audience learns psychological information by a pathway that superficially appears to have more to do with apples than me.

your linguistic behaviour, but this is a political act and evaluable as such.<sup>9</sup> Personal identity has borderline cases, but whether you count an old man as *the same person* as the young war criminal sixty years may involve imprisoning or freeing an individual. Again, we can adopt and convey plans, but we can also evaluate them. While these particular examples aren't squarely in MacFarlane's target zone (recall, his thesis is restricted to certain kinds of gradable adjectives) there are normatively significant predicates to which his account directly applies: *being good* is a gradable adjective (the positive form of the comparative *better*).<sup>10</sup>

Even in low-stakes cases there are significant evaluations of plans to count things as F as okay or not, though often they will have a hypothetical rather than categorical character: we'll evaluate a plan as good *relative to specified practical interests or goals* in play in the situation. In MacFarlane's apple sorting plant, it's not okay to count an apple of the wrong size as large. Doing so would mess up the system. When making meringues, there are certainly points where it's just *wrong* to count egg whites as just stiff. The standards are here set by whether starting to add sugar at that point would produce tasty results.

Against a backdrop where standing moral, political or more local (institutional, culinary) norms can make certain counting plans *not okay*, it's also significant to register cases where it is permissible to count something is F, or count it as not F, on a whim. In a paint shop, where we just have to find some sensible way to shelve our items, it really is a matter of indifference whether we count something as red or as not red, for shelving purposes. But while it may be permissible in the shop context to shelve orangey-reddish paint either on the orange shelf or the red shelf, it's not permissible in context to shelve it with the greens.

#### 4. Practical gapping, glutting and scaling, and MacFarlane's localism.

What we have seen is that there is a significant question to ask of any counting plan, of whether it's (all things considered) permissible or not. It's in terms of this evaluation that a practical analogue of the cognitive role question arises. Concentrate on a case of an agent with relevantly complete factual information about Patchy—knowing exactly what shade of colour he is. Is it correct, given this information, to count Patchy as red? More generally, holding fixed relevant subvening colour-information, how should we treat Patchy, redness-wise?

There are practical analogues of gappers, glutters and scalers. There will be those who say that when it is borderline whether Patchy is red, we should neither count Patchy as red, nor count Patchy as not-red. There will be those who say that in such circumstances, we should count Patchy as red, and also count Patchy as not-red. And there will be those who say that, depending on the exact shade of colour, we should *partially* count Patchy as red, treating Patchy *as red to a greater extent* the redder Patchy is.

Now, practical scaling is fully specified once we have a grip on what *partially counting as F* amount to, and that notion was provided for in section 2. Practical glutting and gapping require not only the notion of *counting x as F* but also *counting x as not-F*. That raises a new range of questions. For example, if there's a behavioural script for F, might there be another separate script for not-F? Can these two scripts be compatible, so that we can enact both on the same item, thus counting it as both F and not-F? On a parasitic model, similar questions arise, substituting patterns of behaviour toward paradigm Fs/not-Fs for the scripts.<sup>11</sup>

---

<sup>9</sup> For discussion of how we should approach the politically loaded question of whether intersex persons and trans women are or should be counted as women, see e.g. Saul (2007). This paper is just the tip of a very large iceberg of feminist thought and writing on such issues!

<sup>10</sup> A poll of Facebook produced the following examples of politically significant (and so arguably high-stakes!) examples of gradable adjectives: provincial, hard-working, free, deserving, fair, proper, aspirational, capable, safe, threatening, healthy, privileged, happy, poor, posh. Some of these are high-stakes because they pick out kinds or properties that are socially important (e.g. terms for classes, such as posh, or for distributional properties, such as fair). Others acquire significance through a more path-dependent role they acquire in political discourse (hardworking, deserving).

<sup>11</sup> You might think that the parasitic model at least answers one question. Paradigm Fs are presumably not labelled as "not F", whereas paradigm not-Fs are, and so if we're to treat a borderline case like a paradigm F, then you might think we can't deploy all the behaviours appropriate to paradigm Fs and to paradigm not-Fs. However, note that paradigm Fs will be labelled, inter alia, as "paradigm Fs", "clear Fs", etc., but we would not want to transfer this aspect to borderline cases. So a fully developed parasitic understanding of counting-as-F will involve some filtering of the sorts of behaviours transferred.

If one assumes that to count *x* as not-*F* is (inter alia) *to fail* to count *x* as *F* then practical glutting in particular, will look incoherent. But that assumption is contestable, and I for one regard practical glutting and gapping as live theoretical options.<sup>12</sup>

MacFarlane will say, however, that *none* of glutting, gapping or scaling is right for the case he focuses on. The recommendations of the practical gappers, glutters and scalers do not fit with the pattern of behaviour he takes to be appropriate in his apple-sorting plant. And I think—for that particular case—he is right. I would expect this to be the general rule for low-stakes vague concepts. Insofar as we look to local interests to resolve a generally low-stakes classification issue, a counting plan seems good just in case it helps further some local institutional goals; and the local convention that emerges idiosyncratically in an apple-sorting plant may well be all that bears on whether or not the counting plan adopted is good or not. Such local conventions can, as in MacFarlane's case, count some borderline cases one way, others another. In other plants, a different local practice of counting-as-large might emerge, drawing the local boundaries differently. MacFarlane's practical *localism* is a rival to practical gapping, glutting and scaling, and we should expect it to be pretty widespread.

To say that practical localism is how *one* instance of the practical role question plays out, is not to say that it's how *every* case plays out, however. Consider again the case of personal identity. Imagine a future independent Wales evolving its own local conventions about the punishment of fission-products of criminal citizens. A consensus evolves to punish exactly one fission product--the one caught by the authorities first, ties broken by a coin flip. Independent Scotland evolves a different set of conventions--perhaps here, it is the closest continuer who is punished for the crimes of the pre-fission individual, again with coin flips breaking ties. When questions of freedom or imprisonment are at stake—and so when assessing whether we should indulge in the practical behaviour constitutive of counting Lefty as the same person as Orri—moral considerations have lexical priority over considerations of according with local custom. For this reason, the kind of local consensus that would close the question of whether we should count a patch as red, leaves open the question of whether we should count Lefty as the same person as Orri. We are able to step back and evaluate whether it is *morally right* to imprison Lefty, for example. How the open question here resolves itself is a matter for substantive first-order moral discussion, but one can envisage defenders of any of these three positions: (i) it's wrong to imprison either Lefty or Righty; (ii) it's right to imprison both; (iii) it's right to punish each so that the sum total of punishment is the same. It's hard to take seriously *substantive* discussion of the analogous questions for counting as red, counting as large, etc.<sup>13</sup>

It is a striking feature of the philosophical literature on vagueness that so much of it concerns itself with a range of predicates that one would expect to figure in relatively low-stakes situations: bald, heap, red, large, tall, etc. There is nothing *wrong* with this, so long as we are careful not to mistake common features of *low-stakes* profile of vague properties for features of *every* vague property. And so while I'd be unsurprised if practical localism was the right thing to say about the common run of low-stakes cases, this leaves wide open the practical role question for high-stakes but vague properties---personal identity, permissibility itself, racial and gender classifications, etc.

Disputes like this over the practical role of borderline cases of vague properties are perfectly consistent with endorsing MacFarlane-style expressivism. We should distinguish: (1) expressivism as a general framework for theorizing about borderline cases, one that prioritizes the practical role over the cognitive role question

---

<sup>12</sup> A semantically glutty treatment of fission linked to decision-theory is argued for in Williams (2014b) and can form the basis for building a model of practical glutting.

<sup>13</sup> Tax law provides some nice examples of high-stakes disputes of counting-as, even when the adjective doesn't strike us as context-independently high-stakes. For example: whether a second-floor apartment counts as a building matters for the tax liability of a development of an apartment block. Whether a Jaffa Cake counts as a cake or a biscuit matters for the rate at which it is subject to VAT. A lot of money can ride on the judgement-calls made in this area. But they are essentially more like the apple-sorting-plant than the personal identity case. In such cases, interested parties dispute over the best fit with precedent and with tax code as set down by legislators. It may be hard to determine what the best-fitting verdict is, but this is not so different from someone really trying to figure out whether the rules of the apple sorting plant and past practice really require one to count an apple as large. But what we find in the case of personal identity is a *different* set of conventions, independent of local conventions, that seem to normatively *override* such considerations.

and provides a distinctive account of communication using vague language; (2) the *particular* way the practical role question gets answered in the example on which MacFarlane chooses to focus. Criticizing (2), or worrying that it may be unrepresentative, are not thereby attacks on (1).

## 5. Belief, Truth and Quasi-realism.

By taking the expressivist route, it looks at this point that we've eliminated the cognitive role question as originally posed. But (as Elstein emphasizes and advocates) there's a well-trodden quasi-realist route for reintroducing such questions in this general framework.<sup>14</sup>

Here's my favoured way of understanding the point. The quasi-realist about Gs agrees with the realist about Gs that we should describe people as having (genuine) beliefs about Gs, and that whether something is an G is a matter that can be true or false. But she imagines that these are fairly superficial characterizations of what is going on. For although realist and quasi-realist will both agree with the biconditional that a belief that x is G is correct iff it's true that x is G, they disagree on the explanatory order. The realist will think that truth explains correct belief. But the quasi-realist thinks that facts about correct belief tell us what it is for the claim to be true. As the dialectic is supposed to play out, in order to prosecute her strategy, the realist needs a belief-free story about ("realist") truth. And equally in order to prosecute her strategy, the quasi-realist needs a truth-free story about correct belief. But here the expressivist machinery comes into play: the belief that x is G is correct, says the quasi-realist, when the associated pro-attitude is the correct one to have.<sup>15</sup>

The *semantically driven* approach to the cognitive role question is thus characteristically realist, since that approach *builds in* the idea that that it is the *alethic status* of borderline proposition p that *explains* the correctness of the relevant cognitive attitude towards p. This is just a generalization, allowing for a richer range of the cognitive attitudes and alethic statuses, of the realist thought about the explanatory order between correct belief and truth.

The pro-attitude associated with a putative belief that apple A is red is a plan to count apple A as red (I'll use that phrase as shorthand for a complex and probably *mixed* state of *planning to count all those apples of shade S as red, where S is a shade that it's doxastically possible that apple A possesses*). So at a first pass, a quasi-realist will say that a belief that apple A is red will be correct just when it's correct to plan to count it as red.<sup>16</sup>

MacFarlane, I suggested, can be read as endorsing varying *local* correctness norms on counting. If an apple is within a certain apple-sorting plant, where a rule has been institutionalised that something is to count as a large apple iff it is over 15cm in diameter, then (since other things are equal) this settles that the correct plan is to follow that rule. That makes believing an apple to be large iff it is over 15cm in diameter correct. And by the original biconditional, the proposition that an apple is large is true iff it is over 15cm in diameter—so long as it is in that particular situation. Transfer it to a different plant with different counting conventions, and the apple may no longer be large.<sup>17</sup>

---

<sup>14</sup> For quasi-realism, see Blackburn (1993). Gibbard too endorses quasi-realism, and building upon this, MacFarlane himself is comfortable with object-language truth-talk.

<sup>15</sup> I've formulated these explanatory "what it is" claims connecting *correctness* of beliefs on the one hand and *correctness* of counting on the other. Quasi-realists may prefer to start in the first instance with the claims of psychological grounding: that *what it is* to believe that x is red, is to count it as red. From this claim about the metaphysics of psychological/behavioural states, it's plausible that the claim about the metaphysics of correctness that I rely on flows. But it's not obvious that we couldn't, in a fashion more like moral externalism, endorse the normative explanatory identity while denying the psychological principle. I'll set that issue aside for now, however.

<sup>16</sup> I'm ignoring for now the complexity that MacFarlane-style expressivist about vagueness may also be an expressivist about normative claims (in the fashion of Gibbard, whose framework is being built upon here). Perhaps this should complicate the discussion, but it's not *obvious* how it'd interact, and enough is going on that it's better to set it aside pro tem!

<sup>17</sup> You might think that largeness is an extrinsic property on the account just sketched, assuming that the *hyperplans* that I am open to recommend *differential* behaviour in different apple-sorting plants. Such hyperplans would make true the following counterfactuals "if I were to take this apple and shift it into the plant down the road, it would no longer be large". In the terminology to be introduced below, this is a *convergent* hyperplan, but one where the cut-off is sensitive to the plant one is in. A different way of developing the account would have it that while in a single plant, one plans to count apples as large or not *unrestrictedly*, and what happens if you get transferred to another factory is the



We will examine how the rehabilitated semantic ascriptions play out with some of the rival answers to the practical role question flagged earlier. In the case where A is a borderline red apple, the practical gapper will say that planning to count it as red is incorrect, as is planning to count it as not red. By the plan-to-belief link, believing that it is red is incorrect, as is believing that it is not red. Our original biconditional said that it was correct to believe that p iff p is true. It follows that the proposition that the apple is red is *not true*, and the proposition that the apple is not red is *not true*. Thus from practical gaps and the quasi-realist bridge principles, we have derived alethic gaps. By similar reasoning, the quasi-realist practical glitter is an alethic glitter.

(Many quasi-realists make equivalence principles (e.g. it is true that p iff p) central to their account of truth-talk. As has become familiar, such principles together with the assertion of an alethic gap will entail a contradiction. Some conclude from this observation that alethic gaps are impossible. Others conclude from the same observation that the equivalence principles have exceptions. It is the latter course that I am recommending here.<sup>18</sup>)

The alethic scalar will not recognize the ideology of truth *simpliciter*, demanding instead to work with a notion of truth to a degree. Her analogue of the original biconditional is that a credence of d in p is (alethically) correct iff the proposition that p is true to degree d.<sup>19</sup> Just as before, there is a realist reading of this, but also the quasi-realist reading whereby we explain truth to degree d in terms of (alethic) correctness of partial belief. For that, we need to identify a pro-attitude to be associated with certain credences. The practical scalar has a candidate for us: *partially treating x as F*. A credence of d that x is F is (alethically) correct iff it is correct to count x as F to degree d.<sup>20</sup> With those bridge principles in place, we can see that the practical scalar will view as correct certain credences, which in turn project degrees of truth onto the propositions in question. So here too we find a quasi-realist gloss on a nonclassical semantics.

We should allow that the norms on counting plans leave us leeway—different practical plans are equally permissible. That raises a problem for the discussion thus far, insofar as talk of the “correct” planning-states appears to assume that there’s but one permissible (and so required) planning state to adopt. Relaxing this faulty assumption causes trouble. Two candidate plan-belief bridges suggest themselves:

- \* A belief that x is F is alethically required iff planning to count x as F is practically required.
- \* A belief that x is F is alethically permissible iff planning to count x as F is practically permissible.

By themselves these are plausible, but the biconditional linking alethic norms on belief to truth value also needs to be reconstructed in our new terms, and the trouble is that both candidates are problematic:

---

*abandonment* of all those plans and replacement by a different set. That is more like a contextualist story, and would not vindicate the counterfactual (since shifting the world parameter does not vindicate the hyperplan parameter). John MacFarlane tells me that the latter is what he intends, though as far as I can see, nothing else in his account *forces* him to take this line.

<sup>18</sup> See Williamson (1994) for a nice discussion of the relation between the T-scheme and truth value gaps. Can’t we at least *introduce* a truth-operator, or predicate, which satisfies the equivalence schema? And if so, wouldn’t this have a *better* claim to *be* TRUTH than whatever is involved in the assertion of alethic gaps? I won’t take a stance on the first, which in the case of a truth-predicate at least, turns on a successful solution to the semantic paradoxes; but argumentation on the second point is crucial point, and is often distressingly thin. For this gambit, see Williamson (1994, Ch. 5). A classic discussion of languages containing multiple truth predicates, including one that satisfies the equivalence schema and relations and one that does not, is McGee and McLaughlin (1994).

<sup>19</sup> This is *compatible* with the original biconditional. For example, on a crude Lockean view about flat-out belief on which it is identified with having a credence above a certain threshold, the original biconditional will be satisfied so long as we adopt an equally crude view about the relation between flat-out truth and truth to a degree—that flat out truth is truth above a certain threshold.

<sup>20</sup> I’m here assuming that there’s a *single* number that measures the extent to which it’s correct to count x as F. But recall the earlier observation that the extent to which we count something as F might vary from respect to respect. There are various possible ways of getting from a profile of numbers to an overall figure—averaging, taking the minimum or maximum. But which of these recipes is the right way to developing scaling quasi-realism deserves further investigation. Thanks to Thomas Brouwer for raising this point and for discussion of the issue.

\* A belief that x is F is alethically required iff it's true that x is F.

\* A belief that x is F is alethically permissible iff it's true that x is F.

In the situation where counting x as F and counting it as not-F are both practically permissible, the belief that x is F is not required but permissible; and likewise for the belief that x is not-F. The first biconditional then will say that x is F is neither true nor false. And the second biconditional will say that x is F is both true and false. These would be spurious gaps and gluts, making no internal sense. For example, on the latter, we'd be committed to saying that it was permissible to disbelieve p, when p was true. So I don't think we should go this way.

In order to block the derivation, I suggest that the bridge principles we may assert are limited to the following:

\* A belief that x is F is alethically required only if it's true that x is F.

\* A belief that x is F is alethically permissible if it's true that x is F.

But this fails to generate guidance on the alethic status of the crucial cases of interest. What we get from the quasi-realist strategy in situations where there is a sphere of permissible counting plans are *partial* distributions of truth value, distributions that are *incomplete* exactly at those points at which multiple counting plans go different ways.<sup>21</sup>

## 6. Normatively defective and normatively contested terms

The quasi-realist as presented here explains our evaluations of beliefs (that an apple is large, red, or whatever) in terms of all-things-considered practical evaluation of plans (to count the apple as large, red, or whatever). There's a question of whether this equation can be maintained---whether those things that we *should* count-as-F are exactly those things we should *believe* to be F.

Suppose there were *clear* cases x satisfying a gradable adjective F that we definitely *should not* count as F. Then it looks like x would be a counterexample to the biconditionals I attributed to the quasi-realist. For here we would have an F that all things considered should not be counted as an F. So it'd be (alethically) right to believe that it to be an F, but not correct to plan to count it as an F.

I earlier drew a divide between low and high-stakes examples of vague predicates. I don't expect to find instances of the above scheme for low-stakes predicates. In low-stakes cases, the *general and systematic* normative considerations will likely be dominated by our interests in coordinating with others in our community---so if there's a prevailing convention to count x as F, that gives each of us reason to count x as F. And only where there are special local circumstances (as in the apple-sorting plant) would we expect this to be overridable—and I think it'll be hard to use those to generate a case where there is consensus that a clear F shouldn't be counted as F.

But higher-stakes adjectives are a richer source of potential examples, because when there's more at stake in deploying a term than just coordinating with your neighbours, then intuitions about what *actually* falls under "F" may align with the conventions for counting people as "F", whereas considerations of whether it's *right* for you (or anyone) to count people as F turn on what's at stake beyond coordination.

One cluster of candidate instances are *normatively defective terms*: predicates F with non-empty extensions, where it's wrong to count *anyone* as F. Perhaps pejoratives are an instance of this---insofar as I understand what it is to count an individual as nerdy, it seems to be mildly *wrong* (because belittling) to adopt that attitude, but that's not (yet) to deny that there are individuals meeting that description.<sup>22</sup> Thick ethical

<sup>21</sup> There's an interesting analogy between this sort of verdict, and that developed in my (2014a), but I won't explore it further here. Note that the *incompleteness* here is a sort of *silence*. We are not given license to "close off" and *deny* that proposition is true, when it's permissible to believe it, and permissible to disbelieve it. Rather, there's no general *external* license to say that it's true, or that it is false.

<sup>22</sup> Elizabeth Barnes pointed out how many objectionable misogynistic gradable adjectives there are. As applied to women, consider the following: trashy/trashier; easy/easier; skanky/skankier; bitchy/bitchier; loose/looser.

gradable adjectives that embody wrongheaded normative stances are another candidate case, as counting someone as e.g. *lewd* might require acceptance of outdated sexual mores.<sup>23</sup>

A different cluster of candidate cases are *normatively contested* terms. Sally Haslanger (2000, 2005) distinguishes “ameliorative analysis” of a term from conceptual or descriptive projects. The latter ask what subjects had in mind or were tracking when deploying “F” all along; but the former asks what pattern of deploying “F” would best serve our (legitimate) purposes. There’s no reason in general to think that the way we currently count people as Fs is the best---it may be that an alternative counting practice (future counting plans) would be better.

If we take somewhat high-stakes terms---for example, the politically significant class-based adjective *posh*---and ask the *normative* question about who it’s best to count as posh, then it may be that we get some surprising results. For example, someone from a paradigmatically working-class background, living a modest lifestyle, might acquire features---e.g. speaking with received pronunciation, familiarity with high culture---on the basis of which they commonly receive the dimensions of social privilege associated with upper-class people in UK society. I can well imagine a case being made that such a person really *should* be counted as posh for political purposes (whether that case is convincing is first-order normative question).<sup>24</sup> But this would I think be revisionary: by ordinary standards, the background and lifestyle would lead us to count the person as *non-posh*.

How should the quasi-realist about gradable adjectives view these cases? An interesting line on politically contested cases would be to maintain that the *best* extension is the *actual* extension. This would be to *deny* that these provide instances of the problematic schema: Fs which we shouldn’t count as Fs. Though ameliorative analysis can form the basis for a critique of *current practice* of counting things as F, it is not in the business of revising the extension.<sup>25</sup> The (morally, politically, institutionally) *best meanings* for our terms act as “reference magnets”. The extension of this strategy for examples like “nerd” and “lewd” would be to again treat normative verdicts as authoritative---if it’s really right that nobody should be counted as a nerd, then there’s nothing in the extension of “nerd”, and *mutatis mutandis* for “lewd”. The effect would be that the right meaning for a term doesn’t just need to be (morally, politically, institutionally) *best* among available candidates, but also *good enough* so that it’s okay to count the Fs as Fs. If the only candidate meanings for a term are (morally, politically, institutionally) *bad*, we get reference-failure. It’s not obvious that this thesis will prove compatible with our best theory of the semantics and pragmatics of pejoratives and thick evaluative terms, and that should be worrying for a quasi-realist. But it is not obvious what *else* they’re going to say here, and the view that they are forced to seems to me to have independent interest.<sup>26</sup>

---

<sup>23</sup> I’ll just mention two candidate accounts that make them (arguably) normatively but not semantically defective: Williamson (2009)’s treatment of pejoratives as conventional implicatures, and Vayrynen’s (2012) treatment of the normative content of thick adjectives as pragmatic implicatures of some kind. The disentangling of semantic and pragmatic phenomenon is one major challenge for the quasi-realist approach here.

<sup>24</sup> Here is a Haslangerian-style analysis of “posh” that might lead to such results: x is *posh* iff (i) x is regularly and for the most part observed or imagined to have a certain features presumed to be evidence of a certain kind K of upbringing; (ii) x having these features marks x within the context of the background ideology of their society as appropriately occupying certain kinds of privileged social position; and (iii) *along some dimension* S’s social position is privileged, and S’s satisfying (i) and (ii) plays a role in that dimension of privilege. (This requires some further elaboration of the kind K relevant to anchor the analysis to poshness, rather than other privileged social class terms).

<sup>25</sup> The idea of the ameliorative (there “analytic”) project as putting forward new meanings is explicit in Haslanger (2000), though see Haslanger (2007) and Saul (2007) for a reconsideration of the issue. Compare again the discussion of “conceptual ethics” and metalinguistic negotiation in Plunkett and Sundell (2013,2014).

<sup>26</sup> This issue arises in the first place for the quasi-realist due to a kind of purity in their conception of the task. Suppose we started from a semantics for vague predicates on which they were equipped with an extension and an anti-extension---borderline cases of the term being those that fall into neither category. We might think that insofar as a speaker is competent with, and prepared to work with, the term “F” in the first place, they must plan to count all the cases in the extension as F, and count all the cases in the anti-extension as not-F. That leaves the borderline cases up for grabs, and here we might say that there is no general *alethic* norm governing counting-as-F, and other non-alethic factors come in. (compare the normative treatment of truth value gaps in Maudlin 2004). So on this, initially semantically-driven view, you might end up with a quasi-realism but only concerning the borderline cases. Such a mixed view wouldn’t be a pure expressivism, but we’ll see in the next section some reasons that suggest we might have to give up that ambition anyway.

## 7. Convergent and divergent counting plans.

This final section looks in more detail at the mechanics of MacFarlane's approach, and identifies a sticking point. To do this, we need a bit more detail on the table.

On the Kennedy story, the underlying syntactic structure of "large" involves a component that maps objects to a totally ordered set of degrees or largeness. The comparative form "larger than" is true of a pair  $\langle x, y \rangle$  iff  $x$  is assigned a higher degree of largeness than  $y$ . The positive form "large" is true of  $x$  iff  $x$  is higher than a certain point—that point is the *delineation* of largeness.<sup>27</sup>

A non-expressivist semantics will seek to capture (at least one aspect of) the vagueness of "large" by saying something substantive about this delineation component. An epistemicist may position an unknowable fact of the matter concerning which degree the delineation picks out. Fara's shifting-sands position posits systematically evasive delineations, sensitively dependent on what differences are significant in context. A supervaluationist may say that there are a range of admissible delineations, and "N is large" is true iff its true on every admissible delineation. And so on.

Though MacFarlane rejects the idea that an "intended" interpretation is fixed by facts alone, he has an analogue to this assumption—he must assume that the world *together with a hyperplan* determines a delineation. His positive characterization of the way that a hyperplan  $h$  fixes a delineation is the following: "the minimum degree of largeness that  $h$  recommends counting as satisfying the positive form of  $g$ ". Here,  $g$  stands for some function assigning degrees (or largeness, redness, or whatever) to objects, and he tells me that "the positive form of  $g$ " is intended to denote the concept of large, red, etc, when the  $g$ -position is filled with the corresponding function (LargenessOf, RednessOf, etc).<sup>28</sup>

A hyperplan, recall, was "something that rules out or permits ... every course of action in every conceivable circumstance". Here's a perfectly kosher counting plan, adopted as I flip a coin into the air: I plan to count something as large iff it's greater than 15cm in diameter and the coin lands heads, or it's greater than 20cm in diameter and the coin lands tails. I have a plan for what will count as large in each situation in every future situation, but it's very hard to see how to extract a single delineation from this—it looks instead like you're letting the coin flip determine which one of two delineations you are going for.

Call a hyperplan *convergent* with respect to counting-as-F iff there is a minimum degree  $d$  such that in each relevant choice circumstance  $c$ ,  $h$  plans to count  $x$  as large in  $c$  iff its degree of largeness is greater than  $d$ . It's clear for convergent hyperplans that the intended effect of MacFarlane's clause is to return  $d$  as "the minimum degree of largeness that  $h$  plans to count as [large]". The problem with the coin-flip plan is that it is non-convergent. In divergent hyperplans, it's unclear what "the minimum degree of largeness [is] that  $h$  plans to count as large". You *might* suggest that this is the degree of a 15cm diameter item, since this is the smallest that *in some choice situation* is counted as large. If that's the case, then asserting "a 16cm thing is large" will express a plan *compatible* with the divergent hyperplan. But on the contrary, it express an

---

<sup>27</sup> See Kennedy (1999, 2007). Compare the Kennedy-based treatment of vague adjectives of Fara (2001).

<sup>28</sup> I'm not sure how *hyperintentional* entities like concepts will be determined from the function from objects to degrees of largeness, etc, that are input here; and nor am so sure that I understand that I know what it is in arbitrary contexts to count something as satisfying the concept *red* (for example). This last point is where our earlier discussion of "counting as F" enters. But I'll spot MacFarlane these points for the present. Incidentally, it's worth noting that if you wanted an absolute characterization of the extension of "large", or a characterization of when a hyperplan for counting something as F was faithful to the F-ness facts, MacFarlane's metalanguage contains the resources for you to do so: a hyperplan is faithful to the facts iff it recommends counting  $x$  as satisfying the positive form of  $g$  iff  $x$  satisfies the positive form of  $g$ . I'd be tempted to use this to deparameterize the hyperplan index of MacFarlane's semantics to get a classification of claims as true or not (MacFarlane's discussion of the futility of truth-conditional semantics is all very well, but it undermines at best the *interest* of positing truth-conditions for certain kinds of explanation---as far as I can see it does nothing to directly attacks the *existence* of such truth-conditions). It's a special feature of MacFarlane's favoured clause that we can do this. In general, if we had a way for characterizing what it is to "count  $x$  as red", CR $x$ , which doesn't *use* the term red or allow us to construct a synonym (e.g. *satisfies the concept red*), then a hyperplan semantics could say that  $x$  satisfies "red" relative to  $w$ ,  $h$  iff  $h$  recommends CR $x$  at  $w$ ; and we would not, within the restricted vocabulary of the metalanguage, be able to characterize conditions under which  $h$  is faithful to the facts.

*incompatible* plan. Asserting this tells your audience that you plan to count something as large if it is over 15cm, which is not consistent with planning to do so in one contingency and doing something quite different in another! That would be like promising to look after your neighbour's cat, and taking this commitment to be compatible with letting whether kitty gets fed be determined by a coin flip. It is no better to let the hyperplan determine 20cm as the cutoff. For then "a 19cm thing is not large" would express something compatible with the divergent hyperplan—but again and for the same reasons, the hyperplan is ruled out by this.

(You might be tempted to find convergence within the hyperplan by appeal to the *actual* fate of the coin. But this is in effect to make the selected delineation depend, not on the complete *contingency* plan involved in h, but on the plan it recommends for those situations which happen to *actually* arise. These may be extremely sparse! In the actual world, the destruction of the world or the invasion of the colour-eliminating aliens may mean there are no further chances where we choose to count something as red going forward. For example: I adopt a convergent plan to count something as red iff it is above 17cm. The 17cm cutoff is encoded in the *contingency plan* in the resolutions of (what turn out to be) nonactual choice situations. But the resolutions of actual choice situations are too sparse to extract the right delineation.)

This is not an objection to MacFarlane's expressivism *as a project*, nor to the idea that the primary communicative significance of assertions of vague gradable adjectives (at least in the borderline region) is to express one's planning states. What it reflects is the following: for some reason or other, our common ground is stocked with convergent hyperplans for counting-as-F, i.e. those which really *do* determine a delineation. There arises a challenge, not an objection. We need to explain *why* only convergent hyperplans get a look in, and why it'd be *incorrect* to adopt perfectly decent *divergent* plans.

It is hard to see why there should be anything at all wrong with divergent hyperplans *as such*---i.e. considered just as ways of structuring our future behaviour. In a board meeting at MacFarlane's apple-sorting plant, we might resolve to follow the coin-flip plan to determine what local rules to institute.

Such plans do look very peculiar to a quasi-realist, however. Let's suppose that the quasi-realist is a classicist, and so takes it to be either true that a 17cm diameter apple is large, or false that a 17cm diameter apple is large. Our divergent plan has us committing, right now, to count that apple as large (in heads scenarios) and to count it as not large (in tails scenarios). So either it is true that it is large and (in tails scenarios) we count it as not large; or it is false that it is large and (in heads scenarios) we count it as heads. Whatever happens, endorsing the divergent hyperplan commits us to a contingency in which we count it as F when it is not F. We are planning for our behaviour not to be faithful to the facts---at least if the coin falls the wrong way. If our planning state is consistent with a convergent hyperplan, this will never happen, since the delineation determined by that convergent hyperplan makes an apple large whenever our plan recommends we count it as large.

I emphasized earlier the central role of *evaluations* of plans to count-as-F. Call evaluations of counting plans *purely external* when you can reach a total evaluation of counting plans bracketing entirely the question of whether or not something is F, and more generally any considerations of the semantic status of F-claims. Maybe sometimes evaluations are purely external in this way---the case of personal identity may be one. But I suspect that in the case of gradable adjectives, the explanation for why some plans (the divergent ones) are bad will appeal to the reasoning I just attributed to the quasi-realist. The evaluations of counting plans will be a mixture of exogeneous considerations about the norms governing relevant counting behaviour, global or local, and endogenous considerations about what pattern our counting-practice has to produce if it is to fit with a semantics for F of the appropriate kind.

This final moral, I think, illustrates a lacuna in MacFarlane's argument against positing absolute truth-conditions for these vague gradable adjectives. Truth-conditions, and the assumption of a correct delineation, he argues, are not relevant to explaining the conversational dynamics that result from ascriptions of "large" to apples and the like. But there is considerable space between that claim and the conclusion that truth-conditions are theoretically redundant. Truth-conditions may be relevant because they *explain* otherwise

mysterious structural features of our conversations. The story I just ran through purports to use them to explain why the common ground he posits contains only convergent plans.<sup>2930</sup>

## Bibliography

Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.

Caie, Michael. 2012. 'Belief and Indeterminacy'. *Philosophical Review*, 121(1), 1–54.

Elstein, Daniel Y. ms.. Quasi-Epistemicism.

Fara, Delia Graff. 2003. 'Shifting Sands: An Interest-Relative Theory of Vagueness'. *Philosophical topics*, 28(1), 45–81. Originally published under the name 'Delia Graff'.

Field, Hartry H. 2000. 'Indeterminacy, Degree of Belief, and Excluded Middle'. *Nous*, 34, 1–30. Reprinted in Field, *Truth and the Absence of Fact* (Oxford: Oxford University Press, 2001) pp. 278-311.

Field, Hartry H. 2003a. 'No Fact of the Matter'. *Australasian Journal of Philosophy*, 81, 457–480.

Field, Hartry H. 2003b. 'Semantic Paradoxes and the Paradoxes of Vagueness'. Pages 262–311 of: Beall, J. C. (ed.), *Liars and Heaps*. Oxford: Oxford University Press.

Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Haslanger, Sally. 2000. 'Race and Gender: (What) Are They? (What) Do We Want Them to Be?' *Nous*, 34(1), 31–55.

Haslanger, Sally. 2005. 'What Are We Talking About? The Semantics and Politics of Social Kinds'. *Hypatia*, 20(4), 10–26.

Haslanger, Sally. 2007. 'Philosophical Analysis and Social Kinds I'. *Proceedings of the Aristotelian Society Supplementary Volume*, 80(1).

Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.

Kennedy, Chris. 1999. *Projecting the Adjective*. New York: Garland Press.

Kennedy, Chris. 2007. Vagueness and Grammar. *Linguistics and Philosophy* 30(1).

---

<sup>29</sup> For reasons mentioned in a previous footnote, MacFarlane is poorly placed to argue that there are *extra costs* involved in positing absolute truth-conditions: his semantic clauses already contain the resources we need to characterize what it takes for a counting-hyperplan to be faithful to the facts.

<sup>30</sup> Thanks to Elizabeth Barnes, Ross Cameron, Thomas Brouwer, Daniel Elstein and John MacFarlane for discussion of this essay. [The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme \(FP/2007-2013\) / ERC Grant Agreement n. 312938](#).

- Maudlin, Tim. 2004. *Truth and Paradox: Solving the Riddles*. Oxford: Oxford University Press.
- McGee, Vann, & McLaughlin, Brian. 1994. 'Distinctions Without a Difference'. *The Southern Journal of Philosophy*, supp. XXXIII, 203–51.
- Paris, J.B. 2001. 'A Note on the Dutch Book Method'. Pages 301–306 of: *Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications*. Ithaca, NY: Shaker.
- Plunkett, David & Sundell, Tim (2013). 'Disagreement and the Semantics of Normative and Evaluative Terms.' *Philosophers' Imprint* 13(23).
- Plunkett, David & Sundell, Tim (2014). 'Antipositivist Arguments from Legal Thought and Talk: The Metalinguistic Response'. In: Hubbs, Graham & Lind, Douglas, *Pragmatism, Law, and Language*. London: Routledge, pp. 56-75.
- Priest, G. 2006. *In Contradiction: a Study of the Transconsistent*. Oxford: Oxford University Press.
- Smith, Nicholas J. J. 2008. *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.
- Smith, Nicholas J. J. 2010. 'Degree of Belief is Expected Truth Value'. In: Dietz, Richard, & Moruzzi, Sebastiano (eds.), *Cuts and Clouds*. Oxford: Oxford University Press.
- Saul, Jennifer. 2007. Philosophical Analysis and Social Kinds II. *Proceedings of the Aristotelian Society Supplementary Volume*, 80(1), 119–143.
- Saul, Jennifer. 2012. Politically Significant Terms and the Philosophy of Language. Pages 195–217 of: Crasnow, Sharon L., & Superson, Anita M. (eds.), *Out From the Shadows*. Oxford: Oxford University Press.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Vayrynen, Pekka. 2013. *The lewd, the Rude and the Nasty*. Oxford: Oxford University Press.
- Walton, K. 1990. *Mimesis as Make-Believe*. Oxford: Oxford University Press.
- Weatherson, Brian. 2003. 'From Classical to Constructive Probability'. *Notre Dame Journal of Formal Logic*, 44, 111–123.
- Williams, J. Robert G. 2011. 'Degree supervaluational logic'. *Review of Symbolic Logic*, 4(1), 130–149.
- Williams, J. Robert G. 2012a. 'Generalized probabilism: dutch books and accuracy domination'. *Journal of Philosophical Logic*, 41(5), 811–840.
- Williams, J. Robert G. 2012b. 'Gradational accuracy and non-classical semantics'. *Review of Symbolic Logic*, 5(4), 513–537.
- Williams, J. Robert G. 2014a. 'Descision making under indeterminacy'. *Philosophers' Imprint*, 14(4).

Williams, J. Robert G. 2014b. 'Nonclassical minds and indeterminate survival'. *Philosophical Review*.

Williams, J. Robert G. 2016. 'Nonclassical logic and probability'. In: Hajek, Alan, & Hitchcock, Christopher (eds.), *Oxford Handbook of Probability and Philosophy*. Oxford: Oxford University Press.

Williamson, Timothy 2009. "Reference, inference and the semantics of perjoratives" In: *The philosophy of David Kaplan*, J. Almog and P. Leonardi (eds.), Oxford: Oxford University Press.

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.