

HENRY JACKMAN

TEMPORAL EXTERNALISM AND EPISTEMIC THEORIES
OF VAGUENESS*

'Epistemic' theories of vagueness notoriously claim that (despite the appearances to the contrary) all of our vague terms have sharp boundaries; it's just that we can't know what they are. Epistemic theories are typically criticized for failing to explain (1) the source of the ignorance postulated, and (2) how our terms could come to have such precise boundaries. Both of these objections will, however, be shown to rest on certain 'presentist' assumptions about the relation between use and meaning, and once these assumptions are rejected, the possibility of a new sort of 'normative epistemicism' will emerge.

1. EPISTEMIC THEORIES OF VAGUENESS, MEANING AND USE

My friend Robert once had a full head of hair, but he has been gradually losing it, so that now, in his late 30s, he is well on his way to baldness. Because of this, the claim:

(RB) Robert is bald.

was once clearly false, will one day be clearly true, but is now neither clearly true nor clearly false. Robert is currently a 'borderline' case of a bald man. Most popular accounts of vagueness understand borderline cases so that current utterances of (RB) not only *seem* but also *are* neither true nor false. For sentences with vague terms, bivalence need not hold.

In contrast to this popular understanding of vagueness, there is frequently discussed (but rarely endorsed) view according to which bivalence still holds for sentences like (RB). Robert may be neither clearly bald nor clearly not bald, but he still must be one or the other. On such *epistemic* accounts of vagueness, we may be unable to know whether or not a vague expression applies in 'borderline'



cases, but all such cases still either do, or don't, fall within the extension of the expression in question. Vague terms are perfectly determinate, it's just that their precise extensions remain unknown to us.¹

Epistemic accounts of vagueness seem best suited to endorsing the consistent application of full-blooded versions of classical logic to natural language. Since such logics are presupposed by our highly successful mathematical and scientific practices, this connection provides non-trivial support for the epistemic view. Furthermore, the non-classical logics associated with non-epistemic accounts of vagueness not only require us to endorse highly unintuitive inferential norms, but also typically can't account for higher order vagueness.² This essay will not, however, be concerned with the strength of such arguments for epistemicism, but will focus instead on how palatable the resulting view can be made. Most arguments against epistemicism deal not with the arguments offered for it, but focus instead on the apparent problems with the view. This is not surprising, since epistemicism not only strikes many *highly* unintuitive, but also seems to run into serious philosophical difficulties of its own.

First of all, satisfactory explanations of *why* we are ignorant of the extension of our own terms are hard to find among epistemicists.³ In addition to this worry, the most common and compelling complaint against epistemic theories (other than simply that they endorse bivalence in borderline cases) is that they seem to sever the intuitive connection between meaning and use. It is quite natural to think that what we mean by our terms is a function of how we (collectively) use them, but there seems to be nothing in our use of, say, "bald", that would settle whether borderline cases of baldness fall within the term's extension or not. After all, given that we typically can't distinguish between the two cases, it seems completely mysterious how our use of "bald" could conceivably make it the case that someone with, say, 342 hairs with an average length of 5.0001 mm on his head counted as bald while someone with 342 hairs with an average length of 5.0002 mm did not. In the absence of any story about how our use could produce such determinate meanings we have no reason to think that we *do* mean anything precise by our seemingly vague terms.⁴

The claim that our usage determines precise boundaries is incredible enough for familiar terms like “bald” or “heap”, but it is even more unintuitive when we consider that the epistemicist must claim that sharp boundaries are fixed even when we introduce *new* terms into our language. For instance, if one casually coins the term “bink” for a shade of pink often found on houses in Bermuda, epistemicism seems to require that one immediately and unknowingly gave “bink” perfectly precise boundaries with no areas of indefiniteness.⁵ Things get even worse when we consider that terms are occasionally introduced with the *expressed intention* of leaving some cases unsettled. For instance, the Supreme Court could introduce the term “brownrate” to mean “with all deliberate speed” where the lack of precision in the definition is intended by the court to allow rules determining the extension of the term to become progressively more complex as new circumstances come to light. Such incomplete stipulation would seem impossible according to the epistemic view.

It is sometimes suggested that such worries betray a lingering commitment to verificationism or some other brand of anti-realism.⁶ However, the worry seems independent of any demand that we produce some sort of equation between, or reduction from, facts about usage and particular meanings. The issue is that even if we knew *all* of the facts about *all* of our usage, we still can’t begin to imagine *how* or *why* just *one* particular set of meanings should be the *only* one that supervenes upon current use. It’s not just that we don’t know the facts, we can’t even imagine how there could be such facts to be knowable.⁷

The apparent fact that non-epistemic views of vagueness are in tension with very plausible principles of classical logic provides reasons for thinking that it would be *very bad* if our expressions didn’t have determinate extensions, but the fact that it would be very bad if our usage didn’t produce determinate extensions is not a compelling reason for thinking that it does so. As they stand, epistemic views of vagueness seem committed to a type of wishful thinking about the ability of our current usage to produce determinate meanings.⁸ One seems, then, presented with a dilemma. Epistemicism appears to be mandated by our use of classical logic, but endorsing the view requires a type of wishful thinking about how

current usage could produce determinate meanings for terms that manifestly seem to lack them. Fortunately, there is a way to avoid this dilemma, but before seeing how, we need to consider another view that has itself been considered somewhat counterintuitive.

2. TEMPORAL EXTERNALISM

The view in question will be referred to here as “temporal externalism” (or “T-externalism” or “TE”), and according to the temporal externalist, our everyday ascriptions of content legitimately reflect a sensitivity not only to the structure of a speaker’s physical and social environment (including the previous usage of his terms), but also to ‘accidental’ developments in his terms’ usage in the period subsequent to his utterances.⁹


Like other more prominent versions of externalism, TE draws upon fairly well entrenched aspects of our linguistic practice. For example, our use of “gold” in 1650 did not distinguish gold from platinum (both dissolved in *aqua regia*, etc.), but by the time large amounts of platinum were discovered in the 18th century, our chemical knowledge had developed enough for this ‘white gold’ to be distinguished from ‘gold proper’. Consequently, we take utterances from 1650 such as “this ring is gold” to be true only if the ring in question is made of *gold*, and treat the claim as false if made of a platinum ring. However, there are good reasons to think that if large platinum deposits in South Africa had been discovered and mined in the 1650s rather than in the 18th century, the resulting ‘white gold’ would have been treated as ‘gold’, and scientists in the 18th century would have discovered that there were two types of gold (just as they later discovered that Nephrite and Jadeite were two types of jade). Had things developed this way, we would now consider a 1650 utterance of “this ring is gold” as being true if and only if the ring in question were made of *gold or platinum*.¹⁰ Such cases suggest that it is merely a *historical accident* that “gold” has the extension it does, and that when we interpret the past use of both of ourselves and of other speakers, we help ourselves to subsequent specifications that were not determined by the facts available at (or before) the time of utterance. Temporal externalists argue that we should treat such ascriptions as correct.

Since T-externalism gives future behavior a significant role in determining what we presently mean, it should not be surprising that some of the apparatus designed to account for explicitly future directed statements can also help explicate a T-externalistic semantics. In particular, the use of “gold” described earlier can be characterized in terms of the following:¹¹

1. A *moment* (“*m*”) is a spatially complete but instantaneous event: all of nature idealized to a zero temporal thickness.
2. The *causal ordering relation*, $m_1 = m_2$. (This relation is taken to be reflexive, transitive, and anti-symmetrical.)
3. A *history* (“*h*”) is a maximal chain of moments.
4. *Our World*. The set of all moments that are connected to this very moment by means of any zigzag combinations of the causal ordering relation or its converse. Every two moments of Our World have some common historical ancestor. All branching is forward, never backward. Incompatible moments in Our World never have a common upper bound.

A bare moment is not enough for the evaluation of the truth of all of the utterances made at that moment. What one needs is a moment/history pair. The truth-value of many of our utterances, most obviously those about the future, depends upon how history continues after the moment of utterance, and so such histories must be included within the context of evaluation. The truth-value of a sentence is always relative to a history.

When, say, Locke uttered “that ring is gold” in 1650 (while pointing to a platinum ring), a T-externalist could describe the case as follows:



Treating m_1 as the point in 1650 where Locke’s claim that the ring is gold, there are, say, two possible histories, h_1 and h_2 , in which the use of the term develops in different ways. In *our* actual history, h_1 , modern chemistry developed at m_2 , the platinum was discovered at m_4 , and we are discussing Locke’s utterance at m_6 . In the alternative

unactualized history, h_2 , the platinum in South Africa is discovered m_3 while the progress in chemistry is made at m_5 . On h_2 , “gold” denotes all of *gold or platinum*, while on h_1 it denotes *gold*. Locke’s utterance will thus be true in h_2 , and false in h_1 . In such cases, future usage does not make the initial *claim* uttered true or false. Rather, it makes the utterance the claim that it is. What makes the claim false (in our history) is still the fact that the ring isn’t made of gold. Future usage helps to determine an utterance’s truth-value by contributing to its content.

Of course, some terms and sentences have their meanings and truth-values independently of what may happen in the future. Assertions that have their truth-value established at a moment, independently of what happens afterwards, can be described in terms of “settled truth.” Roughly, an assertion is *settled true* at a moment if it is true in *every* history passing through that moment. (For instance, claims like “gold is valuable” or “gold dissolves in *aqua regia*” may be settled true even in 1650.) In much the same way, a term’s extension can be said to be ‘settled’ at a moment if the term has the same extension in every history passing through that moment.¹² Further, even if a term’s extension is not *completely* settled at a given moment, it may still be *largely* settled, with their being settled truths about most items in both its extension and anti-extension.

While it may be the case that, in 1650, it wasn’t *settled* that this history would become the actual one, relative to *this* history, h_1 , the term *always* meant *gold*. (It is now (i.e.: at m_6) settled that the term meant *gold* at m_1 , though it was not settled at m_1 that it meant *gold*.) At m_1 the full context of evaluation was not yet in place. We would look back on Lock’s future tensed utterance of “I will never be discussed by philosophers in the 21st century” and say that it was false, even though, at the time of utterance, it was not settled false. In an analogous fashion, we can look back at his utterances involving “gold” and treat them as referring to *gold*, even though, at the time of utterance, that extension was not yet settled.

3. TEMPORAL EXTERNALISM AND ‘NORMATIVE’ EPISTEMICISM

So what do temporal externalism and epistemicism have to do with each other? Well, epistemic theories of vagueness only need be in tension with the idea that meaning is a function of use if one presupposes a ‘presentist’ conception of use in which *current* usage (or current and past usage) is the only use relevant to what our words currently mean. If one allows that *future* usage can contribute to what we presently mean, then the tension between the two views seems less obvious. Nothing in how we currently use the term “bald” may settle whether or not it applies to borderline cases like Richard, but subsequent developments in the term’s use might make precisely these distinctions. The general worry that we have no idea *how* our usage could *possibly* make precise distinctions disappears, since future usage may include explicit conventions relating to them.

By adopting TE, the epistemicist also becomes considerably more resistant to criticism relating to the introduction of new terms in a language. For instance, one need not assume that sharp boundaries are settled at the moment a new term like “bink” is introduced. Rather, one need only understand the introducer as committed to accepting subsequent refinements to the term’s application. In much the same way, cases of deliberately incomplete stipulation fit comfortably into the T-externalistic framework. Such stipulated terms are introduced into a practice with the intention that they be specified eventually, with such specifications read ‘retroactively’ into the earlier use.¹³

The T-externalistic epistemicist also has a very tidy explanation of *why* we are ignorant in borderline cases. If there are no settled facts to be known pertaining to a question, then there is no mystery as to why we don’t know its answer. Indeed, ignorance is precisely what TE predicts. Further, by not completely collapsing indefiniteness with uncertainty, some questions can be understood as merely *epistemically* uncertain (the facts are out there, but we simply don’t know them), while others are *semantically* or *metaphysically* uncertain (the relevant facts are not settled at this moment).¹⁴

At this point, some might wonder whether this account of vagueness is properly considered an *epistemic* theory at all. Epistemicism

is typically characterized by the claim that bivalence holds for every sentence, and the view outlined above seems ultimately not to entail this. If a particular linguistic practice were to die out while the extensions of its terms were still unsettled, then bivalence would not seem to hold for the sentences in the language associated with that practice. If I claim that Robert is bald today, and the world were to end tomorrow, then it would seem to follow that my claim was neither true nor false, since there would be no future use to give the claim a determinate truth value. To make matters worse, unlike natural kind terms such as “gold”, there is little reason to predict that we will *ever* make paradigmatically vague terms like “bald” or “heap” more precise. The epistemicist who appeals to TE thus seems open to accusations of wishful thinking about how our usage will develop in the future.

However, such objections underestimate just how radical a change the inclusion of future use makes to how we can approach these issues. There is nothing we can do about how we have used our words, and so when we look at past use, we are inclined to take a purely *theoretical* stance towards it and thus see the relation between meaning and use in purely *descriptive* terms. Future use, on the other hand, is still within our control, which allows us to take a more *practical* stance towards it, and thus understand the relation between meaning and use in more *normative* terms. While it seems unlikely that all of the vague terms in our language (which may, indeed, be most of them) *will* be made determinate in the future, it is open to the epistemicist to argue that we are nevertheless *committed* to making them so. While we may have no compelling *evidential* grounds for the bivalence of (RB), we still may be committed to treating it as bivalent.¹⁵ From the point of view of an observer who is standing *outside* of a particular linguistic practice, there may be no reason to think that it will reach a stage of complete determinacy, but for the *participants* in that practice the commitment is still there.¹⁶

The fact that it would be (logically) *very bad* if our terms lacked determinate extensions is thus used to underwrite not the implausible *descriptive* claim that current usage somehow *does* produce determinate meanings, but rather the more defensible *normative* claim that we are (collectively) committed to (someday)

producing such determinate meanings. Since this is a claim about what we *should* do, rather than what we *will* do, it isn't subject to the charge of wishful thinking. Epistemicism does not require us to adopt the pollyannaish position that our current use *somehow* makes our language determinate, rather, it *commits* us to making it more determinate in the future, and *commits* us to such future precisifications when our own current claims are evaluated. Our commitment to our terms having determinate extensions is a *practical* rather than a (merely) *theoretical* one, and the resulting view can thus be described as a *normative* epistemicism.

Some might doubt that there could be such norms related to making our language precise given that a perfectly precise language would probably be *less* useful than a vague one.¹⁷ Rather than being committed to making our terms precise, it might thus seem that we are really committed to keeping them vague. However, even if there is a good sense in which we *should* keep our language vague, it's less clear that we are *committed* to doing so, and even if we were committed to keeping it vague, it doesn't follow from this that we aren't still committed to making it precise. Commitments, especially commitments of different sorts, can often conflict, and our 'logical' commitment to making our concepts precise may be in tension with other more *practical* norms relating to ease of use (just as our commitment to consistency may be in tension with practical norms relating to speed of belief formation).¹⁸ The normative status of a commitment allows that we may never fully live up to it, but our failure to live up to it need not make it any less binding.¹⁹

Further, the commitment to making our language determinate needn't be a commitment to make it determinate any time *soon*. There may be a practical commitment to putting off our logical commitment to precision for as long as possible, but that doesn't, in fact, *conflict* with our logical commitment. Because of this, the manifest fact that we have good reasons for keeping our language vague needn't rule out the possibility that we are committed to its being precise.

But what should one say about practices other than one's own? While the normative epistemicist seems to be committed to bivalence in *her own* language, it might seem that she could take a

more ‘third personal’ stance towards the languages of others, and thus consistently deny that bivalence held of them. For instance, imagine a ‘Twin Earth’ just like ours other than that (1) its inhabitants have been using the word “twald” the way we use “bald”, and (2) it is destroyed by a meteor tomorrow, so that its inhabitants never make the extension of “twald” determinate. Should we claim that my twin’s claim “Robert is twald” (hereafter “(RT)”) is neither true nor false? If so, our commitment to epistemicism seems comparatively weak, since it would allow that others could be well placed to truly deny bivalence for the sentences in our language, so long as they have (as they probably do) good reason to predict that we will never make our terms completely determinate.

However, there is no need to deny bivalence for (RT). Even when we talk about another ‘dead’ language, translating it into ours incorporates it into our practice in a way that our norms of determinacy continue to apply to our talk of it. One can view a Twin-Earthian’s utterance of (RT) from two perspectives. From a ‘participatory’ perspective, we incorporate the ‘dead’ practice into our own, and since ours is still living, we can still attribute unknown determinacy to the utterance.²⁰ From a more purely ‘observational/theoretical’ perspective, of course, the speaker’s practice may be dead, and we may see that there was no determinate class picked out by its use of “twald”. However, from such an observational perspective, one need not (perhaps even should not) treat the speaker’s utterances from the ‘intentional stance’ at all, so the question of the claim’s truth-value (or even the utterance’s status as a *claim*) shouldn’t, strictly speaking, arise at all.

4. CONCLUSION

We can now say a little more about the relation between epistemicism and TE. Epistemicism is a thesis about every term in our language, while temporal externalism, on its own, need not be. Temporal externalism seems more plausible for some terms than for others, and one could count as a temporal externalist even if one only endorsed the view for *some* of the (possibly non-vague) terms in one’s language. Indeed, even if one endorsed TE for vague terms, one still wouldn’t be, thereby, committed to epistemicism. If one is

to be an epistemicist, the reasons for being so will have to be found in writers like Sorensen and Williams, not in anything suggested here. On the other hand, if one is inclined to find the arguments in favor for epistemicism compelling (or at least troubling), then one has a very good reason not only to be a temporal externalist, but also to think that the range of cases to which TE applies is much larger than one (or at least I) might have originally thought. Temporally externalistic intuitions run well for natural kind terms partially because we feel that the extensions of such terms *must* be determinate. However, if the epistemicist's arguments are right, we should have this intuition for *every* term.

TE, of course, will strike many as counterintuitive,²¹ but those sympathetic to epistemicism should consider TE in light of the following remark of Sorensen's:

Change in the web of belief should be made at the most peripheral portion available. Beliefs about how language works are far more peripheral than beliefs about logic . . . Instead of changing logic, we should change our opinions about how language works. (Sorensen, 2001, p. 8)

If we accept Sorensen's point, it remains an open question just *what* commitments about how language works should follow from our commitment to meaning being determinate. Among the possible candidates are the following four:

1. We are committed to our current usage in some mysterious way producing determinate meanings.²²
2. We are committed to our meanings being determinate, though nothing *makes* them so.²³
3. Descriptive TE: Future use can make current meaning determinate, and we have reason to believe that it will.
4. Normative TE: Future use can make current meaning determinate, and we are committed to seeing to it that it does.

The main thrust of this paper is that (4) is the most attractive of these options. Of course, if you don't find the arguments for epistemicism at all compelling, then the choice between these four is not one that you have to make. That said, perhaps the main reason people don't find the arguments for epistemicism compelling relates not to any problem with those arguments themselves, but rather from their taking epistemicism to commit one to (1) or (2). In doing a

sort of conceptual cost/benefit analysis of how one should think of one's language, the epistemicist stresses the considerable costs associated with giving up classical logic. High as these may be, however, adopting (1) or (2) still strikes many as conceptually more expensive. I've hoped to show that these are not costs that the epistemicist need really bear, and that (4) is, at least comparatively, cheap.

In conclusion, epistemicism gives one reason to adopt a T-externalistic semantics by suggesting that we are committed to our language being precise, while TE expands the pool of available usage in a way that neutralizes a number of arguments against epistemicism. However, the appeal to the future doesn't just expand the pool of usage that the epistemicist has to appeal to. It also facilitates a more fundamental change in the way we understand the relation between meaning and use, encouraging us to see the relation in *normative* rather than purely *descriptive* terms.

NOTES

* I'd like to thank Dave Beisecker, Robert Brandom, Garry Ebbs, Steven Gross, Danny Korman, John McFarlane, Tom Osborne, Douglas Paterson, James Sage, Susana Schellenberg, Lionel Shapiro, Paul Teller, Brian Weatherston, James Woodbridge and audience members at the 2002 Mountain Plains Philosophy Conference and 2002 Pacific Division meeting of the APA for comments on earlier versions of this paper.

¹ Such views are defended in Cargile, 1969; Horwich, 1990, 1997; Sorensen, 1988, 2001; Williamson, 1992, 1994.

² For discussions of the first sort of consideration, see Horwich, 2000; McLaughlin, 1997; Williamson, 1992. For discussions of the second sort of considerations, see Cargile, 1969; Williamson, 1992, 1994; Keefe and Smith, 1996b; Sorensen, 1988, 2001; Horwich, 2000.

³ The most notable exception being Williamson (1994), though for criticisms of his proposed explanation, see Gómez-Torrente, 2002; Graff, 2002; Sorensen, 2001; Tye, 1995.

⁴ For a forceful expression of this worry, see Tye, 1997; Schiffer, 1999. Williamson attempts to ease such worries by suggesting that the relations between use and meaning may simply be "unsurveyably chaotic" (Williamson, 1992, p. 275, 1994, p. 209). However, even if this left it possible that our usage *could* produce determinate extensions for our terms, it gives us no reason to think that it *does*.

⁵ For a version of this and the next criticism, see Keefe and Smith, 1996b, p. 21; Tappenden, 1995, p. 196.

⁶ Howich, 1990, 2000; Williamson, 1996; Williamson and Sainsbury, 1997.

⁷ For expressions of this sort of criticism see Horgan, 1994, p. 177, 1997, pp. 231–232; Schiffer, 1999, p. 492.

⁸ This is not the case with Sorensen (2001), who suggests that vagueness gives us reason to deny that use determines meaning. However, in the absence of any explanation of why meaning is determinate, Sorensen's view still seems committed to wishful thinking about the determinacy of our language.

⁹ Versions of this view can be found in Ebbs, 1997, 2000; Jackman, 1996, 1999; Lance and O'Leary-Hawthorn, 1997; Wilson, 2000.

¹⁰ This case is discussed in considerably more detail in Ebbs (2000) and Wilson (2000). Such cases are also discussed in Donnellan, 1983; Lance and O'Leary-Hawthorn, 1997; Sorensen, 1997; Wilson, 1982 (though Wilson (1982), Donnellan, and Sorensen do not endorse the ascriptions involved).

¹¹ This framework draws on Belnap and Green (1994) and is presented in more detail in Jackman (1999).

¹² Or at least in those histories where the meaning of the term doesn't change. In general, whatever story the supervaluationist will have about which 'sharpenings' of a language are not 'admissible' will largely translate into an explanation of which changes in usage should count as changes in meaning for the T-externalist. For discussions of change of meaning within a temporally externalistic framework, see Ebbs, 1997, 2000; Jackman, 1996, 1999; Lance and O'Leary Hawthorn, 1997.

¹³ For a discussion of this aspect of legal practice, see Dworkin, 1977.

¹⁴ The T-externalist also seems well placed to provide an interpretation of our use of 'definitely' that has a number of advantages over its more temporally bound counterparts (for a discussion of this, see, Jackman 2003b).

¹⁵ And thus forbidden from endorsing the sort of reasoning characteristic of the sorites (for a discussion of this, see Jackman, 2003b).

¹⁶ Just as an observer of a chess game may doubt that the rules will always be followed, even though the participants in the game are still, if they are to be playing chess, committed to following all the rules. Logical norms govern the practice of assertion, and a commitment to determinacy can be understood as implicit in such norms. They are part of a system of rules a commitment to which is constitutive of the distinction between assertion and (mere) verbal behavior. For a discussion of this, see Jackman, 2003b.

¹⁷ This is a theme familiar from Wittgenstein (1953), and can also be found in Schiller (1939), Wright (1976), and Edgington (1996). It should, however, be distinguished from the claim that our language needs features like ambiguity and context sensitivity, both of which are compatible with a language being perfectly precise (for a discussion of this see Williamson, 1994).

¹⁸ See Cherniak, 1986.

¹⁹ Any more than ethical claims are less binding simply because they are not (and probably never will be) universally observed.

²⁰ This may relate to the 'normative' account of translation defended in Lance and O'Leary-Hawthorne (1997).

²¹ There are many natural objections to TE, and most of them are discussed (and I hope answered) in Jackman (1999) (see also Jackman, 2003a; Stoneham, 2003).

²² Williamson, 1992, 1994.

²³ The view seems to be suggested in Ebbs, 2000, 2002; Horwich, 2000; McGee and McLaughlin, 1995; Sorensen, 2001; Williamson, 1996.

REFERENCES

- Belnap, N. and Green, M. (1994): 'Indeterminism and the Thin Red Line', in Tomberlin (ed.), *Philosophical Perspectives, 8: Logic and Language* (pp. 365–388), Atascadero: Ridgeview.
- Cargile, J. (1969): 'The Sorites Paradox', *British Journal for the Philosophy of Science* 20, 193–202 (reprinted in Keefe and Smith, 1996a).
- Cherniak, C. (1986): *Minimal Rationality*, Cambridge: MIT.
- Donnellan, K. (1983): 'Kripke and Putnam on Natural Kind Terms', in Ginet and Schoemaker (eds.), *Knowledge and Mind* (pp. 84–104), New York: Oxford University Press.
- Dworkin, R. (1977): *Taking Rights Seriously*, Cambridge: Harvard University Press.
- Ebbs, G. (1997): *Rule Following and Realism*, Cambridge: Harvard University Press.
- Ebbs, G. (2000): 'The Very Idea of Sameness of Extension Across Time', *American Philosophical Quarterly* (July) 37(3).
- Ebbs, G. (2002): 'Learning from Others', *Nous* 36(4), 525–549.
- Edgington, D. (1996): 'Vagueness by Degrees', in Keefe and Smith (1996a).
- Gómez-Torrente, M. (2002): 'Vagueness and Margin for Error Principles', *Philosophy and Phenomenological Research* (January) LXIV(1).
- Graff, D. (2002): 'An Anti-Epistemicist Consequence of Margin for Error Principles', *Philosophy and Phenomenological Research* (January) LXIV(1).
- Horgan, T. (1994): 'Robust Vagueness and the Forced-March Sorites Paradox', *Philosophical Perspectives* 8, 155–188.
- Horgan, T. (1997): 'Deep Ignorance, Brute Supervenience, and the Problem of the Many', *Philosophical Issues* 8, 229–236.
- Horwich, P. (1990): *Truth*, Oxford: Blackwell.
- Horwich, P. (1997): 'The Nature of Vagueness', *Philosophy and Phenomenological Research* (December), LVII(4).
- Horwich, P. (2000): 'The Sharpness of Vague Terms', *Philosophical Topics* (Spring) 28(1).
- Jackman, H. (1996): *Semantic Norms and Temporal Externalism*. Ph.D. Thesis, University of Pittsburgh (available at www.jackman.org).
- Jackman, H. (1999): 'We Live Forwards but Understand Backwards: Linguistic Practices and Future Behavior', *Pacific Philosophical Quarterly* 80, 157–177.

- Jackman, H. (2003a): 'Temporal Externalism and Our Ordinary Linguistic Practice', colloquium presentation at the Central Division meeting of the American Philosophical Association, May (available at www.jackman.org).
- Jackman, H. (2003b): 'Temporal Externalism and Constitutive Norms', MS York University (available at www.jackman.org).
- Keefe, R. and Smith, P. (1996a): *Vagueness: A Reader*, Cambridge: MIT Press.
- Keefe, R. and Smith, P. (1996b): 'Introduction: Theories of Vagueness', in Keefe and Smith (1996a).
- Lance and O'Leary-Hawthorne (1997): *The Grammar of Meaning*, New York: Cambridge University Press.
- McGee, V. and McLaughlin, B. (1995): 'Distinctions without a Difference', *Southern Journal of Philosophy* 33(supplement).
- McLaughlin, B.P. (1997): 'Supervenience, Vagueness, and Determination', *Philosophical Perspectives*, 11, *Mind, Causation, and World*, 209–230.
- Schiffer, S. (1999): 'The Epistemic Conception of Vagueness', *Philosophical Perspectives*, 13: *Epistemology*.
- Schiller, F.S. (1939): 'How is "Exactness" Possible', in his *Our Human Truths*, New York: Columbia University Press.
- Sorensen, R. (1988): *Blindspots*, Oxford: Clarendon Press.
- Sorensen, R. (1997): 'The Metaphysics of Precision and Scientific Language', *Philosophical Perspectives*, 11, *Mind, Causation, and World*, 349–374.
- Sorensen, R. (2001): *Vagueness and Contradiction*, New York: Oxford University Press.
- Stoneham, T. (2003): 'Temporal Externalism', *Philosophical Papers* (March) 32(1), 97–107.
- Tappenden, J. (1995): 'Some Remarks on Vagueness and a Dynamic Conception of Language', *Southern Journal of Philosophy* 33(supplement), 193–201.
- Tye, M. (1995): 'Vagueness: Welcome to the Quicksand', *Southern Journal of Philosophy* 33(supplement), 193–201.
- Tye, M. (1997): 'On the Epistemic View of Vagueness', *Philosophical Issues* 8, 247–254.
- Williamson, T. (1992): 'Vagueness and Ignorance', *Proceedings of the Aristotelian Society* 66(supplement), 145–162 (reprinted in Keefe and Smith, 1996a).
- Williamson, T. (1994): *Vagueness*, London: Routledge.
- Williamson, T. (1996): 'What Makes it a Heap', *Erkenntnis* 44, 327–339.
- Williamson, T. and Sainsbury, R.M. (1997): "Sorites", in Hale and Wright (eds.), *A Companion to the Philosophy of Language*, Malden: Blackwell.
- Wilson, G. (2000): 'Satisfaction Through the Ages', in A. Kanamori (ed.), *Proceedings of the 20th World Congress of Philosophy, Volume 6, Analytic Philosophy and Logic*, Bowling Green, Philosophy Documentation Center.
- Wilson, M. (1982): 'Predicate Meets Property', *The Philosophical Review* 91(4), 549–589.

Wittgenstein, L. (1953): *Philosophical Investigations*, 3rd edn., Oxford: Blackwell.

Wright, C. (1976): 'Language-Mastery and the Sortes Paradox', in G. Evans and McDowell (eds.), *Truth and Meaning* (pp. 223–247), Oxford: Clarendon Press (reprinted in Keefe and Smith 1996a).

Department of Philosophy
York University
4700 Keele St.
North York, Ontario M3J 1P3
Canada
E-mail: hjackman@yorku.ca