



Replies to Armstrong, Dennett, the Schroeters, and Stalnaker

Mark Richard

Philosophy Department, Harvard University, Cambridge, MA, USA

ARTICLE HISTORY Received 15 July 2020; Accepted 15 July 2020

Reply to Josh Armstrong

Meaning evolution

The meaning of a word, I say, is constituted by a linguistic community's interpretive common ground (ICG), commonly shared presuppositions about how speakers expect to be understood.¹ ICG often changes without anything deserving to be called a change in meaning occurring, as when a community preserves pretty much every presupposition it makes about marriage save the presupposition that marriages can only occur between differently sexed individuals. Armstrong agrees that this sort of thing should not be understood as 'marry' losing one meaning and gaining another; we agree that an account of meaning must explain how there can be identity of a word's meaning during an interval in which what constitutes that meaning changes. But Armstrong objects that this sort of change vitiates my analogy *meanings are like species*: '[A]gents' lexical presuppositions are the cumulative result of interactions with a wide range of different individuals over the course of many distinct occasions of use. In such cases, the applicability of parent-offspring meaning relations seems to break down' (Armstrong, [forthcoming](#): 7).

I could not agree more with Armstrong's first sentence. Change in linguistic use and presupposition often occurs very quickly; many examples in Richard (2019) are explicitly ones in which semantic change occurs within less than a (human reproductive) generation.² But how much

CONTACT Mark Richard  richard4@fas.harvard.edu  Philosophy Department, Emerson Hall, Harvard University, Cambridge, MA 02138, USA

¹I abridge here and below the definition of interpretive common ground in Richard (2019), 65–68. The Schroeters quote the core of the definition on (S, 2–3).

²This is obviously true of 'skyline' example discussed in sections 6.1 and 6.2, as well as the socially driven examples of 'marry', 'gay', and 'pasta' in Chapter 4.

does the fact, that such ‘lateral transmission’ of meaning characteristics is an important engine of semantic change, vitiate the analogy between species and meanings?

As Armstrong says, I try to abstract from the two dozen or so notions of biological evolution in currency. In spelling out the analogy, I make use of Peter Godfrey-Smith’s notion of an M-population, where such a population is a diachronic collection of individual things that have some degree of autonomy and a significant number of properties in common, with [many of] those properties coming in families of variants and in which there is variation across time in the distribution of variants, variation which is determined by earlier distributions of variants and the interaction of the population with its environment.³ I say that ‘when I say that we should think of meanings as species, what I mean is that we should think of meanings as forming M-populations’ (Richard 2019, 172).⁴ Word meanings, construed as diachronic sets of lexical entries, are obviously M-populations so defined. And change in variant distribution in an M-population does not require that the determination of variation is invariably or even often a result of something like sexual congress.

That said: I do at many points in Richard (2019) suggest that part of what is illuminating about the species analogy is that it highlights the fact that meaning change involves a sort of inheritance. Should I retract *this* suggestion? Well, as I say at several points in Richard (2019), if we are thinking of meanings as things which vary diachronically because of something like reproduction resulting in things with characteristics which are (more) suited to the environment (than alternatives), then the members of the collection which constitutes a word *w*’s meaning are probably best thought of as something along the lines of particular lexical-entries-for-*w*-at-a-time.⁵ These are things – *w*-slices, call them – which interact with one another. The earlier interactions of *w*-slices are causally responsible for characteristics of later *w*-slices. The way this diachronic evolution occurs certainly makes it appropriate to say that the distribution of variants of a word *w*’s characteristics – variant pronunciations, senses, or other features of a *w* – is determined by earlier distributions of variants and the interactions of the population of *w*-slices which constitute a meaning with one another and their environment. It makes it appropriate, that is, to say that a word’s meaning is an M-population. What’s more: This determination of

³Richard (2019), 171. This gloss on the notion of an M-population partially quotes and partially paraphrases (Godfrey-Smith 2009), with mild departures from Godfrey-Smith occasioned by Sober (1980).

⁴See the discussion of the analogy on p. 172, n. 13.

⁵See the discussion on pp. 168–9.

latter characteristics by earlier ones is a matter of the characteristics of the w-slice of a lexical entry L at time $t + k$ being determined by the characteristics of L's w-slice at t along with the characteristics of those w-slices with which L was interacting between t and $t + k$. Normally, the way earlier w-slices determine the characteristics of later ones is that the latter's characteristics are copies or 'blends' of those of the earlier ones. This is enough, I think, for it to be apt to say that meaning change is like the sort of evolution which occurs through transmission of characteristics from parents to children, with some characteristics flourishing, others not flourishing in virtue of the environment.⁶

At the end of the day, the extent to which you're willing to say that talk about meanings is like talk about species is a matter of how thick you think a metaphor needs to be to be illuminating. That characteristics that are transmitted to others are acquired in a 'Lamarckian' way doesn't seem to me to vitiate the analogy. I did not, after all, embrace the biological species concept as the concept we should use to shape the concept of word meaning. And while an extreme Lamarckianism about biological species looks to be incorrect, that doesn't mean that the extreme Lamarckian is not using a species concept. Whether speaking of meaning change as involving a sort of inheritance is useful turns on how we project parentage from the biological case to the linguistic one. If we say that the players in semantic evolution – i.e. the individuals who diachronically constitute meanings – are lexical-entries-at-a-time, there is no bar to saying both (i) linguistic evolution is well-analogized as a matter of 'parents spawning progeny' while (ii) such evolution is by and large effected by repeated interactions between lexical entries.⁷

Languages and reproductive isolation

Armstrong also objects to the claim that meanings are analogous to biological species because such species are normally understood as reproductively isolated populations which either cannot or have not interbred. According to Armstrong

⁶I grant that if we call the lexical entries existing at t which are directly responsible for the properties of a particular lexical entry existing at $t + 1$ its parents, things can sound pretty kinky. If my entry for 'potamology' interacts with no other lexical entries from t to $t + 1$ – which is surely the norm – then the $t + 1$ entry has only the t entry as its parent. And there will be many cases in which an entry at t has many parents – how many would depend in good part on the length of the temporal intervals we use in constructing our models of meanings.

⁷That this is how I understand the analogy should be clear from the early sections of Richard (2020) and (2021), particularly the material on usage based grammar.

... this is a rather significant mistake. There are no two groups of human agents whose members are intrinsically incapable of coordinating their states of mind with one another in acts of interpersonal communication, and, further, communication across the boundaries of linguistic communities has long been a central source of language change.

If a member of *canis familiaris* finds itself on an island populated with California sea lions (*zalophus californianus*), it is not going to be able to interbreed in a way that produces fertile offspring. In contrast, if a group of monolingual speakers of distinct and distantly related natural languages were collected on an island together, there is no barrier preventing those agents from successfully communicating with one another in ways that give rise to a novel natural language. Indeed, this kind of situation has played out many times in the brutal context of colonization in which a group of enslaved individuals who lack a common language manage to create pidgin communication systems which subsequent language learners use to construct the expressively rich natural language systems often called creoles. (Armstrong, [forthcoming](#): 9)

It is of course true that people who don't share a language can coordinate states of mind and create a novel language. But how is this fact a counterexample to claim, that systems of natural language meaning are well-analogized in terms of reproductive isolation? Lexical linguistic reproduction is a matter of speaker interactions in the production and interpretation of a shared lexical item producing changes in at least one of the later stages of the lexical entry for that item, changes which typically involve an increase in the potential for co-ordination between that lexical entry for the item with others. If this is what reproduction is, then meaning systems are reproductively isolated to the degree that utterances of lexical items which manifest meaning in one system do not produce changes in the lexical entries in the other system.

It is not impossible for an utterance in one language to produce changes in interpretive propensities in a speaker who uses another language. I might hear people in Germany speaking about 'Bürgerer', and decide that the way to get a hamburger is to say 'Bürgerer' with a rising inflection. But this is not a matter of my lexical entry for English's 'burger' being modified; I am trying to break in, as it were, to a foreign idiom, creating a (fairly pathetic) beginning of a lexicon for the language into which I am trying to break.

Of course such things as word borrowing and creolization occur all the time. But word borrowing is a matter of speakers of one language adding a novel lexical entry to an existing lexicon, not of interbreeding by coordinating the lexicon of one language with another. Insofar as creolization is simply a more dramatic case of this sort of thing – in which speakers literally create a new semantic system by fashioning bits and pieces of modifications of

existing ones into a free standing whole – it does not tell against the idea that linguistic change within a particular language is pretty much a matter of an isolated system evolving independently of others. The fact that strangers on an island can, starting with no common tongue, co-ordinate by creating one does not show that modeling public languages as more or less independently evolving systems is misleading.

Conceptual engineering

As Armstrong says, I'm somewhat pessimistic about this being something we can easily bring about. My pessimism stems in part from my profession. I have spent a lifetime watching philosophers trying to change the concepts the philosophical community attaches to its philosophical uses of vocabulary. It seems to me that it is very rare for a philosopher to manage to convince more than a small minority of other philosophers that their analysis of knowledge, causation, personhood, meaning, propositions, belief, etc., etc. is correct. But, as I think most enthusiasts for conceptual engineering agree, the sort of thing that Goldman, Hume, Nozick, Grice, Hintikka, and Marcus present as accounts of these properties and relations should be understood as attempts at conceptual engineering: they are as much proposals for conceptual change as they are attempts to accurately transcribe the concepts we have been using all along.

I don't think conceptual engineering is impossible, and I certainly don't think that we are unable to get a small circle of friends to use a term in a new way. Neither do I think that 'the forces of language change are independent of the actions of individual agents'. (Armstrong, [forthcoming](#): 13) But I do think that we have little direct control over whether those in our linguistic community will take up whatever innovations in usage we adopt or suggest. Suggesting that we should use a word in a certain way will often bring those who hear the suggestion to interpret us in accord with the suggestion. But that is not bringing the others to use the words in the way we suggest, which is what the conceptual engineer is trying to pull off. Much more likely to change the way others use a word are factors like the prestige of a speaker or beliefs about whether a way of using a word marks one as a member of a particular social group.



Reply to Daniel Dennett

Dennett and I agree, I think, on what is right and what is somewhat off the mark in Quine's writings on analyticity, the a priori, and indeterminacy.

That agreement is in good part the result of several decades of fruitful dialogue between Dan, me, and my erstwhile colleague Jody Azzouni. However, we aren't here to acknowledge agreement but to discuss differences.

Beyond concepts?

One difference is that Dennett thinks that we should 'drop the requirement that there are concepts that are (and must be) shared for communication to be possible' (Dennett [forthcoming](#): 3); all we need is something like Ruth Millikan's unicepts. A response needs to be prefaced by something about Millikan's picture of cognition.⁸

Millikan observes that it is crucial for animals to be able to recognize when different bits of information are about the same object or property. She posits unitrackers as mental particulars which do this:

A unitracker is a mechanism or faculty same-tracking something, for recognizing when incoming information concerns it, then linking and storing this information together as information about one and the same thing. A unicept is a structure that forms a stable link, originated by its partner unitracker, between items of knowledge that are about the same. It holds them together so they can work together as, for example, a middle term in inference. (Millikan 2018, 43–4)

Unicepts are something like constituents of the mental states which realize beliefs and desires and stand in relations of inference, something like words in an individual's language of thought.

As Dennett suggests, Millikan's view is that all you need to account for cognition and communication are (unitrackers and) unicepts. She thinks there is nothing like traditional concepts, which she takes to be types of mental particulars which: are like Fregean senses, being something like descriptive ways of thinking which secure reference via something like satisfaction; play the role of public language meanings (and thus are not only shareable but shared); are or determine the content of propositional attitudes. Thinking that it is obvious that there are no such things, Millikan seems to take the primary virtue of theorizing about cognition using unicepts to be that they are completely unlike senses: they are particulars; they do not refer through associated information, but

⁸A further excuse for outlining Millikan's view here is that the Schroeters suggest that they think I should embrace a good deal of it, since like Dennett they think that shared meanings are not necessary for linguistic communication.

'refer directly'; unlike concepts, they can refer to an object or property even though they are underlain by a unitracker which often, even normally, fails to track the thing it is supposed to track – unicepts can be confused (or do a poor job tracking their referent), while a traditional concept can't fail to be a concept of, and thus descriptively adequate to, its referent.⁹

Contrasting my views about meaning with Millikan's raises two questions. One is whether the natural language meanings I posit have the alleged defects Millikan finds in concepts as traditionally conceived. The other is whether there is need in an account of human communication for the sort of commonality in term meaning I say there is. Set aside for the moment the question of whether natural language meaning requires something like shared information. Words, on my view, have two sorts of meaning. One is meaning as that which, when we cognize it, grounds our ability to understand and be understood by others; this is meaning as ICG. The other sort of meaning is what's commonly called semantic value, which is or determines the extensions and intensions of our terms. It's not my view that the first sort of meaning determines the other sort; I am not now, nor have I ever been, a card carrying member of the Fregean party.

On the story I tell, to say that your use and my use of 'dog' have the same meaning – to say that we are both tracking the ICG of 'dog' in a language we share – does not in the least imply that they express the same concept in the sense of 'concept' which offends Millikan. Word meaning in the sense I have tried to elucidate does not and is not supposed to determine reference. Though I have reservations about Millikan's biosemantical account of reference, as far as I can see her story is compatible with the overall thrust of the account of meaning I give.¹⁰ I agree with Millikan that term reference is not determined by description and is for the most part 'direct'. And the meanings I say words have can be deeply confused. The ICG of, say, 'atom' in a population is the set of those claims *p* such that (simplifying) it is common knowledge in the population that users of 'atom' presuppose *p* and expect you know they do. But 'atom' can perfectly well refer to atoms in a population even when it is common knowledge that everyone presupposes that atoms are indivisible; to know the meaning of 'atom' is not to know how to correctly

⁹See the discussion in Chapter 3 of Millikan (2018).

¹⁰Richard (2019) suggests that something like a fusion of a Putnamian story about reference and one or another principle of charity provides the beginnings of an account of reference. But substituting a teleosemantics wouldn't affect any of the things I say about meaning in the sense I am trying to elucidate.

define it, but to have a fix on what everyone (perhaps mistakenly) takes for granted about atoms.

Turn now to the question of whether part of competence in speaking a natural language involves shared meaning in some tolerably rich sense, or can be adequately theorized in terms of idiosyncratic unitrackers and unicepts. Some communication surely can be adequately theorized in terms of unitrackers: the sort of communication that arises in simple and not so simple Lewisian signaling games can be pulled off by co-ordination of perceptual states and dispositions to behave established by coordinated unitrackers. But important facts about human linguistic communication will not be captured if we take it to involve only co-ordination of unitrackers and unicepts. To be able to communicate at the level of your average eight year old one needs to be able to reliably: extract from an utterance what is 'strictly speaking' being said as well as what is 'merely' being implied; evaluate the probable overlap between the inferences speaker and auditor would (and would not) draw from the information imparted; efficiently resolve 'ambiguity' involving polysemous uses of a term; extract information from (such things as) word choice about (what the speaker means to reveal or suggest) about her affect with regard to the topic at hand; etc., etc. Part of the story about how we do such things will appeal to general intelligence which is not 'part of the language faculty'. But a good deal of the explanation of how we all manage to do these sorts of things in pretty much the same way will be in terms of shared meanings: what underlies this ability in humans sharing a natural language is a fairly robust agreement across speakers about what it is normally presupposed (and expected to be recognized as presupposed) by speakers when they use a word. This fairly robust agreement, I suggest, is maintained because pretty much all of us are very good at tracking norms of such presupposition. Given that there are such norms for word use, that we are tracking them, and that information assembled on the basis of this tracking is used in interpretation, it seems more than apt to identify (one sort of) meaning of a shared word with the various claims whose presupposition we are expected to recognize as being made when a word is used.¹¹

¹¹On my view the shared meaning of a word is its continually evolving interpretive common ground. Understanding a word at a time *t* is not a matter of having accurately transcribed its ICG at that time in one's lexicon; it is rather having a more or less accurate take on what the ICG of the word is as a result of one's ongoing attempts to monitor how words are used in one's community. I'll return to this in my reply to the Schroeters.

As I'm replying to Dennett and not Millikan, I should add that Dennett and I are not so far apart here. He writes 'We can each have our own concept (or better, unicept) of DOG and HORSE, and as long as our concepts, like our idiolects, stay in robust consonance with those of our neighbors and associates, we will manage just fine'. (Dennett 4) Of course our idiolects do generally stay in consonance, and communication breakdowns are usually quickly repaired. I take it to be part of the task of an account of public language meaning to explain just why this is the case; my proposal is that the community's collective attempt to track common expectations of speakers and interpreters is central to any such explanation. What makes it the case that we all mean the same by 'dog' is that we are tracking and more or less accurately registering the evolving lowest common denominator of community presuppositions associated with the word's use.

Meanings and viruses

Dennett is not altogether happy with my analogy between public language word meanings and genes. I say that there is a certain amount of 'allelic variation' in the lexical entries for a public language word. By this I mean that the information which is associated with a word as its meaning often 'clumps' within a population, so that there are several kinds of informationally overlapping lexical entries for the word in the population. The sort of clumping I have in mind is well displayed in cases in which a meaning is contested in a population, as the meaning of 'marry' in the US currently is: pretty much everyone's entry for the word includes claims like *married people usually live together* and *married people are not supposed to have intimate relations with non-partners*; but one variety of entry includes *same sex marriage is impossible*, while another does not.

'No', Dennett says, 'words are not like genes; they are more like viruses and bacteria'. (Dennett [forthcoming](#): 5) How so? Because in the case of ebola and E. coli 'there is really no useful distinction between mutation and the acquisition of new traits'. (Dennett [forthcoming](#): 5) In some ways, this complaint is of a piece with some of Armstrong's worries about the analogy: changes in meaning, Armstrong and to some extent Dennett think, have little to do with anything like selection of variants generated by the sort of reproduction in which fish, falcons, and the French engage; variants arise because one lexicon manifests itself in a way that makes other lexicons modify themselves

(Armstrong), or due to something well analogized by viral mutation (Dennett).

I did my best in replying to Armstrong to explain why I think the analogy I draw is not misleading, and I don't propose to repeat what I said there. But linguistic change at the level of an idiolect is not well thought of as a kind of mutation. Change in the lexicon is in good part produced by interactions in which the properties of language speakers' mental states are shaped by perception (and sometimes misperception) of how people expect to be understood. This is not mutation. It is diachronic change in the distribution of variants determined by earlier distribution and the interaction of the population with one another and the environment.

Reply to François and Laura Schroeter

Meaning and communication

The Schroeters agree that communication requires 'roughly similar ways of understanding' words. But they complain that 'On Richard's account, successful communication requires speakers and hearers to both grasp the same conventional ICG for particular words', and object that this is not a necessary condition for communication. (Schroeter and Schroeter forthcoming: 7–8).

This makes it sound as if I hold that successful communication occurs only if speaker and hearer have 'gotten ahold' of the claims in the ICG of the speaker's words and recorded them in their lexical entries for those words. That is certainly not my view. The most extended discussion of communication, grasp of meaning, and successful interpretation in Richard (2019) occurs in section 4.8, which discusses what it is for conversation to be fluid, and what it is for fluidity to be stymied. I suggest that: (a) word meaning is what grounds linguistic competence, as what makes one a competent speaker of a group's language is the ability to track the ICG of its terms¹²; (b) successful interpretation involves 'more or less correctly picking up on' what the speaker expects the interpreter to

¹²What grounds the ability to understand uses of [a word] in a group is, in part, one's ongoing monitoring of the ICG that surrounds the use of the word in the group ... the normal route [for understanding a word as used in a group G] involves something like this: one tracks the ICG surrounding the word in G more or less reliably, registers the ways it changes, and makes use of the information this monitoring provides in interpretation ... what constitutes understanding – "knowledge of meaning" – is in part a persisting ability to track a process, in part a fluctuating store of information garnered by the ability. It is not having some particular nugget of knowledge – a small chest of propositions or an immutable rule ...'. Richard (2019), 122.



recognize the speaker as presupposing. The norm in a community is that speakers understand one another because in conversation ‘there is pretty strong overlap’ between what the speaker expects the audience to recognize as presupposed in a term’s use, what the audience takes to be presupposed, and the ICG of the word (Richard 2019, 123). This is not a view on which communication succeeds only if both speaker and hearer match on what they take to be ICG. It is a view on which the ongoing and sustained ability of members of a group to communicate with a language is undergirded by those members tracking communal expectations about use. It is a view on which meaning is what the competent speaker is supposed to be tracking, and internalizing (with modification as use evolves) in order to be and remain competent. What makes the ICG the public meaning of a word is not that all competent speakers have engraved it in their hearts and minds, using it in a wooden way to interpret. ICG is meaning because, as the Schroeters put it, ‘ICG is a sort of interpretive center of gravity within a community at a time, which is imperfectly reflected in individual’s speakers’ idiolects and which can gradually evolve ...’.

Coordination and communication

Richard (2019) suggests that a group’s ability to successfully communicate with a language depends not only on its members coordinating on the ICG of its words but on those members standing to one another in what I call relations of linguistic co-ordination. Linguistic coordination is (very roughly) a matter of speakers coordinating their lexical entries so that the lexical entry X uses to interpret Y’s use of expression e is the entry X manifests in uttering e and vice versa. The Schroeters argue that linguistic co-ordination alone is sufficient to explain how understanding is achieved: in interpreting

P. Pork sausage reminds me of many things that are good,

You can simply plug others’ use of a term directly into your own encyclopedia entry for the word[s in the sentence] to arrive at a reasonable interpretation. In effect, the coordination relations linking different individuals’ lexicons amounts to a tacit presupposition of coreference. (Schroeter and Schroeter [forthcoming](#): 9)

Suppose I utter P and you ‘plug its words’ into your encyclopedia entries. In the case of ‘good’ you are going to have to make a decision. You know that it’s common knowledge that ‘good’ has many uses. It is, for example,

used to ascribe: moral approval; being safe to eat; being delicious; being apt for the purposes at hand. Each such use is associated with a large store of putative knowledge, but to decide which store might be relevant, you must use your knowledge of the various ways the word is used – your knowledge of meaning as ICG – to suss my intentions and figure out into which socket the word should be plugged.

The Schroeters will reply that even if ICG plays a role in interpretation in the case of ambiguity and polysemy, it is not plausible ‘that there are conventional patterns of understanding or ‘folk theories’ shared by our whole linguistic community for words like ‘Obama’ or ‘morally right’ or ‘water’. For most proper names, there is no dominant conception of the reference that’s mutually obvious to all competent users. Similarly, theoretically ambitious terms like ‘morally right’ or ‘water’ can be highly contested and variable’. (Schroeter and Chroeter forthcoming: 5–6) They add that

this point is familiar from the work of externalists.

I must protest. First of all, the point familiar from the externalists is that it is implausible that there is a community wide dominant conception that *determines reference*. But I say repeatedly that meaning as ICG is meaning in the sense of the ground of linguistic competence, what one needs to more or less reliably track to be competent in a language; ICG (often containing large amounts of misinformation) is not supposed to, and is not up to the task of, determining reference.¹³

What, though, of the Schroeters’ remarks about proper names and ‘theoretically ambitious terms’ such as ‘water’ and ‘right’? Is it even plausible that there is something like a dominant fix on the expectations of users about such terms?

¹³This is relevant to the Schroeters’ allegation that I’m ‘committed to a version of the a priori conceptual analysis favored by many neo-descriptivists’. (SCH, 7) They argue that since ICG is common knowledge, one can reason in the armchair to ‘a priori knowledge of ... ICG’, which they take to be reasoning to a conceptual analysis.

Concepts traditionally were taken to be determinants of reference, and so conceptual analysis was traditionally supposed to provide us an articulation of what our thought and talk is about. If that is what a correct analysis provides, and a priori ratiocination allows me to discover that my concept of knowledge is the concept of justified true belief, then I can know a priori that that is what knowledge is. But the fact that ICG is common knowledge does not yield this sort of knowledge. A claim *p* is in the ICG of ‘knows’ because it’s common ground that speakers *assume p* when using ‘know’ and expect the audience to recognize this. Suppose that B.G. (Before Gettier), it was common knowledge that speakers assumed that knowledge was justified true belief. Even if this means we could know a priori that speakers generally *assumed* that knowledge was justified true belief, that’s a long way from knowing what knowledge in fact was.

On top of this, it’s hard to see how any of this is a priori, as I need extensive experience with the ways of the community in order to know what a word’s ICG is therein.



In the case of terms like ‘water’ or ‘morally right’ the answer is that obviously there is. Pretty much every adult speaker of English knows that (it is common knowledge that) users of ‘water’ expect their audience to know that they presuppose that ‘water’ is used to pick out something which: we drink; is found in rivers, lakes, and the oceans; is normally transparent when liquid; evaporates when heated; is used to cool drinks when it is frozen, etc., etc. Whether this litany determines a unique reference is neither here nor there for such knowledge to ground an ability to understand speakers. (And if reference determines a modal profile, then as Putnam observed, the litany obviously does not determine reference.) Of course what is and is not water is contested: Is coffee water? Is coke? Is a plasma created from a ‘normal’ sample of water water? But contestation does not imply that there is not a substantial consensus on what we ‘have to know’ about ‘water’ to count as competent.¹⁴

Something cognate is obviously true of phrases like ‘morally obligatory’.¹⁵ It is common knowledge that users of the phrase assume (that audiences recognize that they assume) that the phrase is used for acts which: one has a compelling reason to perform; are taken by most people to include such things as refraining from killing, not stealing, supporting one’s children, not lying about how many votes you got; one is often criticized for not performing; etc.¹⁶

Proper names are tricky, and a discussion of them needs to be more nuanced than the discussion in Richard (2019), where I casually suggested that part of the ICG (in 2016) of ‘Barrack Obama’ was the claim that he was president. I stand by the claim that given the definition of ICG I employ, the claim that ‘Obama’ names the president was part of the name’s ICG in the U.S. in 2016. What of the Schroeters’ claim that it was not then

¹⁴This is the place to address the Schroeters’ worry that I do not give an account of where to draw the line between claims which are part of a term’s ICG and those which are just parts of a user’s encyclopedia entry.

I distinguish two things: an individual’s take on a term’s ICG, the claims she thinks are commonly presupposed, and what a term’s ICG in a population in fact is. ICG is the latter. If in some cases this means that ICG is somewhat ‘encyclopedic’, so be it. There’s nothing wrong that I can see with the idea that when we all presuppose that we use expression *e* to talk about things that are A, B, C, ... Z, ... and expect audiences to know this, knowing that is part of what makes us a competent speaker.

More discussion of this occurs in the response to Stalnaker.

¹⁵I’ve switched from ‘right’ to ‘obligatory’, though I don’t see it affects the point here.

¹⁶Relevant here is that chapter 6 of Richard (2019) discusses the fact that many terms, particularly ones with a normative dimension, will be contested. The suggestion I make there is roughly that in such cases the ICG of a term *t* will be of the form: there are ways $w_{1...}$, w_k of using *t*; when using it in way w_1 , speakers presuppose p_1, \dots , and p_j ; when using it in way w_2 , they presuppose q_1, \dots , and q_j ; ...; when using it in way w_k they presuppose t_1, \dots , and t_n .

necessary to know that people presupposed that ‘Obama’ named the president in order to understand an utterance of ‘Obama authorized drone strikes’?

Well, of course that claim is correct. Suppose that A speaks English but is ignorant of all facts about US politics and that B, pointing in turn to Obama, Biden, and Brennan, says to A ‘Those three people are Barrack Obama, Joe Biden, and John Brennan. Obama authorized drone strikes’. A will understand B.

But the issue is not what invariably accompanies any successful interpretation of a use of a proper name. The issue is what underlies our competence to understand uses of proper names that are not deitically anchored in the context of use. I say that the answer is co-ordination on their ICG. In interpretation we begin by identifying a proper name token as a token of a phonetic / orthographic proper name type, and thus as a generic name. As generic, that type will have a conventional use as a name of any number of people. In general, people who routinely use generic names (e.g. tokens of types like ‘David’, ‘Rose’, or ‘Francis Bacon’) as names of particular individuals seem to co-ordinate their lexicons in terms of common assumptions which can differentiate amongst different possible referents. This is witnessed by conversational hiccups like ‘Uh, do you mean Francis Bacon the painter or Francis Bacon the philosopher?’ Generic names being generic, we need to be prepared to discern – I would say by delimiting the speaker’s presuppositions in use – who the speaker is talking about. Sometimes the speaker’s presuppositions are one-off. But groups often effectively conventionalize the indexing of a generic name with disambiguating information. When this happens, the group has attached a shared meaning to the name – a meaning that need not determine reference but instead is accessed by the competent speaker in order to guide utterance interpretation.

Referential indeterminacy

The Schroeters are unhappy with what I say about referential indeterminacy. (Schroeter and Schroeter [forthcoming](#): 8) I say that without recourse to our interests in interpretation, there are a variety of maximally good ways to assign reference. Does this mean that I hold that there is no objective fact about a word’s reference? That would be a fairly uncharitable interpretation. Assignments of referents, I say, need to respect such things as the paradigms associated with a term, patterns of deference about how we are to project reference from those paradigms, and one

or another principle of charity (Richard 2019, 4.4–4.5). This, I think, suffices to make it an objective fact that Nancy and Mark’s canine pets Simone and Ajax fall under ‘dog’, as do pretty much all the bearers of dog licenses in Newton.¹⁷ It’s also an objective fact that no cans of Dr. Pepper are in the term’s extension. What I urged was that independent of the interests of the inquirer there won’t be an objective answers to questions like, Are dog-fox hybrids dogs?

The Schroeters are ‘happy to concede’ something like this is true of ‘everyday communication’. But they say it is not true in ‘rational inquiry into the nature of a familiar but imperfectly understood topic’. (SCH, 8) This is because: (a) this sort inquiry presupposes that we have been keeping track of a single topic over time; (b) we start the inquiry with a presupposition of diachronically stable reference; (c) the goal of such inquiry is ‘to arrive at a coherent interpretation that best integrates and vindicates the different perspectives and interests within our coordinated practice’ of using the term at the focus of the inquiry.

I pretty much agree with all of this. As for (a): Most of us who are to some extent invested in the project of conceptual engineering think that preservation of topic is consistent with change in extension. Evidence for this is the fact that we will ascribe assertion disquotationally (if *x* makes an assertion with *A*, we will – adjusting for whatever demonstrative or indexical elements occur in *A* – say that *x* said that *A*) even when we inclined to say the extensions of predicates in *A* shift between *x*’s and our context.¹⁸ So I agree that if we inquire, say, about what species are, we begin by assuming that we, Martin Nowak, Darwin and (perhaps) Linnaeus are all ‘talking about the same thing’. As for (b): ‘talking about the same thing’ is a tricky locution. It could be understood as a matter of stability of topic, of stability of Kaplanian character, or stability of temporal (or possible worlds) intension. I agree with (b) in so far as talk of stable reference is understood in one of the first two ways. But that is consistent with extension variation and indeterminacy of best interpretation. And I agree with (c): in inquiry we are trying to arrive at a coherent interpretation of past and present talk. But when we investigate, we may conclude that there are different ways of assigning extensions to ‘species’ and nothing that makes one assignment the best independent of our interests. There are, after all, some two dozen candidates for interpreting

¹⁷For the record, Simone’s full name is Simone de Bauer.

¹⁸See, for example, the discussion of consistency of topic and diachronic talk of belief preservation in Cappelen (2018), Richard (2019) and Richard (2020). See also the discussion of ‘know’ and gradeable adjectives in Chapter 4 of Richard (2008).

current and past talk about species, many of which are reasonable if imperfect glosses on much biological talk.

In short: I agree that that we typically begin inquiry assuming that we are talking about the same thing as our peers and ancestors. But it just doesn't follow that there is only one interpretation of our talk. Even given that the best way of interpreting us has us at the end of the day talking about the same thing, it needn't be the case that one interpretation is better than all others. And it certainly doesn't follow from the facts, that we begin inquiry by assuming identity of topic and that inquiry's goal is to arrive at a single interpretation of our talk, that we when we begin inquiry we are all talking about the same thing.¹⁹

Reply to Robert Stalnaker

Meaning and reference

Richard (2019)'s project is to contribute to an explanation of linguistic competence. To be a competent speaker of a language is to be able to express thoughts using it in a way which other speakers can reliably (but not infallibly) interpret, and to be able to reliably (but not infallibly) interpret thought expression by others. Thought is by and large representational, and so is about the world, about its objects and their properties. One might conclude, as Stalnaker seems to, that my project involves or should at least contribute to solving the problem of linguistic intentionality – explaining what makes the members of group united by a language able to think and talk about the world using that language. Stalnaker writes

... to be competent with a language, or with particular words in it, is to have the ability to use the words to communicate information about things in the world, which suggests that what we must be in cognitive contact with [to be competent speakers] involves a relation between the words we use and things they refer to, and thoughts they express. (Stalnaker forthcoming: 2)

If Richard's conception of meaning is to be a conception of what we have to "grasp" or connect cognitively with to be in a position to engage in communication, then I think it should provide us with an explanation of its role in explaining how our use of words and other linguistic devices allows us to connect with

¹⁹In the last section of their paper, the Schroeters ask whether I think that meanings as I construe them – as sequences of populations of consisting of lexical entries – are 'like segments of [biological] lineages, demarcated by causal-historical relations of coordination and descent'. I do.



things in the world, and to have the capacity to say things about them. (Stalnaker forthcoming: 3)

As I understand this, Stalnaker thinks that, since I am providing an account of meaning in the sense of what one needs to know in order to communicate, I need to explain how having this sort of meaning makes, or at least contributes to, our use of the words which have it refer.

If I were providing an account of the truth conditions of sentences or of the propositions that are determined by complement clauses, then what Stalnaker is asking for would be provided by the account. For it would tell us that 'Obama' refers to Obama, 'is a dude' is true of the dudes, and thus 'Obama is a dude' is true iff Obama is among the dudes. Or it would tell us that word, phrase, and sentence have the intensions they do, assuming we take propositions to be intensions. Along with the observation – assertively using a sentence with truth conditions C (expressing proposition p) is representing the world as being a place in which C obtains (asserting p) – an account of truth conditions or propositions expressed tells us how our use of our words connects us with the world.

But I am not providing such an account, and the sort of meaning I am trying to elucidate is not elucidated by giving an account of phrase reference.²⁰ In some sense, meaning of the sort I am trying to elucidate is posterior to meaning in the sense of (the determinant of) reference. Knowledge of such meaning is knowledge of social facts about language use. Of course, that we share such knowledge helps insure that we are all talking about the same things. It plays a role 'in explaining how our use of words and other linguistic devices allows us to connect with things in the world', as regularities in use are a part of the complex network of environmental and social relations connecting the speakers of a language (and thus connecting their language) to one another and to the world. But these regularities are only a part of what determines reference; pointing out that knowledge of them makes mutual understanding possible is not trying to give an account of how it ends up being the case that 'young dude' is true of all the young dudes.

²⁰Which is not to say that presuppositions about reference are not part of it. It is common ground that we all presuppose and expect to be recognized as presupposing that: 'cousin' is a word for cousins, 'I' refers to its user, and so forth. So these claims are part of the ICG of the relevant words. And it is an easy, but non-deductive, inference from such things and pretty obvious facts about how sentences compose that 'I have two cousins' as used by me is true iff I have two cousins. So participating in common ground is knowing things we all *assume* about reference. Since we usually have a tolerably good idea what we are talking about, it is a(n imperfect) guide to what we are talking about.

What is in ICG?

Stalnaker asks ‘exactly what propositions are in the ICG for a given expression’. Richard (2019) provides a definition which answers this: p is part of the conventional meaning of e in G to the extent that it’s common knowledge in G that users of e presuppose p and expect auditors to recognize that.²¹ What claims are in the ICG of a word in population P is thus an empirical matter – it depends on the linguistic *habitus* of the group, the field of dispositions users of the word have, to make presuppositions and have expectations, that are commonly recognized within P . That said, a rough, partial indication of when a claim expressed by using or mentioning a word is in its ICG is that speakers generally take the claim to be beyond obvious, reacting to its statement with rejoinders like ‘But that’s just what w means’. I did, after all, motivate the project of Richard (2019) by saying that while I agree with Quine’s objections to the notions of analyticity and the *a priori*, I also think that intuitions about meaning, synonymy, and analyticity are tracking something important. The fact that everyone would agree that (everyone would agree that) the truth of ‘vixens are foxes’ is guaranteed by what its words mean is a sign that, at the time of such agreement, the claim that vixens are foxes is in the ICG of ‘vixen’.²²

The analogy of meanings as species is relevant to questions about what is in ICG. Species typically consist of dispersed populations. There is often variation in a species’s members indexed by such dispersal, variation induced by such things as a tendency of dispersed sub-populations to be relatively reproductively isolated from one another. And while sub-populations of a species are not themselves species, they may be distinctive enough (for example, in the environmental conditions they can and cannot tolerate) that dispersed populations tend to evolve in distinctive ways while still being capable of such things as interbreeding. So members of a species may be, and often are, members of multiple sub-populations with different if connected evolutionary trajectories.²³

²¹The ICG of a term in a population P will be vague. For claim p to be in the ICG of term t , certain generic claims about P must be true. But generic claims, being claims about what is the norm in some situation or population, can be more or less accurate. (See Richard 2019, 3.4–3.5) So the line between being a part of conventional meaning and having a weaker status is bound to be somewhat vague.

²²There is work to be done here. I wouldn’t say the fact that everyone thinks that “vixens are foxes” is true because of its meaning’ implies that the claim that vixens are foxes is in the ICG ‘foxes’ (or of the verb ‘be’). Presumably there is some notion of sentence topic, one on which vixens are a topic of the sentence while foxes are not which would help explain this.

²³For a discussion of variation in dispersed populations see Calosi et al. (2009).

And something similar is true of individuals who share a public language. They will be members of (multiple) sub-groups whose overall linguistic behavior within those subgroups can be, and often is, subject to different pressures because of differences in what is more or less conventional about word use within the group. To the extent that we individuate languages (in part) in terms of their conventions for word use, this means that in some sense a speaker of a public language will speak any number of related ‘dialects’ indexed by the (kinds of) groups in which she finds herself repeatedly interpreting and speaking. So when we speak of the meaning of an expression as used by a speaker, or of the language of a speaker, we will need to recognize that such talk must be indexed to one or another group. A presupposition which is part of conventional meaning relative to one group or dialect – think here of terms with both a quasi-technical and a common use, or ones whose meaning is contested – may not be a part of conventional meaning relative to another, even if that presupposition is often made by speakers inside the second group.

Pragmatics and semantics

Stalnaker’s request for an account of exactly what claims are in ICG is motivated at least in part by the observation that the line between presuppositions which are part of conventional meaning and those invoked for pragmatic communicative purposes is a line continually subject to modification. Part of what Stalnaker wants, I take it, is a criterion for deciding, when a presupposition is (more or less) routinely present and expected to be present, on which side of the line it falls. This is a reasonable request; let me try to give Stalnaker some of what he wants.²⁴

As I see it – and here I follow a line of thought in usage based grammar, construction grammar, and allied views²⁵ – our syntactic and semantic knowledge undergoes continual modification as a result of our repeated encounters with particular uses of phrases. We collect exemplars of use in a variety of contexts and generalize over them, extracting rules of thumb for interpretation of the form *in contexts of such and such a kind, one way phrase e is used is way w*. Having such a rule of thumb does not by itself make expectations used in interpretation part of conventional meaning. I

²⁴What follows is developed in a bit more detail in Richard (2021).

²⁵See, for example, Bybee (2010), Geerearts (2006), and Goldberg (2006).

know that one way my wife uses the phrase ‘stupidest person I have ever met’ in discussing academic administrators is to express mild disdain for a person’s views of university policy. But I do not think this is part of its conventional meaning, even when we consider the dialect spoken only by my wife and I.

What, then, makes an interpretive rule of thumb a part of conventional meaning? Partisans of UBG suggest that over time a ‘mere’ rule of thumb (e.g. *people use ‘will’ to make predictions about the future*) for interpreting a phrase may become ‘mechanized’ as an interpretive strategy, so that it is automatically, and not just inferentially, available for use in interpretation. As such mechanization occurs quite generally in a population, it may become common knowledge that one way people use the phrase is with the expectation that the audience will (recognize that they expect them to) use the relevant rule in interpretation. And at this point the rule has crossed the line between pragmatics and semantics, becoming part of conventional meaning.²⁶

Referentialism

Intuitively, this is the thesis that reference is essential to meaning. More precisely, it is the view that the character of an expression – the function associated with the expression by convention which maps a context of use to an intension – is essential to meaning. I tried to be cautious in what I said about this view in Richard (2019): I said that I was ‘inclined to think it is a bad way to think about conceptual continuity’ (Richard 2019, 97); my discussion of it is signposted as one which ‘express[es] a good deal of skepticism about’ it (Richard 2019, 121). My view of referentialism was and still is unsettled.

That said, I do suggest there are putative counterexamples to the thesis. Stalnaker discusses one involving the meaning of ‘marry’: it is plausible that the relation ascribed by uses of the verb has changed in the last one hundred years (from a relation in which those of the same sex cannot stand to one in which they can); it is implausible to say that the word does not mean what it did one hundred years ago, as the change in what is common ground about speaker expectations about

²⁶Stalnaker’s question is more than reasonable, and more than what’s said here is needed for a full response. As my answer suggests, I think a full answer to the question would have to be based on reasonable conjectures about the mental structures underlying interpretation. Much of the literature on ‘radical pragmatics’ is relevant here – Carston (2021), for example, has a worthwhile discussion of how the line between conventional meaning and pragmatically inferred meaning might be drawn.

assumptions about ‘marry’ is actually quite minimal (Richard 2019, 107). Stalnaker replies that, assuming that I have the same view of the relation between character and context as Kaplan, this is not a counterexample to a reasonable version of referentialism. As I understand the response, it goes as follows: Context, should be understood more broadly than Kaplan understands it. It includes what is common ground among conversants, since this is (pretty much always) relevant to extension. So even if the relation ascribed by the use of the verb shifted between 1921 and 2021 (because common ground about marriage shifted), that does not imply that there was a shift in the function which is the character of the verb at both times: there is a single function from character to extension which is associated with ‘marry’ from 1921 to 2021.²⁷

I agree that my putative counterexample to referentialism is at best incomplete. But I don’t think Stalnaker’s response to the putative counterexample is terribly convincing. On the picture I am working with, a phrase’s reference is identified with a function from context to the phrase’s contribution to truth conditions or to proposition expressed. I agree with Stalnaker that because the relation determined by the use of gradable adjectives like ‘rich’ and ‘rancid’ pretty much always turns on what is mutually obvious in a conversation about presuppositions about wealth and being rotten, some aspects of conversational common ground need to be a part of the contexts on which character operates. A dependency on conversational standards of application is baked into the meanings of such adjectives and so is part of their meaning in the sense determinant semantic value. (And thus it is also part of their meaning in the sense of ground of competence.). But as I see it, idiosyncratic conversational common ground plays a limited role in the determination of conventional reference. What is determinative of conventionally determined reference is in the first instance are such things as the history of the use of the term, associated paradigms, deference to ‘experts’ and social institutions like religions, and the web of causal/social connections between speakers and the world. For many terms it is these which do the bulk of the work in reference determination; ICG and common ground among speakers is more an imperfect reflection of what determines reference than a determiner.

This ‘meta-semantical’ picture is part of what motivates me to say that there is a reference shift in the putative counterexample. I don’t see a

²⁷Stalnaker also remarks that he doesn’t think that there was a shift in relation ascribed, though he concedes that there might be cases in which there is such a shift without the sort of shift which we would call a change of meaning.



strong case for the idea that it is common ground per se that determines the reference of 'marry' at a particular time in a particular socio-linguistic context; what does the determining are patterns of use in the past and present, paradigms, patterns of deference and other socio-cultural factors. It seems to me that, given that it is these factors which are doing is determining reference, there is a good case to be made that the verb's reference in 1920 is a constant function mapping contexts to a relation without same-sex pairs, while the reference in 2020 is a different one.²⁸

That said, I see a case for the verb's reference being constant from 1920 to 2020, since it's not unreasonable to say that a term's future use can be relevant to determining its reference in the past. I think that even accepting this, there are variants of the counter-example in question which suggest that reference isn't essential to meaning, one of which is developed in Richard (2020). My attitude towards referentialism continues to be, as it was in Richard (2019), a skeptical but not dismissive one.

Acknowledgements

I'm grateful to Josh Armstrong, Dan Dennett, Laura and François Schroeter and Bob Stalnaker for their care in reading *Meanings as Species* and their comments and their hard questions about it. I hope my responses at least approach being as good as their comments and questions.

Disclosure statement


No potential conflict of interest was reported by the author(s).

References

- Armstrong, Josh. *forthcoming*. "Meanings Without Species." *Inquiry*.
- Bybee, Joan. 2010. *Language, Usage, and Cognition*. Cambridge University Press.
- Calosi, P., D. Bilton, J. Spicer, S. Votier, and A. Atfield. 2009. "What Determines a Species' Geographical Range? Thermal Biology and Latitudinal Range Size Relationships in European Diving Beetles (Coleoptera: Dytiscidae)." *Journal of Animal Ecology*.
- Cappelen, Herman. 2018. *Fixing Language*. Oxford: Oxford University Press.
- Carston, Robin. 2021. "Polysemy: Pragmatics and Sense Conventions." *Mind and Language*.

²⁸There are interesting questions about how this view interacts with our intuitions about ascriptions of propositional attitudes to our forebears. Chapters 4 and 5 of Richard (2019) attempt to say something useful about this.



- Dennett, Daniel. [forthcoming](#). "A Fond Farewell to Cartesian Linguistics." *Inquiry*. Geerearts, Dirk, ed. 2006. *Readings in Cognitive Linguistics*. Mouton de Gruyter.
- Godfrey-Smith, Peter. 2009. *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.
- Goldberg, Adele. 2006. *Constructions at Work*. Oxford: Oxford University Press.
- Millikan, Ruth. 2018. *Beyond Concepts*. Oxford: Oxford University Press.
- Richard, Mark. 2008. *When Truth Gives Out*. Oxford University Press.
- Richard, Mark. 2019. *Meanings as Species*. Oxford: Oxford University Press.
- Richard, Mark. 2020. "Is Reference Essential to Meaning?" *Metaphysics* 3: 68–80.
- Richard, Mark. 2021. "Superman and Clark Walk into a Phone Booth." In *Oxford Studies in the Philosophy of Language*, edited by Ernest Lepore, and David Sosa. Oxford: Oxford University Press. 
- Schroeter, Francois, and Laura Schroeter. [forthcoming](#). "Meanings as Species in Communication and Inquiry." *Inquiry*.
- Sober, Eliot. 1980. *The Nature of Selection*. Cambridge: MIT Press.
- Stalnaker, Robert. 20  "On Meanings as Species." *Inquiry*.