



Meaning change and changing meaning

Allison Koslow¹ 

Received: 19 February 2021 / Accepted: 18 December 2021 / Published online: 14 March 2022

This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2022

Abstract

Is conceptual engineering feasible? Answering that question requires a theory of semantic change, which is sometimes thought elusive. Fortunately, much is known about semantic change as it occurs in the wild. While usage is chaotic and complex, changes in a word's use can produce changes in its meaning. There are several under-appreciated empirical constraints on how meanings change that stem from the following observation: word use finely reflects equilibrium between various communicative pressures (just as, say, product sales do between various market pressures). Much of the relevant work in linguistics has employed the methods of empirical pragmatics and diachronic semantics. In this way, the study of meaning change can be brought to bear on the conceptual engineer's normative project. The picture that emerges tells against the sorts of engineering projects most likely to appeal to philosophers. Some may stand to succeed, but they have significantly different contours than the typical ones.

Keywords Conceptual engineering · Amelioration · Feasibility · Implementation · Semantic change · Historical linguistics · Metasemantics

1 Introduction

Paradigmatically, a conceptual engineer identifies a word—say, “woman”—that expresses a concept—WOMAN—and proposes that we use the word to express a new but related concept in some area of discourse. The target area of discourse may be exclusive to a cabal of philosophers during working hours, but often it is broader. The reasons for conceptual engineering are potentially many. They include facilitating a

✉ Allison Koslow
akoslow@uci.edu

¹ UC Irvine, Irvine, CA, USA

better account of some phenomenon, furthering some malevolent aim, or making for better puns.¹

Conceptual engineering has most often commanded recent philosophical attention when proposed as a tool for social justice and ideological reform. In that guise, it is sometimes called “conceptual amelioration.”² Where some activists may correct our implicit theory of, say, race or addiction, conceptual engineers may advocate changing our concept RACE or ADDICTION. The case for correcting mistaken theories hardly needs stating; the case for adjusting our concepts is less clear.

The suggestion that we can feasibly change the world by revising our cognitive and linguistic technology is undeniably exciting. Is it also correct? The question can be factored into two components. First, would the revision promote its alleged goal (whatever it may be), when deployed in the target area of discourse?³ Second, is revising our concepts in the way proposed achievable? If a conceptual revision is justified by the goal it promotes, then this second question is paramount.⁴ This paper sketches a method for answering it, for assessing the feasibility of proposed revisions (whatever their aims).

Conceptual engineering requires a theory of what sorts of meaning changes can, as a practical matter, be induced. If you want to move a sand dune, it helps to know how sand dunes move about in the normal course. Fortunately, much is known about meaning change as it occurs in the wild. Much of the relevant work in the linguistic realm has employed the methods of empirical pragmatics and diachronic semantics, including the study of linguistic innovation. I bring this work to bear on the engineer’s project.

The most exciting engineering projects aim to shape collective, rather than individual, linguistic behaviors, because they aim to shape collective, rather than individual, predicaments.⁵ Identifying worthwhile projects involves identifying what proposed changes to group linguistic behavior are more, or less, likely to take hold. Our linguistic behavior, like other complex and chaotic systems, is not to be studied from the armchair. It’s surprising then that discussions of whether conceptual revision is achievable have tended to focus on the degree to which *individuals* can, or cannot,

¹ Scharp (2013), for instance, suggests *logicians* use a revised truth rule to facilitate a coherent logic of truth. But some reactionary political movements may also involve conceptual engineering. See Srinivasan (2019) and Queloz 2021 for a related point.

² Not all conceptual ameliorations have such aims, and not all projects that go by the name “amelioration” advocate conceptual change. See Haslanger (2020) for discussion of kinds of ameliorative projects. Amelioration also relates to recent work on hermeneutical injustice (Manne 2017). To the extent these don’t also involve claims about what concepts we should use, they are beyond the scope of this paper.

³ See Saul (2006).

⁴ The second question is sometimes called the feasibility question (Cappelen 2018), or implementation challenge (Cappelen and Plunkett 2020).

⁵ A limiting case are proposed revisions that, like Scharp (2013)’s proposed truth rule, aim only to facilitate theoretical projects in expert discourse. To the extent they argue, e.g., *logicians should* adopt the revision, they too target the linguistic behavior of *a group*, albeit a small group (e.g., with a liar-shaped predicament).

enact it, given some account of concepts, and what determines their contours.⁶ This paper takes a different tack.

Questions about the feasibility of conceptual revision are like questions about whether a particular product will succeed in a new market, or whether a nonnative species will thrive in our ecosystem. We can study the likelihood of one expression becoming more popular than another in an area of discourse, rather like we can study the likelihood of one frog species outcompeting another in an area of the rainforest. While an expression's usage may be inexhaustible, it is also shaped by various communicative pressures that exert themselves in foreseeable ways. Past data suggests predicative generalizations, whether or not these have the status of exceptionless laws—so too with forecasting frog propagation, pedestrian behavior, and widget sales.

The picture that emerges tells against the sorts of revisionary projects most likely to appeal to philosophers. As we'll see, some engineering projects may stand to succeed, but they have significantly different contours than the typical ones.

2 Engineering, and semantic drift

Conceptual revision is only one sort of (one sort of) engineering project. Distinguish three sorts: introduction, elimination, and replacement. *Introduction*: experts and ordinary speakers alike frequently introduce words and phrases to make useful distinctions. Sometimes they are neologistic, like “laser”.⁷ Sometimes they are borrowed from ordinary language, and sometimes have related meanings (like “safety,” “sexual harassment”). *Elimination*: speakers may also actively discourage the use of certain expressions (slurs, for instance, or failed theoretical terms like “miasma”). Mostly, however, expressions drop from favor as a byproduct of changing fads, interests and technology. (The vocabulary of falconry has gone the way of the birds, and so has “typewriter.”) Occasionally, an expression gradually fades from public consciousness (like “apricity”).⁸ *Replacement*: sometimes the introduction of one expression results in the elimination of another, as has (nearly) happened by rebranding prunes as “dried plums.” In effect, use of one word-meaning pair <W, C> replaces use of

⁶ For example, supposing the objects of revision are semantic values in a language with an externalist metasemantics, Cappelen (2018) is skeptical that conceptual revision is possible at all. Koch (2021a, b) is more sanguine, arguing that individuals can enact long-range control over the meanings of their expressions. Deutsch (2020, 2021), disagrees, arguing that the implementation problem is insurmountable for any “metasemantic view that requires more ... than just the intention on the part of some group of speakers to use the relevant term as if it had that very semantic meaning and reference” (17). Pollock (2021) and Nimtz (2021) appeal to internalist metasemantic principles to argue that conceptual engineering is feasible: individual speakers can enact it. Fischer (2020) empirically investigates individual's ability to reason with engineered concepts. Pinder (2021) argues that conceptual revision is better understood as targeting speaker-meanings, and so may be easily enacted by individuals. Machery (2017) approaches engineering as a psychological, rather than semantic, project and is optimistic: concepts are revisable insofar as they are particular bodies of information retrievable by default from long-term memory for tasks, like reasoning and categorization. Machery (2021) considers the limits of revision, given well-established features of our cognitive architecture.

⁷ “Laser” originated in the scientific community as an acronym for “light amplification by stimulated emission of radiation” and gained currency with the broader public in the 1960s.

⁸ The term, now obsolete, meant the warmth of the winter sun.

another. These may share the meaning-component (as do “dried plums” *-dried plums* and “prunes”-*dried plums*), or share the word component, or neither. (Some replacement happens by mistake: as nautical terms become less familiar, use of “tack” in stock metaphors, as in the previous section, is sometimes replaced by use of “tact”).

Conceptual revision as commonly advocated by philosophers is one kind of replacement project. One word-meaning pair $\langle W, C_{\text{NEW}} \rangle$ is introduced to replace use of another, $\langle W^*, C_{\text{OLD}} \rangle$, where $C_{\text{NEW}} \neq C_{\text{OLD}}$. Replacement, rather than simply introduction, is part of the picture because not only is C_{NEW} deemed desirable, but C_{OLD} is deemed undesirable. We can call this *meaning replacement* since an expression with one meaning replaces uses of an expression with another meaning. When suitably related, C_{NEW} may be described as *revising*, rather than merely *replacing*, C_{OLD} , but this distinction may be set aside for the purposes of this paper. Typically, $W = W^*$. Retaining the familiar word for the new concept seems a natural, though unforced, choice.⁹

This paper mostly concerns meaning replacements, in particular ones that extend outside of purely expert discourse, but many of the points generalize. Roughly, engineering projects propose some change (introduction, elimination, or replacement) to part of the total mosaic of usage to advance some goals. If there’s a strong reason to think the proposed change will not take place to do so sufficiently, the proposal is undermined with respect to those goals. So identifying feasible revisions involves identifying which word-meaning pairs are likely to catch on in which areas of discourse, in which parts of the total mosaic of usage. We may construe usage broadly to include spoken and written applications of a particular expression, thoughts involving it, and inferences, presuppositions and associations triggered by its application. Some parts of the mosaic are in seminar rooms, some on the beach.

Before we go any further, some clarificatory remarks are in order. “Concept” is used in multifarious ways. Some philosophers identify concepts with meanings or semantic values, while others identify them instead with mental representations (and others still with anchors of semantic competence, inferential norms, sets of capacities).¹⁰ Just how use relates to meaning, and how linguistic meaning relates to cognitive architecture, is fraught. Fortunately, there is no need to wade into the attendant controversies here. While different paradigmatic engineering projects may operate on different targets,

⁹ Thus, in the case of the concept WOMAN, the conceptual engineer urges us to use another concept instead of the existing one. The obvious suggestion is then to lexicalize the new concept in the time-honored way, as “woman” (rather than, say, “femina”). Speakers are antecedently likely to continue carving up the world with habitual expressions. These are the expressions caught up in our verbal dispositions, written into our laws and on our signs. They trigger social obligations and figure in instinctive self-descriptions. Short of rewriting everything, the best strategy would seem to be to keep the words but change their meaning (to one suitably related). See Haslanger (2000, 2010), and Cappelen (2018) on so-called lexical effects. But why advocate usage of $\langle W_{\text{OLD}}, C_{\text{NEW}} \rangle$, over usage of $\langle W_{\text{NEW}}, C_{\text{NEW}} \rangle$, or convincing the relevant speakers of some new theory of what the W_{OLD} s are like? Doing so might seem harder on account of involving an extra step: convincing speakers to use an old word in a new way, *and* that doing so would be good. Presumably, the answer is that it isn’t harder, and has a greater likelihood of success. If that’s right—if the reason to undertake conceptual revision, as opposed to related strategies to achieve the same ends, is that it’s more doable—then a strategy for assessing how likely these interventions, and others, are to succeed is crucial. Perhaps there are reasons to employ more than one simultaneously, especially when they support one another, and the opportunity costs, and risks, are small.

¹⁰ For discussion of engineering concepts, understood as sets of capacities, see Haslanger (2020).

conceptual revision may be described as involving changes in the use of words with particular meanings, in an uncontroversial sense.

Quinean skepticism aside, there is an intuitive notion of a word “changing in meaning” that figures in lexicography and diachronic linguistics. Without it, these legitimate empirical disciplines would not exist. No controversial conceptions of concepts, or disputed theories of linguistic meaning, or its determiners, are required to observe that in Chaucer’s day “girl” was correctly applied to children of any sex and in Queen Elizabeth I’s day “one myriad” was ten thousand.¹¹ These changes in application cannot be explained by changes in background belief. Chaucer knew all about the birds and bees when he wrote of “girle knights,” and we have not forgotten any math since the Elizabethan era.¹²

Conceptual revision involves changing the meaning of words in the mundane sense of “meaning” in which it is true that “myriad” and “girl” are used today with different meanings than they once were. Real life examples abound. “Shore,” for instance, originally meant the tidal zone where boats run aground. As the beach changed from a place of business to leisure, it was convenient and clear enough to call the place beachgoers went the “shore” instead of reaching for a new word. Repeated usage eventually resulted in “shore” acquiring a new meaning, the whole beach. Gradually uses of “shore”-meaning-the whole beach have replaced some but not all uses of “shore”-meaning-tidal zone. Conceptual revision, in effect, involves use of one expression replacing use of another, as happened with “shore”.

Two examples help illustrate. First, Jenkins (2016) argues that thought and talk involving “woman” should express the concept PERSON WITH A FEMALE GENDER IDENTITY rather than our present concept WOMAN. Arguably “woman” as traditionally understood fails to apply to transgender women and applies to transgender men. “Woman” understood with Jenkins’ meaning is intended to have neither of these supposed defects. To this extent, Jenkin’s proposal does not merely offer a new theory of what women are, or advocate for new behavior and dispositions to infer around the classification of women. It appears to involve “woman” acquiring an additional meaning, of the sort that might be registered in a comprehensive lexicon, more or less as Jenkins describes it: *person with a female gender identity*.

The second example comes from Barnes (2016)’s ameliorative account of disability. Admittedly, her own description of her project does not involve changing the meaning of anything:

I want to figure out what disability is. This is a project in social metaphysics—I’m not investigating what our word ‘disability’ means, nor trying to give a theory of our folk concept of disability. I’m asking *what it is* for something to be a disability. (2016, p. 10)

¹¹ See the OED entries for “myriad” and “girl” (<http://www.oed.com>).

¹² Meaning change in this sense is compatible with various ways linguistic meanings have been theorized: as external objects and properties, as Fregean senses corresponding to particular mental representations, or what must be known for competent use. It is also compatible with an externalist metasemantics, or a metasemantics where idiolects are pervasive, and even locally modulated. Talk of an expression “changing its meaning” may even be compatible with natural language lacking a genuine analytic/synthetic distinction: an expression may be said to “change its meaning” when there is some change in the canonical delineation of statements true in virtue of its meaning.

She summarizes her theory of physical disability as follows: “disability just is whatever the disability rights movement is promoting justice for” (43); it is not, as is commonly supposed, a matter of diminished bodily capacity. One might well wonder whether Barnes’s account of disability is faithful to the ordinary meaning of “disability.” Insofar as it isn’t, her account may be recast as conceptual revision, without too much loss of fidelity. As she says:

I am using ‘disabled’ rather than a replacement term like ‘differently-abled’. Words are hard to replace. I think it’s easier to shift meanings. (6)

The new meaning is preferable partly because it involves “a social category people have found useful when organising themselves in a civil rights struggle” (41). The hope is that using it more broadly (in place of “disability” with its unshifted meaning) will advance the disability rights movement. Other examples in the literature may be characterized in a similar spirit without too much loss of fidelity, including Dembroff (2016) on “sexual orientation”, Haslanger (2000, 2010) on gender and race terms, and Manne (2017) on “misogyny”.

Supposing this is right, these projects involve pushing the organic process of meaning change in the preferred direction. Successful implementation is not just a matter of whether, e.g., “woman”, acquires a particular secondary meaning. If it did, but, say, no one used $\langle \text{Woman}, C_{\text{Jenkins}} \rangle$ in the target areas of discourse, the project achieves its aims little better than if the word did not acquire the new meaning at all. (For this reason, focusing on whether, and when, we can cause a word to acquire a new meaning, is only a small part of the story required to address the feasibility question.) Success involves alterations in the mosaic of use, as happened with “shore”-meaning-*the whole beach* and “shore”-meaning-*the tidal zone* (and substantial enough alterations to achieve what the project set out to do).

Conceptual engineering, as advocated in the literature, acts on various targets. These include a word (or words) associated with a particular: semantic value (Cappelen, 2018; Koch, 2021a), Fregean sense, mental representation (Machery, 2017, 2021), speaker-meanings (Jorem, 2021; Pinder, 2021), valence, set of competence conditions, set of capacities it triggers and figures in (Haslanger, 2020), or particular linguistic function (Simion & Kelp, 2019). Discussions of feasibility has tended to constellate around the ease, or difficulty, of causing a word to be associated with a new meaning, or concept (in one of the senses above), given metasemantic principles.

I am suggesting this approach may be fruitfully set aside in favor of investigating how word-meaning pairs are, and are not, likely to be deployed. Likewise, to the extent, say, revising a mental representation involves the associated word changing its meaning in the uncontroversial sense “shore” and “myriad” have, then many of the complexities associated with mental representations, and linguistic meanings, may also be set aside for the present purpose. The feasibility of a project is partly a matter of how the distribution of certain word-meaning pairs is, or is not, likely to change across the total mosaic of usage (in the uncontroversial sense of meaning). This point generalizes to other targets of conceptual engineering.

For example, consider Pinder (2021)’s position that conceptual engineering involves the project of designing, and deploying, new speaker meanings, and, so, individual

speakers face few barriers to implementing “conceptual revisions”. Speaker meanings that become commonplace, rather like systematic variations in valence or force, tend to be recognized colloquially in the way a dictionary entry might. Theorists can dispute whether, say, entry 1a and entry 1b are truly distinct meanings, or whether this and that utterance are synonymous, but that needn’t be resolved to appreciate some of the constraints on how words tend to change their meaning in the sense in which “shore” and “myriad” have. To the extent the aims of an engineering project require promoting shifts in the use of a word paired with a particular speaker-meaning, shifts commonplace enough to be describable as a new secondary meaning, the feasibility question remains unanswered. Identifying what word-speaker-meaning pairs are likely to catch on in usage can be approached via what word-meaning pairs (in the intuitive sense of meaning just discussed) are likely to catch on, and replace which others.

More generally, consider various $\langle W, X \rangle$ pairs, where X is a semantic value related to W by an externalist metasemantics, or perhaps an internalist one, or a semantic value subject to frequent modulation (Ludlow, 2014), or a Fregean sense corresponding to a mental representation with a particular structure, an ideation, a particular force, a particular valence, or a package of dispositions, or capacities, or linguistic roles. So long as the project promotes shifts in the use of some $\langle W, X \rangle$ on the scale describable as W being deployed with a new (or new-ish) secondary meaning (in the sense above), the feasibility of the project may be approached via identifying which sorts of word-meaning pairs are likely, or unlikely, to catch on in the relevant ways. And that, I will now argue, can be done—at least partly.

3 Feasibility and rippling effects

Conceptual engineering is not impossible, because it has actually happened, at least in the sense that there have been premeditated changes to English that have stuck. Talk of the “disabled” has mostly replaced talk of the “handicapped.” Words like “housewife” have gradually given way to less gendered alternatives. The use of “queer” has lost in many contexts its negative connotations. Guatemalans and Nicaraguans are less likely nowadays to be called “Mexicans” by white Americans; they are likelier to be called “Latinos,” and increasingly likely in some circles to be called “Latinx”. No committee decided to implement these changes, but they were motivated, not accidental. Speaker behavior is difficult to regulate, and highly complex in many areas of discourse. But this gives us no reason to believe that more coordinated attempts at engineering would not succeed.

We need to understand the mechanisms by which small interventions may propagate. Introducing a few cane toads into Australia in 1935 created an environmental catastrophe decades later. Introducing a few penguins would not have had that effect. What’s needed is a strategy to assess which engineering projects are worth attempting, given their comparative odds of success and the ends they stand to promote. A felt need for linguistic revision, by itself, is not enough. For instance, despite a long-recognized need for an English epicene pronoun, only recently has singular “they” gained purchase. Dozens of candidates were proposed in the nineteenth century, including “hiser”

and “thon”, with little success. Horace Greeley (founder of the Herald Tribune) reputedly offered a cash prize to whoever identified one people would actually use. Singular “they”, of course, has been tripping off tongues for centuries, including Shakespeare’s and Jane Austen’s, but social priorities and grammatical pieties had to shift before it achieved widespread use. Sporadic idiosyncratic uses do not tend to reproduce themselves, any more than most stipulative ones do. Only sometimes they do.

The contemporary linguistics literature on semantic change is small but growing.¹³ It is driven in part by the booming interest in grammaticalization in cognitive linguistics (Sweetser, 1990) and Neo-Gricean approaches to the semanticization of pragmatics, and inference (Horn, 1984, 2011; Horn & Kleinedler, 2000). Other attempts to theorize semantic change loom large in sociolinguistics (Eckert, 2006; Labov, 1994), game theoretic semantics (Deo, 2015), and lexical semantics (Geeraerts, 2009).¹⁴ Some efforts have been made to extend relevance theory and optimality pragmatics to the study of semantic change.¹⁵

Drawing on some of this literature, I’ll now suggest a way to approach meaning replacement that allows the identification of projectable, even illuminating, generalizations. Then I’ll elaborate on some generalizations of particular relevance to conceptual engineers.

4 Generalizing about semantic change

A common sense understanding of development in a language has served as a theoretical model for other sorts of developments. Seeds of Adam Smith’s theory of the economy’s invisible hand may be found in his earlier work on linguistic change (1761/1983). In *The Descent of Man*, Charles Darwin approvingly cites the linguist Friedrich Müller as having remarked:

“A struggle for life is constantly going on amongst the words and grammatical forms in each language. The better, the shorter, the easier forms are constantly gaining the upper hand, and they owe their success to their own inherent virtue.”
To these more important causes of the survival of certain words, mere novelty may, I think, be added; for there is in the mind of man a strong love for slight

¹³ There is a substantial older literature focused primarily on tracing changes associated with word forms stemming from German and French linguistics in the nineteenth century, e.g., Bréal (1900), and Paul (1888/1891). For contemporary work on historical and comparative linguistics, see Blank and Koch (1999), and Janda and Joseph (2003).

¹⁴ Another relevant strand of research involves iterated learning frameworks (Kirby 2001).

¹⁵ Optimality pragmatics formalizes and extends the Gricean principles of cooperative communicative behavior found in Horn (1984) and Levinson (2000). For example, a principle of strength generates a preference for readings that are informationally stronger, a principle of consistency generates a preference for interpretations that do not conflict with the extant context, a principle of faithfulness generates a preference for interpretations of the utterance that do not leave out any of what the speaker says. The interaction of such constraints, founded on such heuristics, explains how the hearer arrives at the intended interpretation. At the same time, this model can be regarded as producing default, presumed interpretations. For applications of such an approach to explaining semantic change, see Falkum (2007), Sperber and Wilson (1986), and Wilson and Carston (2006).

changes in all things. The survival or preservation of certain favoured words in the struggle for existence is natural selection. (1871, pp. 60–61)¹⁶

In economics and evolutionary biology, we have systems that are chaotic at the micro level. The behavior of individual buyers and beetles is unpredictable in part because they are subject to unsurveyably many influences. Nevertheless there are stable patterns at the macro level. Other things being equal, inexpensive products tend to sell better than more expensive ones, and well camouflaged beetle species tend to outcompete worse camouflaged ones. Economics and evolutionary biology are largely concerned with these stable macro patterns, in some cases with an eye towards influencing micro behavior. Taxing a product, for instance, is one way of inducing buyers to prefer an untaxed competitor.

Our use of words is undeniably chaotic at the micro level, the more so if usage is construed broadly to include spoken applications of a particular word, thoughts involving it, as well as inferences, presuppositions and associations triggered by the word's application. Usage is sensitive to personal factors, including what a speaker believes about the meaning of words and her interests. It is also sensitive to her physical and social environment, including when the word in question was last used, how it was used, by whom, with what intonation, and under what circumstances.

But this is no obstacle to identifying stable patterns at the macro level. Are there laws of semantic change, akin to what we find in biology? There are certainly generalizations that have some predictive power and obvious partial explanations.

Lass (1980) issued a last gasp of skepticism that semantic change was simply too chaotic to be studied. He argued, roughly, that, because there are not exception-less laws of history or psychology, projectable regularities in semantic change cannot be identified, only retrospectively described. But he has since given up this position (Lass, 1997). What remaining controversy there is among linguistics is mostly beyond the scope of relevance to conceptual engineers.¹⁷ Phonology too was once thought too chaotic to be studied, but has since become a scientific discipline. Even if semantic

¹⁶ Müller was vehemently anti-Darwinian in the sense that he thought no selection process could have produced articulate humans out of inarticulate apish proto-humans. That said, he did think that a selection process of sorts had operated among the roots themselves, winnowing down an initially large set to the smaller number that in turn formed the basis of the major language groups. Darwin used evidence of selectional change among languages to defend his view that language use promoted the selection of humans over primates, and, debatably, that so-called primitive human races should speak evolutionarily lower languages than so-called civilized races (Radick 2002).

¹⁷ Linguists have painstakingly accumulated much data about—and proposed explanations for—cross-linguistically attested pathways of semantic change. [Blank (1999) offers a useful summary from the perspective of historical linguistics.] Usually these pertain to the descriptive semantics of a particular group of words, not the properties of form-meaning pairs that make them more or less prone to semantic changes. For example, body part terms tend to acquire spatial means, and explanations proposed for the development of body-part terms into spatial terms cannot necessarily be generalized to words of other semantic classes. Apart from disagreements about why particular generalizations obtain, there is disagreement about whether some such tendencies hold universally, or are “unidirectional”, and, if so, why. Some candidates involve perceptual vocabulary and deontic modals. For instance, it's cross-linguistically attested that perceptual vocabulary tends to acquire broader attitudinal meanings (“I see your point, feel pleased, am touched”), and deontic modals tend to acquire epistemic meanings (“could” and “must”). But it's disputed how much our cognitive architecture—including our intuitive feeling that “see” (as opposed to “kick” or “smell”) is suited to figure in metaphors for knowledge—is responsible. See Sweetser (1990), and Urban (2015).

change is not nearly as regular, or predictable, as phonological change, there are still useful generalizations to be had about what sorts of semantic changes are (and aren't) likely.

Let's dwell on this for a moment. Why is the use of "niggardly" in decline? This is clearly not a coincidence. Speakers tend to avoid expressions that are easily confused with slurs (something far more likely for "niggardly" now than 150 years ago). Likewise, it is not surprising that "myriad" acquired the meaning *quite a lot*. Terms for large quantities often go imprecise in this way after becoming popular stock metaphors (see, "a ton" and "a lot", once the quantity of something, e.g., hay, contained on a parcel, or lot, of land). It stands to reason that easily pronounced abbreviations would replace over time the longer terms that they abbreviate ("car" came from "carriage"). Language associated with youth culture tends to become popular. Speakers tend to speak in ways that fit with how they want to be seen: funny, hip, sophisticated or of the people.

Certain types of expressions have a tendency to pejorate (acquire more negative secondary meanings) rather than ameliorate (acquire positive secondary meanings). Euphemisms, for instance, are likely to pejorate (see the chronology tracked in the OED entry for "toilet"). But so are terms for the average ("mediocre," "mean," "common," "vulgar") and the inexpensive ("cheap" but also "vile" and "shoddy"). When an expression acquires a more negative meaning, often it stops being used with its original meaning. For instance, the use of "notorious" to mean *famous or well-known* fell off quickly when the word acquired scandalous overtones (in the seventeenth century). Why? It could be as simple as this: we try not to speak in ways that are easily misinterpreted, particularly to untoward effect.

Generalizations of the foregoing sort are well attested. But they cannot be expected to hold universally. For example, expressions tend on the whole to acquire broader rather than narrower meanings, and more abstract rather than more concrete ones. But sometimes it works in reverse. ("Business" first meant *anxiety*, then *purposeful activity*. Only relatively recently did it acquire the meaning of *occupation or trade*. Or take "segue.")

How then are we to extract generalizations relevant to those hoping to promote meaning replacement? One strategy is experimental. Linguistic innovations often originate and attain stable use in sub-communities prior to spread to the wider public. One sober-minded approach involves studying sub-communities for innovations and promoting those that seem promising. Another involves beta-testing several proposals by seeding them in small communities and promoting more widely those that catch on. Perhaps philosophers should see themselves as promoting experiments in new usage within counter publics. But more can be said.

The fact that we speak so as to achieve certain goals is critical to understanding language change. These goals are sometimes shared—coordinating our plans—and sometimes not—social advancement. In pursuit of these goals, speakers observe various rough rules. The most familiar of them, at least to philosophers, are the Gricean maxims. Grice made a few observations about meaning change, e.g., that habitual implicatures are apt to find their way into lexical meaning, though that was not his main focus. Others, including Neo-Griceans, have tried to locate the causes of language change in the opposing interests of speaker and hearer. They have appealed,

for instance, to souped-up versions of the maxim of relevance (say no more than you must!) and the maxim of quality (say as much as you can!).¹⁸

Conversational goals are served not only by advancing the right propositions, but expressing them via the right words. A word-meaning pair is helpful to the extent that the speech behavior it enables is, for example,

- Convenient (easily done),
- Efficient (good ratio of energy expended to the value of the message conveyed),
- Clear (unlikely to be misunderstood),
- Smooth (unlikely to raise irrelevant issues or distract from the point),
- Transparent (makes clear what question is being addressed),
- Flattering (it casts the speaker in some desired light),
- Concordant (it fits in with the speech of one's chosen group).

A new word-meaning pair meeting these desiderata is likelier to gain a foothold than one that flouts them. There are of course exceptions—the behavior of individual speakers is unpredictable—but these are the types of linguistic innovations we'd expect from a macro-perspective to stand a chance of succeeding. Conversely, innovations are less likely to succeed in areas of discourse where their use is unclear, inefficient, discordant, distracting (or so on). Thus, our speaking in accord with the above desiderata generates patterns at the macro level.¹⁹

For instance, our tendency to clear expressions (enshrined in Grice's Maxim of Manner) is partly responsible for the fact that languages have a well-documented tendency to avoid homonymy, in particular confusing homonymy. In English there are a mere 1600–2000 homonyms and far fewer where audiences are unlikely to be able to recover the speaker's intended meaning from context.²⁰ A beloved example of homonymy avoidance occurred in twelfth century Gascon French: the terms “cattus”

¹⁸ This idea is developed in many ways. It is found in Zipf (1949) and Martinet (1952). More recently, Horn (1984), and then other Neo-Griceans have attempted to derive Grice's maxims from the hearer-oriented Quality-principle (what isn't said isn't, modulo the Relevance-Principle) and the speaker-oriented Relevance-principle (speak with minimal effort, modulo the Quality-principle). Cognitive linguists, like Geeraerts (1990), have attempted to explain the phenomenon of prototypicality as reflecting an equilibrium between these two opposing forces. Bidirectional optimality theorists, like Blutner (2000), model optimization of the linguistic output against a system of ranked constraints that evaluates form-meaning pairs. These ranked constraints include expressive optimality and interpretative optimality. The particular theoretical commitments of such approaches are irrelevant to the present purpose.

¹⁹ Limited work in this direction can be found in Neo-Gricean theories (including Wilson 2003; Wilson and Carston 2007) and cognitive theories of semantic change, metaphor and conventional inference (including Sweetser 1990). Both emphasize the role of the structure of the lexicon in explaining semantic change. For example, it's often observed that changes in words' meanings are due to a tendency of languages to avoid ambiguous form-meaning pairings, such as homonymy, synonymy, and polysemy. On the other hand, when related words are examined together, it has been observed that one word's change of meaning often “drags along” other words in the same semantic field, leading to parallel change (Lightfoot 1999). These seemingly contradictory patterns of change lead to the conclusion that if ambiguity avoidance is indeed a reason for semantic change, its role is more complex than initially assumed.

²⁰ This problem only arises when there is identity of sound, category (e.g., animal), subcategory (e.g., rooster), and register between the homonyms. Notice that a variety of linguistic and extralinguistic factors conspire to alleviate potential homonymy or polysemy, including grammatical gender (as in French *le foie*, *liver*, vs. *la foi*, *faith*), inflection (brothers/brethren, hanged/hung), word order (Fr. *pauvre homme* *unfortunate man* vs. *homme pauvre indigent man*), and orthography (*draft/draught*, *metal/mettle*, *knight/night*),

(*cat*) and “gallus” (*rooster*) merged to “gat.” The resulting ambiguity was highly inconvenient especially in a farming context. Two things quickly happened: speakers withheld “gat” from roosters and three other words acquired *rooster* as a secondary meaning—*aza* (*pheasant*), *begey* (*vicar*) and *put* (*chick*). When languages do tolerate homonymous expressions, their uses rarely clash on account of occupying disjoint regions of discourse. (Consider the differential interpretations of “sheet” on a yacht, at the printer’s, or at a linens shop).

Some macro patterns, including homonymy avoidance, are of particular relevance to engineering projects. The study of what meaning replacements are achievable is closely related to the study of what meaning replacements are likely to eventuate in the normal course. Both require investigating how convenient, clear, inconspicuous, hip, distracting, etc. uses of various word-meanings are likely to be in the linguistic environment and how their introduction is likely to shift in the linguistic environment. One sort of empirical question is how to shift the environment to favor of some engineering proposal, say by exposing the public to it in regular but inconspicuous ways, or alternatively in memorable and politically charged ones.²¹ Another sort of empirical question is, holding fixed how the public is exposed to a word-meaning pair, whether the drivers of use put any additional constraints on what pairs are likely to become popular. The question is what may be said—what macro generalizations may be identified—that directly bear on which meaning replacements are achievable. It turns out quite a lot.

Now that we’ve seen *how* to approach the topic and *that* there are results to be had, let’s turn to these more specific conjectures that bear on the project of engineering. Section 5 elaborates some generalizations, including a few we have already met—expressions are unlikely to be replaced by more fraught or opaque ones, usage tends to be efficient and avoid unclear homonymy—and applies them to engineering projects. There is some bad news. Section 6 offers some positive but highly circumscribed suggestions in view of the fact that meaning replacement predictably involves expressions being replaced by slightly broader ones or, occasionally, slightly narrower ones. Section 7 returns to the negative: more dramatic meaning replacements are usually the product of successive more subtle replacements.

5 Negative applications

Successful engineering requires introducing $\langle W, C_{NEW} \rangle$, thus eventually driving out $\langle W, C_{OLD} \rangle$. This faces (at least) four obstacles. I’ll use the proposals described earlier to illustrate, but the obstacles generalize to others.

The first obstacle we have already met: *homonymy avoidance*. It presents a problem for engineering because once a word-meaning pair is stably in use, a homonym, particularly a clashing homonym, is unlikely to replace it—even one that would otherwise be more desirable. But this is precisely what conceptual engineers most often propose

Footnote 20 continued

pronunciation (human/e). When all else fails, speakers use modifiers or other repair strategies to disambiguate underspecified items, as observed: *fair-sized/fair-minded/fair-haired*, “Funny how? Funny-strange or funny ha-ha?”).

²¹ Nimitz (2021) suggests implementing conceptual engineering by shifting social norms.

to do. As we see in the marketplace, it is difficult to destabilize a product that already stably occupies a large market share (say Hershey's chocolate bars) even by introducing a more desirable competitor. It may be that engineers are stacking the deck against themselves by endorsing homonymy-style proposals rather than neologistic ones. If their proposed homonyms do find stable use, it will likely be within a restricted area of discourse (which may be sufficient for some purposes).

The problem of homonymy avoidance suggests one of three outcomes: either the new word-meaning pair will not spread, or it will only attain use in restricted bands of discourse where that use is unambiguously clear, or the new word-meaning pair will spread but use of the old pair will be replaced by a non-homonymous synonym.

For instance, consider proposals to conceptually engineer "woman" like those found in Jenkins (2016) and Haslanger (2000). The phenomenon of homonymy avoidance suggests that <"femina", WOMAN_{NEW}> is more likely to replace <"woman", WOMAN> than is <"woman", WOMAN_{NEW}> in areas of discourse where the homonymous use would read as ambiguous.²² To the extent "woman"-has a female gender identity is already gaining a hold in usage, the phenomenon of homonymy avoidance predicts that it will be confined to areas of discourse where the usage is *not* read as ambiguous. Such areas may be circumscribed. For instance, they may include some avowals and debates about who counts as a woman, without including generic discourse about, say, what teas women tend to order.

The second problem is the problem of *enduring communicative desires*. Meanings that serve communicative ends persist even when the *words* used to express them do not. This is a problem in part because some of our communicative desires are orthogonal to social projects. In the late nineteenth century, a shift in American slang made mere mention of the traditional term ("cock") for unneutered male chickens unacceptably awkward. But the desire to discuss them was unaffected and so "rooster"—originally a term for *neutered* male chickens—was soon recruited. In sixteenth century England, a revolution in cheese-making technology expanded the kinds of cheese beyond what was traditionally available (a farm cheese). These new sorts of cheese were also called "cheese." Several new expressions for farm cheese quickly came to be used ("farm cheese," "cottage cheese," "unripe cheese," "greene cheese" as in what the moon is made of...). Speakers had a continuing desire to single out the soft traditional cheese in an efficient way.

If the goal is to have certain *meanings* expressed less, then what is needed may be not conceptual engineering, but something that addresses the desires, something which eliminates part of what drives expression of the problematic meaning. For the conceptual engineer this problem arises in both parts of meaning replacement. A new word-meaning pair that fails to serve our communicative agenda is unlikely to be used, and a familiar meaning that better serves our communicative agenda is likely to continue being expressed.

Consider a proposal like Dembroff (2016)'s: if implemented, talk of "sexual orientation" will undergo meaning replacement and expressions like "lesbian" and "straight" will drop from usage, replaced by novel orientation vocabulary, including expressions

²² Sterken (2020) argues that creating linguistic ambiguity can disrupt communication in useful ways. Whether or not this is so, for the present purposes it's enough to note that confusingly ambiguous expressions tend to be replaced by those less so.

for those *attracted to men who were assigned “male” at birth* and those *attracted to men who were assigned “female” at birth*. So long as queer women do not want straight men to fill their events, they will find some phrase that excludes them, likely one synonymous with *women who are attracted to women*. Likewise, so long as homophobia persists, the phrase “because so-and-so is gay” will sometimes be more succinctly explanatory of harms done than some more specific delineation of so-and-so and his proclivities. Sexual orientation vocabulary may not be ideal for the reasons Dembroff identifies, but the way forward is unobvious.

The problem of communicative desires is very general. It is that use of language is finely calibrated to what we desire to talk and think about *and* how we desire to do so (e.g., conveniently, as opposed to, say, grandly). That is, our use of language is efficient to our communicative desires: it reflects a good ratio of energy expended to the value of the total message conveyed. When speakers share communicative desires to a high degree—as among experts desiring explanations of agreed upon phenomena, or among friends who enjoy nicknaming—it may still be tricky to identify an expression that satisfies them better than those already in wide use. But it happens, all the time. The problem is that only some of our communicative desires are tied up in the projects the would-be engineer aims to promote (socially significant, or otherwise). Supplanting usage of one word-meaning pair with usage of another, whose meaning component better serves these projects, may be even more difficult than it appears for somewhat subtle reasons. This is particularly so for non-expert discourse.

To the extent conceptual engineering does not address the underlying forces, communicative and otherwise, that give rise to the target linguistic behavior, there’s reason to worry that circumstances engineers hope to change will continue to reproduce themselves, perhaps with the aid of other words. This is partly why Nimtz (2021) advocates shifting social norms as a means to conceptual revision. Machery (2021) argues that features of human cognition make revising certain concepts particularly difficult: the mind tends to think with them.

Let’s briefly return to “woman.” Plausibly, part of what drives use of “woman”-*woman* (whatever the contours of its meaning may be) is a need to describe or acknowledge someone quickly, on the basis of their presentational gestalt, in a way that requires few epistemic commitments. Pointing down the street, I might say, “The car is parked way down by that woman” or, “That woman, she’s a good doctor” and be well understood. Presumably, *having a female gender identity* and *being a person subordinated on the basis of perceived reproductive capacity* are highly correlated with someone’s overall look and behavior. But it seems unlikely that expressions for these properties may become more reflexively trigger-able by presentational gestalt than “woman”-*woman*, without loss of clarity. Suppose use of “woman”-*woman* drops suddenly (say, it is made illegal, or becomes taboo, or a recent homonym renders its use confusing). Synonyms that serve the same communicative desires are likely to enter usage swiftly, like expressions for roosters or farm cheese once did.

This leads us to the third problem: *loading*. Loaded words—words that are potentially offensive, uncomfortable, awkward, merely distracting, and especially words that are taboo—are avoided in favor of expressions that are less loaded. Speakers avoid expressions that cast them in the wrong light or set the wrong tone in a conversation. The phenomenon of taboo avoidance (and euphemistic substitution) is well

documented. Words may also exhibit a kind of guilt by association.²³ Where once we spoke of haycocks and stopcocks we speak of haystacks and faucets.²⁴ Uses of “mensuration” – once a common term for computing surface area—fell off as “menstruation” entered the public consciousness. (Computation and the horror of periods were unaffected.) Use of “gender” sometimes replaces use of “sex,” which has more explicit and clinical connotations.

Similarly, words that have pejorated—acquired more negative meanings—may tend to be avoided in favor of those that have not. When “notorious” acquired the meaning *notorious*, “notorious”-*well-known* was avoided in favor of other expressions. Likewise, words that have acquired a socially controversial meaning tend to be avoided in favor of those that have not. As “liberal” and “conservative” have become more prominent as political terms, their non-political uses have fallen off outside of fairly set formulae. (For instance, directions to “sprinkle liberally,” or talk of “conservative dressers”.)

“Queer” provides a particularly clear example. The precise origins of the term are unknown, but beginning in the sixteenth century it meant something like *peculiar*. By the seventeenth century it acquired stable use as a euphemism for *unwell* (itself a euphemism) and ultimately assumed that as its literal meaning. By the nineteenth century “queer”-*unwell* began to function as a euphemism for *drunk*, and soon acquired that meaning as well. Sometime around the turn of the twentieth century, “queer” began to function in euphemisms associated with homosexuality in select circles, eventually acquiring a string of related meanings. As these gained prominence in the later quarter of the twentieth century, these other uses of “queer” dropped off. This is partially attributable to homonymy avoidance, but not entirely. “He seems queer” threatens to be ambiguous, but “what queer weather” can only mean one thing. What may be going on here is that usages are avoided that evince a cluelessness about the present moment, or otherwise call to mind a charged topic when doing so may distract the audience.²⁵ Loading may also partly explain the phenomenon of reclamation, whereby in-group members begin to self-identify using a derogatory expression. After “queer” pejorated, it was reclaimed in activist circles where its, and load, affirmed shared identity. Yet

²³ The entry in the first edition of the OED in 1893 read: “the current name among the people, but, pudoris causa, not admissible in polite speech or literature; in scientific language the Latin is used.”

²⁴ Cicero (45BCE/1982) describes the phenomenon in a letter:

What you in your letter call by its own name [“mentula”] he with more reserve calls *penis*; but because so many people use it so, it has become as obscene as the word you used. [...] *Ruta* [path] and *menta* [mint]—we use both words without impropriety. I want to use the diminutive of *menta*, as one might say *rutula* [lit. little path]; it is not done [*non licet*]. (IX.xxii).

Examples need not be so phallogentric. See the OED entry for “coney” (once a word for juvenile rabbits like “kitten” is for cats). Today, thanks to the rise of a piece of vulgar slang, it survives only in “Coney Island” (with altered pronunciation).

²⁵ “Gay” has a related trajectory. It entered English meaning *merry*, acquired via euphemistic substitution the meaning *lascivious* (“gay halls” were brothels staffed mainly by women). Further euphemistic substitution resulted in acquiring the meaning *homosexual*. Once that usage was well known, use of “gay” with these other meanings dropped precipitously.

reclaimed, and derogatory, uses alike were thin outside these circles, where the various loads were unwelcome.²⁶

Loading afflicts some otherwise attractive meaning replacements. Insofar as they reveal hard or unexpected truths about the forces sustaining some classificatory practice, they threaten to distract the audience, or make the conversation unnecessarily awkward.

For example, PETA's "meat"-*murdered animal* is unlikely to spread beyond protests, or debates about vegetarian ethics, to school lunches, shipping contracts, or, deciding what to bring to a potluck. Why do people use the word "meat"? Sometimes they use it to say what they want (at a store, for dinner, etc.). But a word highly evocative of cruelty is likely to be avoided. To the extent PETA's "meat"-*murdered animal* colors "meat" with uncomfortable associations, euphemistic substitution seems a likely outcome. (How much shopping behavior will be affected along the way, is not well researched from the armchair.)

Proposals like Haslanger's for race and gender (e.g., "black"-*subordinated on the basis of perceive ancestry*; "woman"-*subordinated on the basis of perceived reproductive capacity*) may win the day in expert discourse concerned with intersectionality, but speakers are likely to avoid her proposed meaning reassignments at the water cooler and children's birthday parties. Well-wishers will not want to express thoughts like the following: "I hope you grow up to be a strong and powerful person subordinated on the basis of your perceived reproductive capacity."

The phenomenon of loading may suggest the following strategy. Introduce two expressions, where the first—say, a phrase that rhymes with "woman" but which is un-useable in most company—is a means to curtail an objectionable expression and the second is a replacement for the objectionable expression. But the problem of communicative desires tells against this strategy. If "woman" is driven out of use, the likely outcome is that a synonym replaces it, not an artificially supplied expression, let alone a homonym.

The fourth problem for conceptual engineering is the problem of *opacity*. Words that are unclear to audiences, such as technical terms and ones whose meanings are particularly difficult to recall, tend to be replaced by those that are clearer. That is, the primary meaning of a word tends to contract around the subject it is used most often to discuss. Take expressions ripe for folk etymological speculation, like "begging the question" and "inflammable." (We encounter the prefix "in" and negate what follows—along the model of "inactive" rather than "inhabit."²⁷) Such expressions tend to acquire a secondary meaning in line with popular belief ("begging the question" may already mean *raises the question*), if they are not replaced by a synonym (like "flammable").

Another kind of opaque term is a technical one that has migrated to ordinary discourse. Take the vocabulary originating in the wildly popular theory of the four humors developed by Hippocrates of Kos (460–370 BCE): "sanguine," "choleric," "phlegmatic," "melancholy," and "humor" itself. A 1728 compendium of arts and sciences

²⁶ Whether reclamation projects involve meaning replacement partly depends on the analysis of derogatory speech. See Nonberg (2018).

²⁷ So much so the American Fire Safety organization has discouraged labeling products as "inflammable" since the 1920s.

records that “Sanguine Constitutions require a frequent Use of Phlebotomy,” due to an excess (relative to the other three humors) of blood. But “sanguine” (and the others) as of the early fourteenth century acquired two non-technical secondary meanings each associated with a salient symptom of sanguinity: *ruddy* (as in “a sanguine complexion”) and *optimistic*.²⁸ As popular knowledge of the theory declined, use of “sanguine,” the medical term, further faded from ordinary discourse.

“Begging the question” and “sanguine” persist with their original meanings in specialized discourse, discourse where the original meanings are well known and transparent in conversation. These considerations suggest reasons to be hesitant about proposals to introduce theoretically rich word-meanings. Imagine an activist who believes the revolution will be hastened by causing the public to use “worker”-*exploited in virtue of selling his or her labor* largely in place of “worker”-*worker*. This word-meaning pair is unlikely to be adopted outside of Marxist discourse where the background theory is not only well understood but presumed true and conversationally relevant. To the extent it does gain everyday currency, it will be crowded out by more transparent word-meaning pairs unencumbered by Marxist theory.

A Barnes-style proposal for “disability” may fare well with respect to the problem of loading, but less so for the problem of opacity.²⁹ The revision is theoretically sophisticated, and so not conversationally transparent to the uninitiated in most areas of discourse. Even if the revision takes off, activists may need to constantly remind speakers what talk of “disability” is all about. Understandable folk-etymological speculation may lead audiences to negate what follows the “dis-” and take “disability” to mean a lack of relevant abilities. (Like words for what’s average, “disabled” also has a long history of pejoration). Perhaps a better alternative is to lexicalize *WHATEVER THE DISABILITY RIGHTS MOVEMENT IS PROMOTING JUSTICE FOR* with an acronym (along the lines of “LGBTQ”) and promote its widespread usage.

To the extent that the word “disability” is used most frequently among the disabled and those expressly concerned with disability rights and access, a proposal like Barnes’s could come to predominate in disability-focused subcultures (as Barnes suggests it may already have). It may even gain some foothold among analytic philosophers, but it is unlikely to *predominate* in theoretical discourse about disability, since one salient issue is precisely what disability *is*.

Opacity may partly explain the success of reclamations. Return to “queer”. Plausibly this well-publicized reclamation resulted in fewer derogatory uses of “queer” (though perhaps not of related expressions). One explanation is that public and confident uses of “queer” destabilized the belief among homophobes that statements involving the word “queer” would land with the right derogatory zing. Reclaimed uses of “queer”

²⁸ Sheridan in a 1735 letter to Swift: “Do not think me sanguine in this; for more unlikely and less reasonable favours have been granted.” See the OED entry for “sanguine.”

²⁹ See also Dougherty (2020)’s discussion of celebratory-based, rather than solidarity-based, ameliorative projects for disability.

changed the linguistic environment in a way that reduced derogatory uses of “queer.” (To the extent the political dimension of “queer” has since flattened out, this change may also be attributable to opacity). Perhaps the phenomenon of reclamation may shed light on whether “disability”-*whatever the disability pride movement is celebrating* is more likely than a Barnes-style proposal to gain currency in usage. The former might fare better with respect to loading considerations (unless it proves divisive).

The foregoing problems suggest some hurdles for the would-be conceptual engineer. They tell against proposals that involve introducing homonyms of stably used expressions. They tell against proposals that involve replacing a usage sustained by communicative desires by one that does not better satisfy those desires. They tell against proposals that involve broad use of a word-meaning pair that will paint someone in a political light that some people are keen to be seen in and others are loath to be seen in. Finally, they tell against proposals that require speakers to use more epistemic effort to successfully apply an expression. But the news isn’t all bad: there are patterns of meaning change that the engineer can exploit.

6 Positive applications

Often a word will acquire a broader or generalized meaning. For instance, use of generic “xerox” has mostly replaced use of “Xerox™” and “copier.” Occasionally a word’s meaning narrows. “Accident”, originally a neutral term for any happenstance, acquired the narrower meaning of *unfortunate happenstance*. This narrowed term replaced some negative uses of the original term (as the original term became loaded, it was avoided in favor of less loaded alternatives, like “coincidence,” “occurrence”). Broadening replacement mostly occurs in broader areas of discourse (the make of a copier is mostly irrelevant), while narrowing replacement mostly occurs in narrower areas of discourse (to some salient exemplar).

Narrowing and broadening sometimes operate in tandem. In Old English “dogca” referred to a particular breed of dog (thought to be a hunting dog rather like a mastiff). “Hound” denoted the entire kind *dog* (as German “hund” continues to do). Sometime in the fourteenth century, when Chaucer’s warning “It is nought good a slepyng hound to wake” was turning into Heywood’s “It is evyll waking of a slepyng dog,” “dog” presumably had both a narrow and a general meaning. Eventually “hound” (broad) was displaced by “dog” (broad). But a narrower use of “hound” remained in use among hunters who took them to be ideal representatives of the species, dogs par excellence. Today what we call “hounds” are very like what English speakers once called “dogges” and before that “dogca”—large slobbery dogs traditionally used for hunting. We began with a term for dogs, and a more specific term for slobbery hunting dogs, and still have them today, although what we then called “hounds” we now call “dogs,” and vice versa.

Cases of broadening and narrowing often involve some culturally salient subordinate or superordinate category and do not re-carve ordinary felt distinctions. For instance, semantic changes did not redraw the happenstance/bad-happenstance distinction.

But there are two strategies in this area that might be of help to the conceptual engineer. Broadening-type replacement *can* result in the expression of a novel meaning. Female-inclusive “guys” has replaced some uses of the female-exclusive term, but not all (“He’s going out with the guys tonight”). “Aunt” once exclusively referred to the sister of the father. Broadening of this sort may occur naturally in response to shifting social realities. “Selfie” originated as Internet slang for a self-portrait using a web camera. As cell phone technology improved, it broadened to include large group shots. Broadening may also be spurred on by collective effort, as when female-exclusive uses of “actor” began to be replaced by female-inclusive ones.³⁰

It might be feasible in a similar spirit to broaden “hat” to include earmuffs. This would be convenient and not particularly confusing (suppose in winter I say “everybody have their hats?” and you brought earmuffs). Earmuff-exclusive “hat” might persist among milliners and in settings where the distinction between traditional hats and earmuffs is relevant, but in ordinary settings the hat/not-hat distinction will be effectively redrawn.

If there are significant engineering projects that involve gradual broadening of a similarly convenient sort, that is a mark in their favor. There may well be. Today those with a Jewish father, but not a Jewish mother, are now often considered “Jewish.” To the extent the primary meaning of “family” excluded non-heterosexual-bionormative-nuclear families, but no longer does, “family” has broadened along these lines.³¹

But few engineering proposals with a revolutionary flavor will involve broadening or narrowing in this sense. Consider a proposal like Manne’s ameliorative account of misogyny, one that proposes “misogyny”-meaning-*the law enforcement arm of the patriarchy* replace some uses of “misogyny” with its ordinary meaning. This goes beyond simply endowing “misogyny” with a broader meaning in the sense discussed above. “Misogynist” (like “racist”) has *already* broadened in that way. Today we may say a university is misogynist when its policies are *as if* made and applied by misogynists—those, in the narrower sense, who have bigoted attitudes or animus towards women. (One might conjecture this occurred as the desire to address institutional sexism grew and the perceived gap between perpetuating certain harmful practices and being culpable for doing so shrank.) “Rewinding” has broadened in a similar way, though for different reasons. Where once it meant winding back physical tape, now it means whatever turns back the recording. If Manne’s proposal is offered only to those interested in theorizing about patriarchy, it may well succeed, but her ambitions seem to be greater.

A second strategy for the conceptual engineer takes a particular sort of narrowing as a model. When “bitch”-*female dog* entered English, there was no sex specific mate.

³⁰ The original meaning of “actor” was gender-neutral, although almost all players were male. “Actress” (and most other gendered occupational terms) entered English in the eighteenth century as part of American efforts to standardize English along Latinate lines.

³¹ Haslanger (2020), by contrast, suggests the changes undergone are consistent with a pattern of dual character concepts: terms that have one meaning associated with some activity, another with some virtue or underlying characteristic that disposes individuals to participate in that activity. Whether anything so sophisticated is required to describe what happened to “family” is unclear. Outside of some highly contentious circumstances, it was clear and, above all, convenient and polite to call non-HBNFs “families,” more so than any alternative. Exclusive “family” was replaced by subtly broader inclusive “family.”

As a result the general term “dog” developed narrower meaning designating the males of the species. (These sex-specific terms have since pejorated). The existence of the more informative term (“bitch”) together with the choice by an informed speaker to employ a less informative term (“dog”) in a context where the additional information would have been relevant licenses the inference that the speaker was not in position to employ the more informative term. The conventionalizing of this inference is thought to have contributed to “dog” acquiring the meaning *dog* (excluding bitches). Some uses of the broader “dog” were replaced by the sex-specific term, notably in the limited areas of discourse where the sex of dogs was at issue.³²

It may be the case that the best way to get “woman” to mean what one wants it to mean is indirect. It may involve introducing some other term that means the complement of the target meaning of “woman.” For instance, “lady”-*benefitted from her perceived reproductive capacity*. Then “woman” may acquire a complementary meaning in the way “dog” did. But “woman” with its ordinary meaning is likely to dominate use.

In a similar spirit, perhaps an environmental activist might promote use of “recyclables” and “compostables” as a means to narrow the meaning of “trash” and “garbage” (to exclude recyclables and compostables). To the extent calling a recyclable or compostable “trash” is roughly as uncomfortable as calling a thumb “a finger” (that is, somewhat in some situations), these terms may replace some uses of “trash.” If calling something “recyclable” conduces to actually filing it in the recycling bin, a project like this might enhance garbage-sorting compliance.

Diachronic semantic phenomena may be of interest to social justice minded projects other than engineering. Take the well-documented phenomenon of synonymy differentiation. Just as languages tend to avoid homonyms, they tend to avoid true synonyms. Often this avoidance is attributed to hearer-based effects: use is sensitive to what audiences are likely to understand, which is conditioned by what sorts of speech they have recently been exposed to. This in turn foments synonyms accreting increasingly disparate associations before diverging. The size acceptance movement’s attempts to reclaim “fat” (in part to increase value-neutral vocabulary for body size) may succeed in promoting further semantic differentiation between “fat” and, say, “overweight.” The latter is more clinical and more susceptible to folk etymology, and so more likely to pejorate after a period of euphemistic substitution.

7 Meaning replacement: predictably gradual or unpredictably dramatic

Some meaning replacements are not a matter of narrowing or generalization. But these are unlike what engineers hope to achieve, or else are unpredictable. “Shore” acquired a new meaning because our communicative desires shifted to involve the whole beach, but that shift was driven by unforeseeable technological changes. Sometimes an

³² There is no shortage of similar cases. When “hue” entered English, “color” acquired a hue-exclusive meaning. When “thumb” entered English, “finger” acquired a thumb-exclusive meaning. See also “cow” (excluding bulls), “rectangle” (excluding squares), “gay” (excluding lesbians).

expression figuring in a stock metaphor or metonymy is replaced by one that has acquired the related literal meaning. (Calling a glass drinking-vessel “a glass” was once metonymic). But the replacements that conceptual engineers tend to favor are not obviously achievable via stock figuration.³³

Dramatic meaning replacements are often unpredictable precisely because they involve a series of successive more predictable narrowings, broadenings, or literalized figurations. Take the case of “livid.” Today’s livid bruises are flushed purple-red, a vivid or angry sort of color. Some, but not all, may have a similar cast to “a bruise quite livid” in 1501 (see the OED entry for “livid”), when “livid” meant *black-and-blue*.³⁴ Plausibly some modern uses of “livid” have replaced some uses of expressions meaning black-and-blue. Our reflexive classification of bruises have shifted, but only somewhat.

But this rather subtle effect is the product of the semantic profile of “livid” undergoing a series of changes, not one. When “livid” entered English, it functioned more or less as a synonym of “black-and-blue.”³⁵ Webster’s Word Histories (1989) recounts:

A slight extension of meaning had by the end of the eighteenth century given it the sense of “ashen” or “pallid,” as in describing the appearance of a corpse. “Livid” eventually came to be used in this sense to characterize the complexion of a person pale with anger; such as, “livid with rage.” In the twentieth century...because of association with words like *lurid* and *vivid*, and in part because an angry person is at least as likely to be red-faced as pallid, *livid* has acquired the sense “reddish.” Its frequent occurrence in phrases like “livid with fury” has also given rise to a sense entirely unrelated to color, with *livid* now commonly functioning simply as a synonym of *furious* or *enraged*.

Described from 30,000 feet, the sorts of engineering proposals most attractive to philosophers require dramatic semantic changes to which there is no predictable path via a series of smaller changes. On the whole, a term is unlikely to acquire a secondary meaning broader than its original meaning in some ways, but narrower in others—at least not in one fell swoop. For example, perhaps “woman” and “man”, as deployed in wide swaths of discourse, will first acquire (or already have acquired) secondary meanings that are broader in some respects, and not mutually exclusive (where some people are both men and women) before acquiring further secondary meanings that are also narrower in other respects (and re-carve the man/woman distinction in the way some trans activists seek).

³³ It is a striking fact that native speakers tend to broadly agree about which figurations are salient. To the extent the proposed word-meaning does not *appear* to mean literally what some salient use of that word with its ordinary meaning may be said to mean figuratively, a meaning replacement via figuration is unlikely to occur.

³⁴ A more socially significant case might be that of “usury.”

³⁵ The word entered English as a loan from the Middle French “livide”, itself deriving from the Latin adjective *lividus* (meaning *dull, greyish, or leaden blue*).

8 Conclusion

The engineering projects that are advertised as producing the best results face serious obstacles to implementation. As we saw in Sect. 5, such projects face the problems of homonymy avoidance, enduring communicative desires, loading and opacity. As we saw in Sects. 6 and 7, meaning replacement is predictably gradual and unpredictably dramatic. All is not lost: meaning replacements that involve broadening or indirect narrowing (or figurative extension) may be feasible. Even so, these lack the revolutionary feel, and theoretical import, likely to appeal to philosophers.

Those serious about meaning replacement as a tool, say, for effecting social change, must recognize that the devil is in the (empirical) details. There is an applied science of how to increase a widget's share of the market. So, too, there could be an applied science of how to increase a word's share of use. It would be rooted in facts about how word usage evolves over time, identifying more fine-grained generalizations than I have here. Whether developing such a science is ultimately of interest to philosophers partly depends on how seriously would-be conceptual engineers and ameliorators take their own projects.

Nothing I've said tells against engineering proposals catching on among a small band of like-minded speakers with unified aims, and homogenous circumstances. Carnapian explication projects, undertaken for circumscribed theoretical projects, often do succeed. But conceptual engineers may seem to want more. For example, many of the social conditions one may hope to change are facts about the broader community. Pay gaps, equity in domestic labor, and the number of accessibility ramps are unlikely to be affected by usage in an erudite enclave, or confined areas of discourses. Even so, some significant social conditions—for instance, how respected, valued and included some members of a community feel—are affected by such usage. It's up to engineering enthusiasts to calibrate the revision proposed in an area of discourse and the aims that allegedly justify it (whatever they may be). One question is how many of the considerations that matter to activists supervene on the concepts deployed as opposed to the broader dispositions of those deploying them.

There may be reasons to go in for innovating new concepts, apart from their likelihood of catching on and promoting substantial change. Proposing novel concepts may raise consciousness, provoke dialogue, change our dispositions. It may perform a gestalt shift on audience members. But if engineering is worthwhile for these reasons only, it stands accused of false advertising. *Bringing about* conceptual revision was supposed to be a means to a better world, when, in fact, it is *proposing* conceptual revision that is the means to a better world. (Concepts, after all, are not typically analyzed as sets of dispositions, or gestalts.)

Studying which words with what meanings are likely to thrive in a better world may be worthwhile, but it risks being like ideal biology, studying the species that would thrive in a better ecosystem, or like science fiction. Jules Verne, while prescient, didn't invent the solar sail even if he inspired it. An expression's potential use in a brighter future is not by itself a reason it can realize that potential today. This observation generalizes to expressions promoted as tools to facilitate theorizing.³⁶ Identifying a

³⁶ Say, on the basis of their expected explanatory utility, see Carballo (2020).

replacement that *would* figure in fruitful explanations, if used, is not yet a reason the replacement is feasible. One reason a replacement is not feasible is if speakers, on the ground, experience it as subject changing, diverting from their concerns. Studying what word-meanings should be used in broader discourse, never mind constraints driving that discourse, is perhaps rather like studying what the equilibrium pricing of milk should be, never mind the supply and demand. It may be the lead question is not even well formed.

Things do not look good for ambitious conceptual engineering projects of the sort hopes most ride on. They risk being like semantic perpetual motion machines, or blueprints for a bridge elegant in theory but unbuildable in practice. Correcting our vocabulary might have beneficial effects but is vastly harder than is usually presumed.

The difficulty does not arise from the Wittgensteinian observation that usage is unsystematic, and obscurely related to meaning. It also does not arise from the observation that usage is difficult to regulate, or the nature of metasemantics. While meaning change is somewhat unpredictable and uncontrollable, the main difficulty for conceptual engineering arises from the observation that a word's usage is shaped by various pressures which exert themselves in foreseeable ways. More may be said about how a term, e.g., "woman", is likely to be used, beyond application to women, disquoted. Because meaning follows usage in some way, meaning change is somewhat predictable.

Collective action can, and does, work. Strikes do improve working conditions. Perhaps volitional collective changes to linguistic behavior, including conceptual engineering, can be effective. But linguistic behavior is not akin to striking. It is not what improves working conditions, credibly threatening the boss's economic position does that. Successful strikes shape workers' material circumstances by incentivizing those with the power to change them. On its face, changing linguistic behavior does not change many of the circumstances engineers may hope to shift, nor does it change the underlying web of circumstances, and pressures, that give rise to them. Deploying concepts inconvenient for defending abusive working conditions does not shift the incentives that maintain those conditions, or inhibit defenses of them—words will be found. Even so, conceptual engineering may well be a useful tool for adjusting collective circumstances, when deployed with an eye towards opportunity costs, and what changes to linguistic behavior are likely to catch on.

Acknowledgements For conversation, and comments, thanks to Alex Byrne, Michael Della Rocca, Cosmo Grant, Sally Haslanger, Smriti Kanal, Jon Koslow, Katy Meadows, Erin Miller, Rebecca Millsop, Kate Ritchie, Susan Robbins, Kieran Setiya, Jack Spencer, Michael Thornton, Jonathan Vogel, Mallory Webber, and Steve Yablo, as well as audiences at NYU Foundations of Conceptual Engineering Conference 2018, MIT/Michigan Social Philosophy Workshop 2018, and WOGAP 2019.

Authors' contributions In this limited literature on this topic, it is often supposed that because use is inexhaustible, and its relation to meaning unsurveyable complex, it, along with meaning change, is unpredictable and uncontrollable. Sometimes this observation is seen as telling in favor of the feasibility of conceptual revision, sometimes against, and often contingent on the preferred metasemantic theory. This paper gives reasons for thinking both moves are too quick. As far as I know, there is no work in any philosophical journal that tries to directly answer which engineering proposals are feasible or draws extensively on historical linguistics and pragmatics.

Funding N/A.

Availability of data and materials N/A.

Code availability N/A.

Declarations

Conflict of interest Not applicable.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Barnes, E. (2016). *The minority body: A theory of disability*. Oxford University Press.
- Blank, A. (1999). Why do new meanings occur? In A. Blank & P. Koch (Eds.), *Historical semantics and cognition*. Mouton de Gruyter.
- Blank, A., & Koch, P. (1999). *Historical semantics and cognition*. Mouton De Gruyter.
- Blutner, R. (2000). Some aspects of optimality in natural language interpretation. *Journal of Pragmatics*, 17(3), 189–216.
- Bréal, M. (1900). *Semantics: Studies in the science of meaning*. Translated from original by N. Cust. London: William Heinemann.
- Cappelen, H. (2018). *Fixing language: An essay on conceptual engineering*. Oxford University Press.
- Cappelen, H., & Plunkett, D. (2020). Introduction: a guided tour of conceptual engineering and conceptual ethics. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Carballo, A. P. (2020). Conceptual evaluation: epistemic. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Cicero, M. T. (45BCE/1988). *Epistulae ad familiares*. In W. Watt (Ed.), Oxford University Press.
- Darwin, C. (1871). *The descent of man*. J. Murray.
- Dembroff, R. (2016). What is sexual orientation? *Philosopher's Imprint*, 16(3), 1–27.
- Deo, A. (2015). Diachronic semantics. *The Annual Review of Linguistics*, 1(1), 179–195.
- Deutsch, M. (2020). Speaker's reference, stipulation, and a dilemma for conceptual engineers. *Philosophical Studies*, 177(12), 3935–3957.
- Deutsch, M. (2021). Still the same dilemma for conceptual engineers: Reply to Koch. *Philosophical Studies*, 178(1), 1–12.
- Dougherty, T. (2020). Disability as solidarity: Political not (only) metaphysical. *Phenomenological and Philosophical Research*. <https://doi.org/10.1111/phpr.12666>
- Eckert, P. (2006). Variation, meaning, and social change. In N. Coupland (Ed.), *Sociolinguistics: Theoretical debates*. Cambridge University Press.

- Falkum, I. (2007). A relevance-theoretic analysis of concept narrowing and broadening in English and Norwegian original texts and translations. *Languages in Contrast*, 7(2), 119–142.
- Fischer, E. (2020). Conceptual control: On the feasibility of conceptual engineering. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1773309>
- Geeraerts, D. (1990). Homonymy, iconicity, and prototypicality. In D. Geeraerts (Ed.), *Diachronic semantics*. John Benjamins Publishing Company.
- Geeraerts, D. (2009). *Theories of Lexical semantics*. Oxford University Press.
- Haslanger, S. (2000). Race and gender: (What) are they? (Do) we want them to be? *Noûs*, 34(1), 31–55.
- Haslanger, S. (2010). Language, politics, and “the folk”: Looking for “the meaning” of ‘race.’ *The Monist*, 93(2), 169–187.
- Haslanger, S. (2020). Going on, not in the same way. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Horn, L. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, form, and use in context: Linguistic applications*. Georgetown University Press.
- Horn, L. (2011). Etymology and taboo. In archive of conference talks, International Society for the Linguistics of English. <http://www.bu.edu/isle/isle-2-archive/>.
- Horn, L., & Kleinedler, S. (2000). Parasitic reference vs. R-based narrowing: lexical pragmatics meets he-man. *Annual Meeting of the Linguistic Society of America*. https://ling.yale.edu/sites/default/files/files/horn/HornKleinedler2000_he-man.pdf.
- Janda, R., & Joseph, B. (2003). On language, change and language change—or, of history, linguistics, and historical linguistics. In R. Janda & B. Joseph (Eds.), *The handbook of historical linguistics*. Blackwell.
- Jenkins, K. (2016). Amelioration and inclusion: Gender identity and the concept of *Woman*. *Ethics*, 126(2), 394–421.
- Jorem, S. (2021). Conceptual engineering and the implementation problem. *Inquiry*, 64(1–2), 186–211.
- Kirby, S. (2001). Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computations*, 5(2), 102–110.
- Koch, S. (2021a). The externalist challenge to conceptual engineering. *Synthese*, 198, 327–348. <https://doi.org/10.1007/s11229-018-02007-6>
- Koch, S. (2021b). There is no dilemma for conceptual engineering. Reply to Max Deutsch. *Philosophical Studies*, 178, 2279–2291.
- Labov, W. (1994). *The principles of language change*. Blackwell.
- Lass, R. (1980). *On explaining linguistic change*. Cambridge University Press.
- Lass, R. (1997). *Historical linguistics and language change*. Cambridge University.
- Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.
- Lightfoot, D. (1999). *The development of language: Acquisition, change, and evolution*. Blackwell.
- Ludlow, P. (2014). *Living words: Underdetermination and the dynamic lexicon*. Oxford University Press.
- Machery, E. (2017). *Philosophy within its proper bounds*. Oxford University Press.
- Machery, E. (2021). A new challenge to conceptual engineering. *Inquiry*. <https://doi.org/10.1080/0020174X.2021.1967190>
- Manne, K. (2017). *Down Girl: The logic of misogyny*. Oxford University Press.
- Martinet, A. (1952). Function, structure, and sound change. *Word*, 8(1), 1–32.
- Merriam-Webster, Inc. (1989). *Webster's word histories*. Merriam-Webster.
- Nimt, C. (2021). Engineering concepts by engineering social norms: Solving the implementation challenge. *Inquiry*. <https://doi.org/10.1080/0020174X.2021.1956368>
- Nonberg, G. (2018). The social life of Slurs. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts*. Oxford University Press.
- Paul, H. (1891). *Principles of the history of language*. Translated from 2nd Ed. by H. A. Strong. Longmans, Green Co. Original edition, 1886.
- Pinder, M. (2021). Conceptual engineering, metasemantic externalism and speaker-meaning. *Mind*, 130(517), 141–163.
- Pollock, J. (2021). Content internalism and conceptual engineering. *Synthese*, 198, 11587–11605.
- Queloz, M. (2021). Conceptual engineering and the politics of implementation. *Pacific Philosophical Quarterly*. <https://doi.org/10.1111/papq.12394>
- Radick, G. (2002). Darwin on language and selection. *Selection*, 3(1), 7–16.
- Sharp, K. (2013). *Replacing truth*. Oxford University Press.

- Saul, J. (2006). Gender and race. *Aristotelian Society Supplementary*, 80(1), 119–143.
- Simion, M., & Kelp, C. (2019). Conceptual innovation, function first. *Noûs*, 54(4), 985–1002.
- Smith, A. (1761/1983). Considerations concerning the first formation of languages, and the different genius of original and compounded languages. In J. Bryce (Ed.), *Smith's lectures on rhetoric and belles lettres*. Oxford University Press.
- Sperber, D., & Wilson, D. (1986). Loose Talk. *Proceedings of the Aristotelian Society Supplementary*, 86(1), 153–172.
- Srinivasan, A. (2019). Genealogy, epistemology and worldmaking. *Proceedings of the Aristotelian Society*, 119(2), 127–156.
- Sterken, R. (2020). Linguistic interventions and transformative communicative disruption. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Sweetser, E. (1990). *From etymology to pragmatics: metaphorical and cultural aspects of semantic structure*. Cambridge University Press.
- Urban, M. (2015). Lexical semantic change and semantic reconstruction. In C. Bowers & B. Evans (Eds.), *Routledge handbook of historical linguistics*. London: Routledge.
- Wilson, D. (2003). Relevance theory and lexical pragmatics. *Italian Journal of Linguistics*, 15(2), 273–291.
- Wilson, D., & Carston, R. (2006). Metaphor, Relevance and the “emergent property” issue. *Mind & Language*, 21(3), 404–433.
- Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics*. Palgrave-Macmillan.
- Zipf, G. (1949). *Human behavior and the principle of least effort*. Addison-Wesley Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.