

**Meaning Change, in Theory and in Practice**

by

Allison Robbins Koslow

B.A. Philosophy  
Yale University, 2010

Submitted to the Department of Linguistics and Philosophy  
on September 8, 2020 in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Philosophy  
at the  
Massachusetts Institute of Technology

September 2020

© Massachusetts Institute of Technology 2020. All rights reserved

Signature of Author: \_\_\_\_\_

Department of Linguistics and Philosophy  
September 9, 2020

Certified by: \_\_\_\_\_

Alex Byrne  
Professor of Philosophy  
Thesis Supervisor

Accepted by: \_\_\_\_\_

Bradford Skow  
Laurance S. Rockefeller Professor of Philosophy  
Chair of the Committee on Graduate Students

# **Meaning Change, in Theory and in Practice**

by

Allison Robbins Koslow

Submitted to the Department of Linguistics and Philosophy  
on September 8, 2020 in Partial Fulfillment of the Requirements  
for the Degree of Doctor of Philosophy in  
Philosophy

## **ABSTRACT**

Semantic change is of interest in its own right in the philosophy of language. For instance, it sheds light on the relationship between the meaning of a word and its use. It also plays a role in ideology critique. Philosophy of language often focuses on explaining synchronic features related to truth, entailment, and implicature. It is presumed an account of language's diachronic features will follow. But the tools developed are imperfectly suited, and mask phenomena of interest.

Chapter 1 concerns volitional meaning change of the sort advocated in conceptual engineering and amelioration. The question addressed is not what our concept, say, *WOMAN*, is, but what our concept should be. I bring to bear underappreciated empirical constraints on this normative project. Usage finely reflects equilibrium between communicative pressures (just as sales do between market pressures). Revising concepts is not an impossible task, but has significantly different contours than its proponents, and opponents, believe.

Chapter 2 concerns a family of cases that resemble those found in the mid-20<sup>th</sup> century. Austin's famous example is of a bird that looks for all the world like a goldfinch and then blows up like a grenade. Is it a goldfinch? A more realistic case involves the word "food": are meal replacement capsules food? Hard to say. There could be a separation, where British speakers affirm they are, and Americans deny it, without either appearing to misapply "food." Is this kind of open texture reducible to familiar linguistic phenomena, like vagueness? I argue not. It is a *sui generis* feature of meaning worthy of its own theory.

Chapter 3 addresses the charge that radical conceptual analyses, or revisions, change the subject. On a standard picture of meaning, they do. But some puzzles about diachronic synonymy suggest otherwise. I defend a radical position: truth-conditions are projected. They are partly a function of how, in the interpreter's community, it is reasonable to go on with the expression in question. For example, from our present perspective, "meal replacement capsules are food" may be false even as that very utterance is true from the perspective of future (or subaltern) speakers.

## **Dissertation Committee**

Alex Byrne (Chair), Professor of Philosophy, Department Chair  
Sally Haslanger, Ford Professor of Philosophy and Women's and Gender Studies  
Jack Spencer, Associate Professor of Philosophy  
Stephen Yablo, David W. Skinner Professor of Philosophy

**Acknowledgements**

I would like to thank Alex, my chair, as well as the rest of my committee Sally, Steve, and Jack. Thank you also to Allison Balin, Dylan Bianchi, Cheryl Chen, Annalisa Coliva, Cosmo Grant, Jon Koslow, Katy Meadows, Erin Miller, Rebecca Millsop, Sofia Hinojosa Ortiz, Kevin Richardson, Kate Ritchie, Michael Della Rocca, Kieran Setiya, Bob Stalnaker, Kirsi Teppo, Michael Thornton, Jonathan Vogel, Ian Walling, and Mallory Webber. Finally, I am eternally grateful to my parents, Susan and Jon, for their endless patience and support.

## Chapter 1: Conceptual Engineering: Problems and Prospects

Conceptual analysis concerns the concepts we do use. Can some concepts be defined in other terms? Conceptual engineering is the normative successor to conceptual analysis. It asks what concepts we should use. Should we use concept  $C_{\text{NEW}}$  in place of concept  $C_{\text{OLD}}$ ? While conceptual analysis has fallen into disfavor largely because of pessimism about its prospects for success, enthusiasm for conceptual engineering is on the rise.<sup>1</sup>

The idea that some cognitive and linguistic technology is more desirable than others is familiar from the early days of analytic philosophy. Frege, in the *Begriffsschrift*, promotes his concept notation as an instrument to “break the domination of words over the human mind” (1879/1967: 7). Carnap, in *Meaning and Necessity*, recommends the task of explicating, or taking an inexact “concept used in everyday life or in an earlier stage of scientific or logical development” and “replacing it by a newly constructed, more exact concept” (1947/1956: 8).<sup>2</sup> In *Individuals*, P. F. Strawson (himself no fan of explication) famously contrasts descriptive metaphysics, which “is content to describe the actual structure of our thought about the world” with revisionary metaphysics, which “is concerned to produce a better structure” (1959: 9). The latter, he suggests, is “essentially an instrument of conceptual change, a means of furthering or registering new directions or styles of thought” (10).

Conceptual analysis reached concepts through language; so does conceptual engineering. 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. (For the purposes of this paper, concepts are taken to be word meanings, or semantic values.) The target area of discourse may be exclusive to a cabal of philosophers, but often it is broader. The reasons for conceptual engineering are potentially many. They include facilitating a better account of some phenomenon, furthering some malevolent aim or making for better puns.

---

<sup>1</sup> See, for instance, Bach 2019, Brun 2016, Burgess 2019, Cappelen and Plunkett 2019, Brigandt and Rosario forthcoming, Cantalamessa 2019, Diaz-Leon 2019, Fassio and McKenna 2015, Goetze 2019, Koch 2019, Lutz 2018, Nado 2019, Pinder 2017 and forthcoming, Plunkett 2015 and 2016, Queloz 2019, Shea 2019, Sterken 2019, and Tanswell 2018.

<sup>2</sup> For a discussion of explication and conceptual engineering see Brun 2020, Novaes 2017, Kitsik 2017

But mostly conceptual engineering has commanded recent philosophical attention when proposed as a tool for social justice and ideological reform.<sup>3</sup> In that guise it is often called “conceptual amelioration.” Where traditional activists may correct our implicit theory of, say, race or addiction, conceptual engineers advocate changing our concept RACE or ADDICTION. For instance, one might argue that “person” should express an agential concept, rather than a species concept, to promote more appropriate treatment of animals like chimpanzees. (Activist lawyers may be said to engage in a related project when they use habeas corpus litigation to persuade courts to count chimpanzees as legal persons.<sup>4</sup>) The case for correcting mistaken theories hardly needs stating; the case for adjusting our concepts is less clear.<sup>5</sup>

The suggestion that we can feasibly improve the world by revising our cognitive and linguistic technology is undeniably exciting. This paper asks whether it is correct. The question can be factored into two components. First, would it actually promote social justice if we were to revise our concepts in the way the conceptual engineer proposes?<sup>6</sup> Second, is revising our concepts in the proposed way an achievable goal? If conceptual engineering is to be justified by the good it promotes, then this second question is paramount. But curiously little attention has been paid to it.<sup>7</sup> This paper makes a start.

Conceptual engineering requires a theory of what sorts of semantic changes can, as a practical matter, be induced. Fortunately, much is known about semantic change as it occurs in the wild. Questions about the feasibility of conceptual engineering projects are like questions about whether a particular product will succeed in a new market, or whether a nonnative species will

---

<sup>3</sup> It sometimes isn't: Scharp 2013 proposes that logicians use a revised truth rule to facilitate a coherent logic of truth.

<sup>4</sup> See [www.nonhumanrights.org/blog/first-habeas-corporus-order-happy](http://www.nonhumanrights.org/blog/first-habeas-corporus-order-happy) and [aldf.org/article/though-denied-by-new-york-court-of-appeals-habeas-corporus-claim-for-chimpanzees-prompts-reflection/](http://aldf.org/article/though-denied-by-new-york-court-of-appeals-habeas-corporus-claim-for-chimpanzees-prompts-reflection/)

<sup>5</sup> Amelioration relates to recent work on hermeneutical injustice. See Fricker 2007, Manne 2017. Not all projects that go by the name “amelioration” advocate conceptual change. Some are purely epistemic, more in the mold of good old-fashioned conceptual analysis: they set out to clarify what, say, women are, what we mean by talk of, say, “women,” what we are (unreflectively) doing with such thought and talk, and how it contributes to creating the unjust social reality it represents. See Haslanger 2005, 2010. To the extent these do not also involve claims about what concepts we should use, they are beyond the scope of this paper. See Haslanger 2019 for a discussion of different kinds of ameliorative projects.

<sup>6</sup> See Saul's contribution to Haslanger and Saul 2006.

<sup>7</sup> For instance, the attention paid to whether we should conceptual engineer in the event doing so may lead to false beliefs, or epistemic losses, may seem premature. We have yet to determine when and whether conceptual engineering can even succeed in areas of discourse where such problems may result. For discussion on the epistemic value of conceptual engineering, see Simion 2017, 2018, McKenna 2018, and Podosky 2020.

thrive in our ecosystem. They are not questions to be settled from the armchair. 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 conceptual engineering projects.

This paper concerns the hope that revising our linguistic classifications may make that world morally better. I suggest a method for assessing the feasibility of the proposed revisions. The picture that emerges undermines the sorts of engineering projects most likely to appeal to philosophers. As we'll see, some projects may stand to succeed, but they have significantly different contours than the typical ones.

### 1. Conceptual Engineering and Semantic Replacement

For the purposes of this paper, projects to revise or change our concepts may be thought of as projects to change the distribution of expressions with particular meanings.

Distinguish three sorts of engineering project: introduction, elimination, and replacement. All three are familiar. *Introduction*: experts and ordinary speakers alike frequently introduce words and phrases to make useful distinctions. Sometimes they are neologistic (“mansplaining”, “laser”<sup>8</sup>). Sometimes they are borrowed from ordinary language, and sometimes have related meanings (“contract,” “safety,” “sexual harassment”, “woke”). *Elimination*: speakers may also actively discourage the use of certain expressions (slurs, for instance, or failed theoretical terms like “phlogiston”). Mostly 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.”) Others, like “apricity”, gradually fade from public consciousness.<sup>9</sup> *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>

---

<sup>8</sup> The expression “mansplaining” was inspired by Rebecca Solnit’s essay “Men Explain Things to Me: Facts Didn’t Get in Their Way” (2014), but its exact origin is unknown. “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.

<sup>9</sup> The term, now obsolete, meant the warmth of the winter sun.

replaces use of another. These may share the meaning-component (as do “dried plums” -*dried plums* and “prunes”-*dried plums*), or share the word component, or neither.<sup>10</sup>

Conceptual engineering as commonly advocated by philosophers is a 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  $C_{\text{NEW}}$  deemed good, but  $C_{\text{OLD}}$  is deemed bad. We can call this *semantic replacement* since an expression with one meaning replaces uses of an expression with another meaning. Typically,  $W=W^*$ . Retaining the familiar word for the new concept seems a natural, though unforced, choice.

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.<sup>11</sup>

Some clarificatory remarks before we go any further. “Concept” is used in multifarious ways.<sup>12</sup> If (as assumed here) concepts are meanings or semantic values, then it may seem that conceptual engineering as previously explained changes the subject—at least absent some account of how expressions may concern the same subject despite having different semantic values.<sup>13</sup> Moreover, some conceptual engineers resist identifying concepts with meanings and think of their project not as changing our concepts, but as keeping them while making them better.<sup>14</sup> Fortunately there is no need to wade into these controversies here.

---

<sup>10</sup> It’s easy to see how these projects stand to be socially significant. Rebranding “prunes” increased sales. Coining the phrase “sexual harassment” helped us attend to and track its instances as well as confer about them in ways we did not do beforehand. To the extent, use of “sexual harassment” replaced some uses of, say, “overly flirtatious,” it’s easy to see how replacement projects may be socially significant.

<sup>11</sup> For an extended discussion, see Haslanger 2000.

<sup>12</sup> See Machery 2011, Fodor and Lepore 2002, Laurence and Margolis 1999, 2015.

<sup>13</sup> For discussion of the so-called discontinuity objection, see Prinzing 2017, Pinder 2020, Koch 2020, Riggs 2019. For a discussion of preserving the topic, see also Cappelen 2018, chs. 9-10.

<sup>14</sup> For instance, in Haslanger (2019)’s terminology, the concept associated with a term is a set of practical capacities that roughly correspond to its social role. What I call a concept is much closer to what Haslanger calls “the informational content” associated with a term (where that may be understood as something like a semantic value). Thus the informational content of, say, “family” may change while remains conceptually continuous in Haslanger’s

Quinean skepticism aside, there is an intuitive notion of a word “changing in meaning” that figures in lexicography and diachronic linguistics.<sup>15</sup> Without it, these legitimate empirical disciplines would not exist. No controversial conceptions of concepts, or disputed theories of linguistic meaning, 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.<sup>16</sup> 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. The conceptual engineer proposes to change 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 of semantic replacement 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. The conceptual engineer’s project is to push this organic process of meaning change in the preferred direction.

## 2. Five Examples

Having clarified the kind of conceptual engineering at issue, I will now turn to five illustrative examples of it. The most widely discussed comes from Haslanger (2000) and concerns race and gender. She describes her project as follows:

---

sense: “Our concept of family has evolved, improved, at least in part, through the political work done by LGBTQ and adoption activists in conceptual engineering with slogans such as ‘we are families too’” (2019: 21). Asking whether we may induce use of one word-meaning pair to replace use of another is thus roughly analogous to asking whether the concept associated with a word, in Haslanger’s sense, is likely to evolve with the result that the word come to be associated with an informational content of a particular sort. For discussion of the significance of concepts in amelioration, see Sawyer 2019.

<sup>15</sup> This is compatible with meanings being Fregean senses, external objects and properties, or what must be known for competent use. It is also compatible with an externalist metasemantics, or a metasemantics where idiolects are pervasive. 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 their meaning.”

<sup>16</sup> See the OED entries for “myriad” and “girl” (<http://www.oed.com>).



The questions “what is gender?” or “what is race?” require us to consider what work we want these concepts to do for us; why do we need them at all. The responsibility is ours to define them for our purposes...My priority in this inquiry is not to capture what we do mean, but how we might usefully revise what we mean for certain theoretical and practical purposes. (2000: 33-4)

Haslanger argues that those concerned with intersectional theoretical projects and activism should use “woman”/“man” so that they correctly apply to those subordinated/privileged on the basis of perceived reproductive capacity, and “race” so that it correctly applies to those subordinated (or privileged) on the basis of perceived ancestry. It is worth noting that while Haslanger thinks that using “woman” for the new concept is desirable, she does not think it essential:

[T]here are rhetorical advantages to using the terms ‘gender’, ‘man’ and ‘woman,’ and ‘race’ for the concepts I’ve defined, but if someone else is determined to have those terms, I’ll use different ones. (52)

Unlike a conceptual analysis of, say, “woman”—an account of what women are—ameliorative projects of this sort are not subject to traditional counterexamples:

[T]here could be females who aren’t women in the sense I’ve defined, but these individuals (or possible individuals) are not counterexamples to the analysis. The analysis is intended to capture a meaningful political category for critical feminist efforts, and non-oppressed females do not fall within that category. (46)

Haslanger thus offers an account of what “woman” should mean with the goal of

empowering critical social agents. However, whether the terminological shift I’m suggesting is politically useful will depend on the contexts in which it is employed and the individuals employing it. The point is not to legislate what terms to use in all contexts, but to offer resources that should be used judiciously. (48)

Jenkins (2016) provides our second example. She argues that the word “woman” should express the concept PERSON WITH A FEMALE GENDER IDENTITY rather than our present concept WOMAN, or the one proposed by Haslanger. She writes:

[T]he adoption of [Haslanger’s WOMAN] concept would exacerbate the existing and illegitimate marginalization of trans women within feminist discourse [in or out of the academy]. (2016: 396)

Arguably “woman” as traditionally understood fails to apply to transgender women and applies to transgender men; Haslanger’s “woman” at least fails to apply to some transgender women and applies to some transgender men. “Woman” understood with Jenkins’ meaning is intended to have neither of these supposed defects.

Ameliorative analysis is not bound to comply with our ordinary understanding or use of a concept: the target concept may be revisionary, provided that it furthers the goals guiding the analysis. (395)

In this instance the goals guiding Jenkins' project include facilitating feminist theorizing, respecting the gender identities of trans individuals, and combatting transphobia.<sup>17</sup>

The third example concerns the ordinary concept SEXUAL ORIENTATION, which pertains to match or mismatch between a person's sex and the sex of those they are attracted to. It is unclear how SEXUAL ORIENTATION, and associated concepts like LESBIAN and STRAIGHT, apply to those outside the gender binary. Dembroff (2016) "sets out to engineer a revised concept of sexual orientation" (2) such that:

the sexual orientation one has is grounded in what sex[es] and gender[s] of the persons one is disposed to sexually engage, *without reference to the sex or gender of the person so disposed*. (18, emphasis added)

This is supposed to have beneficial effects:

we can think of the engineering project as one that sets out to elucidate and possibly revise or replace our everyday concepts *in light of the impact we would like them to have*. (3, emphasis added)

In this framework, lesbians and straight men exclusively attracted to women would have the same "sexual orientation," an orientation shared with some who are gender non-binary or transgendered. Thus, if successfully implemented, the project "implies a new taxonomy of sexual orientation" (4) that stands to

eliminate and replace the everyday concepts associated with the taxonomy of sexual orientation. (4)

If Dembroff's new concepts replaced LESBIAN, STRAIGHT, and so on, the divide between "cisheterosexuality" (heterosexuality among non-transgender people) and queer sexual orientations (e.g. trans woman attracted to other trans women) would disappear from view. The conceptual revision is supposed to provide:

a taxonomic schema capable of recognizing persons outside the gender or sex binary...[and] conducive for establishing legal and social protections for persons who have queer sexual orientations. (19)

---

<sup>17</sup> Jenkins writes:

Failure to respect the gender identifications of trans people is a serious harm and is conceptually linked to forms of transphobic oppression and even violence. It follows from this that an important desideratum of a feminist analysis of gender concepts is that it respect these identifications by including trans people within the gender categories with which they identify and not including them within any categories with which they do not identify. (396)

And to reduce or eliminate:

the presumption that cisheterosexuality is the normatively standard sexual orientation and all queer sexual orientations are normatively deviant. (19)

The fourth example comes from Barnes (2016). 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: 10)

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. Disability is “a social category people have found useful when organising themselves in a civil rights struggle” (41). However, she also thinks of her project as ameliorative in something like Haslanger's sense:

When evaluating theories of social categories, the ameliorative project asks us to consider what role these categories have to play in social progress—with the background assumption that understanding and explaining social injustice is part of what will help us to address it. (41)

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 engineering, 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 fifth and final example is from Manne (2018). We have words for animus or bigoted attitudes towards women—one of them is “misogyny.” What is of theoretical interest, Manne argues, does not lie in malevolent attitudes towards women and girls, but rather in whatever produces certain bad outcomes for women and girls. In particular:

Misogyny ought to be understood as the system that operates within a patriarchal social order to police and enforce women's subordination and to uphold male dominance (2017: 33)

[It] primarily targets women because they are women in a man's world (i.e., a historically patriarchal one, among other things), rather than because they are women in a man's mind, where that man is a misogynist... the relevant hostilities may be manifested by individual agents, collective (or “mob”) activity, or purely structural mechanisms ... and [target] particular women and particular kinds of women (64)

Manne's project is advertised as an ameliorative one, and she acknowledges that “ameliorative projects are partly stipulative in nature” (62). Using “misogynist” with this new meaning is

supposed to “highlight the political dimensions” of particular phenomena and “render...[them]... psychologically more explicable” (34). She suggests her “analysis also yields an extension for the term ‘misogyny’ that dovetails nicely with usage patterns among feminists” (34).

Are engineering projects like these five at all feasible? Cappelen (2018) shares my doubts about the feasibility of conceptual engineering projects, but not my reasons for them. His case against conceptual engineering needs to be examined first.

### 3. Cappelen Against Conceptual Engineering

Cappelen offers two main arguments. Here’s the first:

[T]o effectively make a change in the extension and intension of [a] term, you would need to understand the mechanisms of reference change. These mechanisms are also not known to any of us and might in effect be unknowable. ... That makes it an illusion to think that we can be in a position to effectively predict and implement changes. (2018: 73-4)

This argument is unsuccessful. It assumes that if the A facts are somehow constituted by or supervene on the B facts, then instrumental reasoning about how to change the A facts requires knowing how they depend on the B facts. But that cannot be right. I know how to change your mind without knowing how the mental relates to the physical (I can talk with you) or how to age a photograph without knowing how pixels arranged thus-and-so give rise to a particular image of a face (I can draw in some smile lines).

Cappelen’s second argument is that even if one could acquire such knowledge, semantic changes would be impossible to implement.

Even if we were perfectly coordinated as a group (something we are decidedly not), [that] would not give the group control [of the reference fixing facts] because the actions and intentions of groups have at best a messy and unpredictable effect on our semantic values. We can of course try to influence other speakers and experts—but that will hardly ever amount to more than a drop in the ocean. (74)

This is akin to arguing that individual voters cannot control who is governor, because no individual’s vote is decisive. Of course, there may be something to this, but it is not impossible to influence electoral outcomes. If it were, people would not work so hard to do so. There are ways to amplify one’s voice. Some involve persuasive campaigning. Some involve organizing. And, of

course, it doesn't hurt to select a candidate that stands, given the right kind of organizing and campaigning, a decent chance of success.<sup>18</sup>

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 for the most part 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.” No committee decided to implement these changes, but they were motivated, not accidental. Cappelen's second argument gives us no reason to believe that more coordinated attempts at conceptual engineering would not succeed.

A drop in the ocean can have momentous consequences. 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. We need to understand the mechanisms by which a small intervention may propagate. What's needed is a strategy to assess which conceptual engineering projects—understood as semantic replacements—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.<sup>19</sup>

The contemporary linguistics literature on semantic change is small but growing, driven in part by the booming interest in Neo-Gricean approaches to the semanticization of pragmatics, inference and grammaticalization.<sup>20</sup> Other attempts to develop theories of semantic change have

---

<sup>18</sup> Cappelen offers a third argument:

...Even if we had all the information about the metasemantics of a term (about the use patterns, the histories, the sources of information, the interaction between the experts, etc.), it would appeal to factors that are in large part out of our control. For example, past facts play a role in determining the meaning of terms, but we can't change the past. (73-4)

This is no more persuasive than the first two. Since we need not know a term's metasemantics to affect its meaning, it is irrelevant whether having that information would help us do so. Even if we cannot now control some of the factors that determine meaning, it does not follow we cannot now control some of the factors that will do so going forward.

<sup>19</sup> Despite a long-recognized need for an English epicene pronoun, only recently has singular “they” gained purchase. Dozens of candidates were proposed in the 19<sup>th</sup> 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.

<sup>20</sup> Most notably Horn 1984, 1999, 2000a, 2000b. For a discussion of semantic change and grammaticalization for generativists, see von Stechow 2014.

loomed large in sociolinguistics,<sup>21</sup> cognitive linguistics,<sup>22</sup> and game theoretic semantics.<sup>23</sup> Some efforts have been made to extend relevance theory and optimality pragmatics to the study of semantic change.<sup>24</sup>

Drawing on some this literature, I'll now suggest a way to approach semantic replacement that allows the identification of projectable, even illuminating, generalizations. Then, in the following three sections, 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.<sup>25</sup> 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 changes in all things. The survival or preservation of certain favoured words in the struggle for existence is natural selection. (1871: 60-1)<sup>26</sup>

---

<sup>21</sup> See Labov 1994, 2012, and Eckert 2006.

<sup>22</sup> See Blank and Koch 1999, Bybee 2009, Sweetser 2006, and Traugott 1999.

<sup>23</sup> See Deo 2015 and 2014. Another relevant strand of research involves iterated learning frameworks. See Kirby 1999 and Smith et al. 2003. Some linguists are turning to corpus data to identify generalizations about words liable to undergo semantic change. <https://www.aclweb.org/anthology/W11-0134>.

<sup>24</sup> Optimality pragmatics may be thought of as formalizing and extending the Gricean principles of cooperative communicative behavior as found in Horn 1984 and Levinson 1995, 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 faithful interpretation generates a preference for interpretations of the utterance without leaving out any aspect of what the speaker says. The interaction of such constraints, founded on Levinson's 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 this sort of approach to explaining semantic change, see Falkum 2007, Carston 2002, Sperber and Wilson 1985/1991, Wilson 2003, Wilson and Carston 2006, 2007.

<sup>25</sup> See Smith 1759/1976, 1761/1983 and Otteson 2002.

<sup>26</sup> 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

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 we construe use broadly to include spoken applications of a particular word, thoughts involving it, and 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.<sup>27</sup> 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. Neither is it 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").<sup>28</sup>

---

primitive human races should speak evolutionarily lower languages than so-called civilized races). For discussion and controversy, see Radick 2001 and Alter 1999.

<sup>27</sup> Traugott 2017. See Norde 2009, Geeraerts 1990. Lass (1980) issued the 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, see Lass 1997. Of course, phonology was once thought too chaotic to be studied, but has since become a scientific discipline. Even if semantic change is not as regular, or predictable, as phonological change, there are still useful generalizations to be had about what sorts of semantic changes are and are not likely.

<sup>28</sup> Efforts are underway to create algorithms to extract generalizations about semantic change from corpus data (like that in the Corpus of Contemporary American English and the British National Corpus). Preliminary work suggests

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 17<sup>th</sup> 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.<sup>29</sup> (“Business” first meant *anxiety*, then *purposeful activity*; only relatively recently did it acquire the meaning of *occupation or trade*. Or, see, “segue.”)

How then are we to go about extracting generalizations relevant to those hoping to promote semantic 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.

---

that frequently used and shorter words are more likely to acquire secondary meanings. See Tang 2018, and Feltgen, Fagard, and Nadal 2017.

<sup>29</sup> Setting aside disagreement about why a particular generalization obtains, there is also controversy 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 is cross-linguistically attested that perceptual vocabulary tends to acquire broader attitudinal meanings and deontic modals tend to acquire epistemic meanings (“I see your point, feel pleased, am touched”; “could” and “must”). But it is 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 Grossman and Noveck 2015, Sweetser 1988, Traugott 2017, Traugott and Dasher 2002, Urban 2015.



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!).<sup>30</sup>

Conversational goals are served not only by advancing the right propositions, but expressing them via the right the 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

---

<sup>30</sup> This idea is developed in many ways. It is found in Zipf 1949 and Martinet 1952 as well as earlier sources like Paul 1888 and Bréal 1900. More recently, Horn 1985, 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 2010 have attempted to explain the phenomenon of prototypicality as reflecting an equilibrium between these two opposing forces. Bidirectional optimality theorists (see, Blutner 2000, Blutner and Zeevat 2004) 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.

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.<sup>31</sup>

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.<sup>32</sup> A beloved example of homonymy avoidance occurred in 12<sup>th</sup> century Gascon French: the terms "cattus" (*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 a printer's shop or at a linens shop).

Some macro patterns, including homonymy avoidance, are of particular relevance to engineering projects. The study of what semantic replacements are achievable is closely related to the study of what semantic replacements are likely to eventuate in the normal course. Both require

---

<sup>31</sup> Linguists have painstakingly accumulated much data about—and proposed explanations for—cross-linguistically attested pathways of semantic change. 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. Limited work in this direction can be found in Neo-Gricean theories (for instance in the work of Wilson 2003, Carston 2002, Wilson and Carston 2007) and cognitive theories of semantic change, metaphor and conventional inference (Sweetser 1990, Traugott 1999). Both emphasize the role of the structure of the lexicon in explaining semantic change. For example, it is 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 (Horn 1985, Anttila 1989). 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 (Lehrer 1985). 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.

<sup>32</sup> That is a problem that arises only when there is identity of sound, category (e.g., animal), subcategory (e.g., rooster), and register between the homophonous lexical items. 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, rang/ringed), word order (Fr. *pauvre homme unfortunate man* vs. *homme pauvre indigent man*), and orthography (draft/draught, metal/mettle, knight/night), 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?").

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. 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 semantic replacements are achievable. It turns out quite a lot.

So now we have seen *how* to approach the topic and *that* there are results to be had. Let's now turn to these more specific conjectures that bear on the project of conceptual engineering. Section 5 elaborates some of the generalizations we have identified—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 semantic replacement predictably involves expressions being replaced by slightly broader ones or, occasionally, slightly narrower ones. Section 7 returns to the negative: more dramatic semantic replacements are usually the product of successive more subtle replacements.

## 5. Negative Applications

Successful conceptual engineering requires introducing  $\langle W, C_{\text{NEW}} \rangle$ , thus eventually driving out  $\langle W, C_{\text{OLD}} \rangle$ . 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 conceptual 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 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 conceptual 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<sub>NEW</sub>> is more likely to replace <“woman”, WOMAN> than is <“woman”, WOMAN<sub>NEW</sub>> in areas of discourse where the homonymous use would read as ambiguous.<sup>33</sup> 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 will be tightly 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 19<sup>th</sup> 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 16<sup>th</sup> 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

---

<sup>33</sup> On the whole, ambiguous speech is avoided; whether such confusion would or would not be productive is irrelevant. (For an argument that such confusion would be productive, see Sterken 2019).

of what drives expression of the problematic meaning. For the conceptual engineer this problem arises in both parts of semantic 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's: if implemented, talk of "sexual orientation" will undergo semantic replacement and expressions like "lesbian" and "straight" will drop from usage, replaced by novel orientation vocabulary, including expressions for those *attracted to men who were assigned "male" at birth* and those *attracted 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.<sup>34</sup> Sexual orientation vocabulary may not be ideal for the reasons Dembroff identifies, but the way forward is unobvious.<sup>35</sup>

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. The problem is that only some of our communicative desires are tied up in socially significant projects. Supplanting usage of one word-meaning pair with usage of another, whose meaning component better serves these social projects, may be even more difficult than it appears for somewhat subtle reasons.

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*

---

<sup>34</sup> See Stock 2019.

<sup>35</sup> Sometimes expressions that find stable use within a subculture enter the mainstream. Novel orientation vocabulary may originate this way.

*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.<sup>36</sup> Where once we spoke of haycocks and stopcocks we speak of haystacks and faucets.<sup>37</sup> 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

---

<sup>36</sup> 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.”

<sup>37</sup> Cicero (45 BCE) recognized the phenomenon:

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*]. (Epistulae ad Familiares, 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).

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 16<sup>th</sup> century it meant something like *peculiar*. By the 17<sup>th</sup> century it acquired stable use as a euphemism for *unwell* (itself a euphemism) and ultimately assumed that as its literal meaning. By the 19<sup>th</sup> century “queer”-*unwell* began to function as a euphemism for *drunk*, and soon acquired that meaning as well. Sometime around the turn of the 20<sup>th</sup> 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 20<sup>th</sup> 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.<sup>38</sup>

Loading afflicts some otherwise attractive semantic 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.

Haslanger’s proposals for race and gender may win the day in expert discourse concerned with intersectionality, but speakers are likely to avoid her proposed meaning reassignments at the

---

<sup>38</sup> “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.

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."<sup>39</sup>) 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 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 14<sup>th</sup> century acquired two non-technical secondary meanings each associated with a salient symptom of sanguinity: *ruddy* (as in "a

---

<sup>39</sup> So much so the American Fire Safety organization has discouraged labeling products as "inflammable" since the 1920s.



sanguine complexion”) and *optimistic*.<sup>40</sup> 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.<sup>41</sup> 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*.

---

<sup>40</sup> 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.”

<sup>41</sup> For a discussion of celebratory-based, rather than solidarity-based, ameliorative projects for disability, see Dougherty 2019.

*Opacity* may partly explain the success of reclamations, whereby in-group members begin to self-identify using a derogatory expression. “Queer” is a recent example. 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” changed the linguistic environment in a way that reduced derogatory uses of “queer.”<sup>42</sup> 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).<sup>43</sup>

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 conceptual 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<sup>TM</sup>” 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

---

<sup>42</sup> Whether reclamation projects like this are well understood as involving semantic replacement depends on whether the derogatoriness of derogatory statements resides in their meaning, their implicatures or something else. See Sosa 2018.

<sup>43</sup> Dougherty 2019.

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 14th 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.<sup>44</sup>

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

---

<sup>44</sup> 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 18<sup>th</sup> century as part of American efforts to standardize English along Latinate lines.

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 socially significant conceptual 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.<sup>45</sup>

But few conceptual engineering proposals with a revolutionary flavor will involve broadening or narrowing in this sense. Consider Manne’s proposal, that “misogyny”—*the law enforcement arm of the patriarchy* should 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. 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

---

<sup>45</sup> Haslanger (2019) offers a somewhat different analysis of this case (assuming “family” indeed has undergone meaning change in the last 50 years, which, again, is not obvious). She suggests the changes it has 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.”

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.<sup>46</sup>

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 conceptual 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.

---

<sup>46</sup> There is no shortage of cases like this one. 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).

## 7. Semantic Replacement: Predictably Predictable or Unpredictably Dramatic

Some semantic 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 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.<sup>47</sup>

Dramatic semantic 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*.<sup>48</sup> 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.”<sup>49</sup> 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 *furiosus* or *enraged*.

---

<sup>47</sup> 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 semantic replacement via figuration is unlikely to occur. Sweetser 1990.

<sup>48</sup> A socially more significant case of semantic replacement might be that of “usury.”

<sup>49</sup> 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*).

Described from 30,000 feet, the sorts of conceptual engineering proposals most attractive to philosophers require dramatic semantic changes to which there is no predictable path via a series of smaller changes.

## 8. Conclusion

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

Those serious about semantic replacement as a tool 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 activist-minded philosophers with unified aims, and homogenous circumstances. Carnapian explication projects, undertaken for circumscribed theoretical projects, often do succeed. But conceptual engineers typically want more. 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. Even so, some significant social conditions—for instance, how respected, valued and included some members of a community feel—are affected by usage in that community.

There may be other reasons to go in for innovating new concepts, even if they are unlikely to catch on. Proposing novel concepts may raise consciousness, provoke dialogue, or perform a gestalt shift on audience members. But if conceptual engineering is worthwhile for these reasons only, it stands accused of false advertising. Bringing about semantic replacement was supposed to

be a means to a better world, when, in fact, it is *proposing* semantic replacement that is the means to a better world.

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.<sup>50</sup> Identifying a 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 nevermind its relationship to 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 that is elegant in theory but unbuildable in practice. Correcting our vocabulary might have beneficial effects, but is vastly harder than is usually presumed. Activists taking the old-fashioned route of correcting our mistakes and trying to change hearts and minds are likely to have better luck.

## **References**

- Alter, Stephen G. (1999). *Darwin and the Linguistic Image: Language, Race and Natural Theology in the Nineteenth Century*. Baltimore: Johns Hopkins University Press.
- Antilla, Raimo. (1989). *Historical and Comparative Linguistics*, Current Issues in Linguistic Theory. Philadelphia: John Benjamins Publishing Company.
- Barnes, Elizabeth. (2016). *The Minority Body: A Theory of Disability*. Oxford: Oxford University Press
- Blank, Andreas. (1999). "Why Do New Meanings Occur?" in Blank and Koch, eds. *Historical Semantics and Cognition* Berlin: de Gruyter.

---

<sup>50</sup> Say, on the basis of their expected explanatory utility, see Carballo 2019.



- Blank, Andreas, and Peter Koch. (1999). *Historical Semantics and Cognition*. Berlin: Mouton De Gruyter.
- Blutner, Reinhard. (2000). "Some Aspects of Optimality in Natural Language Interpretation," *Journal of Pragmatics*, vol. 17: 189-216.
- Blutner, Reinhard, and Henk Zeevat. (2004). "Editors' Introduction: Pragmatics in Optimality Theory" in Blutner and Zeevat, eds. *Optimality Theory and Pragmatics* London: Palgrave Macmillan.
- Bogardus, Tomas. (2020). "Some Internal Problems with Revisionary Gender Concepts," *Philosophia*.
- Breal, Michel. (1900). *Notebooks on Language*. London: Heinemann.
- Brigrandt, Ingo and Esther Rosario. *Strategic Conceptual Engineering for Epistemic and Social Aims*. Edited by Burgess, Conceptual Engineering and Ethics. Oxford: Oxford University Press.
- Brun, Georg. (2016). "Conceptual Re-Engineering: From Explication to Reflective Equilibrium," *Erkenntnis*, vol. 81: 1211–41.
- Burgess, Alexis. "Never Say 'Never Say 'Never'" in Burgess, Cappelen and Plunkett, eds. *Conceptual Engineering and Conceptual Ethics* Oxford: Oxford University.
- Burgess, Alexis and David Plunkett. (2013). "Conceptual Ethics I," *Philosophy Compass*, vol. 8: 1091-101.
- . (2013). "Conceptual Ethics II," *Philosophy Compass*, vol. 8: 1102-10.
- Bybee, Joan. (2009). *Language, Usage, and Cognition*. Cambridge: Cambridge University Press.
- Cantalamesa, Elizabeth Amber. "Disability Studies, Conceptual Engineering, and Conceptual Activism," *Inquiry*.
- Cappelen, Herman. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Cappelen, Herman, and David Plunkett. (2019). "Introduction: A Guided Tour of Conceptual Engineering and Conceptual Ethics" in Burgess, Herman Cappelen and Plunkett, eds. *Conceptual Engineering and Conceptual Ethics* Oxford: Oxford University Press.
- Carballo, Alejandro Pérez. "Conceptual Evaluation: Epistemic" in Burgess, Cappelen and Plunkett, eds. *Conceptual Ethics and Conceptual Engineering* Oxford: Oxford University Press.
- Carnap, Rudolf. (1947/1956). *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Carston, Robyn. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Berlin: Blackwell.
- Cicero, Marcus Tullius. (45 BCE/1988). "Epistulae Ad Familiares." in Watt, ed. *Loeb Classics Series* Cambridge: Harvard University Press.
- Darwin, Charles. (1871). *The Descent of Man*. London: J. Murray.
- Dembroff, Robin. (2016). "What Is Sexual Orientation?," *Philosopher's Imprint*, vol. 16: 1-27.
- Deo, Ashwini. (2014). "Formal Semantics and Language Change" in Bowerman and Evans, eds. *Routledge Handbook of Historical Linguistics* London: Routledge Press.
- . (2015). "Diachronic Semantics," *The Annual Review of Linguistics*, vol. 1: 179-95.
- Diaz-Leon, Esa. "Descriptive Vs. Ameliorative Projects: The Role of Normative Considerations" in Burgess, Cappelen and Plunkett, eds. *Conceptual Engineering and Conceptual Ethics* Oxford: Oxford University Press.

- Dougherty, Tom. "Disability as Solidarity: Political Not (Only) Metaphysical," *Phenomenological and Philosophical Research*.
- Eckert, Penelope. (2006). "Variation, Meaning, and Social Change" in Coupland, ed. *Sociolinguistics: Theoretical Debates* Cambridge: Cambridge University Press.
- Falkum, Ingrid Lossius. (2007). "A Relevance-Theoretic Analysis of Concept Narrowing and Broadening in English and Norwegian Original Texts and Translations," *Languages in Contrast*, vol. 7: 119-42.
- Fassio, Davide, and Robin McKenna. (2015). "Revisionary Epistemology," *Inquiry: An Interdisciplinary Journal of Philosophy*, vol. 58: 755-79.
- Fodor, Jerry, and Ernie Lepore. (2002). *The Compositionality Papers*. Oxford: Oxford University Press.
- Fortson, Benjamin W. IV. (2003). "An Approach to Semantic Change" in Janda, ed. *The Handbook of Historical Linguistics* Oxford: Blackwell.
- Frege, Gottlieb. (1879/1967). *Begriffsschrift: Eine Der Arithmetischen Nachgebildete Formelsprache Des Reinen Denkens*. Translated by in v. Heijenoort. Edited by vanHeijenoort, From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931. Cambridge: Harvard University Press.
- Fricker, Miranda. (2007). "Hermeneutical Injustice" *Epistemic Injustice: Power and the Ethics of Knowing* Oxford: Oxford University Press.
- Geeraerts, Dirk. (1990). "Homonymy, Iconicity, and Prototypicality" in Geeraerts, ed. *Diachronic Semantics* Philadelphia: John Benjamins Publishing Company.
- . (2010). *Lexical Semantics*. Philadelphia: John Benjamins Publishing Company.
- Goetze, Trystan. "Conceptual Responsibility," *Inquiry*.
- Haslanger, Sally. "Going on, Not in the Same Way" in Burgess, Cappelen and Plunkett, eds. *Conceptual Ethics and Conceptual Engineering* Oxford: Oxford University Press.
- . (1999). "What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology," *Philosophical Perspectives*, vol. 13: 459-80.
- . (2000). "Race and Gender: (What) Are They? (Do) We Want Them to Be?," *Nous*, vol. 34: 31-55.
- . (2005). "What Are We Talking About? The Semantics and Politics of Social Kinds," *Hypatia*, vol. 20: 10-26.
- . (2010). "Language, Politics, and 'the Folk': Looking for 'the Meaning' of 'Race'," *The Monist*, vol. 93: 169-87.
- Haslanger, Sally & Jennifer Saul. (2006). "Philosophical Analysis and Social Kinds," *Proceedings of the Aristotelian Society*, vol. 106: 89-118.
- Horn, Laurence R. (1984). "Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature" in Schiffrin, ed. *Meaning, Form, and Use in Context: Linguistic Applications* Washington, D.C.: Georgetown University Press.
- . (1999). "Etymology and Taboo." *Annual Meeting of the Linguistic Society of America* Boston.
- . (2000). "From If to Iff: Conditional Perfection as Pragmatic Strengthening," *Journal of Pragmatics*, vol. 32: 289-326.
- Horn, Laurence R., and Steven R. Kleinedler. (2000). "Parasitic Reference Vs. R-Based Narrowing: Lexical Pragmatics Meets He-Man." *Annual Meeting of the Linguistic Society of America* Chicago.

- Janda, Richard D., and Brian D. Joseph. (2003). "On Language, Change and Language Change - or, of History, Linguistics and Historical Linguistics" in Janda and Brian D. Joseph, eds. *The Handbook of Historical Linguistics* Oxford: Blackwell.
- Jenkins, Katharine. (2016). "Amelioration and Inclusion: Gender Identity and the Concept of Woman," *Ethics*, vol. 126: 394-421.
- Kitsik, Eve. "Explication as a Strategy for Revisionary Philosophy," *Synthese*.
- Koch, Steffen. (2019). "Carnapian Explication, Experimental Philosophy, and Fruitful Concepts," *Inquiry*, vol. 62: 700-17.
- . (2019). "The Externalist Challenge to Conceptual Engineering," *Synthese*.
- Labov, William. (1994). *The Principles of Language Change*. London: Blackwell.
- . (2012). *Dialect Diversity in America: The Politics of Language Change*. Charlottesville: University of Virginia Press.
- Lass, Roger. (1980). *On Explaining Linguistic Change*. Cambridge: Cambridge University Press.
- . (1997). *Historical Linguistics and Language Change*. Cambridge: Cambridge University Press.
- Laurence, Stephen, and Eric Margolis. (1999). *Concepts*. Cambridge: MIT Press.
- . (2015). "The Conceptual Mind: New Directions in the Study of Concepts." Cambridge: MIT Press.
- Levinson, S. C. (1995). "Three Levels of Meaning" in Palmer, ed. *Grammar and Meaning: Essays in Honour of Sir John Lyons* Cambridge: Cambridge University Press.
- . (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge: MIT Press.
- Lightfoot, David. (1999). *The Development of Language: Acquisition, Change, and Evolution*. London: Blackwell.
- Lutz, Sebastian. "Armchair Philosophy Naturalized," *Synthese*.
- Machery, Edouard. (2011). *Doing without Concepts*. Oxford: Oxford University Press.
- . (2017). *Philosophy within Its Proper Bounds*. Oxford: Oxford University Press.
- Manne, Kate. (2018). *Down Girl: The Logic of Misogyny*. New York: Oxford University Press.
- McKenna, Robin. (2018). "No Epistemic Trouble for Engineering 'Woman'," *Logos and Episteme*, vol. 9: 335-42.
- Merriam-Webster. (1989). "New Book of Webster's Word Histories." *Merriam Webster* Springfield, MA: Merriam-Webster Inc., Publishers.
- Miranda, Rogelio. "Carnapian Explication and the Canberra Plan's Conceptual Analysis: A Comparison and Critique," *Logical Analysis and History of Philosophy*.
- Nado, Jennifer. "Conceptual Engineering Via Experimental Philosophy," *Inquiry*: 1-21.
- . "Conceptual Engineering, Truth, and Efficacy," *Synthese*.
- Novaes, Catarina Dutilh (2018). "Carnapian Explication and Ameliorative Analysis: A Systematic Comparison," *Synthese*: 1-24.
- Otteson, James. (2002). "Adam Smith's First Market: The Development of Language," *History of Philosophy Quarterly*, vol. 19: 65-86.
- Pinder, Mark. "The Austerity Framework and Semantic Normativity," *Inquiry*: 1-19.
- . (2017). "On Strawson's Critique of Explication as a Method in Philosophy," *Synthese*.
- Plunkett, David. (2015). "Which Concepts Should We Use?: Metalinguistic Negotiations and the Methodology of Philosophy," *Inquiry*, vol. 58: 828-74.

- . (2016). "Conceptual History, Conceptual Ethics, and the Aims of Inquiry: A Framework for Thinking About the Relevance of the History/Geology of Concepts to Normative Inquiry," *Ergo*, vol. 3: 27-64.
- Podosky, Paul-Mikhail Catapang. "Ideology and Normativity: Constraints on Conceptual Engineering," *Inquiry*.
- Prinzling, Michael. (2018). "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection," *Inquiry*, vol. 61: 854-80.
- Queloz, Matthieu. (2019). "The Points of Concepts: Their Types, Tensions, and Connections," *Canadian Journal of Philosophy*, vol. 49: 1122-45.
- Radick, Gregory. (2002). "Darwin on Language and Selection," *Selection*, vol. 3: 7-16.
- Riggs, Jared. "Conceptual Engineers Shouldn't Worry About Semantic Externalism," *Inquiry*.
- Sawyer, Sarah. "The Role of Concepts in Fixing Language," *Canadian Journal of Philosophy*.
- . "Talk and Thought" in Burgess, Herman Cappelen and David Plunkett, eds. *Conceptual Engineering and Conceptual Ethics* Oxford: Oxford University Press.
- . "Truth and Objectivity in Conceptual Engineering," *Inquiry*.
- . (2018). "The Importance of Concepts," *Proceedings of the Aristotelian Society*, vol. 118: 127-47.
- Scharp, Kevin. (2013). *Replacing Truth*. Oxford: Oxford University Press.
- Schupbach, Jonah. (2017). "Experimental Explication," *Philosophy and Phenomenological Research*, vol. 94.
- Shea, Nicholas. "Concept-Metacognition," *Mind and Language*.
- Simion, Mona. (2017). "The 'Should' in Conceptual Engineering," *Inquiry*, vol. 61: 914-28.
- . (2018). "Epistemic Trouble for Engineering 'Woman'," *Logos and Episteme*, vol. 9: 91-98.
- Smith, Adam. (1759/1976). *The Theory of Moral Sentiments*. Oxford: Oxford University Press.
- . (1761/1983). "Considerations Concerning the First Formation of Languages, and the Different Genius of Original and Compounded Languages" in Bryce, ed. *Smith's Lectures on Rhetoric and Belles Lettres* Oxford: Oxford University Press.
- Sosa, David. (2018). *Bad Words: Philosophical Perspectives on Slurs*. Edited by Sosa. Oxford: Oxford University Press.
- Sperber, Dan, and Deirdre Wilson. (1985/1991). "Loose Talk" in Davis, ed. *Pragmatics: A Reader* Oxford: Oxford University Press.
- Sterken, Rachel Katharine. "Linguistic Interventions and Transformative Communicative Disruption" in Cappelen, Plunkett and Burgess, eds. *Conceptual Engineering and Conceptual Ethics* Oxford: Oxford University Press.
- Stock, Kathleen. (2019). "Sexual Orientation: What Is It?," *Proceedings of the Aristotelian Society*, vol. 119: 295-319.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.
- Sweetser, Eve. (1988). "Grammaticalization and Semantic Bleaching." in Axmaker, Jaisser and Singmaster, eds. *Berkeley Linguistics Society*.
- Sweetser, Eve E. (1990). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Vol. 54, Cambridge Studies in Linguistics. Cambridge: Cambridge University Press.
- Tang, Xu. (2018). "A State-of-the-Art of Semantic Change Computation," *Natural Language Engineering*, vol. 24: 881-913.

- Tanswell, Fenner Stanley. (2018). "Conceptual Engineering for Mathematical Concepts," *Inquiry*, vol. 61: 881-913.
- Taylor, John R. (1999). "Cognitive Semantics and Structural Semantics" in Blank and Koch, eds. *Historical Semantics and Cognition* Berlin: de Gruyter.
- Trask, R. L. (1994). *Language Change*. London: Routledge.
- Traugott, Elizabeth Closs. (1999). "The Rhetoric of Counter-Expectation in Semantic Change: A Study in Subjectification" in Blank and Koch, eds. *Historical Semantics and Cognition* Berlin: de Gruyter.
- . (2017). "Semantic Change." in Aronoff, ed. *Linguistics Oxford Research Encyclopedia* Oxford: Oxford University Press.
- Traugott, Elizabeth Closs, and Richard Dasher. (2002). *Regularity in Semantic Change*, Cambridge Studies in Linguistics. Cambridge: Cambridge University Press.
- Urban, Matthias. (2015). "Lexical Semantic Change and Semantic Reconstruction" in Bower and Evans, eds. *Routledge Handbook of Historical Linguistics* Abingdon: Routledge.
- Wikforss, Asa Maria, and Teresa Marques. "Shifting Concepts: The Philosophy and Psychology of Conceptual Variability." Oxford: Oxford University Press.
- Wilson, Deirdre. (2003). "Relevance Theory and Lexical Pragmatics," *Rivista di Linguistica*, vol. 15: 273-91.
- Wilson, Deirdre and Robyn Carston. (2006). "Metaphor, Relevance and the 'Emergent Property' Issue," *Mind & Language*, vol. 21: 404-33.
- Wilson, Deirdre, and Robyn Carston. (2007). "A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts " *Mind & Language*, vol. 21: 404-33.
- Zipf, George. (1949). *Human Behavior and the Principle of Least Effort*. Cambridge: Addison-Wesley.

## Chapter 3: Open Cases and Open Texture

### 1. Introduction

Is bitcoin money?<sup>51</sup> Are gestational surrogates mothers? Are glass walls windows? Cases like these are easy to identify, and even easier to imagine. (Suppose we invent meal replacement capsules: are they food?) Identify an invention, or a discovery, or a changing interest—real or imagined—that forces a decision about how to apply a familiar term, though nothing in past usage, our dispositions, patterns of deference, or the social and physical environment, apparently settles the matter. When such cases arise, we settle on *something* eventually. Life must go on. But neither appears entailed by the meaning of our expressions, or to change them. Call these “open cases”.

I’ll discuss a number of them without any pretense that they are clear examples of the very same phenomenon. Some may be amenable to standard treatments in the semantic medicine cabinet, but some are not. My conclusion is that these reveal a distinct phenomenon, worthy of its own account. While the predicates in question certainly have familiar features—like most terms, most are vague—these features are not responsible for the openness of the cases. I’ll cycle through various ways of analyzing this openness, starting with vagueness. I then consider other semantic phenomena that might be used a model (independently of any vagueness), including partial definition, context-sensitivity, indeterminate reference, and underdetermination.

### 2. Open Texture

The cases mentioned at this paper’s start call to mind a group of cases familiar from mid-20<sup>th</sup> century literature objecting to the verification criterion of meaning, and the search for definition. Let’s start by discussing these. They identify a vagueness-like phenomenon that is distinct from vagueness as we think of it today. In Philosophical Investigations, Wittgenstein imagines his chair disappearing:

I say, “There is a chair”. What if I go up to it, meaning to fetch it, and it suddenly disappears from sight? —“so it wasn’t a chair, but some kind of illusion”. —But in a few moments we see it again and are able to touch it and so on. —’so the chair was there after all and its disappearance was some kind of illusion”. —But suppose that after a time it disappears again-or seems to disappear. What are we to say now? Have you rules ready for such cases

---

<sup>51</sup> <https://www.nber.org/papers/w19747>  
<https://cointelegraph.com/news/defining-bitcoin-money-currency-or-store-of-value>

—rules saying whether one may use the word “chair” to include this kind of thing? But do we miss them when we use the word “chair”; and are we to say that we do not really attach any meaning to this word, because we are not equipped with rules for every possible application of it? (1953/1968: 80)<sup>52</sup>

Likewise, in his discussion of games, he notes that the rules of tennis say nothing about how high or hard one may throw the ball, so the concept of a legitimate move in tennis is unregulated in a certain respect (1953: 68). It is not that we know that one meter is permissible while one hundred are not: this height has never been a factor in the game and there are no rules concerning it. By contrast, if we make a gesture towards a region of roughly one meter in diameter and say, “Stay roughly here” (71), then this determines, although not sharply, the boundaries of a region.<sup>53</sup> Austin, discussing the criterial theory of knowledge, gives another such example:

If we have made sure it’s a goldfinch, and a real goldfinch, and then in the future it does something outrageous (explodes, quotes Mrs. Woolf, or what not), we don’t say we were wrong to say that it was a goldfinch, we don’t know what to say. Words literally fail us: “What would you have said?” “What are we to say now?” “What would you say?” ... It seems a serious mistake to suppose that language (or most language, language about real things) is “predictive” in such a way that the future can always prove it wrong. What the future can always do, is to make us revise our ideas about goldfinches or real goldfinches or anything else. (1946: 160)

Again, the situation seems different from vagueness. For one, we know what a borderline goldfinch would be like: it resembles paradigmatic goldfinches but can also mate with, say, cardinals. We know to say: it’s not definitely a goldfinch, not definitely not a goldfinch. Austin’s case brings out two phenomena. The first we might call “semantic dumbfounding”. As he wrote some years earlier:

Ordinary language breaks down in extraordinary cases. (In such cases, the cause of the breakdown is semantical.) Now no doubt an *ideal* language would *not* break down, whatever happened. In doing physics, for example, where our language is tightened up in order precisely to describe complicated and unusual cases concisely, we *prepare linguistically for the worst*. In ordinary language we do not: *words fail us*. (1940: 36–37)

---

<sup>52</sup> Kripke 1984/2000 identifies a more radical lesson in PI’s discussion of rules, namely that no considerations are sufficient to uniquely determine the applicability of a term, even mathematical expressions like “+”. But even granting some facts are sufficient to determine the meaning of “+”, an, say, the application of “window” to typical household ones, it may seem open or unsettled whether “window” is truly applicable to glass walls.

<sup>53</sup> This is the way, claims Wittgenstein, in which we use the concepts of number. When the complex numbers and the transfinite cardinals were invented (not discovered), what we already counted as numbers did not determine whether these new mathematical entities should also be considered as such. It is not that they were on some vague borderline separating numbers from other mathematical creatures: they constituted a new mathematical dimension, and our use had to be extended towards it. (e.g., 1956: II.22-23).

The second phenomenon we might call “semantic dogmatism”:

If there’s a *lusus naturae*, a miracle... that would not mean I was wrong [about their being goldfinches] (163)

On the one hand, there’s the inclination that, even when experience takes a strange turn, there are some things we cannot be wrong about (they are definitely goldfinches). Alongside it, there are competing inclinations to say that we were *completely* wrong (they are definitely not).<sup>54</sup> We might imagine a heated argument. Vagueness, by contrast, tends to elicit equivocal reactions, and agreement that equivocality is the right response. Suppose we start quarreling about whether skateboards are vehicles. We’re likely to quickly realize that there’s no good question here: I’m a little more inclined to say they are than you, but when we realize that, we see it’s one of those unclear cases. We know what to say, in the skateboard case: it’s not clearly a vehicle and it’s not clearly not a vehicle. So, if the vagueness of “goldfinch” (or “chair”) were at issue in the case, we agree that such and so is not clearly a goldfinch, not clearly not a goldfinch. But that does not appear to be our reaction.<sup>55</sup>

The person who made the most of such goldfinch-like cases was Friedrich Waismann, who, by 1937, was Austin’s colleague, and before that a collaborator of Wittgenstein.<sup>56</sup> (He emigrated from Austria to England officially in 1938). Waismann argued such cases revealed something he called “the open texture of language”:

Open texture is a very fundamental characteristic of most, though not of all, empirical concepts, and it is this texture which prevents us from verifying conclusively most of our empirical statements. Take any material object statement. The terms which occur in it are non-exhaustive; that means that we cannot foresee completely all possible conditions in which they are to be used; there will always remain a possibility, however faint, that we have not taken into account something or other that may be relevant to their usage; and that means that we cannot foresee completely all the possible circumstances in which the

---

<sup>54</sup> Stebbing makes much of this tension, roughly identifying the former—goldfinches wisecrack—with Cambridge analysis and Moorean directional analysis, as opposed to what she calls “postulational analysis”, more or less Carnap’s method of explication. Scientific discovery cannot overturn every day usage, and:

We have to think *about* what we *were* thinking about. The philosopher considers a *given expression*, and analyses it in order to find *another expression* which says *more* clearly what the original expression said *less* clearly. This investigation is not linguistic. We must first *know* what facts are the case before we can fruitfully employ analysis for the purpose of clarifying our thoughts about the world. (1933: 86)

<sup>55</sup> Talk of open texture as it appears in the law more closely applies to cases of vague predicates. So, in legal theory, unclarity about whether skateboards, and electric wheelchairs, are prohibited in the park is symptomatic of the open texture of the law. Waldron 1994, Schauer 2008

<sup>56</sup> See Bunikowski 2016. Waismann emigrated from Austria in 1938.



statement is true or in which it is false. There will always remain a margin of uncertainty. Thus the absence of a conclusive verification is directly due to the open texture of the terms concerned (1945: 3)<sup>57</sup>

To illustrate the idea, Waismann used numerous extravagant examples, including a disappearing friend (2), gold that emits a new sort of radiation (2); a cat-like object that grows to enormous proportions, or speaks Latin, or is resurrects from death (3). In all those extraordinary, unforeseeable circumstances, he insisted, we would be hesitant about what to say, because it is not provided for in the meaning of the statement. Our understanding of “that is a cat” gives us no basis on which to choose between the statement and its negation in such a case.

While he used the phrase “vague” differently than we do today, he distinguished between open texture and the features of paradigmatically vague predicates, like “heap” or “pink”.

Vagueness should be distinguished from open texture. A word which is actually used in a fluctuating way (such as “heap” or “pink”) is said to be vague; a term like “gold”, though its actual use may not be vague, is nonexhaustive or of an open texture in that we can never fill up all the possible gaps through which a doubt may seep in. ... Vagueness can be remedied by giving more accurate rules, open texture cannot. An alternative way of stating this would be to say that definitions of open terms are always corrigible or emendable... [Open texture is the] the permanent possibility of vagueness (1945: 123)

But there is no suggestion, for instance, that the fluctuating use of the color term “pink” would be the fate of chemical “gold” should science turn up a surprise.

Verificationism involved the idea was that the meaning of empirical statements, such as “that is a cat”, might be identified with the kind of experience that conclusively verified it. Metaphysical statements such as “God is love” were dismissed as meaningless on the grounds that they had no method of verification. Since verification is something carried out by speakers of the language, the positivist theory promised to anchor the semantics of the language in the practice of using it, without a spooky metaphysics of meaning. The fatal difficulties for the program are familiar. If the meaning of “that is a cat” is exhausted but the kind of experience that would conclusively verify it, then one might expect it to be translatable into a statement about those experiences. But no one could provide a remotely plausible translation.

---

<sup>57</sup> W.C. Kneale suggested this translation of Waismann’s phrase, “*porosität der Begriffe*”, which more literally translates, “the porosity of language”. Waismann’s discusses open texture in a series of essays from 1949 to 1953 concerning the analytic-synthetic distinction, and, most famously in 1968.

Waismann tried to describe a subtler connection between meaning and verification. On his view, empirical statements are never conclusively verified, for two reasons. First, although a finite number of experiences provide good enough evidence to warrant the assertion of “that is a cat” they are never logically sufficient for its truth (in principle one might always turn out to be the victim of a demon, or an elaborate hoax). Second, more interestingly, suppose that all other relevant experiences, future, as well as past, are somehow given. They may still be logically insufficient to decide the truth or falsity of “that is a cat”, even without assuming the truth-value of the statement can transcend all possible experiential evidence.<sup>58</sup> Thus, according to him, the failure of the verification criterion of meaning—the claim that the content of an expression is exhausted by the evidence for its truth:

is not, as has been suggested, due to the poverty of our language which lacks the vocabulary for describing all the minor details of sense experience, nor is it due to the difficulties inherent in producing an infinite combination of sense-datum statements, in the main it is due to a factor which though it is very important and really quite obvious, has to my knowledge never been noticed—to the “open texture” of most of our empirical concepts. (1968: 118-119)

Sometimes Waismann’s objection is glossed as a claim about the difficulty of supplying complete necessary and sufficient conditions for ignorant or cognitively limited creatures like ourselves. That’s only part of it. Another part is that it is often *in principle* impossible to supply a complete verbal definition, as there are always conditions that can be added, or removed. The claim is not that proper *definitions* are infinitely long and that practical limitations block us from constructing, or identifying, them. The claim is there is some “open” or “porous” feature of meanings themselves.<sup>59</sup>

---

<sup>58</sup> To this end, his point strikes the verificationist criterion of meaning perhaps as the view that meanings are intensions:

How, then, should we formulate the “method of verification” – that is, the connection between a proposition *p* and the statements *s*<sub>1</sub>, *s*<sub>2</sub>...*s*<sub>*n*</sub> which are evidences for it? I propose...the evidences *s*<sub>1</sub>, *s*<sub>2</sub>...*s*<sub>*n*</sub> speak for or against the proposition *p*, that they strengthen or weaken it, which does not mean that they prove or disprove it strictly. (1968: 107)

<sup>59</sup> The idea of open texture was also intended just soften a number of dichotomies, most notably between the analytic and the synthetic. Two claims must be distinguished. There’s the banal claim that “analytic” and “synthetic”, like most words, have borderline cases between them. (Quine’s: all green things are extended; the bride’s emotional father: she was born our youngest daughter). The more contentious claim that open texture in the object language entails open texture in the metalinguistic terms “analytic” and “synthetic” requires more elaborate argument than Waismann supplies. For one of the ways this could go, see Richard 2019 for a discussion of analyticity at a time.

### 3. Modernizing

Verificationism has been gleefully buried many times over.<sup>60</sup> But if rightly so, there is a question about what feature, or features, of meaning Waismann's cases bring out. It is not hard to see why natural language meanings were once trifling: they cannot be kicked, and appear to defy exhaustive statement in other words (unlike some mathematical sentences). But metaphysical sentences have been brought back into the fold since the heyday of verificationism. It is now widely accepted that sentences have observation transcendent truth conditions, including sentences about natural language meaning. Waismann's cases, or ones like them, deserve another look.

Much philosophy of language has gone under the bridge since 1945 and it's hard to say what precisely these earlier figures had in mind. It's sometimes said that they ran together what we have learned to properly distinguish.<sup>61</sup> For example, the cases exemplifying open texture frequently involved what are now called "natural kinds", like cat, water, goldfinch, and gold. Yablo (2000), (2002), argues that the Kripke-Putnam's theory of rigid designation suggests such cases either involve a conceivability illusions (where we imagine discovering something with surface properties very like gold, but which is not), or else a skeptical scenario (where we imagine the underlying features giving rise to gold's characteristic properties are actually very different than evidence presently suggests they are). If things do turn out to be very different than we thought—if the underlying microphysical properties are not at all as believed—then we will rearrange.

---

<sup>60</sup> Wittgenstein 1953, Quine 1951, Kripke 1980.

<sup>61</sup> Austin, master of distinctions, illustrates:

We say, for example, that a certain statement is exaggerated or vague or bald, a description somewhat rough or misleading or not very good, an account rather general or too concise. In cases like these it is pointless to insist on deciding in simple terms whether the statement is "true or false." Is it true or false that Belfast is north of London? That the galaxy is the shape of a fried egg? That Beethoven was a drunkard? That Wellington won the battle of Waterloo? There are various degrees and dimensions of success in making statements: the statements fit the facts always more or less loosely, in different ways on different occasions for different intents and purposes. What may score full marks in a general knowledge test may in other circumstances get a gamma. And even the most adroit of languages may fail to "work" in an abnormal situation or to cope, or cope reasonably simply, with novel discoveries: is it true or false that the dog goes round the cow? What, moreover, of the large class of cases where a statement is not so much false (or true) as out of place, inept ("All the signs of bread" said when the bread is before us)? (1950: 11-12)

By my count, he gestures towards about seven different phenomena, and inexplicably considers a dog running round a cow an unexpected turn (though perhaps slightly more exotic than James' squirrel round a tree). Among these perhaps is vagueness—is Belfast North of London?—which is distinguished from a dog going round a cow, and unexpected turns, however those may inexplicably be related.

Perhaps we will say “that’s not gold!” (At least that’s what we did say when learned to identify pyrite, fool’s gold).<sup>62</sup>

But as things stand, there is plenty of reason to believe the underlying properties of what we call “gold”—being Au—give rise to its distinctive properties, and plenty of reason to believe only they can do that.<sup>63</sup> The truth conditions associated “gold” encode facts about what this actually is, not information about how to apply “gold” taking other ways the world might be as actual. So, on the one hand, what it is reasonable to believe of gold is largely an a posteriori matter. On the other hand, the meaning of “gold” revealed from the armchair will not dictate how to rearrange *if* things go sideways, because meanings do not encode that sort of information. One way to take the argument is this: *meanings* are not open; but what it is reasonable to judge is, especially in the face of surprising circumstances. But in many cases, there remains a question about what those meanings are like. Notice that a Kripkean treatment of “goldfinch” does not explain why people have the reactions they do to Austin’s goldfinch case. At best, it delivers the result that goldfinches cannot explode, or cannot wisecrack, but it doesn’t explain why we are semantically dumbfounded, or else dogmatic.

Take a more straightforward case. I drink, and bathe in, water, which is necessarily H<sub>2</sub>O, but I do not know that. If the chemists announced that water is composed of this substance they call XYZ, I would defer to them, or, express faint interest, and go on reporting that I drink, and bathe in, water. Words would not fail me. Austin’s example is supposedly one where even the experts say, “My God, what would you say?”, and ordinary speakers are supposed to also have the reaction that they have no idea what the expert would say, or that it is totally meaningless, or that something else has gone badly wrong. This is the dumbfounded response. Assume there is this difference when it comes to the reference of natural kind terms like “goldfinch” and “water” and non-natural kind terms (and I’m somewhat skeptical there is). Still, rigid designation doesn’t seem like it can explain the difference in our reaction to Austin’s case and the water case.<sup>64</sup> The same

---

<sup>62</sup> Kripke 1980

<sup>63</sup> See Yablo 2000, 2002. He identifies a related thought in Waismann, in developing arguments against anti-physicalism and two-dimensional approaches to conceptual analysis.

<sup>64</sup> The same may be said for children’s stories where mice marry. What’s fictional is that they have a townhall, and a sense of propriety, not that, *somehow*, the marriage relation can obtain between rodents. *Stuart Little* is not *conceptual* fiction. If Peter Rabbit is not a rabbit, the reason is that fictions are not rabbits, or fictions do not exist. The reason is *not* that it is an a posteriori necessity that rabbits lack the mechanisms to speak (in contrast to Farmer

goes for dogmatic responses to the case, say, that, it is perfectly obvious that, then, goldfinches wisecrack and explode.<sup>65</sup>

Being stymied as to what to say, may be chocked up to a hesitance to generate the implicatures assertion would herald, rather than a feature of the meaning of terms involved (“Ow! Why didn’t you warn me about the chair?”). There are many reasons for an utterance to be inappropriate, and being untrue just one of them.<sup>66</sup> Perhaps semantic dogmatism is not so much a semantic phenomenon, as a pragmatic one: speakers communicate diagonal propositions as a means of narrowing down live options for ways the world might be in the context of the conversation. They do not utter true sentences (and may be ignorant about the meaning of “goldfinch”), but they do adjust the context set correctly.<sup>67</sup>

But neither response obviously catches the full the story. In 1899, Rutherford made the shocking discovery that atoms were mostly empty space. Years later he remarked: “It was almost as incredible as if you fired a 15-inch shell at a piece of tissue paper and it came back and hit you”. A classic example of semantic dogmatism followed in the exchange between Sir Arthur Eddington (science popularizer and physicist) and L. Susan Stebbing about how to describe the discovery. Eddington’s position was that science shows us that floorboards are not solid (1927/1958). In his so-called “two tables problem”, he wonders how to reconcile the table of science, which is not solid, with the table of everyday experience, which is. Stebbing infamously maintained, in direct

---

MacGregor). That seems right, whatever else we may say about the truth value of sentences in fictions. To hive off this phenomenon to a theory of *fiction* strikes me as misconstruing ordinary language in a fairly deep way.

<sup>65</sup> Also notice science fiction-like speculation mostly comes out as intelligible. As it turns out, it is not physically possible to make capsules nutrient dense enough to replace meals, and it is physically possible for a woman to bear a child without donating any gametes. Does it follow that it is a posteriori necessary that meal replacement capsules are not food *on account of this fact*? Surely, at one point speakers would have avowed that “mother” (with its biological meaning) rigidly designated the woman who gestated a child, and bore those risks, and any suggestion otherwise was subject changing of the highest order. One reaction is that, in the case of “mother”, we were mistaken, and that the evidence can underdetermine theory in surprising ways. The rise of gestational surrogacy has prompted further collective revising of our theory of motherhood. But however we go on does not appear to have been *antecedently* baked into the meaning of “mother” (“food”), or to herald changing the subject from *mothers* (*food*) to something else.

<sup>66</sup> See Grice 1989.

<sup>67</sup> See Stalnaker 1984, 1999a, 1999b. The semantic content of a sentence on an occasion of use is distinguished from the content asserted by an utterance of that sentence on that occasion. In general, the default option for the assertoric content is the semantic content of the utterances, but on this picture, the mechanisms of conversation—and peculiar revelation—sometimes force the two apart. In this connection, he outlines some principles for asserted content that systemize cases of divergence. Central to his framework is the notion of a context of communication, which is roughly the set of live options at each point in a conversation and a conception of the role of assertion as that of cutting down the live options and so narrowing the context.

reply, that most obviously floorboards, and tables, are solid, and science cannot overturn commonsensical usage (1937). Whatever else we may say about the exchange, neither was confused about the discovery, and neither were their readers. The dispute concerned the application of terms in ordinary English descriptions, and plausibly not what they implicated. There was certainly no deep difference in the readers of Eddington and Stebbings, or their context. And, again, the Kripkean strategy does little to explain the strength of conviction on either side.

There are many cases structurally similar to Waismann's, et al.'s, but that are at least not obviously well-resolved by appeal to the machinery of rigid designation. For one, they do not involve science fiction. For another, they do not include terms that, at the time, were good candidates for being rigid designators (how many of these there are depends on who you ask).

#### 4. Cases and Modernizing

Five cases will be helpful. I will call the phenomena they exemplify "open texture" as a nod to Waismann, but without more pretense of doing exegetical work. I am in good company. Waismann's evocative phrase, and cases of the sort described, have been recruited in support of a wide range of philosophical positions, if rarely in a starring role.<sup>68</sup> The cases may exemplify more than one phenomenon. That's fine. The rest of the paper will argue the balance of these cases are not well analyzed using familiar semantic tools.

We have already met the first and third cases.<sup>69</sup> Each apparently involves a linguistic choice about what applications of a predicate should serve as precedent, and which similarities should count. The choice is triggered by learning that criteria which in the past we explicitly, or implicitly, believed were jointly satisfied, and would continue to be, are not. Some of these choices passed unnoticed, and some involve reassessing what we once took applications of the term to entail.

**"Window":** Are glass walls windows? It is hard to say. Yes, insofar as windows are views to the world. No, insofar as windows are lodged into walls. Our judgements as competent speakers pull in different directions. Mostly we settle on calling glass walls "windows", but our hand does not seem *forced* by the meaning of "window". For instance, if there were a readily available term for both windows and glass walls, we might have been more disposed to deny glass walls are windows.

---

<sup>68</sup> Legal theory is a notable exception. See Hart 1963, Schauer 2008.

<sup>69</sup> Many more cases are to be found in Green 2011, LaPorte Wilson 1980, 2006. See also Brown 2000, Jackman 1999, 2005, LaPorte 2000.

**“Blackbirds”:** The first English-speaking colonists arrived in the New World with a stock of classificatory predicates that had unproblematic applications in the Old World. These included “blackbird”, as well as “rabbit”, “oriole”, “robin”, and others. Their common linguistic training disposed them to apply many of these predicates to New World objects. “Blackbird”, for example, was applied to 26 of the species in the family *Icterids*, completely different from the *Turdus merula*, members of the *Turdus*, or thrush, family called “blackbirds” in Europe. They had in common that they were mostly black (though the New World birds had bright batches of color, brighter than many Old World black non-blackbirds, like starlings). Today, we distinguish Old World and New World blackbirds (the former being *Turdus merula*, the latter 26 species of icterids that do not correspond to a formal taxon).

**“Solid”:** Eddington (1927/1958) told the public that science reveals that a floorboard “has no solidity of substance” (39) because it is composed of elementary particles that lie comparatively far apart. Stebbing (1937) maintains Eddington abuses non-scientific language: “The point is that the common use of language enables us to attribute a meaning to the phrase ‘a solid plank’; but there is no common usage of language that provides a meaning for the word ‘solid’ that would make sense to say that the plank on which I stand is not solid” (51-52).<sup>70</sup> The heart of her complaint is that a predicate must possess “some usage”, or something like a fixed core meaning, from which its subsidiary uses must descend. She suggests, these satellite uses of “solid” cannot completely overturn the parent discriminations from which they derive.<sup>71</sup> Whatever the (de) merits of her points, Stebbing and Eddington appear to disagree about how to describe recent scientific discoveries using the ordinary term “solid.”

**“Athlete”:** In 1999, ESPN ranked Secretariat the 35<sup>th</sup> of the 100 Greatest North American Athletes of the 20<sup>th</sup> Century. This occasioned incredulous fat chewing by sportscasters.<sup>72</sup> Can non-human animals even be athletes?

**“Planet”:** In 2006, the International Astronomical Union (IAU) voted to augment the definition of “planet” in a way that foreseeably resulted in Pluto’s demotion: it included a clause limiting the category of planet to those bodies that have cleared their orbit of debris. This appears to be not so much a discovery as a decision. The impetus for the vote was that the existing definition of “planet”—a large round body that orbits a sun—applied to the recently discovered Eris, along with perhaps some 200 smaller trans-Neptunian objects in the Kuiper belt. The vote itself was highly contentious among participants, and deliberations took into account everyday considerations, e.g., about public sentiment about planets and the costs of updating maps, as well as scientific ones. The underlying features of planets

---

<sup>70</sup> The expanded quote is sometimes cited as a *locus classicus* of the paradigm case argument.

<sup>71</sup> Her overriding purpose is to argue that the deliverances of science must err (or be regarded purely instrumentally) whenever they reverse the opinions of everyday common sense, rather than supplement them with new vocabulary. Part of her objection to Eddington (unquoted) stems from her apparent skepticism that “solid” is multi-valued. She appears to assume that because the original usages didn’t apparently *entail* the empty space in floorboards, Eddington’s usages are false.

<sup>72</sup> See Ludlow 2014. He recruits cases like this in his arguments that a language is best thought of, not as a stable abstract object, but a series of fleeting microlanguages.

turned out to be less significant to us than maintaining a certain sort of size and unity among the objects so classified.

Or course, contemporary linguists do not set out to define words like “window” and “athlete”, and every philosopher is ready with an argument that conceptual analysis is hopeless. Never mind the project of exhaustively stating what athletes or windows are in other words. What apparently endemic feature of meaning gives rise to cases like these?

The suspects usually thought to determine reference—patterns of deference to experts and the natural world, usage in the broader community, speakers’ dispositions, paradigms and patterns of association—do not settle how to proceed with the predicate. It is unclear, for instance, what aspect of the world, or which experts, to defer to.<sup>73</sup>

The meaning of “solid” apparently did not legislate how to describe the discovery. The meaning of “blackbird” as used in the Old World did not destine it to correctly apply to New World creatures, nor prohibited it. Speakers, for example, did not appear to defer to some original baptism, as when Europeans applied Marco Polo’s name for Mozambique mistakenly to the island off its coast. While “Madagascar” eventually acquired a new meaning, New World use of “blackbird” appeared to mean as it did in the Old World.<sup>74</sup> That the editors at ESPN, the sportscasters were experts on athletes and athletics does not appear to advance the question of whether Secretariat is an athlete. Likewise, that the broader public is not first disposed to take him to be. Non-human animals are certainly not paradigmatic athletes, but ranking calls into question the dispositions of many competent speakers to affirm being an athlete entails being human. It may seem, on reflection, that nothing in the meaning of “athlete” forces a verdict, despite how the expression had been used to date. But perhaps you vehemently disagree. That is, the case of “athlete” may trigger semantic dumbfounding in some speakers, but semantic dogmatism in others.

The IAU resolved the linguistic choice concerning “planet” by fiat. That it is now false to say that Pluto is a planet speaks to how strong a trendsetter it is (as compared, say, to the editors at ESPN). But mostly resolution is piecemeal, facilitated by confluence in the pressures that drive usage (of which bureaucratic fiat is one). Some are internal to meaning, involving the properties regulating competent use. Others are external, involving our cognitive limitations, interests in

---

<sup>73</sup> See Putnam 1975, Burge 1979, Kripke 1980.

<sup>74</sup> See Evans 1985 for a discussion of “Madagascar”.



coordination, matters of convenience and salience, the order of events, and so on. If Rutherford's discoveries had upended a widely endorsed theory about the high density of atoms, and been more controversial, perhaps today we would deny that tables are solid.

What about utterances of "Pluto is a planet" made prior to 2006? One way of assigning truth-values to utterances tells us that pre-2006 utterances of "Pluto is a planet" are true. But this is not obviously correct: textbooks were rewritten, the general feeling is earlier scientists armed with partial information had erred, or else that Pluto has been demoted, and unfairly. These considerations do not point to "planet" acquiring a new meaning—in effect, to vocabulary change—so much as changed beliefs about *planets*. But, insofar as planethood is as the IAU says it is, the fact that the IAU could have voted the other way does not so much reveal the underdetermination of theory by data, as something about the way meanings are ordinarily conceived: we can go on in more than one way with a predicate without sinning against it, or changing vocabulary.

In many ways we are creatures like others. Plausibly investigating the semantics and metaseantics for parrot thought involves investigating what properties they are thinking about. I take it there is little temptation to say that these properties are fully determinate and also little temptation to say that when a parrot stares at a puddle it is contemplating H<sub>2</sub>O. But we can happily attempt to look at the structure of parrot concepts and their deployment. So, what's different between parrots and us? We have scientific discourse, and they do not. Even so mostly *we* are not engaging in scientific discourse. It is a banal observation that we, like parrots, have not anticipated all of the cases in which we might be puzzled as to whether a particular word or concept applies. It is a less banal, but no less correct, observation that our dispositions, worldly knowledge, and environmental relations do not come close to determining what our reactions to novel cases will or should be. We can imagine maintaining this even upon an alleged discovery that these factors do associate occult intensions with our expressions, in accord with some absolute and undiscoverable metaseantic law, or metaphysical fact.<sup>75</sup> Appealing to these positions as treatments of our open cases (as opposed to the reference of natural kind terms, or vague predicates) sounds on par with Leibniz's commitment that this is the best of all possible worlds.

---

<sup>75</sup> See Dasgupta 2018.

Just as it seems out of place to appeal to some such law, or metric of naturalness of a property, or of a thought, to render the parrot's thoughts determinate, so too our own. If there is something that distinguishes us most obviously from parrots, it is that we self-correct and seek coherence in our statements and judgments—of all sorts—over time.

What then is the feature of linguistic meaning that gives rise to the appearance of openness, or choice? At first glance, the cases do not turn on under-specificity. “Between 9 and 5” means something, and it means something very precise. But it is not as informative as “between 9 and 10”. “Window”, for example, is under specific in the sense that its truth conditions do not distinguish portholes from bay windows. But our question is whether applications of “window” distinguishes those options from glass walls. Neither do our cases obviously involve the phenomena associated with figuration, or loose speech, in the way of “France is a hexagon” or “The eyes are the window on the soul”. Our open cases can be framed as existence questions, where the answer we give is ontologically committal (“Do non-human athletes exist?”). Typically, figurative speech is not seen as ontologically committal.

But why would we lack knowledge of such things even if we found out all the underlying facts about Secretariats, what sports experts think, what the public thinks, and so on? What kind of status is this, where the question is not only something we cannot answer, but where we cannot conceive of getting in a position to answer it? It may seem as though there is no fact of the matter about whether Secretariat is an athlete, and that we can feel puzzled, ignorant, is a spandrel, generated by load bearing features of the expressions involved, but of no other significance.

## 5. Vagueness

Sometimes inquiry resistance of this sort is thought to be a hallmark of vagueness. Many of the predicates that figure in the open cases are vague, like so many predicates of natural language. But the openness of open cases do not appear to turn on their vagueness, but rather on something else. While a neutral characterize of vagueness is hard to come by, vaguely, it involves unclarity at boundaries.<sup>76</sup> We have no independent handle on what it is for a predicate to be vague in advance of talking about borderline cases, and the notion of being borderline is explained by pointing to a certain sort of case. “Bald” is paradigmatically vague: it is unclear to competent speakers where,

---

<sup>76</sup> Raffman 2014

say, the bald ends, and the not-bald begins. Roughly, Jim's scalp is a borderline case of "bald" because he is neither clearly bald nor clearly non-bald.<sup>77</sup>

The same goes for open texture. A predicate has it, is open, precisely because there are some open cases involving that predicate. These have some characteristic features, including semantic dumbfounding and semantic dogmatism, the revising of judgments about paradigm cases, and what is entailed by a term's application, a feeling that a linguistic choice is being made. Not all cases have each of these features. Maybe it's not totally clear there's one phenomenon. Even so, these open cases have several features that borderline cases lack.

As already noted cases of vagueness do not elicit semantic dumbfounding, or dogmatism, as open cases do. Instead they elicit equivocality. Competent speakers may agree it is unclear whether the man is tall (when apprised of the relevant comparison classes for tallness and facts about the height of any relevant populations). Likewise they may agree it is unclear whether the plaster is solid. Perhaps doing so is relevant to competence. But agreement seems unavailable in Stebbing/Eddington case of "solid": some speakers hold there is no answer as to whether tables are solid, others that they clearly are, others that they clearly are not. Were they to have come upon some crumbling plaster, Stebbing and Eddington would presumably agree that it is was not clearly solid, but also not clearly not solid. That is a reason to floorboards, unlike crumbling plaster, are not borderline solid.

Consideration of some open cases may undermine judgments about apparently core cases—say, that floorboards are solid—but consideration of borderline cases does not undermine apparently core cases. (Whether the plaster is solid does not tell against the tables being, because the reasons associated with whether the plaster is solid, or not, are irrelevant to the assessment of table. Resolving borderlines does not tell against the core reasons competent speakers apply "solid." Consideration of some open cases may undermine established judgments about entailments—say, whether being an athlete entails being human—but consideration of borderline cases also does not appear to do so.

---

<sup>77</sup> Sometimes soriticality is thought to be sufficient for vagueness (dubious, but see Bueno and Colyvan 2012), or at least diagnostic. But open cases do not appear to involve soritical reasoning. Sometimes gradeability is treated as diagnostic of vagueness, see Shapiro 2006, Raffman 2014, Fara 2000 for discussion, as well as, collectively, arguments that none of these are necessary or sufficient for vagueness. Our open cases also mostly do not involve gradeable expressions. While "solid" has a gradeable meaning, that is not the one at issue in the case.

Again, I'm not denying the open predicates are vague. They are. The point is their vagueness does not explain what is open about the open cases. Being open has importantly different features than being borderline. For example, we can contrive a case for a paradigmatically vague word—"bald"—where there is semantic dumfounding. Imagine a guy with two heads, one of which is as hairy as any, the other of which as bald as any. It is indeterminate whether to call this person bald or not, but it has nothing to do with borderline cases, or incremental reasoning: the heads are individually clearly either bald or not-bald, there is no borderline in-between state involved. Is he bald? We throw up our hands. "Bald" is open, but *not* because of its having borderline cases (or its soriticality). It is open because we can imagine an "odd case" outside the usual domain of application.

Another distinction between open texture and vagueness is this. Vagueness involves unclarity about the boundary in orderings of individuals, say, from the red to the not-red (in the blue direction), the red to the not-red (in the green direction), and so on. Maybe there is unclarity about whether an individual is red that falls in thus and so way on one ordering, thus and so way on another, and so forth on a third ordering. One way an expression can be vague is if there is unclarity about how to relate multiple partial orders.

Plausibly, open texture is unclarity about *which* orderings are relevant in the first place, independently of where to draw the boundaries on any ordering, or how to aggregate them. If you have only 10 hairs, you are definitely bald, but if you have only 10 hairs and they are an inch thick, 8-feet long and braided atop your head, you are what exactly? That is, it is unclear whether orderings from the thin to the thick, and the short to the long, are relevant orderings for determinations of baldness at all (in addition to orderings from the few hairs to the many)? For open predicates there may be a question about which orderings, in contrast to a question about the weighting of the extant parameters associated with the term. It is not, as it is in cases of vagueness, a question of boundaries on a linear ordering or a multi-linear ordering.

Vagueness, for example, differs from the kind of indeterminacy found, say, in the claim that Marlowe liked magic acts (Joseph Conrad did not say). In principle this kind of unclarity could be removed without changing the meanings of any of the worlds involved. (Conrad *could* have said). In contrast, vagueness is a permanent unclarity whose removal would seem to change the meaning, by threatening the word's communicative utility in surprising ways. (It would be difficult

to express, for instance, common dating preferences—e.g., *she likes dating tall men*—without overselling the point and our powers of discrimination).<sup>78</sup>

Open predicates exhibit a kind of unclarity for certain outré fact patterns – where when the unclarity is removed, and application is settled for those fact patterns, it does not *seem* the expression has changed its meaning. “Solid”, and “blackbird”, do not *appear* to have changed their meanings. Many people would say “planet” hasn’t either; it just turns out, or was decided, that Pluto isn’t one. (We report early speakers homophonically, and disquote for counterfactual reports, e.g., if Newton and Carl Sagan were here, we all talk about the planets). Surely the changes these expressions have undergone are larger than the changes “bald” would undergo if a fixed number of hairs was settled on. As choice points for a word are encountered and resolved, the meaning of a term appears to be developed and enriched. By contrast, the vagueness of a term is a lasting feature of its meaning and does not indicate a similar place for enrichment.

The “window” case is the most likely of our cases to turn on vagueness, but even here it’s not clear. Glass walls might seem like borderline windows: they are sort of windows, sort of not. Considering whether they are doesn’t obviously elicit semantic dumbfounding or dogmatism, though we do tend to side on calling them “windows”, though it may have the feel of loose speech. Considering whether glass walls are windows does not threaten to undermine core judgments, but it does threaten to undermine an established entailment (e.g., if there’s a window on that side, part of that is opaque; if my hand is on the one window, it is not on the wall). Removing the unclarity about glass walls might threaten the communicative utility of “window” in a straightforward way because we contrast windows and walls. This last is a reason to think it’s open whether glass walls are windows, not vague.

So is this. We can construct several partial orderings for “window”, ranging from the translucent to the opaque, the openable to the unopenable, the minute to the very large, the non-load bearing to the load-bearing. Locating a prospective window on each ordering may yield an inconclusive verdict on whether, or not, a tiny cracked pane in a cracked wall is a window—it’s a borderline window. Arguably, the wall situation is different. There uncertainty about which orderings matter, not how placement on some given orderings do. Perhaps it is vague whether

---

<sup>78</sup> See Schiffer 2013, Wright 2003.

glass walls are windows, but I tend to think it is open. Regardless, the above considerations are far more decisive for the other open cases. It is not vague whether Eris was a planet, or floorboards were solid, or merulas were blackbirds, or Secretariat is an athlete.

So, open texture is not vagueness, although open predicates may be vague (at least many are). We can generate open cases from paradigmatically vague expressions like “bald” but which are not themselves borderline cases (or higher order borderlines). While open texture elicits semantic dumbfounding or dogmatism, vagueness elicits equivocality distinctive of borderline cases. While precisifying a vague predicate appears to change its meaning, the same is not clearly so for the feature that gives rise to openness.

## 6. Partial Definition

We’ve dispensed with vagueness as a model of open texture. Open texture is not vagueness. Now we’ll consider some other models of it, some of which themselves have been used as a model for vagueness. Partial definition is one, as is context-sensitivity and supervaluation (which I’ll consider the next two sections). A natural thought is that open predicates are semantically incomplete, only partly defined. It is as if we, The Definer, or the metasemantic laws, forgot to map, e.g., “is a blackbird” to a truth value for as yet encountered lepora. Take Soames’ example of “smidget” (1999). It is stipulatively defined such that:

- (i) any adult human being under three feet in height is a smidget
- (ii) any adult human being over four feet in height is not a smidget
- (iii) anything that is not an adult human being is not a smidget.

This predicate is not vague; its conditions of application are perfectly precise. But it contains a perfectly precise gap: for some objects “smidget” is silent. Confronted with Bill, who is three feet six inches tall, we should accept neither, “Bill is a smidget”, nor “Bill is not a smidget”. But then according to the account of truth which Soames (following Kripke) proposes, neither should we accept, “Bill is a smidget” is true, nor “Bill is a smidget” is false. So, any actual utterance of ‘Bill is a smidget’ will fail to get a truth value. Soames argues if we know that “smidget” is undefined

for an object, we should reject both “*a* is a smidget” and its negation (where *a* names the object).<sup>79</sup> He maintains that “applies to” is also partially defined, and so we must reject both “smidget” applies to *a*”, and its negation. I’ll argue this is not a good model for open texture because in open cases are a) unforeseeable and b) there are considerations that bear on their resolution (apart from it being unclear whether there are any semantically incomplete predicates). Neither of these applies to semantically incomplete predicates (if there are any).

Was “blackbird” circa 1620 like “smidget”? It is easy to get into the spirit. Consider some candidates the meaning of “blackbird” in 1600: “Blackbird” applies to *x* if and only if: (1) *x* falls under OLD, (2) *x* falls under NEW, (3) *x* falls under MIXED where OLD, NEW, and MIXED determine respectively {*x*|*x* is a *Turdus merula*}, {*x*|*x* is an *Icterid* of one of 26 species}, and {*x*|*x* is in the Old World and *x* is a *Turdus*, or *x* is not in the Old World and *x* is an *Icterid* of one of 26 species}. These assignments are pairwise mutually exclusive; severally each is problematic. (1) makes false most of their utterances about New World blackbirds (“Blackbirds ate the pumpkins”). (2) makes false most of their standing sentences about Old World blackbirds (“The blackbirds do back home are numerous”). Such assignments are likely to disconnect their language from the world so that it will be hard to make sense of the practical activities they use language to engage in. (3) seems strained since they were able to distinguish Old World blackbirds from Old World black birds, as both are to be found, and were linguistically distinguished as “blackbird” and an array of other terms for other black birds, in the Old World. Moreover, claiming they meant MIXED by “blackbird” all along amounts to claiming rather mysteriously that they were—at least their predicate was—prepared for an eventuality that was completely outside their ken, including the distinguishing 26 regional species from the more numerous and sometimes blacker members of *Icterid*. Appealing to the way of the natural world as a determinant of extension fall short of answering the question: replace the concepts OLD, NEW, and the relation falling under a concept in conditions (1) and (2) with the relevant Linnaean kinds, and the relation being a member of a kind, respectively, and repeat the argument. But if the colonists were prone to this kind of error, what makes us think we aren’t? These are the considerations in favor of thinking that in 1600 the

---

<sup>79</sup> Soames employs an account along these lines, invoking the apparatus of Kripke’s theory of truth, to justify the rejection of the instance of Excluded Middle and figures in his modified contextualism about vagueness.

extension of “blackbird” is semantically incomplete: (4)  $\{x|x \text{ is one of the Old World } Turdus \text{ merula } \dots\}$

But there are problems. First, it is unclear whether *any* predicates are partially defined, as opposed to completely defined partial-predicates (let alone any predicates of natural language).<sup>80</sup> For example, suppose one stipulates all positive integers map to truth, and all negative ones to falsity, but makes no stipulation for zero. On a traditional picture, one fails to specify a concept. Some conditions have been laid down, but they are met both by a concept mapping 0 to truth and by one mapping it falsity (two distinct natural number concepts). It would be wrong to say one has specified a *partial* concept, a concept undefined for some objects. To do that, one would need to stipulate that the concept be undefined for zero, but one has not even got that far. To fail to stipulate a value is not to stipulate that there no value. Frege, for example, insists that some stipulation be made even for non-numerical cases, such as the moon. Unless a stipulation has been made for every case, nothing has been defined.

Seen this way the metaphor of “forgetting to define a truth value” looks suspect. Suppose the various metasemantic forces that collectively determine a truth value fail to do so for “blackbird” as applied to some members of *Icterid*. Nothing has been *forgotten*: there is *indeterminacy* among fully determinate extensions. There is familiar indeterminacy among fully defined concepts, where such indeterminacy is representable as determinacy (see Section 10). A predicate where some terms are *defined* as undefined is not partial, *silent* as to some objects; rather “true” is three-valued. Mathematicians sometimes define a function to be “undefined” in certain cases; this is quite different from not making a stipulation for those cases. Likewise, stipulating that “smidget” be thus and so, and otherwise undefined. There is nothing particularly *partial* about this.

The second problem with analyzing “blackbird” as a case of semantic incompleteness is phenomenologically inapt. The case didn’t lead to semantic dumbfounding, so much as a sort of dogmatism: New World birds were called “blackbirds” almost immediately upon settlers’ arrival. There were strong considerations that weigh in favor of going this way or that (and that in fact did). The definition of “smidget” gives no guidance as to whether someone between the heights of

---

<sup>80</sup> Frege 1884/1974, Williamson 1994.



the tallest A member and the shortest B member is a smidget. There are, by stipulation, no considerations that can bear on whether Mr. Smallman is one. If the “blackbird” case turns on its semantic incompleteness, we should expect an absence of considerations bearing on its application to New World icterids. But this is not so. We know, or at least have a good sense, of what sorts of considerations matter: overall look, behavior, habitat, tastiness, comparison to other locals. Competence has something to do with meaning, and we may suppose there are no confounding factors at play.

The third problem is familiar from the last section: “smidget” is semantically incomplete in an entirely foreseeable way because “smidget” is defined over an ordering of heights (of non-child humans) beginning with shortest person in A and ending with the tallest person in B. The situation with the rabbits and blackbirds, George, is not like this. There are a variety of partial orderings involving the considerations that bear on the application of “blackbird”. The question is how these relate.

There may be items where application of “smidget” to them seems open in the way application of “blackbird” to an icterid is. For instance, suppose someone taller than the tallest member of A but shorter than the shortest member of B is perpetually stooped such that, stooped, they are shorter than the tallest member of A. Are they a smidget? The issue of partial definition does not appear to be what makes our cases vexing.

## 7. Context Sensitivity

We’ve left behind partial definition, and vagueness, as models of open texture. Perhaps context-sensitivity (apart from these) will better. It does not. Analyzing openness as context-sensitivity is to give up any hope of illuminating the phenomenon.

Open cases are suggestive of questions that seem defective. Take Eddington: “Are tables *really* solid? Science describes the tables as not solid, even though everyday talk does.” Or the IAU: “Is Pluto *really* planet?” These questions don’t have the same defect as asking whether or not swatch #376 really is red, or whether Mr. Smallman is a smidget. They also do not have the same defect as asking whether the ball is *really* on the left, never mind the viewer’s perspective. Knowing whether some icterid is a blackbird, or wall is a window, or horse is an athlete, does not appear to be a matter of knowing the context of utterance. Many expressions are indexical, like “I”, or otherwise contain unarticulated constituents, like “left”. The truth conditions of, “My pants

are on fire”, or “The ball is to the left”, vary with some semantically encoded parameter whose value is determined by the context of utterance (if not facts about it).

Of course, plenty of open predicates are context-sensitive. “Solid” likely has unarticulated constituents (one corresponding to relevant applied force: cakes are solid at a soda shop, but, at the hambershop, cakes are not solid). The ordinary celestial term “planet” does too: Arkas is a planet, just not in this solar system. But the openness of open cases does not determine on the context sensitivity.

For one, arguments that such and so expression is context sensitive turns on robust agreement among competent speakers about utterance truth that is taken to be straightforwardly indicative of meaning. But this is precisely what open cases lack, or, when they have it, it does not appear demanded by their meaning.<sup>81</sup> That speakers may come to agree suggests that meanings may be enriched in ways that persist across contexts. But unlike context-sensitive variation, linguistic choices govern usage from that point forward, across contexts. The mere fact that we can generate situations in which it is appropriate to say that “this horse is an athlete” (as opposed to that incorrigibly lazy one over there) is not by itself a reason to suppose the utterance is *true* as opposed to triggers a true implicature (which was the point of making the utterance in the first place).

Let’s say “planet” is context sensitive in that its content varies with whether the context of utterance is before or after the IAU vote. As used in 1960, Pluto is in the extension of “planet”. As used in 2020, it is not. This is not a change in character, and so isn’t a change in meaning. A proposal like this can be understood in various ways depending on what’s meant by talk of properties. A property might be something with a modal profile—it is, or determines, a function that takes a world at a time to an extension. Or it might simply be a function from times to extensions. My own feeling is that if we adopt either proposal, or anything much like them, we have thrown in the towel. It is not clear why we should adopt such a proposal about “planet” unless we’re prepared to say the same sort of thing about every other seemingly non-context sensitive term. If we say this about the word “planet”, we are, in effect, saying that the word “planet” can’t change its meaning. For what we are saying is that speakers at some point started using “planet”

---

<sup>81</sup> See Fodor 1964.

to pick out a property accompanying the use of the word with a set of common paradigms, presuppositions, criteria of application, whathaveyou. These latter changed—by conscious regulation at least twice in the last 100 years, and several times before that—and those changes may or may not cause changes in the property picked out. But neither change in them nor in the property picked out counts as a change of meaning. In this way, the truth conditions of *planet*-talk may vary as the truth conditions of *being to the left of*-talk do, in a diachronic way (in addition to relevant solar system way).

You *can* say that, but it is a just a sneaky way of saying that the notion of meaning has no descriptive, or explanatory, heft.<sup>82</sup> Either the character of “planet” is highly intricate, a content corresponding to each future IAU vote, or all possible ones, or ...There’s certainly no formula from ways the world is to which contents describe the character. (Contextualism about “know” fares much, much better). Moreover, this strategy seems to suppose there is some antecedent fact about what the contours of the character of “planet” are. And, as in the discussion of partial definition, is this is precisely what seems to be missing, yet there are reasons that bear on the application of the term going forward.

Saying that “planet” has an extendable character does not advance the issue. At present, in the very same context, I can use “planet”-1960 and “planet”-2020. So, “planet” does not appear context-sensitive in this way. Talk of extendable character just repackages what was originally puzzling about the case.

Now sometimes the right semantics for an expression may well be indexing the function associated with an expression to the time of utterance. Pace Richard 2019, “pasta” may be a good candidate. Uttered in 1800, the sentence “pasta can be made of bacon agar” may well be false, but utterances of that sentence in 2020 are true. “Pasta” has not changed its character, although, at different times, what is true of *pasta* has changed. The kind pasta, as it were, has grown up. Genres (culinary, literary, musical, and so on), institutions (like college), works in progress (*my thesis*) may be particularly amenable to this treatment. But many expressions do not pick out kinds that change their contours. Perhaps it may be argued that the possibility of creating open cases reveal precisely the opposite: that terms like “planet”, and “blackbird” do, in fact, pick out a single kind

---

<sup>82</sup> That’s part of why I take it in the early days of contextualism about knowledge it was taken to be a kind of deep concession to musty ordinary language philosophy. For discussion, and anecdotes, see DeRose 2005.

with shifting contours, and how those contours shift depends on speakers in some way. But this is highly unappealing. If someone says: “Pluto is not a planet anymore, but it used to be, just like merulas weren’t blackbirds but they are now.” The right reaction, plausibly, is: what happened to *Pluto*? Did the birds interbreed? Did they get darker? (Perhaps *athlete* is somewhat plausible candidate for a version of this treatment).

More to the point, there is little reason to suppose that, all along, our terms pick out kinds that “grow up.” Suppose “caterpillars” came into use before their life cycle was well known. When it was learned, surely speakers could have gone on to talk of caterpillars getting their wings. Caterpillars grow up to be butterflies, and “caterpillar” “grew up” to refer to a phase in the life cycle. But there is, I submit, little sense in claiming all along its meaning destined this result. Is there *more* sense to saying that, say, “book” all along referred to a kind that would “grow up” to include eBooks?<sup>83</sup>

Some predicates that figure in open cases are context sensitive, but not in a way that explains the openness of those cases

## 8. Underdetermination

We have dispatched analyzing open texture as a sort of vagueness, partial definition, or context-sensitivity. We may wonder whether the apparent openness of the cases owes to semantic underdetermination. Perhaps our so-called open cases are open because they lack the pragmatic modulations that would supply the utterances truth evaluable content. So-called truth conditional pragmatics (TCP) holds that linguistic meaning of an (ordinary, non-indexical) expressions *need*

---

<sup>83</sup> Here is a real case with similar contours. Parasites often have bizarre and complicated life cycles. It often takes a lot of detailed research to link a parasite species at one stage with the same species at another stage. Sometimes independently identified and named parasites turn out to be different stages of a single species. Take, *Taenia multiceps*, a revolting but otherwise relatively harmless tapeworm found in the guts of rural dogs. Likewise gid (a neural disease of sheep) was known to be caused by *Cenuris cerebralis*, a cystic brain parasite. What should have happened when it turned out that *C. cerebralis* was the larval stage of *T. multiceps*? Both names were introduced by an implicit ‘that kind of parasite’ demonstration. Both in-the-head conception and prior expertise defer to hidden indexicality, and we would have discovered that “*T. multiceps*” and “*C. cerebralis*” referred to the same kind and therefore meant the same thing all along. Upon that discovery, would we decide which synonym to keep as the official species name, and the other would fall into disuse? This isn’t quite what happened. *C. Cerebralis* has gone on to be used for the larval stage of *T. multiceps* (which has gone on to be the name for the organism). What should make it that *T. multiceps*, but not “butterfly”, refers to the organism rather than a phase stage? Surely all along there was no fact of the matter about which term referred to which property.

*not be* what the expression contributes to propositional content. The speech act contributes aspects of truth-conditional content that are not derivable from the linguistic properties of the sentence (even taking on board the mandatory assignment of contextual values to indexicals and free variables). This is compatible with the idea that the linguistic properties of the uttered sentence (plus saturation) do determine a minimal, truth-evaluable content, which may or may not coincide with the content of the locutionary act.<sup>84</sup>

It is an empirical matter how many expressions are subject to pragmatic modulation and enrichment (the process by which an interpreter identifies a more particular relation to take as the meaning of an expression on a particular occasion of use). Perhaps only a few expressions, like “ready,” are although defenders suggest many more. It may be argued that they are open precisely because the circumstances supplied lacks the pragmatic considerations that would “swoop in” and determine a truth evaluable content

An example illustrates the contrast between TCP and truth conditional semantics (TCS). Travis 1975 and Kripke 2011 each analyze a case they attribute to Austin but to different effect. The case runs as follows. A parent takes a child to a toy store where the toys are plastic models of various animals. The child asks, “Is that a goose?” The parent says, “No, that’s a duck.” What can we conclude about the meaning of “duck” from such an example? Kripke argues that to conclude much of anything commits the “toy duck fallacy”: what the parent says is literally false (although what is implicated is true).<sup>85</sup> TCP unabashedly commits the “toy duck fallacy”: the parent’s utterance is true, and true just in case it is a duck, in this case a toy one. The restriction to the toys is the result, not of actual inferencing, but of pragmatic processing of the utterance in an under the hood sort of way. Which prediction is right? According to TCP, the minimal proposition it’s a duck is not even available to language users. If asked, any ordinary speakers who has just produced

---

<sup>84</sup> A stronger claim—sometimes called “radical contextualism”—takes content to be primarily a speech act notion: sentences possess contents only insofar as they inherit the content of the speech acts, they are used to perform. It follows that sentences by themselves do not carry (even minimal) truth-evaluable contents (Travis). They carry schematic meanings which only determine truth-evaluable contents in the context of a speech act, and how this goes is also not derivable from linguistic properties of the sentence. It is not semantically encoded. Radical Contextualism entails Truth Conditional Pragmatics, but not the other way around.

<sup>85</sup> Kripke writes: “One should not conclude that something is a legitimate use of language, something that ought to be recognized by a dictionary or semantic theory, even though there are appropriate circumstances in which one would say it” (345).

an utterance like yours will—so it is claimed—deny having said any such thing. (“What? A real duck? Who said anything about that?”), as opposed to demurring.

Perhaps then “Secretariat is an athlete”, or “glass walls are windows”, is rather like “it’s a duck”: its truth conditions are underdetermined by what is semantically encoded. If some utterances of the former are underdetermined or sound clearly false, that’s more evidence that the circumstance is impoverished of the pragmatic considerations necessary to yield a truth value, or that we import default pragmatic considerations. (Though, by the same token, this analysis may seem to commit the toy duck fallacy).

Arguments that some expression is subject to pragmatic modulation tends to rely on strong agreement about truth conditions associated with utterances. But this is missing. Open cases like those canvassed at the outset of this paper figure in key arguments for truth conditional pragmatics (and radical contextualism), although they usually focus on schematic predicates, subject to fluctuating use. (The reason for this will become clearer in a moment.) These arguments usually take the form of underdetermination arguments, holding fixed the context of utterance, it appears the truth conditions of an utterances of a type vary across circumstance.

Searle (1978) enquires into the truth-conditions of “The cat is on the mat”. It is easy to describe the sort of state of affairs that would make the sentence true (with respect to a particular assignment of values to indexical expressions). But once we have described such a state of affairs, we can embed it within an extraordinary situation:

Suppose that the cat and the mat are in exactly the relations depicted only they are both floating freely in outer space, perhaps outside the Milky Way galaxy altogether... Is the cat still on the mat? And was the earth’s gravitational field one of the things depicted...? What I think is correct to say as a first approximation in answer to these questions is that the notion of the literal meaning of the sentence “The cat is on the mat” does not have a clear application, unless we make some further assumptions, in the case of cats and mats floating freely in outer space (211).

Searle argues that when we describe an empirical situation, e.g., “the cat is on the mat”, we make certain features explicit, but an indefinite number of other features remain implicit and constitute a sort of hidden “background”, e.g., there are gravitational forces acting on both. The applicability of a term to novel situations depends on their similarity to the source situations, i.e. to the situations by association with which the term has acquired the meaning it has. For the term to be (clearly) applicable, the target situation must be similar to the source situations not only with respect to those features which easily come to mind and constitute the ‘explicit’ definition of the term, but

also with respect to the hidden background. If the two situations considerably diverge with respect to the latter (the cat and mat are both in outer space), it is unclear whether or not the term will be applicable, even though the explicit conditions of satisfaction are satisfied. On this picture, the contribution that the Background makes to the content of representations cannot be represented truth conditionally because any representation of the contribution of the Background (a sentence in a semantic theory, for example) itself depends on the Background for its content.

But, even granting that there are not any representations that have their content intrinsically (without some contribution from the Background), Searle's argument assumes (without argument) that truth conditional semantics cannot rely on context-sensitive representations to represent how the Background affects the content of certain representations *as they actually arise*. Contrary to what he argues, the so-called open cases are not a compelling objection to truth conditional semantic approaches. (And one may suspect he wrongly concludes that a sentence S lacks context-independent truth conditions from the fact that there are circumstances in which the truth value of S would be indeterminate). A similar response applies to Recanati's version of this argument: it assumes that TCS is committed to giving a reductive account of the content of expressions. As long as TCS does not accept such a reductive account of content, then Recanati's argument does not get a grip.<sup>86</sup>

---

<sup>86</sup> Underdetermination arguments against TCS, and for radical contextualism, proceed in somewhat similar terms, relying on so-called Travis Cases. TCS is the project of determining a way of assigning truth conditions to sentences based on A) the extension of their constituents and B) their syntactic mode of combination. A Travis Case: ... consider the words "The leaf is green", speaking of a given leaf, and its condition at a given time, used so as to mean what they do mean in English. How many distinct things might be said in words with all that true of them? Many. ... Suppose a Japanese maple leaf, turned brown, was painted green for a decoration. In sorting leaves by colour, one might truly call this one green. In describing leaves to help identify their species, it might, for all the paint, be false to call it that. So, words may have all the stipulated features while saying something true, but while also saying something false. (1994: 171-2)

Travis describes two contexts, one in which we are sorting leaves by color and the other in which we are sorting leaves by species. In the two contexts factors A) and B) are the same. Yet the use of "The leaf is green" in the decoration-sorting context is true, but the use of it in the species-sorting context is false. Truth-conditional semantics is the project of determining a way of assigning truth-conditions to sentences based on A) the extension of their constituents and B) their syntactic mode of combination. Travis's argument appears to show factors A) and B) under-determine the truth-conditions of utterances of "The leaf is green". Travis concludes some information available in the context is relevant for determining the truth-conditions of the utterance, but which information is not semantically encoded. To the extent that Travis cases can be imagined all over the place, he takes this as a reason to endorse radical contextualism (although he might as well have endorsed the weaker TCP). See, for instance, Travis 1989.

Travis Cases, no matter how many can be constructed, plausibly only represent a threat to truth conditional semantics only insofar as it aims to give a reductive account of sentence meaning in terms of bivalent truth. If it does

There is much to say about truth conditional pragmatics, and radical contextualism. We have touched on three questions. First: are our cases well analyzed by appeal to them? Second: do our cases really motivate one of these positions. Third: are RC and TCP are coherent, and are they viable competitors to truth conditional semantics? The answer to the first in some sense depends on our answer to the third, and this is a very complicated question. But it seems open cases do not provide a reason to prefer truth conditional pragmatics over truth conditional semantics insofar as the latter does not aspire to reductive account of meaning. Truth conditional pragmatics also leave significant questions unanswered.

According to Recanati 2004, the correct interpretation of a modulated utterance is what it would be reasonable for the interpreter to take the speaker to have intended. One objection is that it ignores the fact that the primary basis for our knowledge of what other people have in mind is what they *tell* us. If from your saying, “the best currency is bitcoin”, I learn you believe bitcoin is the best currency, then that’s because there is something about his use of those words that allows us to make the inference. I can infer you believe bitcoin is currency, because I interpret your words in a particular way. So, we must be able to say what makes your words have that meaning on the occasion of your uttering them, other than that that is the meaning that we can reasonably assume you have in mind.<sup>87</sup>

There is reluctance among linguistic pragmatists to treat seemingly polysemous expressions as genuinely ambiguous. Many of the predicates they analyze as undergoing modulation are schematic, and are another man’s polysemy (“use”, “lemon”). One source of this reluctance is Grice’s “Modified Occam’s Razor”, ‘senses are not to be multiplied beyond necessity’ (1989: 47). This is usually construed as requiring that we not posit a new sense *wherever there is a pragmatic derivation* of the message from an uncontroversial old sense.

---

not – which it need not — these objections have little bite. They might seem to suggest TCS cannot account for open cases (whether or not we endorse radical contextualism) since they can be constructed for open cases in ways that elicit stable truth judgments. But more is needed: an argument that predicates in open cases resemble stable cases like “lemon”, or “ready”, cases TCS would analyze as polysemy, or unarticulated constituents, respectively, and which “window”, as applied to glass walls, does not resemble on this score.

<sup>87</sup> Another objection: it assumes that thoughts cannot be open in just the way that spoken words can be. If mental representations can be open as well, though on the occasion of any particular tokening they have a more specific meaning, then we cannot equate that more specific meaning of a thought with the meaning of some other thought underlying it.



Everyone agrees that some expressions are genuinely polysemous, e.g., what generates ambiguity in “his eyes and his soul are blue”. The Razor, thus construed, cannot be right because it would make metaphors, and stock loose speech, immortal. The metaphorical meaning of a word, or what is communicated by loose speech, is derived from its conventional meaning. Over time, a metaphorical meaning often becomes regularized and conventional: the metaphor “dies”; the expression is now polysemous/ambiguous. Yet a derivation of what is now a new conventional meaning from the old conventional meaning is still available. That derivation will be center stage in the diachronic linguistic explanation of the presence of this new meaning in the language. But what is sought is not a theory that secures communication but an account of how language helps. Thus, appeals to TCP do not further the issue.

Truth conditional pragmatists, more generally, eschew treating polysemous expressions as ambiguous. Recanati (2010) again: “it does not seem that there is a *discrete* list of senses available but, rather, a continuum of possible senses to which one can creatively add in an open-ended manner” (18).<sup>88</sup> But the creative use of expressions is not at odds with the view he opposes. It accepts that an expression can have a speaker meaning other than its conventional meaning – in a metaphor, for example. And, in time, such a speaker meaning can become a new conventional meaning of the expression. Indeed, this is the story of polysemy (2010: 70). Everyone agrees that languages change. What then is the line between predicates like “get” and “use” subject to fluctuating use (plausibly not ambiguous), and an expression like “blackbirds”, which is plausibly ambiguous: “she eats, but won’t wear blackbirds” (said not at a food fight, but at the feather boa competition). Returning to our open cases: are we to think the meaning of “blackbird”, “window” underdetermined their truth conditions as applied to new birds, or new technology, in a way that permitted fluctuating use, as a schematic verb might? It is unclear why we should think this, as opposed to thinking that sometimes implicatures are semanticized.

## 9. Non-Vagueness Indeterminacy

We’ve dismissed vagueness, partial definition, context-sensitivity, and underdetermination as a models for open texture. Last one. Some indeterminacy is thought to be supervalueable. Sometimes vagueness is thought to be this sort of indeterminacy. Set vagueness aside. Maybe open texture is.

---

<sup>88</sup> See also Recanati 2004, p. 134, discussion of “get”.

It isn't. One trouble is that too many sentences come out only partly true. A more interesting trouble concerns *which* expressions are indeterminate in this way.

Suppose "athlete" is undecided with respect to non-human animals, and this is not because the term has an indeterminate meaning, but because we have failed to decide which precise meaning to give it. *That we haven't yet decided is consistent with the idea that the decisions that we have made can make some sentences true.* Say a precisification of a language is a way of assigning a precise meaning to every word in the language consistent with the decisions that we have already made. So if we have decided that any object that is F satisfies "is an athlete," and no object that is inanimate does, then on all precisifications, "is an athlete" will only be satisfied by objects that are F, and animate. But if we have not decided whether any objects that are non-human satisfy this predicate, then there will be some precisifications on which they do, and some on which they do not. The distinctive supervaluationist claims are

- (a) there are many precisifications of language, because semantic indecision is pervasive;
- (b) a sentence is true if it is true on all these precisifications;
- (c) a sentence is false if it is false on all these precisifications; and
- (d) if a sentence is true on some precisifications and false on others, then it is neither true nor false.<sup>89</sup>

The indeterminacy in one term may be constrained by the indeterminacy in other terms. For example, two extensions A and B may each result from legitimate sharpenings of respectively 'red' and 'orange', but there may be no legitimate joint sharpening of 'red' and 'orange' that yields these extensions (e.g., if A and B overlap). Admissible interpretations of a language are those assignments of entities where each term is assigned a partial denotation in a manner that respects the constraints. Roughly, there is one global sharpening of all terms that yields the assignments made by the interpretation. Now truth can be defined via Bas van Fraassen's supervaluation: a sentence is true (false) if and only if it is true (false) under all admissible interpretations; where truth (falsity) relative to an admissible interpretation receives the usual Tarskian treatment.<sup>90</sup>

Not everyone uses this machinery for vagueness. Field (1973) proposes that Newtonian "mass" partially denotes (since nothing in pre-Newtonian usage settled the matter of whether (say)

---

<sup>89</sup> Though this is something of an exaggeration: correspondence theorists tend to reject (d).

<sup>90</sup> Field 1973 famously argues that scientific terms, in particular, "mass" as used in Newtonian physics does not determinately denote, and instead partially denotes.

an ant on Earth is heavier than an elephant in space space). He proposes an account of partial designation on which we can say that it partially picks out having greater weight than and partially picks out having greater mass than. Given our stipulations about the case, some sort of partiality or indeterminacy seems to be mandated if we assume that facts about usage (intentions, dispositions, etc.) at or before the time of utterance are what determines meaning.

A few things. First, applied to “solid”, rather than “mass”, implies that very few older statements about solid things were supertrue. That’s a major cost. Terms have partial denotations related to each linguistic choice that is eventually faced and linguistic choices that could possibly be faced.

The more open predicates are, the more partially they refer. We now know very many of our older terms to refer indeterminately, beyond those of scientific discourse. That’s reason to think many terms partially refer in ways that today they do not seem to. At least, with vague predicates, the indeterminacy is of the foreseeable sort. The more axes on which linguistic choices may arise, the more “acceptable precisifications. Supervaluationism for vagueness predicts non-supertruth/falsity in the borderline. Supervaluationism for open texture predicts it more broadly. Some of our cases involve linguistic choices that herald belief revision, e.g., for “tables are solid”, whereas precisification around the borderline typically do not. So, in 1600, utterances of “tables are solid” is not supertrue, even if after we precisified similar utterances are.

Second, it is not obvious that this account can make sense of semantic decision. For example, one might imagine various Newtonians debating the question of whether the ant or the elephant is heavier, giving reasons and arguments (e.g., the two objects could be lifted with equal ease, but the elephant would exert greater force). We can easily imagine one party backing down, retracting statements, and so on. We can imagine all this happening inside Newton’s head. This is how thought experiments work, and also interpreting the result of actual experiments. Assuming thought experiments can refine concepts, produce semantic decision there is a problem here. If it is indeterminate whether “heavier than” picks out *having greater weight than* or *having greater mass than*, no past users were determinately correct, and neither party to the debate is (determinately) correct; and if designation in this case is partial, both sides to such a debate would be partially wrong. A view like this does poorly capturing deliberation. There deploying one precisification, then deploying another. Likewise, there is no retraction. If we engage in

suppositional reasoning and come to different conclusions, it appears we use different terminology.<sup>91</sup>

It seems the possibility of making linguistic choices stands in a complex relationship to the possibility of defeat of ordinary beliefs. Let's go back to 1600. There is no determinate fact of the matter that tables are solid (as used in 1600). There is a question about how to apportion our credences. Should we be less sure that tables are solid than that tables contain a fair amount of empty space? One reason to find this analysis unappealing is that, arguably, we cognitively imperfect creatures use imprecise linguistic language in part *because* being responsive to rational requirements associated with highly precise thought and talk is untenable.

It is doubtful that the notion of partial denotation can be explained independently of the notion of truth. There is plenty of reason to think that the reference of *some* terms is prior to the sentences we take to be true. But this is not always so. The partial denotations of some terms (e.g., those introduced in certain theories) is settled in part by the (basic) sentences that we accept about them – or come to accept. The same may be said about many terms from ordinary discourse. So, truth does not appear to be reducible to partial denotation. I think this matters for reasons that will become clear.

Fourth, the more partial reference there is, the more partial reference relation starts to look like the plain old reference relation coupled with epistemic tentativeness on our part. What makes it the case that “F” indeterminately, or partially, refer to X and to Y? They are acceptable precisifications. Presumably, that means X and Y are suitably precise, and they are acceptable because it is not determinately the case that X is F, and it is not determinately the case that X is not F (and so on for Y). So, explaining indeterminacy by appeal to having acceptable precisifications is not particularly clarifying. If the locution “not determinately the case that X is F” concerns a state of affairs, then there is no determinate fact of the matter about whether something is 1600-solid. Plausibly the distinctions we care about is as much between determinate/indeterminate facts, as *determinate* facts and *determinate* not-facts.

Notice it seems unappealing to analyze in/determinacy in terms of ignorance/knowledge. Sometimes we can be ignorant when there is *no* fact of the matter to be ignorant of. I do not know

---

<sup>91</sup> Chalmers 2011 fails to draw a credible account of analyticity from Bayesian updating for similar reasons.

what I want for dinner, and, when I come to know, it is not as though my doing so results from unclocking a fact that was there all along. It is not indeterminate whether I want Indian for dinner, in the sense I do not know I want it (but there is a fact of the matter about my wants), or in the sense it is indeterminate (as opposed to determinate) fact of the matter.

There is a question about how to express indeterminacy in the home language, when we are not dealing with vagueness. We can say: future scientific theories, ones more sophisticated than our own, will have two terms, which they hold are non-synonymous, and will look back at our own use of “F” and say for practical purposes it translates into the future language as “X” and other purposes as “Y”, and there is no further fact of the matter.<sup>92</sup> This is unconvincing. Just because a future theory regards our term “F” as imperfectly translatable into “X” and “Y”, it doesn’t follow that *our* term is referentially indeterminate. Maybe “F” better carves nature at its joints, and future speakers have made a sort of unnatural distinction or are responsive to features we are not and reflectively do not care to be. Or maybe it is simply a different distinction. Suppose “F” is “sweater” and “X” is “redsweater” and “Y” is “greensweater”, and these future speakers have dropped “sweater” as hopelessly indeterminate. If characterizations of referential indeterminacy this way are to be extensionally adequate, the future theory concerning “X” and “Y” must in some sense be better than ours.

The problem is that, by hypothesis, there aren’t resources for explaining what “better” comes to. We cannot say the future theory contains truths concerning “X” and “Y” that our current theory cannot express, at least not if we are disquotationalists about truth (and so cannot ascribe truth to conceptually alien theories). Likewise, we cannot say the future theories “carves at the joints” better than the current one. Because “X” and “Y” doing so is a matter of them referring to objects or properties that are special in some way. If the future theory is just “better” in some epistemic sense it had better not be one that presupposes the theory gets things more right in some sense that presupposes truth or joint-carving-ness. As far as I can tell, they must say the future theory is “better” because it exhibits more theoretical virtues than ours, simplicity, explanatory depth, coherence, and the rest. But that seems to put the cart before the horse. Theoretical virtues are supposed to be virtues because they are somehow indicative of the theory being true, getting

---

<sup>92</sup> See Field 2018, Leeds 1997.

the joints right. That's why understanding referential indeterminacy in terms of translations might seem apt in the first place: "F" is referentially indeterminate iff "F" would be translated into multiple terms of a truer-and-more-joint-carving theory. But if we stripe away the epistemically better theory of these features, then again, this proposal seems extensionally inadequate. Imagine that "F" is translated into multiple terms "X" and "Y" of a future theory that exhibits more theoretical virtue but in fact is not any more true or revelatory of nature's structure than our own. Is this sufficient for referential indeterminacy of "F"? It seems that it isn't. All the more so when we are dealing with ordinary discourse, with all its plurality of vocabulary, rather than a discourse and subpart of language, which by stipulation is restricted to those that capture the scientific world.

We *could* say "F" is referentially indeterminate if there is more than one eligible candidate referent in nature, each equally fit to be "F"'s reference as far as our linguistic practice is concerned. But if that's right, how future speakers would translate "F" is beside the point. IF their theories are "better" than ours, then that's partly because their terms carve at the joints better, are less referentially indeterminate. If "carves at the joints" is understood in terms that rule out determinate reference to medium sized dry goods, then we are left to wonder why we speak such a language, if we could do otherwise, or why we cannot do otherwise. If "carves at the joints" is understood in a way that does not, it is hard to see how appealing to indeterminacy helps address our open cases.

## 10. Open Texture is Sui Generis

Well, that was exhausting. Where are we? I've argued that open texture—stipulatively, that feature of expressions that gives rise to open cases—is not vagueness, partial definition, context-sensitivity, supervalueable, or underdetermination. It's not explained by recourse to the machinery of rigid designation, or plausibly a biproduct of finely tuned, but difficult to figure, metasemantic laws. Open texture may be characterized as a sort of linguistic choice point (it may go unnoticed) where there is more than one right way to go on. These choice points are often accompanied by semantic dumbfounding or dogmatism. They may cause reconsideration of core cases, or established entailments. Resolving open texture does not appear to result in meaning change. I think the fact that there is more than right way to go on with a predicate is sui generis. If you disagree—if you think it is analyzable as one or more of the phenomena canvased, then you may well think language does not have an open texture in some further interesting sense. I do.

The open cases set the mood, perhaps, for skepticism that much is semantically encoded, or at least semantically encoded as intensions. It goes without saying predicates have various associates: past applications, disposition to apply, characteristic inferences, projections, paradigms, presuppositions, associations, implicatures. What is encoded might be a general rule of application, reference fixing description, paradigms, beliefs, or facts about the environment or whatever.

Maybe there is a recipe that takes some of these associates to a referent such that changes in the associates sometimes produce changes in referent. But why think there is any such recipe? One reason would be if the correct recipe is ahistorical: each instant has its own recipe. Another would be if there is one recipe (from the associates of the predicate to its reference) which is something like a weighted average: the further away from the present moment, the lower the weight some associate is given. A third would be if the future weighs more heavily, since we acquire more knowledge of what we want to be talking about. A fourth would be if the magic recipe perhaps involves speakers' (counterfactual) judgments to determine a distinguished interval used to calculate reference. Perhaps it is implausible that there is anything about patterns of use or the intentions and beliefs of users that makes one of those, or some other story, the a-contextually correct story about reference.

Discussions of reference often seem to presuppose a picture on which a word's use is something like a collection of different "forces", each pushing the word towards a particular part of the world, and thus a certain extension. To discover a word's extension on this picture, one sums the forces to get a single net force; that total force gives the word a unique momentum through semantic space, thus determining a unique extension. Such discussions seem to assume that the various strands of a predicate's use must compose somehow, that there must be a single recipe for combining all the facts about the way we use a word so as to come up with a best candidate for the word's extension. But why should this be so?

Consider, for example, the word "cat". Suppose that more or less the same "mechanism" underlies all normal adult use of the word in perceptual situations. Given this, someone might suggest that to find the word's extension we need to find the answer to the question "Which properties does this mechanism best register?" It would be silly to suggest that this question is irrelevant to the word's semantics. Someone else may observe that the word "cat" is a member of a family of words that are associated with a kind of concept that has a long evolutionary history.

She may claim that to find the word's extension we need to find the answer to the question "Which properties did this sort of word emergence and continued use intend to register?" Again, it would be odd to suggest that the answer to this question is irrelevant to the word's semantics. There is the use of a predicate, construed broadly, and there is what kind of properties the word is meant to track, and what kinds of properties actually regulate its deployment. Why are we to suppose that there is a single proper way to commensurate their differing effects in computing an extension, as opposed to many ways of abstracting away from the question of just what the extension is to facilitate a theory of composition and validity in natural language?

Perhaps there is nothing to know about an expression's character such that it gives rise to open cases. That's no knock on the utility of talk of character, and type level meanings, or the utility of compositional semantics, or talk of reference. The way I have suggested meanings are open is irrelevant for the core aims of these projects, and so not well represented in their machinery. Individuals are not perfectly described as sets of parts at a time either. But, for some inquiries, doing so is in order. Mostly when we theorize sentence meaning, we are not concerned with what, if any, features of it make some sentences intelligibly contestable in such a way that a) more non-semantic facts will settle it or b) or the facts about what the sentence means would (if only we could know). But if we are, we do not want to theorize meanings in a way that obviates the phenomenon. That's why it's important to recognize that open texture is not some of the more familiar semantic phenomenon, and is not amenable to some familiar (absolute) metasemantic theories.

Observations of this sort are sometimes parlayed into hyperbolic reflections that the aforementioned projects are wrong headed—an odd thing for one theorist to say to another when they are apparently so fruitful.<sup>93</sup> Here are three levels of response:

Extending the scope of the theory of meaning necessitates a new understanding of the formal structures and tools but it will not require us to abandon them altogether

Extending the scope of the theory will lead to a dramatic reformulation and reimagining of the theory of meaning so that previous work will have to be discarded

Extending the scope of anything like a theory of meaning to the phenomena we have discussed is impossible, that these phenomena are resistant to analysis by anything with the structure of a formal theory.

---

<sup>93</sup> Cf. Beavers and Stanley (MS) apparently think otherwise.



To my mind, all three are too hyperbolic. Likewise, cloud cover does not appear in scientific models of predator-prey relations. But there is no sense in which they need tweaking to reflect it, so long as we do not come to believe the model also describes the weather, conclude there is no weather, or that the weather is not of interest for other projects.

## **References**

- Austin, J.L (1946) "Other Minds." *Proceedings of the Aristotlean Society* 20: 148–87.
- . (1961) "On the Meaning of a Word." In *Philosophical Papers*, 23–43. Oxford: Clarendon Press.
- . (1962). *Sense and Sensibilia*. Edited by Warnock, G. Oxford: Oxford University Press.
- Beaver, D., and J. Stanley (2019) "Toward a Non-Ideal Philosophy of Language:" *Graduate Faculty Philosophy Journal* 39 (2): 503–47.
- Ben-Yami, H. (2016) "Vagueness and Family Resemblance." In *A Companion to Wittgenstein*, edited by Glock, H. and J. Hyman. Chichester: John Wiley & Sons, Ltd.
- Bromberger, S. (2012). "Vagueness, Ambiguity, and the 'Sound' of Meaning." In *Analysis and Interpretation in the Exact Sciences*, edited by Frappier, M., D. Brown, and R. DiSalle, 78:75–93. The Western Ontario Series in Philosophy of Science. Dordrecht: Springer Netherlands.
- Bueno, O, and M. Colyvan. (2012). "Just What Is Vagueness?" *Ratio* 25 (1): 19–33.
- Bunikowski, D. (2016). "The Origins of Open Texture in Language and Legal Philosophies in Oxford and Cambridge." *Rechtstheorie* 47 (May).
- Burns, C.D. (1938). "Philosophy and the Physicists. Susan Stebbing." *Ethics* 48 (4): 559–60.
- Byrne, A. (1993). "Truth in Fiction: The Story Continued." *Australasian Journal of Philosophy* 71 (1): 24–35.
- Chalmers, D. (2011). "Revisability and Conceptual Change in 'Two Dogmas of Empiricism.'" *Journal of Philosophy* 108 (8): 387–415.
- Clapp, L. (2012a). "Indexical Color Predicates: Truth Conditional Semantics vs. Truth Conditional Pragmatics." *Canadian Journal of Philosophy* 42 (2): 71–100
- . (2012b). "Three Challenges for Indexicalism: Three Challenges for Indexicalism." *Min & Language* 27 (4): 435–65.
- Cummings, S. (2020). "Semantic Reasons." Personal Website.
- Dasgupta, S. (2018). "Realism and the Absence of Value." *The Philosophical Review* 127 (3): 279–322.
- DeRose, K. (2002). "The Ordinary Language Basis for Contextualism, and the New Invariantism." *The Philosophical Quarterly* 55 (219): 172–98.
- Devitt, M. (2013). "Is There a Place for Truth-Conditional Pragmatics?" *Teorema* 32 (2): 85–102.
- Eddington, Sir Arthur. (1927). *The Nature of the Physical World*. Ann Arbor: University of Michigan Press.

- Evans, G. (1985). *Collected Papers*. Oxford: Clarendon Press.
- Field, H. (1973). "Theory Change and the Indeterminacy of Reference." *Journal of Philosophy* 70 (14): 462-81.
- . (1974). "Quine and the Correspondence Theory." *Philosophical Review* 83 (2): 200-28.
- . (2000). "Indeterminacy, Degree of Belief, and Excluded Middle." *Noûs* 34 (1): 31.
- . (2001). *Truth and the Absence of Fact*. New York: Oxford University Press.
- . (2018) "Egocentric Content." *Noûs* 51 (3): 521–46.
- Fodor, J. (1964). "On Knowing What We Would Say." *Philosophical Review* 73 (2): 198–212.
- Gauker, G. (2017). "Open Texture and Schematicity as Arguments for Non-Referential Semantics." In *Meaning, Context and Methodology*, eds. Conrad, S. and K. Petrus. Berlin, Boston: De Gruyter.
- Graff Fara, D. (2000). "Shifting Sands: An Interest Relative Account of Vagueness." *Philosophical Topics* 28 (1): 45–81.
- Greene, M. (2011). "'Chocolate' and Other Kind Terms: Implications for Semantic Externalism." *The Philosophical Quarterly* 61 (243): 270–92.
- Grice, H.P. (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Hart, H.L.A. (1961). *The Concept of Law*. United Kingdom: Oxford University Press.
- Hawthorne, J. (2006). "Testing for Context-Dependence." *Philosophy and Phenomenological Research* 73 (2): 443–50.
- Hellie, B. (2019). "Semantic Gaps and Protosemantics." In *Quanta and Mind*, edited by Acacide Barros, J., and C. Montemayor. Cham: Springer International Publishing.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.
- . (1984). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Repr. Cambridge: Harvard University Press.
- . (2011). "Unrestricted Exportation and Some Morals for the Philosophy of Language." In *Philosophical Troubles: Collected Papers, Volume 1*. New York: Oxford University Press.
- LaPorte, J. (2000). "Rigidity and Kind." *Philosophical Studies* 97 (3): 293–316.
- Leeds, S. (1997). "Incommensurability and Vagueness." *Nous* 31: 385–407.
- Liston, M. (2005). "Does 'Rabbit' Refer to Rabbits?" 1 (1): 18.
- . (2010). "Rabbits Astray and Significance Awandering: Review Essay on Mark Wilson's *Wandering Significance: Review Essay*." *Philosophy and Phenomenological Research* 81 (3): 809–17.
- Ludlow, P. (2014). *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
- MacFarlane, J. (2016). "Vagueness as Indecision." *Aristotelian Society Supplementar Volume* 90 (1): 255-283.
- Montminy, M. (2010). "Two Contextualist Fallacies." *Synthese* 173 (3): 317–33.
- Quine, W.V. (1951). "Two Dogmas of Empiricism." *The Philosophical Review* 60.
- Raffman, D. (2014). *Unruly Words: A Study of Vague Language*. Oxford: Oxford University Press.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- . (2010). *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Richard, M. (2019). *Meanings as Species*. First edition. Oxford: Oxford University Press.
- Schauer, F. (2008). "A Critical Guide to Vehicles in the Park." *New York University Law Review* 83: 26.

- Shapiro, S. (2006). *Vagueness in Context*. Oxford: Oxford University Press.
- Soames, S. (1999a). "Presupposition." In *Handbook of Philosophical Logic, Vol IV*.
- . (1999b). *Understanding Truth*. New York: Oxford University Press.
- Sosa, E. (2005). *Ontological and Conceptual Relativity and the Self*. Oxford University Press.
- Stalnaker, R. and T. Baldwin. (2001). "On Considering a Possible World as Actual." *Proceedings of the Aristotelian Society, Supplementary Volumes 75*: 141–74.
- Stalnaker, R. (1978). "Assertion." In *Context and Content*, 78–95. Oxford: Oxford University Press.
- . (1999). "Twin Earth Revisited." In *Context and Content*, 210–21. Oxford: Oxford University Press.
- Stebbing, L.S. (1937). *Philosophy and the Physicists*. London: Methuen & Co., Ltd.
- Stephen, S. (2013). "Quandary and Intuitionism: Crispin Wright on Vagueness." In *Essays for Crispin Wright: Logic, Language and Mathematics*, edited by Miller, A. Oxford: Oxford University Press.
- Travis, C. (1975). *Saying and Understanding*. London: Blackwell Publishing.
- . (1989). *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford: Clarendon Press.
- Waismann, F. (1951) "Verifiability." In *Logic and Language*, edited by Flew, A. Oxford: Basil Blackwell.
- . (1968). *How I See Philosophy*. London: Macmillan and Co. Ltd.
- Waldron, J. (1994). "Vagueness in Law and Language: Some Philosophical Issues." *California Law Review* 82 (3): 509–40.
- Williamson, T. (1994). *Vagueness*. Oxford: Oxford University Press.
- Wilson, M. (1982) "Predicate Meets Property." *The Philosophical Review*. 91(4): 549-589.
- . (2006). *Wandering Significance: An Essay on Conceptual Behavior*. New York: Oxford University Press.
- Wittgenstein, L. (1968). *Philosophical Investigations*. Oxford: Basil Blackwell.
- Wright, C. (2003). "Vagueness: A Fifth Column Approach." In *Liars and Heaps*, edited by Beall, J.C. Oxford: Oxford University Press.
- Yablo, S. (2000). "Textbook Kripkeanism and the Open Texture of Concepts." *Pacific Philosophical Quarterly* 81 (1): 98–122.
- . (2002) "Coulda, Woulda, Shoulda." In *Conceivability and Possibility*, edited by Szabo Gendler, T. and J. Hawthorne, 441–92. Oxford: Oxford University Press
- . (2008). "No Fool's Cold: Notes on Illusions of Possibility." In *Thoughts: Papers on Mind, Meaning, and Modality*, 151–70. Oxford: Oxford University Press.
- Zeifert, M. (2020). Rethinking Hart: From Open Texture to Prototype Theory—Analytic Philosophy Meets Cognitive Linguistics." *International Journal for the Semiotics of Law - Revue Internationale de Sémiotique Juridique*, May.

## Chapter 3: Going on in the Same Way

### 1. Introduction

Speakers make, and recommend, novel usages all the time. Sometimes it is tongue in cheek, as when Irving Kristol quipped conservatives are liberals that got mugged by reality. Sometimes it is in view of some discovery about a kind in question, say, about what ulcers are. Sometimes it is in conjunction with conceptual analysis, as when Frankfurt (2005) told us what bullshit is. Sometimes a word is used in a new way in conjunction with a radical theory, as when Clark and Chalmers (1998) told us that the internet houses some of our beliefs. Sometimes the new usage is also in conjunction with activist aims, as in the movement to use “woman” in ways that are inclusive of trans women (and exclusive of trans men).

When a novel usage is received as particularly radical, it may be said to change the subject, say, from bullshit to bullshit\*, or, more significantly, from belief to belief\*, or from women to women\*.<sup>94</sup> Usages that become common within persuasive speech, or a subaltern group, or among experts (and those copying them), may be received as either outright false, or else subject changing, in other areas of discourse, and vice versa.

For example, in 2015, Ta-Nehisi Coates argued that policing practices in predominantly black neighborhoods like Ferguson, MO, constitute a system of plunder.<sup>95</sup> Some celebrated his analyses as obviously correct, or revelatory. Some agreed the practices are unacceptable, but found it obvious that “plunder”, with its standard meaning, does not apply truly, and non-figuratively, to them: either he straightforwardly misapplies “plunder”, or else speaks stipulatively about plunder\*. One strategy is to argue that, although “plunder” may not apply, it should, and that we should revise our concept PLUNDER so that it does. But this proposal may also seem to change the subject from *plunder*, to something else.<sup>96</sup> If the established usage is correct, novel usages associated with revisionary projects, and contentious conceptual analyses, appear to change the subject, or else be clearly false.

---

<sup>94</sup> See Bettcher 2012.

<sup>95</sup> <https://www.theatlantic.com/international/archive/2013/10/when-plunder-becomes-a-system-of-governance/280885/>

<sup>96</sup> Haslanger 2000, 2019 advocates pursues this strategy in response to charges that ameliorative analyses offer false accounts of the ordinary meanings of “woman” and “man”. In the interim, so-called conceptual engineers have followed suit.

That's basically right, at least according to a standard picture of linguistic meaning. Concepts, understood as word meanings, are represented as functions—intensions—from any number of ways the world might be (even ones that boggle the mind) to extensions. A concept, then, can be thought of as a set of possible objects, namely the ones it is true of, and its negation as the rest of the objects. So, an expression cannot change its truth conditions while preserving its meaning. While the exact relationship between meaning and use is obscure, how an expression is used closely figures in determining an absolute fact of the matter about its meaning, what it is true of. So, to the extent established usages are correct, novel usages are either false, or else they concern a subject other than the usual one.

It's well known that radical, or revisionary, novel usages face a worry about continuity. The worry is roughly this: if I ask you a question about the Fs, and your answer concerns not the Fs, but the Gs, then you haven't answered my question. You have changed the subject. Philosophers working in a range of areas express worries in a similar vein, for a range of concepts. Mark Richard (2019) on Haslanger's ameliorative analyses of "race" and "gender": "For that matter, doesn't [Haslanger's revisionary project] cross the line between conceptual therapy and stipulative rebranding?" Peter Ludlow (2005) on anti-contextualists in epistemology: "Any investigation into the nature of knowledge which did not conform to some significant degree with the semantics of the term 'knows' would simply be missing the point." Peter Railton (1989) on his own revisionary naturalism in metaethics: "[If a revisionary metaethical naturalist] wishes to make his case compelling, he must show that his account of a person's good is a rather clear case of tolerable revision, at worst." Most pointedly, Haslanger herself: "Revisionary projects are in danger of providing answers to questions that weren't being asked" (2000: 34). Most recently:

When social change happens, there is likely to be controversy and disagreement about how to extend the concepts we've been using to do the work we now need them to do. Such changes should be acknowledged as such, and should not be held hostage to what we have thought we were doing all along, and how to continue that. (2019: 27)

I think this response is unnecessarily conciliatory to the background picture of language that drives the subject change worry. This paper offers a surprising tweak.

Whether radical and revisionary novel usages change the subject is related to a question in the philosophy of language about meaning over time. There, the main obstacle to conceptual continuity may be framed as a kind of contingency: if we could go on other than we do, we change the subject, however we go on. This can be brought out with a family of puzzles about diachronic

synonymy, involving homophonic cross-temporal and counterfactual reports. I'll sketch a solution where proponents of novel usages can have their cake — truth conditional semantics — and continuity too. It is the view that meaning facts are not absolute, but rather projected by the interpreter's community. I think this picture puts the bump in the rug in the right place when it comes to conceptual contestation and continuity. At a minimum, it does not try to theorize these things by first idealizing away from them.

In the background, I am interested in two different kinds of attitudes towards language. One is a kind of naturalist attitude that says, on the whole, users of the language are right about when they do, and do not, change vocabulary. Philosophers of language may describe, and explain, our linguistic practice, but neither they, nor anyone else, is in a position to correct it. The other attitude is more imperialist. It casts philosophers of language as experts, and part of the job is to go around telling people when they do, and do not, mean the same as others. Whether, for instance, activists are participating in a single social movement concerned with women's rights, or creating a new second movement, concerned with something else (however extensionally similar). There are social reasons not to go around correcting people, but, at bottom, when people talk as though they mean as they always did, experts know better, or in principle could, because meaning facts are robust enough support such judgments.

In some ways it is theoretically easier to have an imperialist standpoint, as we'll see. There are fewer constraints on the theory. But to a great many people that just starts in the wrong place. The reasons for the naturalist attitude are not merely instrumental—not merely a means towards convincing others of activist or political claims, thereby advancing a cause. They are potentially internal to the aims of discourse, including activist discourse, and the theory of meaning over time. This paper sets out to vindicate the naturalistic approach.

I'll start by saying more about the subject change worry and why it matters (section 2). I'll then motivate one version of the puzzle, and introduce some of its variations (sections 3-4), before considering some otherwise promising pictures of conceptual continuity that fail to adequately account for them (sections 5-7). These include temporal externalism, and the view that synonymy is a sort of similarity relation. In the last three sections of the paper, I sketch my preferred solution, and its implication for the subject change worry, and radical or revisionary usages.

## 2. The Subject Change Worry, and its Importance

Let's get things started with an illustration familiar from contemporary discussions about gender and sex within the public sphere. Twitter, that sewer of ordinary usage, is helpful. Take a tweet from The Human Rights Campaign (or HRC, a leading LGBTQ+ rights organization): "trans women are women, and trans men are men" (12/19/19). In the next few days, several users tweeted back. A sample:

PheeberWeaver: I think it's offensive someone male thinks they can redefine woman to suit their purposes.

ForeverMagdalen: No one has the right to redefine Woman without women's consent.

ArlenaRioDell: I am not transphobic, changing the law to redefine woman affects me as a woman immensely. Also, affects my future grandchildren. You sweep our concerns under the rug.

Whatever else may be said about the spirit igniting these comments, they evince a common feeling that the HRC speaks stipulatively: (said in the old vocabulary) it's redefining women to include some men. Thus, it changes the subject, from whatever the ordinary term "woman" refers, to something else. Some so-called conceptual engineers, and ameliorators, have embraced the charge of redefinition—arguing we should revise our concept WOMAN so that it truly applies just this way. But again, there's a worry that the subject is being changed.

Continuity of subject matters to people. For example, it matters to people, regardless of gender identity, that their efforts to articulate who they are in other words have not failed in some appreciable way. Likewise, it's important to some people that in advocating for the rights of women they are advocating for the same group that waves of feminists before them have advocated for, and that they may ask the questions that were asked, and sensibly disagree about the answers given. Say, it is important that they are thinking and talking about the same group that Simone de Beauvoir was, following through on her intentions, and continuing the fight for women that was started more than 75 years ago.

Strawson (1953) gives voice to a more general worry in his exchange with Carnap about role of explication in philosophy. For our purposes, to the extent novel usages are not outright false, they may be thought of as stipulative explications.

[T]o offer formal explanations of key terms of scientific theories to one who seeks philosophical illumination of essential concepts of non-scientific discourse, is to

do something utterly irrelevant—is a sheer misunderstanding, like offering a textbook on physiology to someone who says (with a sigh) that he wished he understood the workings of the human heart. (505)<sup>97</sup>

While explication is consistent with the process of revising hypotheses in scientific inquiry, in philosophical inquiry it may be received as changing the subject. Somewhat less lyrically:

[T]ypical philosophical problems about the concepts used in non-scientific discourse cannot be solved by laying down the rules of use of exact and fruitful concepts in science. To do this last is not to solve the typical philosophical problem, but to change the subject. (506)

Strawson's point concerns a distinction between inquiries that do not require results stated using the very concepts that kicked off the inquiry, and those that seek illumination about those very concepts. Some scientific inquiries fit the profile of the former, some philosophical the latter. When an investigation delivers a result that is taken to problematically change the subject, the inquiry has been stymied. Modernizing a bit, two worries latent in Strawson's point are sides of the same coin:

*Continuity of Inquiry*: if we change the intension/extension of the expressions we use in some inquiry, then we are no longer inquiring about the same thing.

*Verbal Disputes*: if parties to a disagreement use a term central to the debate associated with different extensions or intensions, then they simply talk past one another: they express logically consistent propositions, rather than express substantive dispute.<sup>98</sup>

When we engage in sustained inquiry into some topic—when we investigate or make choices about it, when we change our beliefs, or when we debate it—we have, in the typical case, the ability to be talking about that same subject over time, across theories, contexts, and viewpoints, even as we change our own beliefs or express disagreement with others. Despite these differences and changes, there remains a single topic of our discussion—a sort of stable communicative ground.<sup>99</sup> Novel usages—that purport to offer a revolutionary theory or engage in explication—would seem to threaten this crucial aspect of our inquiry and discourse. If instead of discussing a single topic

---

<sup>97</sup> For a more in-depth, if not uncontroversial, account of this exchange, see Pinder 2020.

<sup>98</sup> Bracketing issues related to variations in the standards relevant for context-sensitive, like “rich”. See Richard 2004.

<sup>99</sup> The integrity of scientific inquiry may not require continuity of inquiry in this sense. See, for example, Chomsky 1995.



over time and across speakers, we instead change the meanings of the terms central to those discussions, proliferate homonyms, then we have lost that stable ground.

Preserving the continuity of inquiry has been invoked as near decisive support for the externalist claim that the meaning of natural kind terms in public language are at least partly determined by the underlying natures that give rise to the surface of phenomenon of interest, e.g., “water” referred to H<sub>2</sub>O even before the discovery that the wet potable stuff was H<sub>2</sub>O. In ordinary parlance, we say that scientists *discovered* that water was made of H<sub>2</sub>O, and learned what our term referred to all along (rather than the discovery heralding a change of vocabulary). Likewise, we take ourselves to be able to intend to do the very thing Socrates once intended to do: have a drink of water. We also take ourselves to be able to think about minds, those same things Descartes thought about, but prefer another theory to his.<sup>100</sup> The ease with which we regulate usage of “water” by pointing out that this or that is (is not) chemically alike to what is in the taps and oceans is taken as evidence about the meaning of “water” all along.

One difference between scientific inquiries, and others, is that the usage of terms in the former are defined, deployed in the service of common goals, and disposable when they hinder progress. Terms of ordinary language—say, “freedom”, or “woman”—that we may hope to clarify by means of philosophic reflection are not deployed in the service of common goals, or subject to uniform patterns of deference, so much as a common interest in coordination and communication. The disparate uses of some terms, say, “free” or “coerced”, are most used tracing areas of contestation in public life. When we seek philosophical illumination of a concept associated with some phrase, we must answer to the present usage that sustains the phrase’s meaning, on pain of

---

<sup>100</sup> What does it take to kill an example? The story goes that Putnam 1975’s Twin Earth thought experiment—where we agree that water-y stuff on Twin Earth, XYZ, is not water—supports the claim that “water” rigidly refers to H<sub>2</sub>O, and that natural kind terms, more generally, rigidly refer to the hidden natures that give rise to surface phenomenon. Whatever else the argument establishes, it seems unlikely it establishes that H<sub>2</sub>O is the referent of our ordinary term “water”. That surely is shorthand. H<sub>6</sub>O<sub>3</sub> has equal claim on being water, and “H<sub>6</sub>O<sub>3</sub>” is not synonymous, does not corefer with, “H<sub>2</sub>O”. (Putnam 1990 makes this point in the context of renouncing externalism, on his way to internal realism). Water is not identical with H<sub>2</sub>O, H<sub>6</sub>O<sub>3</sub> and so on, because ice is not water (neither is steam) but both are similarly related to such samples. Fine, so water is composed of H<sub>2</sub>O, H<sub>6</sub>O<sub>3</sub>, and other like compounds, in a liquid state. In order for water to exhibit cohesion — one of its characteristic features— some reasonable amount of the molecules must actually be H<sup>+</sup> and HO<sup>-</sup>. So, “water”, the presumptively scientific term that rigidly designates, is more like “polka dot” than it might seem. It does not refer to XYZ, but neither does it exclusively refer to H<sub>2</sub>O. The arguments of this paper suggest an asterisk to those who analyze the reference of natural kind terms by way of semantic externalism, since not all terms that go on to function as natural kind terms *begin* that way. Arguably, the real lesson of Putnam’s arguments is *future* temporal externalism. But that is beyond the scope of this paper.

changing the subject.<sup>101</sup> I take part of Strawson's point to be this: a persisting dispute, or question, about a concept, even after empirical information has been supplied, is not evidence that the concept should be explicated, analyzed away, as we might for the purpose of scientific inquiry.<sup>102</sup> We risk abstracting away from the very phenomenon we want an account of.

It is a delicate matter. For example, does ordinary use of "water" share an intension with scientific use? Set aside Twin Earth, and whether "water" refers to, as it did all along, a chemical kind. Oceans are made of water, and so are grimy puddles, but these contain less H<sub>2</sub>O per unit than celery juice.<sup>103</sup> If a new civic filtration system involved steeping water in tea leaves, and what came from the taps tasted of tea, we would still call it "water", and likely stably go on to do so. Right now, if we made it in a pot, we'd call it tea.<sup>104</sup> Suppose we do stably go on to call the filtered tap water "water". Calling tomatoes "vegetables" in culinary contexts is not loose speech, would the situation with the new tap water be any different? Externalists will say that the point is that these things are not XYZ, and never are, and that competent speakers revise their use of "water" on the basis of information about what contains H<sub>2</sub>O (or like compounds) in what quantity. In any event, clarifying the ordinary term by pointing out that tea, club soda, coffee, really are just dirty water may well be received as changing the subject by ordinary speakers.

Appeals to the physical environment only issue an exhaustive account of the ordinary term "water" by, in effect, abstracting away from the differential patterns of use that mark it as distinct from the scientific one. "Watch out—there's a puddle of water on the street!" If it were a puddle of H<sub>2</sub>O, we would not be so concerned. It's not that these sentences aren't about water, per se. It's that they are not quite about a chemical kind, liquid H<sub>2</sub>O. "Oh, stop it with the club soda already—just drink tap water!" Depending on where you are, the club soda is a good shot closer to pure H<sub>2</sub>O. There is much of interest to say here, starting with what is sometimes called The Toy Duck Fallacy, following Kripke 2011. While Kripke cautions against reading too much (into the

---

<sup>101</sup> See Strawson on connective analysis, 1959, 1992.

<sup>102</sup> Strawson defends "imagining ways in which, without things other than ourselves being different from what they are, we might view them through the medium of a different conceptual apparatus" (513).

<sup>103</sup> Some figures for the percentage of H<sub>2</sub>O per unit for various liquids, courtesy of Pietroski 2005. Club Soda 99.9, Diet soda, not cola 99.8, Tea 99.7, Diet Cola 99.54, Stuff from Pietroski's well 99.4, Coffee 99.39, Espresso 97.8, Ocean water, avg. salinity 96.5, Michelob Ultra 95.4, Bud Light 95.0, Distilled vinegar 94.78. See: [ndb.nal.usda.gov/ndb/foods/show/](http://ndb.nal.usda.gov/ndb/foods/show/)

<sup>104</sup> See Chomsky 1995.

meaning) from differential uses, he also says sometimes they are not easily explained away (e.g., as truncation or loose speech, as Austin does). For now, it is enough to say the point may extend to reading too much into meaning from later tendencies to refine, or retrospectively, clarify usage. More than this is beyond the scope of this paper. The point for now is that ongoing controversy about the extension of some term, among competent, and tolerably well-informed, speakers, is significant, if this is what we hope to theorize.

Of course, there is no ongoing controversy about the contours of the ordinary term water, and whether it is distinct from the scientific term. There are not contests of the dominant understanding of *noteworthy* adulteration is (“Club soda is water! Club soda for all!”). But there could be: say, the government distributes packages to purify urine instead of fixing bad tap water (“That is not water! And you know it” “Sure it is: NASA astronauts think so”). Appeals to experts, or chemical kinds, may put some people at ease. Some may even come around quickly to the solution being a source of water. But appeals to the machinery of externalism miss the point.

Likewise, appeals to popular usage, the dictionary, high profile experts, or dominant meanings. These do not settle the dispute since these are precisely what one openly party contests. (And these appeals are likely to generate unwelcome implicatures). Pretheoretically, such a dispute also does not seem verbal, like a dispute about whether banks are where the money is or where the water is, or whether the ball is on the left (when it rolls between us). It seems like a dispute about whether the purified stuff is water, capable of intelligibly persisting among informed speakers. Our ordinary notion is bound up with action, but it need not follow that our dispute is a metalinguistic negotiation serving as proxy for, or byproduct of, a decision about what to do. Whether or not they think it’s water, strapped residents will use it and complain about it.

To respond to the subject change worry for novel usages in non-scientific discourse, we must ask, “what does the continuity of continuous inquiry consist in, and what does the substance of substantive disagreements consist in?” And try to answer it without first idealizing away the phenomenon. This is the big question. That our words mean the same thing is one answer to it— one theory of what forms the stable ground of our conversations and debates. Among non-experts, this happens to be among the most widely accepted answers at the moment. I think rightly so, but for somewhat odd reasons as will become clear.

But many, including revisionary theorists and philosophers of language more generally, reject that theory. Simply observing that, in certain contexts, the sentence “the speakers are talking about the same thing” has *a* true reading in English, or is an acceptable falsehood, doesn’t provide such a response. The demand is not for the vindication of ordinary judgments or claims, but rather for a substantive answer to the big question—for a theory.

But there is substantial pressure on this theory to vindicate ordinary claims, on pain of undercutting the motivations for thinking sentences express propositions in the first place. One good argument for believing that there are meanings, and propositions, is that the best explanation of the fact that I believe what you said is that there is *something* that you said which I believe. Likewise, a good explanation of the fact that “a cat” translates to “un chat” in French is that they may be represented as functions that make the same truth conditional contribution to sentences that they figure in. Part of the work propositions do is figure in a unified explanation of our judgments about content-match, or samesaying.

A natural thought is that competent users of some expression are good at detecting when it is synonymous with another, even those used in a temporally or theoretically remote community, and that we can know when they are. That is, meaning facts seem perspicuous. Mental states belonging to the same agent certainly appear to overlap in content, as when I form the intention to do something I have always wanted to do. Content may also be shared between different agents, as when you finally realize something I never doubted, or you firmly hope for what I shudder to conceive. There is the practice of homophonic translation of older texts: we treat them as testimony.

Content-match is at the bottom of norms of rationality, such as “Don’t intend to do what you don’t think possible” and “Don’t both believe and disbelieve the same thing.” It tells us our apparent knowledge that the present use of some term is synonymous with its uses in the recent past and near future is real knowledge. Intra-subjective perspicuity allows ignorance about what precisely the terms in question mean: that is no bar to being a good detector of synonymy. The rough idea is that an individual is generally in a position to tell when two of their mental states, say an intention and a desire, match in content.

The intersubjective component of the relation grounds norms of interaction, including communication, the aim of which is precisely to transmit content. We care about engaging with the thoughts other people think and express. In many cases a necessary condition for doing that is

samesaying, that we can evaluate what they said or thought. As a result, a kind community with whom we have access to what they think and said emerges. There's a kind of alienation that happens when there is not intersubjective perspicuity. There may be barriers to contemplating what Thales meant when he said the world is composed of just one thing. Perhaps there is nothing we can say that samesays with Thales' dictum. That may be right. At any rate, the situation seems significantly different in the case of my niece who avows that dirt is made of chocolate. I can samesay what she says and also assert that it is false, and know both of these things. Children are not too remote, and neither are speakers in the recent past and future or those in communities that are not too theoretically remote from our own.

Perspicuity is the sort of claim denied only by experts. Semantic externalists often complain that semantic internalism (which ties meaning facts closely to individual variation) predicts that we commonly mean different things by our words, and hence that our reports of one another's speeches are systematically false.<sup>105</sup> Frege suggested people rarely mean the same thing by their words and that meanings are often glimpsed obscurely through a fog (it can take a long time to figure out what a word always meant).<sup>106</sup> But he is also said to have imposed a perspicuity constraint on content,<sup>107</sup> and viewed it has an embarrassment to his view that there should be a common stock of thoughts passed down from earlier days.<sup>108</sup> Similarly, anti-contextualists about a particular word often complain that contextualism about that word predicts that it rarely if ever happens that two utterances of the word contribute exactly the same content, so that speech reports involving that word are far more precarious than we normally suppose.<sup>109</sup> Some epistemicists reject perspicuity on the grounds that small shifts in use induce correspondingly small shifts in reference.<sup>110</sup> Clear-sighted homophonic reports of close neighbors prompt an epistemic worry. When the underlying changes are gradual, there are no tell-tale signs that distinguish cases of recent or imminent semantic change from cases of semantic stasis. How, then, can we know the

---

<sup>105</sup> See Burge 1979, Chalmers 2012.

<sup>106</sup> See Burge 1992 certainly sees perspicuity as an embarrassment to Frege's view, suggesting Frege confuses the notion of sense having to do with cognitive significance with that having to do with meaning in a public language.

<sup>107</sup> See Evans 1985 (first published 1982), 18–19, on Frege 1892/1984

<sup>108</sup> See Abreu Zaveleta forthcoming, Hawthorne and Dorr 2014.

<sup>109</sup> See Cappelen and Lepore 2005.

<sup>110</sup> See Williamson 1994.

propositions expressed by our cross-temporal and counterfactual reports?<sup>111</sup> What has not been fully appreciated is that, once we turn our attention to cross-temporal and counterfactual reporting, analogous problems arise for nearly everyone.

### 3. First puzzle: Non-Branching

Now that we've seen the force of the subject change worry, and some constraints on a satisfying response to it, let's turn from synchronic synonymy to diachronic synonymy. There is a family of puzzles that turns on the observation that, sometimes, a word appears to maintain its meaning despite changes in its application conditions. This section sets out the first of these: the non-branching puzzle. It takes the form of an inconsistent tetrad: four independently plausible statements that are mutually inconsistent. I'll start by motivating a specific instance of it for the predicate "book", before setting it more schematically.

The first claim is this: Two predicates are synonyms only if they have the same intension. It's widely accepted that two predicates are synonyms only if, no matter how things are, they apply to the same things. Somewhat more precisely: for two predicates to be synonymous they must determinately apply to the same things, indeterminately apply to the same things and determinately fail to apply to the same things. What a predicate does, does not, or indeterminately, applies to—its truth conditions—is frequently represented as a possible worlds intension, where an intension is a function from possible worlds to extensions. The intension of the noun "book" is a function that takes whichever world is actual to the set of all actual books. A change of intension is, or produces, a change of meaning.<sup>112</sup>

Of course, the noun "book" can be used to talk about spatiotemporally located things that carry inscriptions of certain contents, or abstract contents that get encoded in many ways and places. That is, "book" is polysemous: it has multiple "sub-senses." Polysemy may be compared to homonymy, a case of distinct words sounding the same. It may be among the most interesting

---

<sup>111</sup> That's one thing when it's a problem for vague predicates – where an expression's actual meaning is surrounded by a vast cloud of slight variants, all of which seem equally well qualified to be possible meanings and there variations are not at issue. But it's another when the variations at issue, larger and more significant. Traditional anti-skeptical maneuvers may be helpful for the former (whether they are helpful enough is a good question), but they are beside the point for the latter.

<sup>112</sup> Perhaps synonymy requires more than sameness of intension. That's fine for my purposes. For discussion of whether synonymy is hyperintensional, see Nolan 2014.

feature of lexical meanings, even if diagnosing examples can be hard.<sup>113</sup> Just how words and meanings are individuated (the difference between polysemy and homonymy) is frequently treated as a sort of housekeeping issue, one that a standard way of thinking about meaning abstracts away from. Uses of the orthographic scrawl b-o-o-k, or pronunciations of the usual sounds, are synonymous just in case they are associated with the same intension, and intensions themselves do not contain lexicographical information. Uses of “book” with its text-sense are associated with one intension. Uses of “book” with its tome-sense are associated with another. From the standpoint of intensions, there is no difference between homonyms and polysemes.<sup>114</sup>

The second claim is that “book” (with its tome meaning) as used in 1750 does not (or does not determinately) apply to eBooks. Think of Samuel Johnson, man of letters and writer of the first English dictionary. Plausibly, he had no disposition with respect to whether eBooks are books, or else conflicting ones. He may even have had a disposition to say otherwise. An eBook is the digital item one buys a license to when one buys a certain electronic file. When apprised of how they work, and what readers in 2020 do with them, we might imagine him pounding the table: “That, sir, is no book!” If we gave him a Kindle loaded with 1000 e-books, would he say, “I hold 1000 books in my hand”? Plausibly not. The Kindle was released in 2007, within a few months the London Times worried that “the slow death of the book may be with us”, and The New Yorker, got out its violin, for an article, “The Twilight of Books.” We may imagine Johnson too would wring his hands about the impending death of the book. At a minimum, it seems improbable that the laws of meaning, such as they are, would fix it that “book” as used in 1750 determinately applied to eBooks in advance of the processes of negotiation and accommodation by which usage is extended when unforeseen, or novel, circumstances arise. Plausibly, “book” as Johnson used it, did not apply to eBooks (or, if you like, did not determinately apply to them).

---

<sup>113</sup> Bromberger 2012, Pietrovski 2005, 2018, Boguraev and Pustejovsky 1996, Vincente 2018, Recanati 2004.

<sup>114</sup> An avenue not treated in this paper is to contrast cases of polysemy with the predicates and uses of interest in this paper. There’s a good case that polysemous terms like “France” have *one* meaning (*France*) but associated with at least two sets of truth conditions (France is hexagonal, France is a republic, but France is not a hexagonal republic). Problems come if instead we try to assimilate names like “France” to entities that satisfy predicates: there are hexagonal republics, or else there are really two Frances. While people refer to things, words don’t. if we adopt the view that theories of understanding are theories of meaning, then we shouldn’t expect identity between meaning and truth conditions (/character). This may be worth exploring more.

The third claim is that eBooks are, in fact, books; that “book” as used in 2020 applies to eBooks. A Google for, “How to read eBooks?” turns up:

Amazon’s eBook store is the largest on the internet. Aside from its huge choice of books, there are plenty of other features which keep customers coming back for more.

For example, there’s the Kindle Unlimited subscription service. For \$9.99/month, you can download and read as many books as you want from the collection of more than 1 million titles. But you won’t find the latest releases or bestsellers on the list.

Prime members also have access to Prime Reading. It’s an ever-changing library of more than 1,000 books, magazines, and comics that you can rent for free.

The application of “book” to eBooks does not appear to be loose or metaphorical speech. Consider this infelicitous tweet from December 24, @FreeDailyeBooks:

Free eBooks and reduced rate eBooks tracker/Crawler that eBooks on daily basis. Books here are \$0.00 to \$0.99 but make sure it is before you buy

The infelicity is not on account of applying “book” to eBook. The blurbs discuss how to acquire a particular sort of book, the sort you can download, buy, rent, read, and that resides in libraries. These are uses of “book” for tomes, only these tomes are digital (as compared to abstract objects). When I read my Kindle, I report that I am reading one of the books in my library. The book I am holding is the same as the book I am reading, in this case, an eBook. It’s a bit odd to the ear to say, while holding only a Kindle “I hold 50 books in my hands”. But it does not seem odd because it is literally false (compare, if I were holding only a duck). Plausibly, it sounds a bit odd because our paradigm book remains a print book. We do not *seem* systematically mistaken that eBooks are books. We can imagine staff meetings in 1990s where engineers report toiling to create new and better books for the 21<sup>st</sup> century and celebrate their success, and a higher up commanding them to do so.

Finally, the fourth claim is that “book” hasn’t changed its intension since 1750. In that year, Johnson penned a scathing review of Milton’s epic: “*Paradise Lost* is a book that, once put down, is very hard to pick up again”. Frankly, I agree with Johnson. I believe what he said, and am laughing at the same joke his contemporaries laughed at. Whether the joke resides in the semantic content of Johnson’s statement, or its implicature, it works precisely because “book” has a tome-meaning, where the tomes in question are of the sort that are heavier the longer the text is. We report Johnson’s verdict homophonically. Disquoting is sufficient to translate him, fully, into the modern vernacular. We report that he talked about books, and intended to write books. We may imagine Johnson bragging, “In the 21<sup>st</sup> century, men will still read my books”, and, as a matter of



fact, they are only read on Kindles. We are inclined to say he spoke truly, although he couldn't have imagined the circumstances ensuring his continued fame. It seems fine to disquote, and infer his book is still a success. If *book* is a poor translation of Johnson's "book", the inference equivocates—which it doesn't appear to do. Take Johnson's imaginary utterance, "eBooks are not books". At least in some moments, we are not inclined to say that such an utterance *was* true, or *is* true. We are most inclined to say it is false, and always was. These reflections form the inconsistent tetrad:

- (1) Two predicates are synonyms only if they have the same intension.
- (2) eBooks at 2020 are not (or are not-determinately) in the intension of "book" at 1750.
- (3) eBooks at 2020 are determinately in the intension of "book" at 2020.
- (4) "Book" at 2020 is synonymous with "book" at 1750.

From (2) and (3), eBooks are in the intension of "book" at 2020, but not the intension of "book" at 1750. By (1), "book" changed its meaning between 1750 and 2020, which contradicts (4). Something has gone wrong. Which statement, or statements, should we reject?

One immediate response is that, really, there is no inconsistency. What makes "book" talk true in 1750 is somewhat different than what makes "book" talk true in 2020, but the character of "book" has not changed. Books have. The situation is comparable to "college": the truths about colleges in 1750 are different than the truths about colleges in 2020, but all along we have been talking about college. It is just that college has changed. Likewise, the tango has changed, but the character of "tango" has not. "Pasta" does not appear to have changed its meaning, although today some bacon agar is classified as pasta, whereas in 1800 it almost analytically would have been excluded. It's not that "pasta" has changed its intension, it's that pasta, as a food genre, has grown up.<sup>115</sup> Perhaps these are the right analyses of kinds that can grow and change, that are in some sense historical particulars. But book is not obviously among them. But books, unlike colleges, are not institutions, and unlike tangos, are not obviously genres. They do not seem to be instances of a kind that grows and changes. Back to the tetrad.

---

<sup>115</sup> Richard 2019 considers the true heir to meaning talk to be the anchors of semantic competence, rather than truth conditional functions. In his sense, and analysis, it is indeterminate whether "pasta" has changed its meaning. But that does not conflict with my point here.

The defense of (2) appealed to the special relationship between meaning and use, or dispositions. Plausibly, speakers lacked dispositions to affirm eBooks are book, even in view of information about eBooks, and how the world would come to be. But the relationship between meaning and use, broadly construed, is obscure. Famously, there's no worked out account of how that's meant to go.

Johnson was learned and capable, but plausibly neither he, nor anyone else, had formed dispositions to affirm conditionals such as “if people in the future invent widgets with such and so contours and call them “books”, then they are books”. Any conditional dispositions he may have had to affirm the sentence “eBooks are books” just in case he is confronted with widgets with such and so contours hardly seem dispositive of the meaning of “book” in advance of being triggered. For example, he might also have had conditional dispositions to affirm “eBooks are not books” just in case he is confronted with widgets with such and so contours, if near a post office.<sup>116</sup>

One popular response is to say that the way of the world determines reference. Some properties are more eligible as referents, more natural, less gerrymandered, and “magnetically attract” the reference of our terms.<sup>117</sup> For instance, *being H<sub>2</sub>O* is just more eligible as the referent of ‘water’ than *being H<sub>2</sub>O and wet, potable, raining down...* and so on. Suppose that is right for some terms. Even so, there does not appear to be an outstandingly eligible property to be the referent of “book”. *Being a readable text in any format* seems no more or less natural than *being a bound and leaved text*. Supposing one is, then there are two mysteries. One concerns how we know reference at all if we are not good detectors of the naturalness facts (suppose SJ knows “eBooks are not books”). The other concerns why we get to know them better and better (suppose he did not, but we know they are books). To the extent “book” in 1750 excludes eBooks, or does not determinately include, them, (2) stands.

Claim (3) also appears secure. At least some uses of “book” in 2020 properly apply to eBooks. The Kindle marketing team, you, and I are not pervasively in error when we systematically refer to eBooks as “books.” We take ourselves, on the whole, to be good judges of the extension of our terms and what expressions are synonymous, when apprised of the general context. If not

---

<sup>116</sup> For argument to another sort of conclusion, see Chalmers 2011, and Grice and Strawson 1956.

<sup>117</sup> See Lewis 1984, McLaughlin 2016.

around the penumbra, say, of decaying unreadable pages, then at least around core cases. Call this consideration Charity. The trouble is that Charity is a reason to endorse (2) and (4) as well as (3).

Charity is not so easily dismissed, at least this is a methodological assumption of this paper. If we start by investigating what evidence we get for synonymy (both synchronic and diachronic samesaying, as well as intra- and inter-subjective), judgments are of primary importance. What other evidence about meaning is there? We can start from a metaphysical stance, asking what the determinants of meaning are, and then investigating which usages stand to be synonymous from there. Mostly, philosophers of language proceed in this way, assuming that, in fact, there is an absolute fact of the matter about what expressions mean, and synonymy facts can be adduced from it. It strikes me that this approach is not well suited to investigating the change of subject worry. We may idealize away our judgments about conceptual continuity, but not when theorizing it. A further consideration is that beginning with the metaphysical approach ushers in the family of problems related to the determinacy of reference. Starting from our practice of gathering evidence about meaning suggests a more fruitful way forward. But a fuller defense of this approach is beyond the scope of this paper.

Now I think there are many instances of the non-branching puzzles, including: pen (and quill), paper (and wood chip paper), slime mold, gene, talking (within earshot as opposed to the phone), rape (and marital rape), pet (and friendly barn loiterers), drug (and alcohol), addicted (to a behavior, not a substance), food (and nutritional capsules), family, witches (and unpopular women), kings (and divine appointees), solid, weight, temperature, ruby/topaz, glass, soup (and smoothie), selfie (with multiple people), B-flat, 5 degrees centigrade, head of household, mother (gestational surrogates), number (and 0), unicorns (and Marco Polo's near "discovery" that they were as real as rhinoceroses).<sup>118</sup> There are even more that we can imagine. These are cases where there is no outstandingly natural property in the vicinity of use, or fully uncontroversial patterns of deference, or a referent that is clearly institutional. The tetrad, more schematically:

- (1) Two predicates are synonyms only if they have the same intension.

---

<sup>118</sup> When Marco Polo arrived in China, he wrote that there were unicorns but not at all like what he expected. For one they had large, low heads, and thick-set bodies. They were rather more gray than white. We can imagine that had circumstances been slightly different, this mixed discovery—that unicorns are real as rhinoceroses—might have resonated through the European use of "unicorn".

- (2) X at t2 is not (or is not-determinately) in the intension of ‘F’ at t1.
- (3) X at t2 is determinately in the intension of ‘F’ at t2.
- (4) ‘F’ at t2 is synonymous with ‘F’ at t1.

For example, when lightning struck sand, Greek-speakers called the product “ὑαλος” as early as the 8<sup>th</sup> century BCE. Speakers plausibly had no dispositions to apply “ὑαλος” to colored windows or firebird cars—having never encountered anything like them. Plausibly they wouldn’t be so disposed even after being apprised of their intrinsic and extrinsic properties. The early philosopher might even have said “ὑαλος” rigidly designated silica dioxide, and so did not properly apply to those things. Today scientists call many compounds “glass” – even some that lack silica dioxide altogether. Colored glass, fiber glass, cd glass are not diluted glass, or fool’s glass; they are just glass. Yet it seems clear enough that the Ancient Greek ὑαλος translates as *glass*. What glass *is* has not changed, even if we now know how to make more of it.

#### 4. Three More Puzzles: Branching, Fusion, Line

Now I’ll introduce three variations on this puzzle. These add constraints to any solution to the branching puzzle.

##### a. *Branching*

Now suppose that things had gone differently: we did not classify eBooks as books, because, say, Barnes and Nobles waged a defamation campaign, rather than hopping on the eReader bandwagon. Plausibly, there is a new non-branching puzzle. On the one hand, in 1750, it is indeterminate whether eBooks are in the intension of “books” as used in 2020. On the other hand, in 2020, it is true that “book” has not changed its intension since 1750, and so true in 1750 that eBooks are determinately not in the intension “book”. That the two non-branching scenarios are intelligible suggests speakers in either world could have gone on with “book” other than they did. Yet, whichever way they went on, “book” does not change its intension.

Now suppose the 18th century use of “book” was the ancestor of two subsequent linguistic communities, one in the US, and one in UK, both of which take themselves to be faithful to the original use. In the US, “eBooks are books” is false in 2020. In the UK, “eBooks are books” is true. If both groups mean by “book” as they meant along, it follows they mean the same as one another in 2020, assuming synonymy is an equivalence relation. Contradiction. We can represent the situation as an inconsistent triad:

- (1) eBooks are not in the intension of “book” as used in the US in 2020; it hasn’t changed its intension since 1750.
- (2) eBooks are in the intension of “book” as used in the UK in 2020; it hasn’t changed its intension since 1750
- (3) “Book” as used in 1750 for tomes (as opposed to texts) is not ambiguous.

Speakers in the US that deny that eBooks are books take themselves to mean as they did. Had either later community not existed, there could still be a branching puzzle.

Let’s return to *Charity*. It says that speakers tend to be good judges of synonymy. We may suppose at least one fragment community speaks truly. Now consider *Parity*. It says that, if statements of one fragment stand in the same relation to past use as the statements of another—if they are comparable similar in respect of naturalness, fitting patterns of deference, attitudes of speakers to prior use, and so on— they are on a par with respect to synonymy to it. So either both fragments’ mean as was meant, or neither do. The US and UK uses of “book” plausibly stand in the same relation to the 18<sup>th</sup> century uses of “book”. Together Parity and Charity suggest both fragment communities speak truly. If only one speaks falsely, that violates Parity. If both do, that violates Charity.

More schematically: there is a community C that splits into two fragments F<sub>1</sub> and F<sub>2</sub> that function independently. As the members of F<sub>1</sub> and F<sub>2</sub> learn more, they express their revised views using their old vocabulary, but find different ways of doing this, with the result that certain terms in F<sub>1</sub> become clearly non-synonymous with the corresponding term in F<sub>2</sub>: were the fragments to start interacting again they would never translate each other homophonically. And yet neither fragment recognizes a change in their own usage from before the fragmentation: each would translate C homophonically.<sup>119</sup>

Fictional examples are a dime a dozen, though real life ones are somewhat harder to come by. Partly this is because when usage diverges, it tends to diverge even more. We can identify some brief time frames that resemble a branching case. For example, British and American uses of “chip” for crispy bits “chipped” off a potato, or New World uses of Old World terms for flora

---

<sup>119</sup> Wilson 1982 (e.g. the example based on *The Island of Lost Women* that opens the article), and in Lance and Hawthorne 1997 (the modified Salem witch story, p. 45). Field 2009 has an example, involving logical vocabulary, in particular negation as used in three different logics.

and fauna. Notice that branching, like non-branching, cases can be constructed more often than speakers judge themselves to change vocabulary. It often seems as though we could have gone on other than we do with a term, but that, if we had done so, we likewise would not have judged ourselves to change vocabulary.

We can imagine many more versions of a branching puzzle: ones with more than two branches, ones where one of the branches splits twice, ones where a branch splits and one of resulting fragments closely resembles an original branch.

Cases of branching usage call to mind the person fission thought experiments.<sup>120</sup> Each branch, taken on its own, appears to be identity preserving, but cannot be in the presence of the other branch, on pain of transitivity failure. Parfit (1984)'s influential conclusion is that the notion of sameness of person relevant in everyday commerce, the one that "matters", does not track numerical identity of persons and is not itself an equivalence relation. Likewise, expressions where branching cases can be constructed must not be identity preserving along either pathway (never mind if branching usage actually occurs). I'll discuss this response to the puzzles in section 5.

#### *b. Fusion*

A third sort of case involves branching, in the opposite direction: instances where non-synonymous subaltern usages come to resemble one another such that at a later time it appears some expression is synonymous with both earlier usages. Call this a fusion case.

"Liber" entered Latin as a name for birch bark, but soon acquired, as a secondary meaning, *scroll* (which were sometimes written on birch bark). When codices were invented in Roman around the year 0, they were a run-away hit. By 300 AD, the number of codices and scrolls were roughly equal. *The codices*-exclusive term, "caudax" (a term for block of wood), dates back as far as uses of "liber" applied to *scrolls and codices*. Let's imagine going somewhat differently. Suppose in the Western Empire, scrolls called "libri" were invented first. Suppose at the same time in the Eastern Empire, codices were invented first, and called "libri" on account of being often made on birch bark. Some years later East-West trade became easier, and each part was pleased to

---

<sup>120</sup> Williams 1970, Parfit 1987.

discover a new style of “liber”, taking themselves to mean all along as they always did by talk of “liber”. This is a fusion case.

- (1) In the East at 0, codices are not (or are not determinately) in the intension of “liber”.
- (2) In the West at 0, codices are in the intension of “liber”.
- (3) At 100 AD, “liber” means as it did in the East at 0, and as it meant at in the West at 0.
- (4) “Liber” as used at 100 AD is not ambiguous.

Some real life fusions do not involve common expressions, or fragmented communities. In the later 19th century, “dipsomaniac” was introduced as a medical term for heavy drinkers displaying highly compulsive behavior and the delirium tremens, among other things. “Drunk” was colloquially applied to those who were often visibly drunk and disorderly. “Dipsomania” and “drunk” as used in the late 19th century are not synonymous. But there’s also good sense to the claim that both are synonymous with “alcoholic” as it is used today: that *being a drunk, being a dipsomaniac, and being an alcoholic* are all the same thing, and that some drunks are rarely visibly drunk, and may no longer drink. Question: “what’s the difference between alcoholics and drunks?” Answer: “Alcoholics go to meetings.” Possibly the joke works precisely because it thwarts the expectation it establishes: the punchline points up a difference in register, rather than extension. “Dipsomania” retains some usage today but for a clinical profile much more restrictive than it once was.

These cases are unlike this one: suppose the samba merges with the mamba, and we call it the “amba”. Genres grow and change. What it is to be an alcoholic, or be a book, or have PTSD, does not seem to be like this.<sup>121</sup> Another response is that we did not realize “Hesperus” and “Phosphorous” were synonymous until we did more astronomy. Likewise, we did not realize “dipsomaniac” and “drunk” were synonymous until we did more medicine and psychology. But this is dis-analogous. The established patterns of deference for astronomical terms were absent for

---

<sup>121</sup> Another: Just fragmented discourse. In the 19<sup>th</sup> century, a particular constellation of symptoms were dubbed “railway spine” after it was determined many of those presenting symptoms had been in railway accidents. After a time, people began presenting with those symptoms who had not been in high speed accidents, and railway spine started to be classified as a kind of hysteria (distinguished from head and neck trauma). A similar constellation of symptoms was dubbed, “battle fatigue”, after it was learned that many with them had recently been in combat. Arguably, the later uses of “railway spine” and “battle fatigue” are enshrined in the diagnosis of “PTSD.” These expressions pick out, and indeed did pick out, one and the same phenomenon. However, “railway spine” and “battle fatigue” do not seem synonymous with one another.

“drunk” and “dipsomaniac”. Facts about human behavior made it likely that these terms would be deployed for a condition we would later call “alcoholism”, but did not mandate it. Use of the one might have been restricted to anti-social behavior, the use of the other to physical symptoms, or outmoded theories of addiction.<sup>122</sup>

*c. Sequential Puzzle*

Another sort of puzzle involves cases of gradual changes in usage, or cases of sequential revision. Take a word with a storied history like “livid.” Its use has evolved gradually over the last few centuries. Initially, it was applied to the purplish/bluish color of a bruise (as the Latin “lividus” was). Today it is applied to angry emotions, and the bright reds of angry faces. The meaning of “livid” is evidently not the same at the beginning and at the end of this history, and there were a number of steps in between:

- (1) at t1, “livid” is applied to grayish/bluish colors of bruises
- (2) at t2, “livid” is applied to ashen color, the color of corpses
- (3) at t3, “livid” is applied to ashen color, the color corpses, and of angry-looking faces
- (4) at t4, “livid” is applied to red/purplish color of angry-looking faces
- (5) at t5, “livid” is applied to *enraged people*.

It might appear to follow that there is at least one second-long interval such that the meaning of “livid” is not constant throughout that interval. But given the gradual character of the underlying changes, all the second-long intervals are extremely similar in relevant respects. This makes it tempting to conclude that semantic changes occur during all of them, or at least during some non-trivial proportion of them.

More generally, we may consider a long series of possible worlds, each only slightly different from its neighbors in microphysical respects, stretching from the actual world to one in which the use of “livid” is very different, and conclude from the fact that there are no non-arbitrary grounds for positing semantic discontinuities at just a few points in the series, all or most of the worlds in the series differ from their neighbors as regards the meaning of “livid”. Never mind the

---

<sup>122</sup> Such cases perhaps suggest analyzing puzzles in terms of the metaphysics of *words*, but that is beyond the paper’s present purpose.



niceties, it seems there must be some sharp cutoffs. “Livid” does change its meaning a few times. Plainly “livid” today does not mean *ashen* (*Charity*).

But it also seems speakers, at any time, rightly aver they mean the same as they did at any adjacent time (even considering some temporal remove). The problem is that if speakers at each time are good judges of synonymy, and take themselves to mean as their immediate neighbors, everyone means the same thing (*Parity*). But that cannot be right.<sup>123</sup> The sequential problem:

- (1) “Livid” as used at t1 is not synonymous with “livid” as used at t5.
- (2) Uses of “livid” are synonymous with uses at adjacent times.

As far as I am concerned, we need not distinguish cases of novel usages due to theoretical revision and those due to changes in practical circumstances. Thus, I think there’s a sequential case for “gene”, and “temperature”, as well as “livid”, and “nice”. We can imagine cases that involve both sequential and branching elements, as perhaps “centigrade”, and “B-Flat” have undergone.

One question concerns how similar the sequential puzzle is to the non-branching, branching, and the fusion cases. The answer in part depends on whether these puzzles have a unified source. It is not entirely clear, but I tend to think it is similar enough to warrant inclusion. The sequential puzzle, like the others, arises from diachronic synonymy judgments, the sort of judgment that seems unproblematic when considered on its own, but not alongside other such judgments. Even so, only the sequential puzzle arguably relies on a tolerance principle.

### 5. Future Temporal Externalism: Deny (2)

The next three sections consider three responses to the puzzles, taking off from the non-branching puzzle. While otherwise appealing, these are unable to account for the other puzzles, or only at substantial cost. Recall the non-branching puzzle:

- (1) Two predicates are synonyms only if they have the same intension.
- (2) eBooks at 2020 are not (or are not-determinately) in the intension of “book” at 1750.
- (3) eBooks at 2020 are determinately in the intension of “book” at 2020.
- (4) “Book” at 2020 is synonymous with “book” at 1750.

---

<sup>123</sup> In effect, this is an argument that meanings are not at the edges of a partition; it’s most straightforwardly an argument against modal robustness, which is beyond the scope of this paper.

(2) looks secure *if* meaning is determined by past and present use. But if the intension of “book” in 1750 depends not only on its usage to date, but on how it will go on to be used, (2) looks less secure. eBooks may well be in its intension in 1750, because of how 21<sup>st</sup> century speakers went on to use it. 18<sup>th</sup> century speakers just did not, and could not, have known.

This view predicts that a world where speakers go on to classify eBooks as books is one where, all along eBooks were books. It also predicts a world where speakers go on to deny that eBooks are books is one where, all along, eBooks are not books. That speakers could have gone with “book” other than they did does not entail that “book” changes its meaning. This looks like a satisfying resolution of the non-branching puzzle.

The claim that present fact depends on future occurrence is not as outlandish as it first sounds. Sometimes the distribution of certain properties depends on future occurrences. Which rendezvous among friends are dates depends on what happens later on. Some are dates, not because of the stated intentions, or secret mental states of the participants, but because of the romance that later bloomed. Likewise, some move in a chess match may count as the winning move, not because it thereby forces a win, but because of the unforced strategy adopted later by the opponent. When we improvise musically today, whether we are playing the first refrain, or another verse, depends on what we play later, just as whether it’s classical, or heavy metal, depends on how loud and how fast we get. Perhaps meaning facts at a time, like date facts, or winning move facts, or genre facts, depend on events that occur later.

It’s popularly thought that facts about linguistic meaning depend on facts that “ain’t in the head”, such as those residing in the speaker’s physical environment, or usage in the broader community.<sup>124</sup> This is the thesis of semantic externalism. It’s also popularly thought that meaning facts depend on facts about how expressions were used in the past. This the thesis of past temporal externalism (PTE).<sup>125</sup>

Future temporal externalism (FTE) is true if the meaning of at least some expressions depends on how they go on to be used at a later time. It has contemporary defenders, and a growing number, as well as some historical sympathy.<sup>126</sup> Wittgenstein (1967):

---

<sup>124</sup> Burge 1979, Putnam 1975.

<sup>125</sup> Kripke 1980.

<sup>126</sup> Jackman 1999, 2005, Ball MS, Sawyer MS.

Let us imagine a god creating a country...which exists for two minutes and is an exact reproduction of a part of England...One of these people is doing exactly what a mathematician in England is doing, who is just doing a calculation. —Ought we to say that this two-minute-man is calculating? Could we for example not imagine a past and a continuation of these two minutes, which would make us call the process something quite different? (336)

While he is not talking about linguistic meaning here, plausibly he suggests that if we can associate with “+” the addition function (rather than say the quaddition function), 2-minute man would count as calculating, in particular he would be adding. The association of “+” with the addition function (or not) depends not only on what happens before the two minutes, but what happens afterwards. But how long afterwards? Donnellan (1983):

An outrageously bizarre view of language—that the extensions of one’s terms may be determined by the psychological quirks of some people several centuries hence (201)

[and] just as bizarre to suppose that the extension of one’s terms should depend upon ... future historical accidents (202)

Never mind the time frame. Conventions take some time to establish. It’s commonly thought meaning is established at least in part by convention. But, if future practice *cannot* determine past meaning, it seems it need not do so at all. (Why not hold that whatever determined the meaning of “+” in advance of some convention taking root determines its meaning afterwards as well?) If the convention one’s usage participates in partly determines the meaning associated with that usage, there may not be some collective agreement beforehand about what that convention is like, nor an antecedent fact of the matter. This is just the observation that usages of an expression helps to establish the conventions that govern it.

To be sure, FTE has some surprising repercussions, and amplifies some odd consequences of past temporal externalism. For instance, PTE predicts we cannot introspect facts about meaning because we cannot introspect how speakers across town use (or used) language, or facts about our physical environment. Meaning facts are like facts about popularity in a large high school: because these facts supervene on attitudes and dispositions of students in other wings, students they may never have met, students cannot introspect who the most popular student is. Future temporal externalism tells us that some of these speakers are in the future (the “other wing” is in the future, as it were). It predicts that meaning facts are not introspectable because we cannot introspect the future. While we can research the physical and social environment, we can only predict the future, and any uncertainty is inherited by beliefs about expression meaning.

FTE links knowing about the future and knowing about meaning. If we did know what we meant, we could know surprising facts about the future. For instance, because I know that meal replacement capsules are not food, I know that a rich company didn't invent them in the future (because if the company were rich, they'd buy ads that would establish a convention of calling them "food"....).<sup>127</sup> Likewise, if I knew the identity of the most popular song in America, I would know all sorts of surprising things about America – for instance, that Americans do not love pop music as one might expect, and so on. But, of course, because we do not know this sort of surprising fact about the way of the world, we don't know what our words mean.

It's well known that a priori knowledge of meaning is difficult to come by for externalist theories.<sup>128</sup> The problem is particularly keen for FTE. Supposing we know FTE is true, meaning facts look undermining. When we flip a coin to decide whether to fragment the linguistic community, just by knowing that "*book*" means *book* I can know whether the coin will land heads. I can think to myself "*book*" means *book* and it wouldn't have if usage branches in the future. Suppose that it's 20 years ago, and I know that in 20 years this prototype will be called a "book," and speakers will say a new kind of book was born on this very day. Do I thereby know that eBooks are books? Plausibly not. Even if a knowledge of meaning is elusive for non-future versions of externalism, facts about utterance truth are still knowable (insofar as we often know we are not in deviant circumstances). Knowledge of the sorts of future contingencies that shape future usage is on much less sure footing.

FTE may appear incompatible with the knowledge norms of assertion, because it ties knowledge of meaning to knowledge of the vagaries of future use. Suppose Stalnaker asserts: "That is either Zsa Zsa Gabor or Elizabeth Anscombe." On his account, he asserts the diagonal proposition. (There are also other possibilities.) The temporal externalist should not accept that the diagonal is what is asserted, because that is not something she stands to know at the time of assertion (potentially). But she can help herself to the idea that the assertion adds the diagonal to the context, and this is enough to let the conversation proceed.

This is not a fully satisfying response. Consider the whole conversation. When we retract, we make it the case that the proposition we asserted was one that we did not know. But we also

---

<sup>127</sup> See Ebbs 2000.

<sup>128</sup> See, for example, Boghossian 1989.

retract the assertion: an appropriate response, given that the norm was violated. And similarly, when we stick to our guns, we make it the case the proposition we asserted is one that we did know. Here there is no violation of the norm, so the re-assertion is appropriate. (Analogy: making a promise that you can keep, but might not.) The knowledge norm of assertion is a regulating target for speakers and as such capturing it is a kind of hurdle for theories of assertion, of knowledge, of metasemantics. That speakers satisfy the knowledge norm is thanks to structural features of the metasemantics, not speakers' epistemic relation to the asserted contents. But that is not the spirit of a norm of action, which is what the knowledge norm of assertion is. A bump in the rug is getting moved, in a way that makes for more, not less trouble for the present purposes.

There are other reasons to be uneasy about FTE. Externalism about “beechness”, “arthritis”, or “sofa”, is compelling because we defer to our linguistic community, and judgments of experts. We defer to the linguistic community because we coordinate with them, and we defer our judgments to those who, if we knew what they did, we would speak as they do (and our peers would too). But these reasons lack analogues that support future temporal externalism. We coordinate with our temporal peers, not future speakers. We defer to matters of fact, the sort of thing experts might now, but which we ourselves do not. We do not defer to what is presently unsettled, in the way that the future is. These are not knockdown arguments, but they are pleas for a more robust defense of FTE.

Whatever FTE's merits, it fails as a solution to the puzzles. While it offers a potentially satisfying treatment of the non-branching puzzle, its treatment of the branching and sequential puzzles reveal that it is unstable: its predictions undercut its motivations (at least for the present purpose).

Start with the branching puzzle. UK speakers go on to affirm “eBooks are books” and US speakers go on to deny it. Temporal externalism predicts, on this basis, that “book” as used in 1750 is either ambiguous or indeterminate. By hypothesis it is not ambiguous: US and UK speakers speak a common language in 1750. But if it was indeterminate, then 18<sup>th</sup> century uses of “book” are not synonymous with later US and UK uses. FTE predicts that meaning change accompanies future branching. What's more it predicts semantic differences where intuitively there are none. Take a world qualitatively alike to the world with branching, but which unceremoniously ends with a bang in 1950. FTE either predicts a) pre-1950 uses of “book” are not synonymous across these worlds (it is indeterminate in the one, but not the other), or else b) they are synonymous, and

that “book” is indeterminate in both worlds. If the latter is right, FTE predicts that whenever branching usage *is possible* there is indeterminacy. This undercuts the solution it appears to offer to the non-branching case: it is not that (2) is false (and eBooks was in the intension of “book” all along), but, rather, (4) is false: “book” as used in 2020 has a different intention, is not synonymous with, “book” as used in 1750.

Let’s go back to the two versions of the non-branching case discussed above (in one Barnes and Nobles wages a defamation campaign). Notice FTE predicts that there is meaning change in *both* cases: it predicts the meanings facts at 1750 are the same in each case, but is unable to explain speakers’ later judgments that they mean as they once did. If FTE instead predicts that “book” means differently in the short-world than in the long-world, speakers come out bad judges, not just of the intension of their term, but of whether they mean the same as one another—since “book” in these worlds would seem to be synonymous, as synonymous as any expression that did not undergo future branching. The reasons to account for semantic errors in terms of deference to usage in the community, or the way of the physical world, do not carry over. We do not coordinate with future speakers, by appeal to them, or by appeal to future events.

Now take a sequential case. It poses a special problem for FTE. Does the meaning of “livid” as used in 1820 supervene on the use of “livid” 50 years after that, or 150 years after that? Suppose speakers are right to judge they samesay with those whose usage only deviates slightly from their own. Speakers at t1 judge that they mean what speakers at t2 do. So, t2 is in the supervenience base for the meaning of “livid” as used at t1. Speakers at t2 judge that they mean what speakers at t1 and t3 do. So, t1 and t3 are in the supervenience base for “livid” at t2. Speakers at t3 judge themselves to mean as speakers at t4 and t2. So, t4 and t2 are in the base for t3. But then t1 is as well, since t2 means as t1 does, and so on. On pain of predicting the meaning of “livid” is constant across all the time slices, FTE predicts meaning supervenes synchronically, or else speakers are incorrect in their samesaying judgments fails. That is, either FTE fails, or else its motivation for the present purposes is undermined. To the extent meaning supervenes on use at, or up to, the time of utterance, FTE is false. If FTE is correct, it does not furnish a satisfying resolution of the puzzles. So, set it aside.<sup>129</sup>

---

<sup>129</sup> Should we find a counter public that uses “computer” and “phone” slightly differently from ourselves but with as much assurance, we might rearrange what we find it *reasonable* to say about what *our* expressions mean (or not). But

## 6. Similarity View: Deny (1)

Let's return to the non-branching puzzle.

- (1) Two predicates are synonyms only if they have the same intension.
- (2) eBooks at 2020 are not (or are not-determinately) in the intension of “book” at 1750.
- (3) eBooks at 2020 are determinately in the intension of “book” at 2020.
- (4) “Book” at 2020 is synonymous with “book” at 1750.

We tried some strategies for denying (2), most recently FTE, to no avail. Another strategy is to deny (1) and affirm instead: two predicates may be synonyms, in the ordinary sense, even if they do *not* have the same intension. What's required for synonymy, or the truth of samesaying judgments, is either *sameness* of something else—topic, say—or else *similarity*, of something, either of intension, of functional role, of the semantic anchors of competence, of sets of intertranslatable expressions, or something else. These views have some intuitive appeal. While, the jargon “having the same linguistic content” and “having the same intentional content” comes down to having an identical intension, we may doubt they capture ordinary samesaying and synonymy. In particular, they may seem too demanding.

There's also appetite for this idea, and it presently appears in a few different incarnations in the literature. Haslanger (2019) argues that concepts may be individuated by functional role, and conceptual continuity is ensured in “F” talk so long as there is continuity in the functional role it occupies in our lives, even if the informational content (intension) associated with “F” changes. Sundell 2020, implicitly following Haslanger, expresses sympathy for this idea. Cappelen 2018 responds to the subject change worry by appeal to sameness of topic, where topics are individuated by ordinary samesaying judgments.<sup>130</sup> Sawyer (MS) attempts to individuate Cappelen-esque topics by reference to future usage. Richard 2019 argues that conceptual continuity consists in continuity in the anchors of semantic competence: in effect, continuity is preserved so long as fluid communication and coordination is.

---

surely what “phone” and “computer” mean around here is no different than what they mean in a world *with such a counter public*. Perhaps the off chance of there being a counter public should not shake my faith in temporal externalism. But first I'd need a reason to support a crazy view like this. It says what “F” means in standard Boston English right now may hinge on what future speakers do, and whether rockslides occur, and how quickly Mennonites in Ohio come around to iPhones.

<sup>130</sup> Perplexingly he recruits only talk about “says”, despite Hawthorne 2006's points about the looseness of says-talk in his eviscerating review of Capellen and Lepore 2005. But never mind.

If synonymy does not consist in sameness of intension, then (2), (3), and (4) are not inconsistent. “Book” as used as 2020 samesays with “book” used in 1750 does, although eBooks are in the intension of one but not the other. An expression may retain its meaning, in this sense, while changing its intension. Novel usages are not subject changing insofar as subjects are not individuated by intensions. That’s the basic insight. The basic cost is that of offering an error theory for ordinary content attributions and talk of samesaying. Doing so leaves the subject change worry otherwise intact (on pain of denying that there is anything *else* for content attributions to consist in). Neither is particularly attractive.

Start here. Notice that the compelling strategy for denying (1) is to deny that, in ordinary content attributions, samesaying is an equivalence relation. Theories of conceptual continuity that rely on *identity* of something are subject to versions of the puzzles. Since identity of anything is, presumably, an equivalence relation, branching and fusion puzzles may be constructed (say, where each fragment community takes itself, but not the other, to preserve, the topic). Non-branching puzzles may be constructed if there are asymmetries in the judgment about whether, say, “book” as used in 1750 concerns the *same* topic, fills the *same* functional role, has the *same* intension, as “book” as used in 2020. Sequential cases may be constructed insofar as speakers are good judges of synonymy (whatever that consists in), and gradual variations in usage do not appear to produce change in topic, functional role, intension. If functional roles and topics are individuated by samesaying judgments, versions of the puzzles likely persist. (Plausibly Johnson would deny that talk of books in 2020 concerns the same topic as his talk of books, though we would affirm it). Cappelen is explicit that topics are individuated by samesaying judgments, and there is nothing more to them. So if crosstemporal and counterfactual reports involving “book”, and samesaying judgments, produce puzzles (as they do) efforts preserve subject matter by retrenching to topics cuts little ice. Likewise, when accusations of subject change are levied against novel usages in ordinary discourse.

Theories of conceptual continuity that rely on *similarity* of something—role, anchors of competence, intension—fare better. Similarity relations are not equivalence relations. For example, A may be similar to B and C without B and C being similar. So, in the branching and fusion puzzles, use of “F” in each fragment community may samesay with use of “F” in the heterogenous community, without being ordinarily synonymous with one another. The claims compromising the sequential puzzle come consistent: temporally adjacent uses may all be



ordinarily synonymous without temporally remote ones being so. Novel usages may still samesay with established uses even if their truth conditions may be distinct.

Even so, we might wonder how these views fare in explaining the source of the puzzles. For example, homophonic translation succeeds in some cross-temporal reports, but not in analogous counterfactual reports. “Johnson extolls the virtues of books” is true, but “If Johnson were here, he would extoll the virtues of books” is not obviously true. The reason is that if Johnson were here, what he extolls is not the referent of our term “book”. Counterfactual reasoning like this is part of the support for (2). Cross-temporal reports are part of the support for (4). The synonymy-as-similarity view does not on its own explain why homophonic reporting should succeed in the one but not the other. For example, suppose samesaying, conceptual continuity, consists in similarity of functional role. It’s unclear how to analyze the difference between the reports. It appears homophonic translation succeeds because the functional role of “book” as used by Johnson is continuous with our own. But, in the nearest world where Johnson is here, it fails because talk of “book” does not fill a functional role continuous with the one filled by our term “book”. It sounds as though it is relevant similarity in intension is doing the work, rather than continuity of functional role. Perhaps the antecedent makes salient the gap between Johnson’s usages, and presumptive dispositions, and our own.<sup>131</sup>

Denying that ordinary attributions of samesaying, ordinary continuity of subject, is an equivalence relationship comes at a cost. If there are intensions, associated with expressions and sameness of intension is what *true* samesaying consists in, then the strategy for denying (2), in effect, offers an error theory for ordinary attributions. Distinguish two versions of the subject change worry. The first concerns cross-temporal reports, counterfactuals, and samesaying claims in ordinary discourse. The synonymy-as-similarity-of-something view does vindicate unmarked samesaying judgments in ordinary discourse. But, plausibly, most of these do involve sameness of intension. (One may wonder about the motivation for a theory of continuity that repudiates complicates the basis of unproblematic ordinary judgments.)

The second subject change worry concerns these phenomena in a more theoretical register. If the intension/extension associated with an expression change, speakers before and after the

---

<sup>131</sup> There are many troubles with treating synonymy as mere similarity, but they are beyond the scope of this paper. See Fodor 2002 (especially ch 5 and 6).

change will semantically express propositions that are about different kinds of things. Similarly, speakers who express disagreement using the respective meanings will semantically express propositions that are logically consistent. So the similarity-as-synonymy-view does not advance a satisfying explanation of the puzzles, or novel usages. For example, in some cases, where semantic change appears to have taken place, inquiry is only continuous *colloquially*, and disagreement only substantive *colloquially*. In disputes, where what is at issue is whether the subject is being changed, this cuts little ice. Entering a more theoretical register is not that hard. You do not need a seminar room. (“Johnson’s word “book” is not *really* synonymous with ours” or “Locke didn’t really mean *property* as we think of it”). If the intensions of the relevant terms had remained constant in these cases, we would have an account of what makes the inquiry continuous and what makes the disagreements substantive—the speakers inquire into the same kinds of things; they express propositions that are logically inconsistent. But if there has been a change of intension in these cases, then there is no way to explain what it is that’s continuous about continuous inquiry and what constitutes the substance of a substantive disagreement. By this token, the synonymy-as-similarity-view offers no way to represent the problem when novel usages are contested, and accused of changing the subject. There is an absolute fact of the matter about whether they do, or do not, and that fact appears to consist in whether the novel usages express the same intension as established ones, or not.

### 7. Semantic Vocabulary is Plastic: Deny (4)

Let’s again return to the non-branching puzzle.

- (1) Two predicates are synonyms only if they have the same intension.
- (2) eBooks at 2020 are not (or are not-determinately) in the intension of “book” at 1750.
- (3) eBooks at 2020 are determinately in the intension of “book” at 2020.
- (4) “Book” at 2020 is synonymous with “book” at 1750.

Strategies for denying (1) and (2) have foundered. Another strategy is to reject (4), on the grounds that, there are in fact, many different synonymy relations, not because ordinary synonymy talk should be analyzed as similarity talk, but because semantic vocabulary is highly sensitive to the use facts, and so, in effect, “synonymous” has one intension at 1750 and another at 2020. Thus, in 1750, (4) is false, but in 2020 (4) is true. At any given time, synonymy consists in sameness of intension.

Dorr and Hawthorne 2013 suggests another sort of response. They are concerned with capturing speech reports and counterfactuals in view of some puzzles given two options about how word meaning supervenes on the world. Assume the semantic facts supervene, at least nomologically, on facts specifiable in some “lower level” vocabulary, such as that of microphysics. Imagine two models for how such supervenience might work for a given such fact. On the first, the semantic facts in question supervene on the base in the same way as the fact that a certain person’s height is between 61 and 62 inches. Tiny changes in the underlying base sometimes produce semantic changes but mostly they don’t. They call this being modally robust. On the second model, the semantic facts (true semantic propositions) vary more like the proposition that four particular point-particles are arranged in a perfect square: whenever it is true, there are arbitrarily tiny microphysical differences which would make it false. They call this being modally plastic.

They argue various considerations suggest that non-semantic vocabulary frequently is robust, while semantic vocabulary is plastic and vague. What variance in the base definitely produces semantic changes is a vague matter resolved by the context of utterance. (It is vague and so frequently contextually determined whether uses of “P” express P or P\* when the supervenience base for “P” is in the boundary region between “P” definitely meaning P and definitely meaning P\*). The argument is that such a view avoids the problem with thinking non-semantic vocabulary is modally plastic: doing so leads to pervasive error in ordinary counterfactual speech reports. A view like Dorr and Hawthorne’s also avoids the problem with thinking non-semantic vocabulary is modally robust: if it is, it might appear some utterers of “P”—those “near” a partition—are mistaken to think those in nearby worlds express what they do by “P.” Dorr and Hawthorne argue that if semantic vocabulary is modally plastic and non-semantic vocabulary is modally robust, speakers will express what their immediate neighbors do by P.

This view is tailor-made to fair well with the non-branching, sequential and fusion case, although one may wonder *what precisely* about usage is such that “F” as used at one time refers, refers\*, refers\*\*, but at any given time semantic vocabulary only expresses one of these relations. But an explanation of *why* in a branching case usage at a time should indicate a relation where earlier and later usages of “F” express the same intension, while at the earlier time semantic vocabulary indicates a relation that does not. One might suspect that an explanation of that is what a more illuminating, less ad hoc solution to the puzzle looks like.

Even so, the plastic picture successfully predicts the expected truth-values for some counterfactuals and speech act reports that struggle if non-semantic vocabulary is plastic or else the partitions sharp. For instance, “if Sally were here, she’d ask for a salad” and “He said Sally asked for a salad” turn out false if “salad” means salad but is modally plastic or on the boundary with meaning salad\*. Sally would ask or asked for salad\*. They analyze these statements *If Sally were here, she’d ask\* for salad* and *He said Sally asked\* for a salad*.

On their picture there are a plurality of relations for every apparent piece of semantic vocabulary. For every variation in the supervenience base – the facts that determine meaning – there is an assertion-like relation, meaning-like relation, a belief-like relation and so on. It predicts, in effect, that homophonic utterances made in communities that vary only slightly with respect to the supervenience base will very frequently express the same proposition. Or rather: one utterance asserts P and the other asserts\*; “P” in one community means P and “P” in other community means\* P.

This picture upholds perspicuity in the sense that when speakers take themselves to inter- or intra-subjectively relate to the same content they often do. It’s just that the relations in question vary very finely. It may be true that “I believe what she believes” but it’s being true involves two distinct belief relations. What the sentence expresses is *I believe what you believe\**. Testing whether this satisfies perspicuity: ask yourself is *that* synonymous with “I believe what you believe.” Plausibly the answer is no. That’s enough to reject their view by my lights.

Dorr and Hawthorne’s view is in part motivated by the thought that some unfortunates near the edge of their cell will be prone to making false counterfactual/ cross-temporal reports. How do we know that we are not like them? We could endorse the *claim* that those who are not too dissimilar from us should be interpreted homophonically. But how could we know *that*? The trouble is that if there are any worlds where homophonic interpretation is appropriate, there must be close worlds w1 and w2 such that homophonic interpretation is appropriate for w1 but not for w2. Philosophers at w1 speak falsely when they endorse the claim above. Treating semantic vocabulary as plastic and boundary regions as contextually-resolvable vagueness goes part way. But the trouble reemerges when it comes to reports of reports, like “I believe what you believe.”

Moreover, a view like theirs will do nothing to address the asymmetries in synonymy judgments I have identified, unless they may be explained as contextual variation on either side of the boundary region for non-semantic vocabulary. The idea would be when one is reporting on a

particular temporally or counterfactually close situation, the vagueness of semantic vocabulary is partly resolved so as to make the report legitimate. That doesn't help for the kinds of cases I care about. Turn to the sequential case. I'll suppose there is boundary vagueness, because they would for that sort of case, but we could as easily suppose there is not.

Suppose that Fred uttered "In August, Barney said that the bruise was livid" once in September and once in October. On the second occasion, the similarity between the usage of "livid" in August and October was salient to him. Because of this, the vagueness of semantic vocabulary was resolved on that occasion in such a way as to make "Barney in August meant the same thing by "livid" that I mean now" legitimately assertible. During the September speech, the similarity between September and August was salient to Fred, but the similarity between October and August was not salient to him, since he was not thinking about October. On the proposed account, this difference in the salience facts means that the vagueness of "say" is resolved differently in Fred's September and October speeches. Thus, it is not definitely the case that there is a single proposition that Fred asserted both in September and in October. So, if you, on some later date, say "Fred has twice said that in August, Barney said that the bruise was livid," your speech will not be definitely true, even if the similarity between all the earlier dates is salient to you on that later occasion. Predictions of this sort are problematic on their own. No minor tweak will help: it is of the essence of the approach they put forward that it buys us the right to discount certain small differences in use as irrelevant to the meaning of non-semantic vocabulary only at the price of making small differences extremely relevant in the case of semantic vocabulary. These maneuvers also make no head way addressing the apparent asymmetry of our synonymy judgments, unless it's by appeal to contextual variation in the meaning of semantic vocabulary. It hard to see how that could help. Plasticity of whatever vocabulary creates problems for reports of use of that vocabulary.

All this is to say, reports of reports require an error theory if semantic vocabulary is highly plastic. But avoidance of error theories is just what drove us to endorse the plasticity of semantic vocabulary in the first place. So, let's set this approach to the puzzles aside.

### **8. Preferred Solution to the Puzzles: Projectivism**

Some otherwise attractive strategies for rejecting claims (1), (2), and (4) of the non-branching puzzle founder. Error theories for ordinary content attributions bring costs, as do appeals to the

behavior of future speakers or to a multiplicity of semantic relations. My preferred strategy, projectivism, avoids these pitfalls. We co-refer when we samesay, and we more or less know when we do. But projectivism denies that there are absolute facts about meaning. For example, projectivism says that, in 2020, the utterance “eBooks are books” does not have a truth value. It only has one in view of a perspective, say from the perspective of 2020 (true), or of 1750 (false). Mostly this makes no difference because the perspective of the assessor and the perspective of the audience/speaker are the same. Mostly, we are right by our lights. The next section makes a conjecture about why this might be. For now, the important point is that in the puzzle cases these perspectives come apart. Let’s *again* return to the non-branching puzzle.

- (1) Two predicates are synonyms only if they have the same intension.
- (2) eBooks at 2020 are not (or are not-determinately) in the intension of “book” at 1750.
- (3) eBooks at 2020 are determinately in the intension of “book” at 2020.
- (4) “Book” at 2020 is synonymous with “book” at 1750.

Projectivism suggests the tetrad is consistent in the sense that each of the statements is true from some perspective, although there is no perspective from which all four are true. Evaluated from the perspective of 1750, (2) is true, but (3), or else (4), is false. Evaluated from the perspective of 2020, (3) and (4) are true, but (2) is false. That is, Johnson, by his community’s lights, is right to deny, “eBooks are books”. We, by our community’s lights, are right affirm, “eBooks are books”, and to affirm that “book” means as it did before the invention of eBooks. From each of our perspectives, we disagree on the truth of a proposition. But what proposition varies with our perspective. It is, we might say, projected by interpreters’ community. Meaning facts vary in this way because how it’s reasonable to go on does. Projectivism is comprised of three claims.

Claim 1: Semantic discourse is assessment sensitive: the truths varies with the perspective of the assessor.

Claim 2: the intensions of at least some expressions vary with the perspective of the assessor.

Claim 3: what an utterance means is a projection of how it is most reasonable to go on with it in the assessor’s community, given how it is used and other facts about the context of utterance.

I’ll give more substance to the view in the next section, but it’s worth spending a moment on its motivations, and the puzzles.

Let's back up. Facts about the meaning of "book" exhibit a peculiar pattern of stability and variability. On the one hand, the intension of "book" seems to vary across time. This gives us claims (2) and (3) of the non-branching puzzle. If Johnson were here, he would not mean *book* by "book." On the other hand, it seems stable across time: Johnson's talk of books, and our own, are about *books*. We report him homophonically. This gives us claim (4). Notice, the temptation is *not* to say, "eBooks are books in 2020, but they *were not* books in 1750." Rather, it is to say that they are books, and always were. What it is to be a book hasn't changed, although technology has. When we take Johnson's perspective, at the context of utterance, "eBooks are not books" seemed true, but that very utterance assessed later on it seems false. We might imagine Johnson being a medical miracle, living down to the present day, and disavowing his earlier utterances, just as, perhaps, some bloggers now disavow their earlier articles decrying that eBooks heralded the death of the (print) book.

Considerations like these—about homophonic third person reports, about retraction, about stability and variability of content in an area of discourse, about genuine but apparently faultless disagreement—are sometimes taken as evidence that truth in that area of discourse is relative to a context of assessment (a perspective).<sup>132</sup> The argument has two stages: first establish the pattern of data, second argue that the best explanation of it takes truth to be relative to the context of assessment (holding fixed the context of utterance, and whatever other parameters).<sup>133</sup> The tools so called assessment sensitivists have developed, notably MacFarlane 2014, provide some machinery to give substance to the claim that the true value of an *utterance* may vary with the perspective taken.

Projectivism about intensions may be understood as the view that semantic vocabulary is assessment sensitive. To the extent arguments for assessment sensitivity in other areas of discourse

---

<sup>132</sup> Areas of discourse include discourse about predicates of personal taste (Kölbel 2004; Lasersohn 2005), moral judgments (Kölbel 2004), epistemic modals (Egan, Hawthorne and Weatherson 2005; Egan 2007; MacFarlane 2011), future contingents (MacFarlane 2003), indicative conditionals (Weatherson 2009), gradable adjectives (Richard 2004), deontic modals (MacFarlane 2014, ch. 11), and knowledge attributions (Richard 2004; MacFarlane 2005, 2010, 2014).

<sup>133</sup> Accordingly, strategies to reject a diagnosis of assessmentism take two broad forms; either they reject the data (retraction or otherwise) as merely apparent, or else they reject that the best explanation of it involves assessmentism. Perhaps the theory of assertion is rejected, or a better explanation is supplied, or unacceptable conclusions are drawn out. See Spencer MSb for all three.

make assessment sensitivity about them a worthy position to consider, the same may be said about the arguments here for projectivism. Actually, I think the case for projectivism is rather stronger.

The argument from faultless disagreement for predicates of personal taste illustrate the assessment sensitivist's strategy well.<sup>134</sup> The basic problem is this: if Sam says "*Paradise Lost* is boring" and John says "*Paradise Lost* is not boring", it seems possible for (a) both sentences simultaneously to be true (relative to their respective speakers), but also (b) that one expresses the very proposition the other denies.<sup>135</sup> The relativist's intuition is that this is right, and truth-relativism is her strategy for reconciling these intuitions: the same proposition (or proposition-like object) is affirmed by Sam that is denied by John, but that proposition is true relative to Sam's context of assessment (circumstance of evaluation, perspective), and false relative to John's. It is then added that two people can overtly disagree—contradict each other—insofar as they affirm (or deny) a single proposition, although its truth varies with the perspective we take (either the speaker, her interlocutor, or perhaps a some third party).

This suggestion is implemented as a variation on Kaplan's view in "Demonstratives" by adding people to the circumstances of evaluation, in addition to the three parameters Kaplan had (worlds, times as locations). Words like "boring", and "tasty", are assigned the same content relative to different individuals, but contextually relativize the assignment of truth values to contents, so that the same content may be assigned different truth values relative to different individuals. This allows for the possibility that two utterances express identical semantic content, but with one of them true and the other one false. Sam says what John denies, though its truth varies with perspective.

Note that the puzzles are structurally like cases of so-called genuine, but faultless, disagreement that are diachronic. For example, in the non-branching case: if SJ affirms "eBooks are not books" and moderns affirm, "eBooks are books", it seems possible for both sentences simultaneously to be true (relative to their respective speakers), but *also* to overtly contradict or disagree with each other, in that one expresses the very proposition the other denies, ones about

---

<sup>134</sup> Lasersohn 2005, Macfarlane 2014, Kölbel 2015

<sup>135</sup> Though Lasersohn 2005 characterizes these objects as proposition, they are not propositions in the standard sense, i.e. they are not functions from worlds to truth-values. See Cappelen 2007.



*books*.<sup>136</sup> The inclinations that give rise to these puzzles are: (a) both we and SJ are right, and we both said something true, and (b) we say that eBooks are not books, SJ denied it, and so we disagree.

Notice that the puzzles also exhibit the hallmark of retraction, which is not so much that a single speaker recants, as that an utterance that appears true from one perspective appears categorically false from another. That is: *Paradise Lost* is boring, and any utterance to the contrary is wrong, even those made by John (from Sam's perspective). We might imagine Sam, after changing his mind about the book, rightly saying: "I was wrong, *Paradise Lost* was never boring." By the same token, it's not that eBooks didn't used to be books, but they are now, or, that, from Johnson's perspective, it's false, but from my perspective, it's true. It's just true (from the modern perspective) that eBooks are books, and any utterance to the contrary, even those made in 1750, were wrong. (We might add: "book" may have a meaning exclusive to print books, but that expression is not synonymous with "book" as it has been used for centuries. That expression is a newcomer, born because technology has made relevant the distinction between digital and print books).

There are two reasons retraction plays a key role in arguments for assessmentism. First, other theories, like contextualism and invariantism, do badly capturing it. The second is deeper: a Brandomian theory of assertion says that assertion involves committing to defend the truth of what is said. So, when we no longer defend some earlier statement, when we retract it, although no further information has falsified it, that is evidence its truth value has changed.<sup>137</sup> Some of the reasons to doubt retraction data obtains, say, for epistemic modality, don't apply to the non-branching case. Suppose I aver, "Bob could be in his office", and, when I later learn the building is shut down for Covid-19, I say, "I take that back, Bob cannot be in his office". My retraction may be explained by my learning a simple fact, that bears uncontroversially on my beliefs about where Bob could be. Before I knew, I spoke falsely, but blamelessly. If I had known earlier what I later learned, I would not have said what I did. But a similar analysis does not apply to the non-

---

<sup>136</sup> p. 662, Lasersohn (2005)) Though Lasersohn characterizes these objects as proposition are not propositions in the standard sense, i.e. they are not functions from world to truth-values. Cappelen against this with PCR

<sup>137</sup> There are good worries about this picture of assertion and its relationship to the objective epistemic ought. See Spencer MSb. I tend to think they do not apply to projectivism about meaning for reasons that relate to what perspectives consist in. But we are not there yet, and this point is largely beyond the scope of this paper.

branching case, or informed disagreements about meaning. We supposed that Johnson knows how speakers went on to use “book”. More to the point, there does not appear to be an *uncontroversial* fact he might have learned which would settle whether they are books. Retraction cannot be recast as learning what is true from every perspective.

It’s worth bearing down on this a bit, since there appears to be a problem. It may be argued: from the perspective of 2020, (1) “book” has not changed its intension since 1750, (2) it is an absolute fact that, from the perspective of 1750, “eBooks are books” is indeterminate, and, therefore, (3) “book” is indeterminate, or else (1) is false. The response on behalf of the projectivist is this: not all facts are accessible from a perspective. From the perspective of 2020, the facts *from the perspective of 1750* are not accessible. In this sense, Johnson entertained thoughts that we cannot, except elliptically as *the thoughts that Johnson entertained*. In this, talk of perspectives, contexts of assessment, are not analogous to talk of perspectives on a coin on a table. That the coin is circular when viewed from above, ovoid when viewed at an angle, are mutually consistent. Whether this is a problem depends on whether you are bothered by projectivism being difficult to state in this way.

In the interests of completeness, let’s apply projectivism to the other puzzles. For the branching and fusion puzzles, it denies there is a single perspective at which the claims comprising the puzzle are true. Consider the branching puzzle:

- (1) eBooks are not in the intension of “book as the US in 2020; it hasn’t changed its intension since 1750.
- (2) eBooks are in the intension of “book” as used in the UK in 2020; it hasn’t changed its intension since 1750
- (3) “Book” as used in 1750 for tomes (as opposed to texts) is not ambiguous.

From the perspective of 1750, (1) or (2), or both, are false. From the perspective of the US in 2020, (2) is false. From the perspective of the UK in 2020, (1) is false. Now, the fusion puzzle:

- (1) In the East at 0, codices are not (or are not determinately) in the intension of “liber”.
- (2) In the West at 0, codices are in the intension of “liber”.
- (3) At 100 AD, “liber” means as it did in the East at 0, and as it meant at in the West at 0.
- (4) “Liber” as used at 100 AD is not ambiguous.

From the perspective of the East, (3) is false. From the perspective of community West, (3) is false. From the perspective of the the heterogenous Empire, (1), (2), or both, are false or indeterminate. Lastly, the sequential puzzle:

- (1) “Livid” as used at t1 is not synonymous with “livid” as used at t5.
- (2) Uses of “livid” are synonymous with uses at adjacent times.

From the perspective of any given time, both (1) and (2) are compatible. At each time slice, from the perspective of that time slice, the temporally adjacent uses of “livid” are synonymous. However, there is no time slice, such that from it, uses of “livid” at every time slice are synonymous with temporally adjacent uses. So there is no puzzle.

We can draw an analogy between meaning facts and friend groups. Take ten well-informed people some of whom are closer than others. Ask each who is in the friend group, and you may receive several different answers. One theory says only some people are right: there’s an absolute fact of the matter about who is, and isn’t, in the group. We can offer an error theory: mostly when we talk about friend groups, or make judgments that seem to be on the basis of friend group facts, we aren’t really. Some other sort of relatedness is what matters. Another theory says everyone is right, but there are several distinct but collocated friend groups. A third theory says everyone is right, and there is a single friend group, but its contours depends on whose perspective you take. I lean towards the third sort of theory—about meaning facts.<sup>138</sup>

## 9. Sketch of Projectivism

Projectivism rests on three claims. The first two we have already met, the third is a conjecture about why they should be so.

Claim 1: Semantic discourse is assessment sensitive: the truths varies with the perspective of the assessor.

The truth value of at least some statements about reference, intension, synonymy vary with perspective of the assessor. That is, there is at least one expression where a statement about its semantic properties has one true-value as assessed from one perspective, but another as assessed

---

<sup>138</sup> The boundaries of English may be like this, perhaps type-clusters of polysemous words are, vowels may be like this?

from another. So, the intensions associated with at least one expression vary with the perspective of the assessor. That is, Claim 1 entails Claim 2:

Claim 2: the intensions of, at least, some expressions vary with the perspective of the assessor.<sup>139</sup>

Take Johnson's (imagined) utterance that eBooks are not books. As assessed from 1750, the intension of "book" as used in 1750 excludes eBooks. As assessed in 2020, it includes them. As Johnson said, "A writer only begins a book. A reader finishes it." There is no absolute fact of the matter about whether eBooks are books, or what "book" means. The perspective of an assessor must be supplied. Claim 1 entails some terms are context sensitive in a special way: they are context-of-assessment-sensitive. The intension they express varies with the context of assessment. Some vocabulary, say "battery" or "electron", may be such that the truth of the statements involving them do not vary with the perspective of the interpreter. Perhaps "electron" is like this. Other vocabulary I have argued is not. We can disagree on which vocabulary is which. Perhaps the 2020 utterance "'Battery' as used by Volta means *battery* as used by denizens of 2020" varies with the perspective of assessor, perhaps it doesn't. One way to adjudicate is to try to construct puzzle cases with the relevant term. That is a way to test what perspectives could be occupied. A way to identify an occupied perspective to identify persisting, informed, intelligible disagreements.

Every perspective, or every occupied perspective, may concur. So, for example, it may be true assessed from every perspective that "villain" as used in 2020 is associated with a different intension than "villain" "villain" as used in the 14<sup>th</sup> century (when it was applied to local peasant farmers, from the "ville"). Likewise, the intensions associated with epistemic and thin moral predicates may be concordant.

There appear to be absolute facts about meaning when there is concordance among all the perspectives, or all the perspectives we deem worthy. Perhaps many more perspectives are possible than are occupied. Perhaps, from some perspective, print books are not in the intension of "book", but eBooks and audiobooks are. But this perspective is not occupied, and so there are not disputes about it. Likewise, there could be people in our lives who would judge the contours of the friend group in ways we find very strange. But there aren't, or we don't recognize them as worthy conversation partners, and so we don't fight about it.

---

<sup>139</sup> A much stronger claim is that every expression is like this. Set that aside.

Shake off the view that possible worlds are like rooms. Instead think of worlds like total stories, as sets of propositions.<sup>140</sup> Projectivism about meanings says there is more than one actual world, one total story. Some parts of the total stories are the same, some are not. Just like in total story at 1750, in the total story at 2020, the pyramids were built in 2500 BCE, murder is wrong, and everyone knows it. In both, Johnson was bored by *Paradise Lost*. In both, expressions are meaningful, and books\* are eBooks, but books\*\* are not eBooks. But there are some differences.

In the actual world at 1750:

“Book” means in 1750 what it does in 2020 (or else it is indeterminate whether it does).

Had Johnson affirmed, “eBooks are not books”, he would have been *right* to do so, and he would have expressed the proposition: *eBooks are not books\**.

Actual expressions are associated with *books\**, not *books\*\**. Propositions about *books\*\** are not entertainable, except by indirect means (e.g., entertaining the proposition 21<sup>st</sup> century speakers entertain by their own lights).

In the actual world at 2020:

“Book” means in 1750 what it does in 2020

Had Johnson affirmed, “eBooks are not books”, he would have been *wrong* to do so, and he would have expressed the proposition: *eBooks are not books\*\**.

Actual expressions are associated with *books\*\**, not *books\**. Propositions about *books\** are not entertainable except by indirect means (e.g., entertaining the proposition Johnson entertained by his own lights).

In both stories there are meaning facts, and discourse about meanings. But in the world at 1750, some utterances are associated with different truth conditions than those very utterances are in the world at 2020. In the world at 1750, Johnson spoke of *books\**. In the world of 2020, he spoke of *books\*\**.<sup>141</sup>

---

<sup>140</sup> Wright 2008 calls this the ternary conception. See also Baghramian and Coliva 2020.

<sup>141</sup> Perhaps fundamentality talk suggests a way to register the variation: facts that vary across occupied perspectives are less fundamental relative to them than those facts that do not. So facts about where the paper is, and where words are, are more fundamental than facts about the meaning of “book” relative to the world at 1750, and to the world at 2020. (Perhaps they are equally fundamental relative to the world at 1749 and 1750, or perhaps those perspective are just the same one). For a development of a related idea, see Spencer MSa.

*Objection:* Surely, saying that the facts about what books are like varies with perspective does not advance the issue. That is like saying you have your facts, and I have mind.

*Reply:* Not quite. From my perspective, our facts are the same and you’re mistaken about them. These aside, many other facts are constant across my perspective and yours, including facts about paper, the distribution of words, and

Claim 3 is a conjecture about why meaning facts vary in this way. It gives substance to talk of a perspective. Projectivism predicts that (by our own lights) more than anyone else we speak truly, and are right about the meaning facts. It's as if manifest destiny that we judge eBooks are books. However, we went on to use "book", it would be the way that was correct all along. This is surprising, given that the patterns of usage that drive semantic change are beset by contingency. Reasonability judgments are also like this. We are reasonable *by our own lights* much more often than those whose judgments about what's reasonable differ from ours. Likewise, holding the facts fixed, our judgments about the most reasonable theory are the ones we should have had all along, at least according to us. But it is no mysterious why. As Lewis 1971 argued, reasonability judgements are immodest because reasonability standards are: we are rationally required to believe our own methods are best, or else to suspend belief (when faced with ties). Claim 3 conjectures that immodest standards explain the variability and stability of meaning facts; the manifest destiny that all along we should mean *book* by "book".

Claim 3: what an utterance means is a projection of how it is most reasonable to go on with it in the *assessor's* community, given how it is used and other facts about the context of utterance.

A perspective of assessment consists in facts about how it is reasonable to go on with the predicate in question in the assessor's community. Where what is reasonable takes into consideration facts about the context of utterance and past usage, as well as facts about the physical and social environment, popular beliefs, and patterns of deference, in the assessor's community, what its members are likely to judge as the referent, and so on.).

Take 1750 uses of "book". The most reasonable way for modern speakers to go on with *Johnson's* word is by applying it to eBooks, by translating it homophonically. The most reasonable way for Johnson to go on with *his* word "book" is by refraining from applying it to eBooks. The difference owes to differences between Johnson's community and our own: a difference in popular beliefs, dispositions to respond to new technology, facts about what distinctions are salient to members of the communities, and much else, rationalizing projecting one use over another. But it is also a mistake to reify a character for "book" such it applies to eBooks when *used* in the context

---

readable technological widgets. What varies is what propositions are associated with particular utterances. On both stories we have a disagreement about the extension of "book".

of 2020, and not when used in the context of 1750. The reason is that there is no perspective from which both such uses are correct. From the perspective of 2020, the 1750 uses are wrong.

Individuals can be wrong about what an expression means, but, at least for some expressions, whole communities cannot be, by their own lights. If the intension of “book” is projected by how it is reasonable to apply the term in the assessor’s community, the intension of “book” may be conditioned by popular beliefs about it.<sup>142</sup>

I am not giving a theory of how to individuate communities, except to say that a good indication that the truth of some utterance varies across them is that we can imagine a puzzle case. We can imagine a community that includes 18<sup>th</sup> century speakers and 21<sup>st</sup> century speakers, and ask how it is reasonable to go on using “book”. That potentially generates a fusion case. Another good indication is the existence of a persisting informed dispute. In this, I am in good company as not much more is said about how to individuate contexts of utterance, or contexts of assessment beyond competent speakers displaying a certain pattern of judgments. In this way, projectivism is descriptive, rather than prescriptive. It is descriptive in that it offers an account of certain puzzles and disagreements. But it does not prescribe how to individuate communities, or how to read the semantic facts off of the non-semantic facts (such that would falsify some utterances once and for all). Projectivism takes our ordinary judgments to describe the meaning facts, as they are, by our own lights.

Claim 3 is a metasemantic claim. It says what a word means is the most reasonable thing for it to mean compatible with its use (broadly construed) and some facts about the assessor’s community. Any metasemantic theory must take use (broadly construed) at the context of utterance seriously on pain of generating a coincidence (we tend to use our words in ways that are true), and a mystery (we know more or less what our terms mean). If we take the appearance of so much assessment sensitivity at face value, the metasemantics must be a function of, at least, use (at the context of utterance) and something stable and variable in the right way. On the standard non-descriptivist metasemantic picture what a word means is a function of use and fit. Temporal externalists have a view about what use matters. Teleosemanticists have one view about what

---

<sup>142</sup> Sometimes reflections to this tune are recruited in a case against compositionality (see Travis 1989), or referential semantics (see Kripkenstein’s skeptical solution 1982, or Brandom 1994 on inferentialism). But a version of them appears to be compatible with truth-relativism about meaning discourse.

constitutes fit. Interpretivists another (it rationalizes a speaker at context of utterance). Lewisians say that fit consists in picking out the most natural properties in the vicinity of the use. To the extent use and the naturalness facts may cut against each other, or there are ties in the naturalness facts, meaning is indeterminate. I am suggesting that fit be understood as the most reasonable way to go on as determined from the context of assessment.

It's helpful to see how the metasemantic theory on offer differs from David Lewis' in "New Work for a Theory of Universals" (1983). Suppose we're in a Lewisian context where everything Lewis thinks about meaning and metasemantics is right—where the most reasonable way to go on with terms is to defer to the most natural properties in the vicinity of past use and, indeed, that is why my terms mean as they do. What my theory adds is that there are other contexts. Perhaps it's indeterminate what "salad" and a whole host of words means in a Lewisian context because use underdetermined meaning and there is no outstandingly natural property in the vicinity to attract the referent of the terms. For Lewis, assuming he is in a Lewisian context of assessment, the meaning of "salad", et al., is indeterminate in all and exactly the same ways he thinks they are. His theory is correct—in a Lewisian context of assessment. To the extent, he said you can't change what you mean by changing your context of assessment, he is right—at a Lewisian context of assessment.

"Food", for example, is indeterminate in a Lewisian context for exactly the reasons projectivism predicts: there is no outstandingly reasonable way to go on with "food" with respect to say nutritional capsules, it just so happens in the Lewisian context the most reasonable way to go on involves reference to the most natural property in the vicinity of use, and there is none here. But terms that are indeterminate as assessed from Lewis' context are perhaps highly determinate as assessed from another context: as assessed from the dystopian and efficient future, perhaps it is true "some foods are nutritional capsules" and as such, from this perspective, is true now, despite its being false from the present one.

For all we know there may even be gruesome contexts of assessment. An interesting question is how gruesome they can be and whether we can identify any context as a gruesome one.<sup>143</sup> According to Projectivism, "food"-utterances that are false or indeterminate in a Lewis

---

<sup>143</sup> I think this is a very interesting question.



context may be true or highly determinate in other contexts, and perhaps vice versa. Then there's the perfectly good question as to whether we're in a Lewisian context, or an interpretivist one, or some other sort. I won't try to answer that.

Consider a term "F" and its various potential truth conditions. My reasonability standards rank them in terms of reasonability given how the terms are used understood broadly. If there is a most reasonable one, that's what "F" means. If there is not, what "F" means is indeterminate among the most reasonable contenders. I am thinking of it like this:

	For Lewis C <sub>a</sub>	For Millikan C <sub>a</sub>	For Epistemicist C <sub>a</sub>
(1) "F" means X : .1	reasonable	.1	.1
(2) "F" means Y : .1		.1	.1111
(3) "F" means Z : .4		.1	.1

At the Lewis C<sub>a</sub>, "F" means Z" is true. At the Millikan C<sub>a</sub>, that statement is indeterminate. At the Epistemicist C<sub>a</sub>, it is false. Projectivism is compatible with people, past or peculiar, being very indeterminate about what they mean by their own lights as well as with their meaning what we mean (by our lights) but speaking falsely. A good question is what sort of metasemantics is operational at our context. Perhaps it might be reasonable to go on by supervaluation and deny that we mean what past speakers do, but I doubt it.

It is not quite right to say that as usage grows, and is more carefully negotiated, there is less assessment based variation in the intensions associated with expressions. Fewer perspectives may be occupied. In a sense it becomes harder to exit a context of assessment, in favor of another. Regardless, an assessment *insensitive* semantics may be warranted across the board, depending on what we are interested in theorizing truth conditionally. But if meaning is determined by usage, and something else, as usage becomes more determinate whatever the other thing is will play a smaller role. Sometimes usage really does crystallize around something in particular (consider, "atom", "galaxy"). No part of my theory is at odds with that possibility.

The puzzles emerge, it seems to me, because there are multiple reasonable ways of going on. I am suggesting that what an expression means, where meaning is understood as reference, varies according to how it is reasonable to go on with it in the assessor's community. If there is a tie, the expression is indeterminate. Conflicting reasonable ways of going on, in effect, occur in

different actual worlds with different meaning facts. They are individuated by ways of going on reasonably.<sup>144</sup>

### **10. Novel Usages, Fusion and the Women**

Here's how we got here. Some novel usages received as particularly radical may be accused of changing the subject. A standard way of thinking about meaning says that either established usage errs, the novel usages do, or else the novel usages change the subject. There is an absolute fact of the matter which of these three options is right. To the extent established usages are correct, novel usages err, or change the subject. But sometimes it seems like a word retains its meaning despite changing its truth conditions. The status of contemporary novel usages is not so different from usages that were once novel, but which become standard. So, we analyzed some puzzles involving diachronic synonymy of this sort. My preferred response to these puzzles is to hold that meaning facts are projected by the speakers' community. There is no absolute fact of the matter about whether this or that is a book. But there is some concordance among perspectives. Whatever the perspective, it may be print books are books, and leopards are not.

Let's return to the case of "woman" in light of projectionism, and start with a fictional branching case. Imagine an island very like the US 50 years ago. Its inhabitants speak a common language. On one half of the island, a trans rights movement gains momentum. As more society starts arranging itself in more equitable and inclusive ways, it comes to be accepted that trans women are women, and trans men are men. On the other half of the island, a free comportment movement takes hold. It comes to be accepted that people may avail themselves of the garb, mannerisms, bodily procedures they see fit, within some standard range. The practices of major insurance companies, and bathroom policies, stay with the times. There's some variation in what people choose, and this is seen as largely unexceptional. Yet, it seems analytic, if anything is, that women are adult human females, and men adult human males. Maybe there's a word for what trans women and cis women have in common, but no one uses it much.

---

<sup>144</sup> Of course another response to the puzzles, one that is not relativist, is to drive apart a theory of meaning and a theory of truth, in the way some do who take polysemy to be a core distinctive feature of natural language, rather than quirky homonymy. On a picture like this meaning constrains both use and truth, in subtle and interesting ways. An account that does justice to natural language will have to accommodate facts which suggest both that (i) use and truth are very complicated, perhaps in many intractable ways, and (ii) meaning is systematic and in many ways theoretically tractable, even for creatures with our limited cognitive powers. Recanati 2010, Wilson and Carston 2007.

Each part of the island thinks “woman” (and “man”) have undergone semantic change on other part of the island, or else are being used in (perhaps inexplicably) incorrectly. Plausibly, both parts of the island have a claim on conceptual continuity, and neither obviously sins against the word “woman”. The result is apparent transitivity failures since “woman” does not mean the same in these different parts of the island, on pain of half of the island speaking in error.

Projectivism analyzes the branching usage as follows. Each of three communities speak truly by its own lights. For instance, each half of the island is right, by its own lights, that they mean by “woman” as they did prior to the movement. But there are no lights by which the other two communities speak truly, and to the same subject one another. So, there is no puzzle. Cross-island quarrels concern what women are, but the truths vary with the perspective taken. From the pre-movement perspective, whether, or not, it’s true that trans women are woman, it is indeterminate whether the post-movement speakers mean as they once did.

We can ask, “what are women *really*?” Projectivism tells us there is no absolute answer. There is only an answer by some lights. There is, for example, an answer to the question, “What does “woman” mean on this island?” We can analyze the speech situation on the *entire* island as a case of fusion, and ask after the meaning facts by the lights of the mixed community. The answer reflects how it is reasonable to go on with “woman”, given its use to date on both halves of the island, and how it is likely to be received going forward, given the social and physical environment, patterns of deference, and the beliefs, associations, and so on, it is likely to trigger in the heterogenous community. If there is more than one outstandingly reasonable way to go on in a community where interlocutors may be from either side of the island, the expression is indeterminate. That matter may evolve as the background considerations driving use shift, including speakers’ dispositions and beliefs about the meaning of the terms being used.

Asking what women are (by our own lights) is a bit like asking what “woman” means in Islandese. There are not uniform patterns of deference to human biology, first personal avowals, activist causes, past classification, presentational gestalt, and so on.<sup>145</sup> There was time when patterns of deference were more uniform, and subaltern uses were far less numerous, if any. But

---

<sup>145</sup> From the perspective of the dominant community, some subaltern uses are incorrect, or concern another subject. From the perspective of the subaltern community, some dominant uses are incorrect, or concern another subject. The mixed community encompasses these subaltern and dominant uses, and whatever else besides.

now there are branching uses of “woman” as applied to trans women. There is the so-called subaltern branch, and the dominant branch, dominant if only because it is supposed most numerous, and taken to be the default interpretation. Some recommend usages of “women” that are received as radical. But by their lights, they are absolutely right to do so, and are not changing the subject. By the lights of those who receive them as radical, they may well be changing the subject, or just straightforwardly false. No one is obviously at risk of changing the subject, by their own lights.

How is it reasonable to go on with our stock of inherited words, given the complex pressures driving usage, has to do with mundane facts about how we tend to speak and what speakers are liable to understand. Sometimes it partly depends on how we answer moral and political questions, about what usages matter, and how much accommodation of one another is expected in view of what is appropriate. The meaning facts from the perspective of either branch community do not settle how to go on with “woman”.

The important point for the present purpose is that the question about what the Fs are, really, when controversy peaks, is not a deep metaphysical question, so much a question about how, given everyone in the community, it’s reasonable to go on applying a term. It is a familiar question about how to talk to, and about, people, and what to do in the face of competing needs and interests. But it is no easier for being familiar. Subaltern usages may be “outvoted,” be too fringe to affect how it is most reasonable way to go on with some term in core ways, given the many pressures that presently drive usage in the heterogenous community. From that perspective, some subaltern uses may turn up false, or subject changing. Reporting this by itself should be as unobjectionable as noting that some candidate won the general election. That is not the same as failing to count to votes, or a vote for mediocrity.

## **11. Conclusions**

If arguments from variability and stability, retraction, faultless disagreement arguments, inference to the best explanation are good enough for relativism about predicates of personal taste, etc., they may be good enough for arguments for truth relativism about semantic vocabulary. Perhaps my modus ponens is your modus tollens, and such arguments come off looking all the worse. Whether or not the arguments succeed for these other areas of discourse, I think similar considerations militate in favor of projectivism.

We got here in an effort to get leverage on a synchronic worry about subject change in ordinary discourse. Perhaps you think diachronic synonymy is not a good model for synchronic synonymy, and that the diachronic puzzles are not a good model for synchronic ones. I think they help hold a standard picture of meaning to account for the many practices that are explained by samesaying. I tend to think we are good judges of synonymy, and that we can, and do, samesay often. I take this as a datum, which is worth accommodating for some theoretical projects, particularly those that turn on speaker judgments about meaning. Whether novel usages—like those that figure in radical analyses and revisionary projects—change the subject is one.

Projectivism says that speakers may recommend novel usage without changing the subject, or speaking falsely, by our own lights – at least sometimes. Continuity of meaning, or lack thereof, may be understood in terms of truth and reference, without the imperialist metaphysics of meaning. Those recommending novel usage can have their cake and eat it too. They can promote intensional change without changing the subject—because what it is being promoted is the reasonable way to go on with the term at a community of assessment of which they are a part.

Insofar as philosophers of language think there is an absolute fact of the matter about what intensions are associated with what expressions—a fact of the matter about what proposition, e.g., “book”-, “woman”-, “chocolate”-, utterances express—there will be trouble making sense of persistent informed controversy. Diagnoses of verballity do not settle disputes about what women are, practically or in spirit. Projectivism suggests that we and Johnson may have a state disagreement about whether, say, eBooks are books, while by our own lights we are each right.

Projectivism is a descriptive view. It does not say how to individuate contexts of assessment, or how many there are. The interesting ones are the occupied ones. A good question is which terms are like “book”, and why suppose that contested concepts, like gender ones are. My answer is deflationary: it is the presence of certain sorts of disagreements. There are not absolute meaning facts that stand to resolve them. But projectivism does suggest how, and why, there might be other concordances across occupied perspectives. While there are not absolute truth conditions associated with meaning discourse, there are considerations that bear on the reasonability of what we say—given how the expression has been used, and facts about the context of assessment—and the meanings facts reflect these. That on its own is not terribly informative, but, I think, it puts the bump in the rug in the right place.

Natural language may not fit the model of a language in which names, including names for properties, are semantically, and absolutely, associated with entities that are satisfiers of predicates. This is a fruitful model that allowed theorists to start accounting for a certain range of facts. And it simplifies the discussion, in harmless ways, when the openness of lexical items is not at issue. But the explanatory value of this model needs attention if we want better models that start to account for ways in which lexicalization, conceptual continuity, contestation, and truth interact. The same of course has been said for lexicalization, composition, and polysemy, or other ways in which natural languages are not-*Begriffsschrift*-y.

## **References**

- Abreu Zavaleta, M. (2019). "Communication and Variance." *Topoi*.
- Baghrarian, M., and A. Coliva. (2020). *Relativism*. London: Routledge.
- Ball, D. (forthcoming). "Relativism, Metasemantics, and the Future." *Inquiry*, 41.
- Bettcher, T. (2012). "Trans Women and the Meaning of 'Woman'." In *Philosophy of Sex: Contemporary Readings*, edited by Power, N., R. Halwani, and A. Soble, Sixth. New York: Rowan & Littlefield.
- Boghossian, P. (1989). "Content and Self-Knowledge," *Philosophical Topics*, 17: 5–26.
- Bromberger, S. (2012). "Vagueness, Ambiguity, and the 'Sound' of Meaning." In *Analysis and Interpretation in the Exact Sciences: Essays in Honor of William Demopoulos*. New York: Springer.
- Brown, J. (2000). "Against Temporal Externalism." *Analysis* 60: 178–88.
- Burge, T. (1979). "Individualism and the Mental." In *Midwest Studies in Philosophy IV: Studies in Metaphysics*, 73–121. Minneapolis: University of Minnesota Press.
- . (1988). "Individualism and Self-Knowledge." *Journal of Philosophy* 85: 649–63.
- . (1992). "Frege on Knowing the Third Realm." *Mind, New Series* 101 (404): 633–50.
- Byrne, A. (2020). "Are Women Adult Human Females?" *Philosophical Studies*, January.
- Cappelen, H. (2008). "Content Relativism and Semantic Blindness." In *Relative Truth*, edited by García-Carpintero, M. and M. Kölbel, 265–86. Oxford University Press.
- Cappelen, H., and E. Lepore. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. London: Blackwell Publishing.
- Chalmers, D. (2012). *Constructing the World*. Oxford: Oxford University Press.
- Chomsky, N. (1995). "Language and Nature." *Mind* 104 (413): 1–61.
- Clark, A., and D. Chalmers. (2000). "The Extended Mind." *Analysis* 58 (1): 7–19.
- Donnellan, K. (1983). "Kripke and Putnam on Natural Kind Terms." In *Knowledge and Mind* edited by Ginet, C., and S. Shoemaker, 84–104. New York: Oxford University Press.
- Dorr, C., and J. Hawthorne. (2014). "Semantic Plasticity and Speech Reports." *Philosophical Review* 123 (3): 281–338.
- Ebbs, G. (2000). "The Very Idea of Sameness of Extension Across Time." *American*

- Philosophical Quarterly* 37 (3): 245–68.
- Egan, A. (2007). “Epistemic Modals, Relativism and Assertion.” *Philosophical Studies* 133 (1): 1–22.
- Egan, A., B. Weatherson, and J. Hawthorne. (2005). “Epistemic Modals in Context.” In *Contextualism in Philosophy*, edited by Gerhard Preyer and Georg Peter. New York: Oxford University Press.
- Field, H. (2001). *Truth and the Absence of Fact*. Oxford University Press.
- . (2018). “Egocentric Content.” *Noûs* 51 (3): 521–46.
- Fodor, J., and E. Lepore. (2002). *The Compositionality Papers*. New York: Oxford University Press.
- Frankfurt, H. (2005). *On Bullshit*. Princeton: Princeton University Press.
- Frege, G. (1884). *Die Grundlagen Der Arithmetik: Eine Logisch-Mathematische Untersuchung Über Den Begriff Der Zahl*. Translated by J.L. Austin. 2nd Revised Edition. Oxford: Blackwell Publishing.
- Gallie, W.B. (1956). “Essentially Contested Concepts.” *Proceedings of the Aristotelian Society New Series* 56: 167–98.
- Grice, H.P., and P.F. Strawson. (1956). “In Defense of a Dogma.” *The Philosophical Review* 65: 141–58.
- Hacking, I. (2007). “Putnam’s Theory of Natural Kinds and Their Names Is Not the Same as Kripke’s.” *Principia* 24 (1).
- Haslanger, S. (2000). “Gender and Race: (What) Are They? (What) Do We Want Them To Be?” *Noûs* 34 (1): 31–55.
- . (2019). “Going on Not in the Same Way.” In *Conceptual Ethics and Conceptual Engineering*, edited by Burgess, A., H. Cappelen, and D. Plunkett. Oxford: Oxford University Press.
- Jackman, H. (1999). “We Live Forwards but Understand Backwards: Linguistic Practices and Future Behavior.” *Pacific Philosophical Quarterly* 80: 157–77.
- . 2005. “Temporal Externalism, Deference, and Our Ordinary Linguistic Practice.” *Pacific Philosophical Quarterly* 86: 365–80.
- Kölbel, M. (2002). *Truth Without Objectivity*. New York: Routledge.
- . (2004). “Faultless Disagreement.” *Proceedings of the Aristotelian Society* 104(1):53-73.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.
- Lance, M., and J. Hawthorne. (1997). *The Grammar of Meaning*. New York: Oxford University Press.
- Lasersohn, P. (2005). “Context Dependence, Disagreement, and Predicates of Personal Taste.” *Linguistics and Philosophy* 28 (6): 643–86.
- . (2009). “Relative Truth, Speaker Commitment, and Control of Implicit Arguments.” *Synthese* 166 (2): 359–74.
- Lewis, D. (1983). “New Work for a Theory of Universals.” *Australasian Journal of Philosophy* 61 (4): 343–77.
- Ludlow, P. (2005). “Contextualism and the New Linguistic Turn.” In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by Preyer, G., and G. Peter, 11–51. Oxford: Oxford University Press.
- MacFarlane, J. (2003). “Future Contingents and Relative Truth.” *Philosophical Quarterly* 53: 321–36.
- . (2011). “Epistemic Modals Are Assessment-Sensitive.” In *Epistemic Modality*, edited

- by Egan, A., and B. Weatherson.
- . (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Nolan, D. (2014). “Hyperintensional Metaphysics.” *Philosophical Studies* 171 (1): 149–60.
- Parfit, D. (1987). *Reasons and Persons*. Oxford: Clarendon Press.
- Pietroski, P. (2005). “Meaning Before Truth.” In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by Preyer, G., and G. Peter, 253–300. Oxford: Oxford University Press.
- Pinder, M. (2020). “On Strawson’s Critique of Explication as a Method in Philosophy.” *Synthese* 197 (3): 955–981.
- Pustejovsky, J., and B. Boguraev. (1996). *Lexical Semantics: The Problem of Polysemy*. London: Clarendon Press.
- Putnam, H. (1975). “The Meaning of ‘Meaning.’” *Minnesota Studies in the Philosophy of Science* 7.
- . (1990). “Is Water Necessarily H<sub>2</sub>O?” In *Realism with a Human Face*, edited by Conant, J., 54–79. Cambridge: Harvard University Press.
- Railton, P. (1989). “Naturalism and Prescriptivity.” *Social Philosophy and Policy* 7 (1): 151–74.
- Richard, M. (2004). “Contextualism and Relativism.” *Philosophical Studies* 119 (1): 215–42.
- . (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- . (2019). *Meanings as Species*. First edition. Oxford: Oxford University Press.
- Sawyer, S. (forthcoming). “Truth and Objectivity in Conceptual Engineering.”
- Spencer, Jack. (MSa). “Is Nonfundamental Existence Relative?” Personal Webpage.
- . (MSb) “Truth Relativism and the Objective Epistemic Ought” Personal Webpage.
- Strawson, P. F. (1959). *Individuals*. London: Routledge.
- . (1963). “Carnap’s Views on Constructed Systems Versus Natural Languages in Analytic Philosophy.” In *The Philosophy of Rudolf Carnap*, edited by Schilpp, P.A. 503–18. Open Court: La Salle.
- . (1992). *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.
- Sundell, T. (2020). “Changing the Subject.” *Canadian Journal of Philosophy* 50 (5): 580–93.
- Vincente, A. (2018). “Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words.” *Philosophical Studies* 175 (4): 947–68.
- Weatherson, B. (2009). “Conditionals and Indexical Relativism.” *Synthese* 166 (2): 333–57.
- Williams, B. (1970). “The Self and the Future.” *The Philosophical Review* 79 (2): 161–80.
- Wilson, D., and R. Carston. (2007). “A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts.” *Mind & Language* 21: 404–33.
- Wittgenstein, L. (1967). *Remarks on the Foundations of Mathematics*. Oxford: Basil Blackwell.
- Wright, C. (2008). “Relativism about Truth Itself: Haphazard Thoughts about the Very Idea.” In *Relative Truth*, edited by García-Carpintero, M. and M. Kölbel, 157–86. Oxford University