Conceptual Engineering Should be Empirical

Ethan Landes

Draft: 3 June 2023

Abstract

Conceptual engineering is a philosophical method that aims to design and spread conceptual and linguistic devices to cause meaningful changes in the world. So far, however, conceptual engineers have struggled to successfully spread the conceptual and linguistic entities they have designed to their target communities. This paper argues that conceptual engineering is far more likely to succeed if it incorporates empirical data and empirical methods. Because the causal factors influencing successful propagation of linguistic or conceptual devices are as complicated and interwoven as they are, proper empirical research will greatly boost the likelihood that propagation is successful. In arguing for the superiority of empirical conceptual engineering over armchair-based conceptual engineering, this paper proposes a framework for understanding the causal forces at play in propagation. This is a three-part framework between the label of a lexical item, the psychological states associated with the lexical item, and the worldly things associated with the lexical item. By understanding the way causal forces affecting propagation play out at these three levels, conceptual engineers can better conceptualize, study, and harness the different causal forces affecting the success of their conceptual engineering projects.

Section 1: Introduction

Despite their apparent differences, conceptual engineering and experimental philosophy are natural bedfellows. Conceptual engineering is the ameliorative method of identifying deficiencies in our conceptual or linguistic repertoire, engineering a better replacement, and propagating said replacement (Isaac, Koch, and Nefdt 2022; Belleri 2021; Chalmers 2020). In contrast, experimental philosophy is the employment of empirical methods to answer traditional philosophical questions and is driven by the conviction that many traditional philosophical questions turn on complex and/or implicit empirical claims (Horvath and Koch 2021; Machery 2017; Pohlhaus 2015). Despite their apparently different aims, both experimental philosophy and conceptual engineering represent alternatives to the methods of mainstream analytic philosophy (Torregrossa 2022), and so unsurprisingly, there has already been considerable work exploring the intersection of the two. We can split this work into three broad camps. There are those who have suggested conceptual engineering can or should be done empirically at some stage in the process, there are those who have used extant empirical data to try to better understand conceptual engineering, and there are those who have undertaken experimentally-informed conceptual engineering.

The largest set of papers examining the relationship between conceptual engineering and empirical methods have argued that conceptual engineering can or should be led by empirical methods to one extent or another. Nado (2019) argues a retooled experimental philosophy can play an important role in functional accounts of conceptual engineering (see Nado 2021) by helping discover the functions of various concepts, words, and so on. Focusing on aesthetics, Torregrossa (2022) argues experimental data can (and already has) revealed defectiveness of aesthetic concepts but remains skeptical that experimental philosophy can help answer all normative issues faced by all conceptual engineers. In conversation with Torregrossa, Andow (2020) argues that normative truths are well within the grasp of experimental methods and maps

out a framework for fully experimental conceptual engineering. In a similar vein, Thomasson defends a conceptual engineering framework built around words as opposed to concepts or meanings and includes such a framework's coherence with empirical linguistics as one of its strengths (2021, 11). In related discussions of Carnapian explication, others have defended both the role of experimental methods in identifying aspects of the explicandum (the precursor concept) (Shepherd and Justus 2015; Schupbach 2017; Koch 2019) and in testing how well the explicatum (the explicated concept) will spread (Pinder 2017). Expanding upon this Carnapian literature, Wakil (2021) argues that conceptual engineers need (in an instrumental sense of need) to use experimental philosophy to check if an engineered concept succeeds at being an improvement over its precursor.

Other papers have actively drawn from empirical methods to better understand conceptual engineering as a method. Fischer (2020) and Machery (2021) both draw from their own past empirical work to argue that some conceptual changes will be more feasible than others. Drawing upon linguistics, Koslow (2022) argues that while meaning and concept change is chaotic when viewing change at the level of particular words, when examining language change at the macro level, predictable patterns emerge that we can study and harness for the purposes of doing conceptual engineering. Looking at the micro level, Landes (forthcoming) examines the propagation of "social distancing" and "coronavirus" early in the COVID-19 pandemic, arguing that the label "social distancing" potentially hindered propagation of the concept SOCIAL DISTANCING.

Third, some papers have actively incorporated empirical methods into conceptual engineering projects. In a series of papers, Machery and collaborators empirically explore the concept of INNATE, identify it as defective, and recommend that it be eliminated (Griffiths and Machery 2008; Griffiths, Machery, and Linquist 2009; Machery 2017; 2021). They thus use the methods recommended by some of the authors above, using experimental findings of inconsistent judgements to motivate conceptual change. Napolitano and Reuter (2021) use experimental

methods to demonstrate that folk find "conspiracy theory" to be negatively evaluative despite many philosophers' insistence that the term is best analysed as neutral and descriptive. Napolitano and Reuter then argue for the introduction of "conspiratorial explanation" as a neutral counterpart to "conspiracy theory" (see also Reuter and Brun (2022) on "truth"). Landes and Reuter (ms) instead look at the propagation stage of conceptual engineering, developing empirical methods to identify whether conceptual change has taken place. Moreover, Landes and Reuter successfully changed, at least at the scale of a few hours to a few days, the default judgements participants made about PLANET to be in line with the IAU's 2006 redefinition (IAU 2006).

This paper belongs to the first category of those that have argued conceptual engineering should involve empirical methods. Closely related to Pinder (2017), this paper argues that adopting empirical methods will improve the odds of successfully propagating a conceptual or linguistic entity. While Pinder focuses on establishing the importance of propagation (what Pinder calls "uptake") on a Carnapian framework, by focusing instead on conceptual engineering, this paper takes the value of propagation for granted. Because conceptual engineering is fundamentally purpose-driven, a conceptual engineering project succeeds or fails based on whether propagation is successful (Wakil 2021, 4–5). This paper instead argues for the value of empirical methods for conceptual engineering through the development of a framework that maps the causal forces affecting propagation. This mapping is itself valuable – it will help conceptual engineers conceptualize, study, and harness the factors affecting propagation. However, mapping these causal forces will also demonstrate why armchair-based conceptual engineering is unnecessarily hobbled compared to conceptual engineering that integrates empirical methods.

This paper proceeds in two parts. Section 2 maps out a framework for thinking about the causal forces at play in conceptual and linguistic propagation. This framework is centred around a single lexical item such as a word or phrase, distinguishing its form, the psychological states associated with the lexical item, and the things in the world associated with the lexical item. Using this framework, Sections 2.1 to 2.4 each look at different types of causation, thereby mapping out

some of the empirical questions that affect propagation. Section 3 explores the consequences of this framework for conceptual engineering. Section 3.1 demonstrates some of the causal forces that need to be considered, using semantic externalist conceptual engineering as an illustration. Section 3.2 argues that the complexity of causal forces means that not only should conceptual engineering be empirical, but that conceptual engineers need to look beyond engineering individual words, meanings, or concepts and instead engineer entire packages of labels, mental states, and extra-mental changes.

Section 2: Labels, Thoughts, and the World

No single paper can do justice to the complex entangled web of causal connections relevant to conceptual engineering. Instead, the goal here is to offer a framework for thinking about that web. Therefore, discussion here will map the key causal joints in conceptual engineering and offer a useful metaphysics – rather than a true metaphysics – for engaging in conceptual engineering, regardless of how conceptual engineering is construed.[1] This metaphysics splits reality into three parts based on its relationship to lexical items, namely, the label the lexical item uses as a vehicle, what psychological states a language user has associated with the lexical item, and what stuff in the world is associated with the lexical item.

The *label* is a lexical item's appearance – the connections of letters and sounds that constitute the vehicle by which a concept or meaning is expressed by a language user. The label is what a lexical item looks and sounds like, which, as will be discussed below, can carry a sizable amount of information for a language user. Saussure has a similar notion of the *signifier*, the vehicle

[1] Andow (2016) serves as the inspiration for this sort of instrumental meta-methodological metaphysics. The metaphysics itself is inspired by the semiotic triangle (Ogden and Richards 1930), although I do not intend to take on any substantive claims from Ogden and Richards.

by which concepts (*the signified*) are communicated. In contrast to Saussure, however, I will be talking about labels as physical manifestations of language whereas Saussure thinks about signifiers as the psychological representations of such manifestations (1916, 66).

Thought will refer to the entities grounded in our own token brain states or mental states. That is, thought is the stuff in our head associated with a lexical item.[2] Since the notion is meant to be theory neutral, how the category of thought is populated will depend on a whole host of theoretical commitments and will differ from reader to reader and conceptual engineering framework to conceptual engineering framework. For example, content internalists will put conceptual content in the category of thought (Machery 2017; Pollock 2021), whereas content externalists (Sawyer 2020a; Scharp 2013) or concept eliminativists (Machery 2009; Cappelen 2018) will not. The same is also true of internalists, externalists, and eliminativists about semantic content. What is not up for philosophical debate, however, is that our mind is structured in ways that are both (fairly) stable and (partially) individualistic. The exact details of how our mind is structured, however, are still being uncovered by neuroscience, cognitive science, philosophy of mind, and other disciplines.

The third category, the *world*, is a catch-all to cover everything external to us, including both physical things and abstracta. Like before, what abstracta belong in the world depends on a whole host of philosophical commitments, and "world" is meant to be a neutral term that can be applied to any set of metaphysical commitments that admits to the existence of a world beyond our minds. There is of course a sense in which thoughts and labels, as the terms are used here, are things in the world. However, the goal here is not to offer a water-tight or even true metaphysical picture of reality but to offer a way of sorting reality that helps conceptual engineers understand, appreciate, and exploit causal relations.

[2] Content externalists in the tradition of Burge (1979) may not like this use of "thought". Unfortunately, I could not find a more neutral term.

Labels, thoughts, and the world, as understood and discussed here, will be talked about as being unified around a singular lexical item. For example, the label "dog" has three letters and is pronounced /dɔg/ or /dɒg/. At the same time, we all have our own thoughts associated with "dog". In my case, "dog" reminds me of my puppy, a dachshund named Fergus. Similarly, the label itself is associated with things in the world. This includes Fergus himself, but also other dogs and our social norms about dogs. Moreover, depending on your theoretical commitments, the worldly entities associated with "dog" may include the natural kind Dog, the semantic value of "dog", the Fregean concept DOG, the causal-historical chain of "dog", and other abstracta.

As this section discusses, labels, thoughts, and the world are causally connected, albeit sometimes indirectly so. While it might seem strange to say the label "dog" is causally connected to the (externalist) semantic value of "dog", they indirectly are. "Dog" is causally related to our cognitive states associated with dogs and such cognitive states are causally connected to the world, including the parts of the world that ground externalist semantic values.[3] There are four directions of causation relevant to conceptual engineering – label to thought, thought to label, thought to world, and world to thought – and each will be looked at in turn.

2.1: Label to Thought

We are used to thinking about the causal connection between language and thought in relation to expression and interpretation. We hear a sentence and in response experience certain mental imagery or form certain beliefs. However, the causal relationship between words and thought is far more complex than just sentential interpretation. In particular, what a label looks

[3] As will become clear below, the labels of lexical items are not indirectly causally related to the grounds of every possible abstracta in theoretical space. This framework discussed in this paper does not have a story to tell for how mathematical or logical terms are related to the grounds of certain platonist understandings of the meaning, referent, or concept of mathematic and logical entities. It does nonetheless offer a way for thinking about how to best engineer the labels of such abstracta to minimize unwanted or false mental representations. Thanks to [Person] for help clearing up this point.

like and sounds like can influence what thoughts we have associated with labels and what mental states we have about the external world. This is particularly clear in capitalized descriptions, such as "the Renaissance", "the Giant's Causeway", and "the Rocky Mountains", that straddle the line between definite descriptions and proper names (Rabern 2015). To see the epistemically available information that can be carried by the form of capitalized descriptions and how such information is different from other types of non-semantic information carried by language, imagine that we are almost entirely ignorant of 20th Century European history and a friend says to us "My grandfather fought in the Spanish Civil War." Even if we have never heard of the Spanish Civil War and have never heard the proper name "the Spanish Civil War", from that sentence we can nonetheless reasonably infer things both about the referent of the proper name and Spanish history more generally.

First, we can use syntactic information to infer things about the name's referent. Indeed, children as young as two years old use syntactic information and the meanings of prepositions to infer semantic information (Waxman et al. 2009; Lidz, White, and Baier 2017; St. Pierre and Johnson 2021). In the example sentence, the preposition "in" tells us the referent of the "Spanish Civil War" is an event the grandfather took part in as opposed to an opponent he fought (compare to: "my grandfather *fought* the Spanish Civil War"). We can also use contextual information to infer things about the name's referent. Given the thing named involves the speaker's grandfather and grandfathers are typically 40 to 80 years older than their grandchildren, we can reasonably infer a general timeframe of the event.

Setting these two inferences aside, there is still non-semantic information that we can reasonably infer from the sentence "my grandfather fought in the Spanish Civil War". Namely, we can infer that Spain had a civil war. This might look like an inference based on the term's semantics or syntax, but capitalized descriptions do not gain their semantics through the composition of the term's component parts (Rabern 2015). If Spain had a second civil war, "the Spanish Civil War" could remain the name of the conflict in the 1930s, but "the Spanish civil war" would be an empty

definite description. Therefore, despite having a syntax that resembles definite descriptions, capitalized descriptions have semantics that work like proper names in that they are rigid and non-compositional (Rabern 2015).

This introduces a bit of puzzle. How can we infer that Spain had a civil war from the term "the Spanish Civil War" if the meaning of "the Spanish Civil War" is non-compositional? The answer is that words often wear their meaning on their sleeves. Even though the relationship between a lexical item's form and its meaning and reference is arbitrary in theory, it is not arbitrary in practice (see Linz and Grote 2003; Dingemanse et al. 2015; Winter et al. 2017). Across our lexicon, there are regularities in the relationship between the labels of lexical items and semantic properties of those lexical items. These regularities can in turn justify language users' semantic and object-level inferences.

In semiotics and linguistics, multiple notions exist to describe such regularities. *Transparency* is the degree to which a language user can infer the meaning of a multimorphemic word from the meaning of its parts (Libben 1998; Libben et al. 2003). The key here is multimorphemic – transparency is a property of compound nouns, terms with affixes, and other multi-part labels. Moreover, transparency is relative to both individuals and meanings. For example, "reader" (meaning a person who reads) is for most English speakers a transparent combination the root *read* and suffix *-er*, whereas "reader" (a job title that some universities use to designate the rank below a professor) is for most *opaque*, i.e., not transparent.

Second, *iconicity* is the resemblance between the form of a linguistic sign and its meaning. Sometimes iconicity is straightforward – the American Sign Language sign for *to drink* looks like someone is taking a drink from a glass with their right hand (see Linz and Grote 2003; Baus, Carreiras, and Emmorey 2013). Another type of iconicity, *ideophones,* are words that resemble sensory experiences related to the meaning (Dingemanse 2012; 2018). These include *onomatopoeias*, such as "whoosh", "whack", and "moo". Ideophones also include iconic *phonesthemes*, or systematic associations between certain sounds and meanings, such as the association of "cr-" in English with

abrupt sounds – e.g., "crack", "crunch", "creak", and "crash" (Mompean, Fregier, and Valenzuela 2020).

Third, *systematicity* are statistical regularities between the form of a word or phrase and its meaning. In contrast to iconicity and transparency, systematicity makes no reference to whether language users perceive the connection or whether there is some principled connection between form and meaning.[4] For example, on average English nouns have more syllables than verbs, are more likely to start with bilabials (e.g., b, m, p) than verbs, and have more vowels as a percentage of total word sounds than verbs (Cassidy and Kelly 1991; 2001; Monaghan, Christiansen, and Chater 2007).

Due to regularities between the form and meaning of lexical items, our ability to infer semantic and object-level information from labels extends beyond capitalized descriptions like "the Spanish Civil War." Other categories of lexical items similarly allow for inferences about their meaning or referent based on the label. To demonstrate, consider five object-level inferences that would be supported the information conveyed by five different labels if the listener was unfamiliar with those labels:

- **Proper name**: "I was visiting *Rožďalovice*."
    - *Conveyed information:* The speaker was in a historically Slavic part of Europe.
- **Multimorphemic adjective**: "The house was *decagonal*"
    - *Conveyed information*: The house had 10 sides.
- **Compound noun**: "I saw a *bluebird*."
    - *Conveyed information*: The speaker saw a bird that was blue.
- **Slightly opaque noun phrase**: "He cooked with *clarified butter*."[5]

[4] Phonesthemes are sometimes discussed as a form of systematicity (e.g., Monaghan et al. 2014).

[5] By *slightly opaque*, I mean that the meaning of the adjective-noun phrase is not immediately clear from its parts. Here is a quick thought experiment to demonstrate that clarified butter is not merely butter that is clarified: Imagine chefs discovered that if you add a small drop of baking soda, the butter becomes perfectly clear but is otherwise unchanged. That butter would be clarified but it would not be clarified butter.

- *Conveyed information*: The speaker cooked with butter that had been cleared up in some way.

- **Monomorphemic word**: "There was a *whack*."
  - *Conveyed information*: There was a sound that ended sharply.

In order to be able to talk about all of the ways a label can carry information, call the property on display in the examples above a label's *informativeness*. A lexical item is informative when its form carries information that a language user is apt to use to infer useful or true information, such as about the lexical item's meaning or referent.[6] Conversely, the information language users are apt to infer from a label can also be harmful or false. For example, here are five such *misleading* labels that may cause false inferences about the world:

- **Proper name**: "I was visiting *Trumpington*."
  - *Misleading information*: The location is not related to the former US president.
- **Multimorphemic adjective**: "The gas is *inflammable*."
  - *Misleading information*: The gas is flammable, not nonflammable (see Koslow 2022, 16).
- **Compound noun**: "I ate a *blackberry*."
  - *Misleading information*: Blackberries are usually purple and are not technically berries.
- **Slightly opaque noun phrase**: "Please *social distance* yourselves."
  - *Misleading information*: The practice primarily involves physical or epidemiological distancing with other people, not distancing that is *social* (see Landes forthcoming).
- **Monomorphemic word**: "Pigs *oink."*

[6] Good, useful, or trueness can come apart. Thanks to [Person] for this point. Note also that this is a newer formulation than in [Author]. This formulation is more framework-neutral and uses "informative" as a more informative label for the notion.

- *Misleading information*: The grumbling noises pigs make sound nothing like "oink".

In the cases of informative or misleading language, the label itself enables inferences, which then impacts the mental representations we have, not just about the language, but about other things as well. In this way, labels can causally influence our beliefs, mental imagery, and other mental states.

Informativeness and misleadingness are therefore a three-way relationship between labels, thoughts, and the world since they describe both the extent to which the appearance of a lexical item is prone to cause certain mental states and the way in which those mental states line up with the world. An informative label will tend to cause mental states that are true, fitting, apt, etc., whereas a misleading label will tend to cause mental states that are false, unfitting, defective, etc., Importantly, however, while the notions' relationship between label and thought is causal, the notions' relationship between thought and the world is not. Moreover, the relationship between thought and world depends on the framework. Functional or pragmatic frameworks (Nado 2021; Isaac 2021; Thomasson 2020) may frame informativeness and misleadingness in terms of usefulness or efficacy, whereas realist frameworks (Scharp 2013; Cappelen 2018; Sawyer 2020b) may frame them in terms of truth or defectiveness.

Label-to-thought causal forces ultimately affect propagation by influencing what inferences people make about novel lexical items (see Majnemer and Meibauer 2023; Mandelbaum and Young forthcoming). Misleading labels will slow down propagation by, for example, confusing people about a novel term's meaning (see Landes forthcoming). Informative labels can combine known ideas in novel and illuminating ways – as was the case with "sexual harassment" (Maitra 2018) and "mansplaining" (Koslow 2022) – in order to help spread the associated concepts and speed up propagation. The epistemic problem faced by conceptual engineers wanting to design an informative label is that what information is drawn from a label will differ from person to person. What is informative to one person might be misleading to most other people. This is especially a risk given that conceptual engineers generally have, at present moment, much higher levels of

education than the general public, so what is informative or misleading to conceptual engineers may not be what is informative or misleading to normal people. Consequently, conceptual engineers will be more likely to pick a label people find informative if they check with the community in which they aim to propagate a cognitive or linguistic item. This requires empirical work. This empirical work could be quantitative such as studies on how labels affect semantic and non-semantic judgements (see Majnemer and Meibauer 2023; Mandelbaum and Young forthcoming), or it could be qualitative work such as focus groups and interviews. Either way, empirical work examining label-to-thought causation puts conceptual engineers on better footing to pick labels that aid in propagation**.**

2.2: Thought to Label

Informativeness and misleadingness are part of a larger feedback loop that exists between a term's form and our mental representations associated with the term. While, as discussed in the previous section, the information carried by a term can cause certain beliefs and associations, our beliefs, associations, and other parts of our psychology influence what labels appear in the lexicon. Notice I am not merely saying that our psychology influences what vocabulary we pick up from other people. While this is true, the point is subtler: our psychology influences what form individual entries in our vocabulary – and the wider lexicon – take.

To discuss this, start with two properties of the relationship between a term's form and our mental representations, a term's fit and stickiness. *Fit* is the felt aptness of a label's form, whereas *stickiness* is the memorability of a label-meaning pair. A term's fit and stickiness are related to informativeness, since an informative term will usually feel like it is apt and be memorable. "The Spanish Civil War" is both a fitting and memorable name for the civil war fought in Spain. Nonetheless, fit and stickiness are both causally and conceptually distinct from informativeness.

Informativeness without fit occurs when a label conveys information, but the way in which it conveys information seems inappropriate, such as calling the Judeo-Christian God "Super-duper-extra-ultimately-strong-smart-and-good-thing". Informativeness without stickiness occurs when a label carries information but is difficult to remember, such as calling the Spanish Civil War, "the 1936 to 1939 War Between Spanish Republicans, Syndicalists, and Communists vs Spanish Nationalists, Falangists, and Monarchists". If someone earnestly tried to introduce it as name for the war, people would either forget it or forgo it for something much easier and more memorable – like the "the Spanish Civil War". At the same time, terms do not need to be informative to fit or to be sticky. Misleading terms are often socially problematic exactly because they are sticky, and companies often have meaningless but memorable names.

Both fit and stickiness are ways our psychology puts causal pressure on what labels are in the lexicon. Fit puts causal pressure on what labels are introduced in the first place. We intuitively understand that some words are better labels for a thing than others and will often hunt around for the right label for a new idea. In contrast, stickiness puts causal pressure on what labels remain in the lexicon once they are introduced (see Monaghan, Christiansen, and Chater 2007). Terms that are not sticky can be introduced, but they will not be readily retained or used by language users. Consequently, they may not enter the lexicon or may die out between generations of language users.

Fit and stickiness are merely illustrative, however, of the larger network of causal pressures our beliefs, desires, associations, etc., put on a lexicon. In drawing upon empirical pragmatics and diachronic semantics to aid in the understanding of conceptual engineering, Koslow (2022) catalogues a number of phenomena that are, on the present framework, thought-to-label causal forces. These include homonymy avoidance, convenience, and efficiency.

Homonymy avoidance is people's tendency to avoid and find ways around ambiguities that cause confusion (Koslow 2022, 94–96). When a label is problematically homonymous or polysemous – e.g., there are often not salient contextual or syntactic clues to disambiguate between

senses – speakers will avoid the label to avoid confusion. Like fit, this will prevent the introduction of certain word-meaning pairs, and like stickiness, this will put pressure on people to avoid the label.

Convenience and *efficiency* are how easy a speech act containing the label is to do and the ratio of energy to communicated information, respectively (Koslow 2022, 94). As demonstrated above in relationship to fit and stickiness, informative labels are often very inconvenient and not efficient at referring to the referent. Convenience and fit come apart when a label seems fitting but is hard to pronounce or spell, and efficiency and fit come apart when a label is short but does not seem fitting (e.g., calling the Spanish civil war "Ba"). The convenience and efficiency of labels will put pressure on both what is introduced as well as what is propagated, as the people who coin language will often strive for snappy language and language users will, all things being equal, prefer easier ways of communicating.

Fundamentally, thought-to-label causal forces affect a word-meaning pair's sticking power. Concepts or meanings that are paired with labels that feel right or labels that have some other sort of aesthetic draw will spread more readily than concepts or meanings paired with awkward labels. The armchair is inferior to empirical methods in its ability to determine how a target community views the aesthetic properties of a label-meaning or label-concept pair. We may be very different demographically from our target community, we may be too close to our own creation to realize what it looks like for the first time, or we may just have idiosyncratic judgements. The most accurate way to know how a target community will react to something is to check with the target community. Empirical methods, whether through quantitative methods like psychological studies or through qualitative methods like focus groups, will thus increase our odds that a designed conceptual or linguistic device has a label that will help it stick.

2.3: Thought to World

Now that the two causal directions between labels and thought have been mapped, it is time to turn to the two causal directions between thought and world. In its most familiar forms, the causal relationship from thought to world is mundane. This morning I was hungry, so I put cereal and milk together in a bowl. My thoughts influenced the world – it reduced the amount of cereal in my cupboard – and the world influenced my thoughts – the availability of cereal led to my belief I did not have to run to the grocery store.

Conceptual engineers are interested in far larger and more interesting changes in the world than how my beliefs and desires influence the contents of my cupboard. Many conceptual engineers are interested in changing social norms as means to some end, such as changing meaning (Nimtz 2021; Thomasson 2021) or ending oppressive power structures (Haslanger 2000; Manne 2018). Because social norms and institutions are grounded in the beliefs, habits, and expectations of individuals, what social norms and institutions exist depend in large part on what we think they are and should be. This includes formal power structures such as law. The legal institution of marriage expanded to include same-sex couples because enough people situated in the right place in society believed the institution should change and were motivated to change it. This interplay between institutions, norms, scientific facts, etc., on one side and beliefs, desires, motivations etc., on the other are, on the present framework, the most relevant causal relationship between thought and world for conceptual engineers.

In order to clarify discussion, I will follow Isaac, Koch, and Nefdt (2022) in distinguishing between the purposes, goals, and targets of conceptual engineering. The *purpose* of conceptual engineering is the final aim of the conceptual engineering project. These are generally things in the world, such as power structures, behaviors, truth-values, and kindhood, whose change would achieve some desired good. What makes conceptual engineering unique compared to other activist activities is that conceptual engineering aims to achieve their purpose through some linguistic or conceptual change. These linguistic or conceptual changes are the *goal* of conceptual engineering, and the linguistic or conceptual entities they hope to change are the *target*. It is by paying attention

to the interplay of a project's target, purpose, and goal that conceptual engineers need to pay attention to the causal forces between thoughts and the world. Thought-to-world causation in particular affects whether the necessary changes (that is, the project's goal) are propagated in the right way to fulfill the project's purpose. Discussion here is complicated by the variety of conceptual engineering projects currently on offer. Some frameworks have goals and targets at the level of thought and some frameworks have goals and targets at the level of the world. These will be discussed in turn.

Frameworks whose target is at the level of thought include frameworks that target individually-grounded understandings of speaker meaning, conceptual content, or meaning (e.g., Plunkett and Sundell 2013; Machery 2017; Pinder 2021). Even though the target of conceptual engineering for these frameworks is at the level of thought, thought-to-world causation still matters to the success of a conceptual engineering project. Conceptual engineering is purpose-driven, and on such accounts the purpose is achieved by spreading speaker meanings, conceptual content, meaning, etc. This only happens if the targeted mental representations (speaker meaning, conceptual content, meaning, etc.) have the desired effect. These cannot backfire or remain epiphenomenal. Therefore, conceptual engineers with targets at the level of thought need to study thought-to-world causation to understand whether people will react to the propagated target in a way that actually aids the project's purpose.

Conceptual engineers with targets at the level of the world – such as norms, concepts, or legal institutions, (Haslanger 2000; Scharp 2013; Cappelen 2018) – similarly need to be aware of how their attempts at revision will be interpreted. While changes in thought are not the target or goal of such projects, changes in thought in response to intentional propagation will affect the grounds of their target. While concepts (understood as mind-independent entities), semantic values (similarly understood), or linguistic practices are not thought-level entities, their grounds often are. Looking at semantic conceptual engineering projects, the grounds of the target might include collective linguistic beliefs, collective linguistic use (Sawyer 2020a), or collective intentions to take

part in some reference chain (Sterken 2020; Riggs 2019, 11). These collective linguistic practices are at the level of world, but the individual instantiations of the grounds –linguistic beliefs, linguistic habits of use, linguistic intentions – are at the level of thought. This is also true with other worldly changes, such as norms and institutions. Norms and institutions exist because they are instantiated, often in extraordinarily complex ways, in the minds of the people taking part in them. Therefore, changes at the world level will most likely require changes at the level of thought, and world-level engineers cannot escape the need for research into how thought and things in the world interact.

Stepping back, there are two broad families of thought-to-world empirical questions facing conceptual engineers. First, how are the entities conceptual engineers hope to change structured? Here we want to understand how individuals and their mental states contribute to the structures that conceptual engineers hope to alter as part of their project's purpose.[7] Second, how will specific changes in mental states change the behaviour, grounds, or norms underlying such entities? Conceptual engineers' ability to answer both sorts of questions will affect the efficacy of their propagation efforts, as it will determine whether the right messages are spread to the right people to fulfil the project's purpose.

From the armchair we can guess what the right messages are and who the right people are. However, our understanding would be limited to our own experience and our personal understanding of the often-opaque factors that, in first question, constitute the ways we contribute to larger social structures, and in the second question, have changed the ways we contribute to larger social structures. Epistemically, conceptual engineers would be much better off drawing from anthropology and sociology, where a variety of empirical tools are employed to understand social structures. For example, methods like social network analysis would allow conceptual

[7] The first question may not be empirical on all frameworks and may instead be answered on some frameworks by metaphysical claims about the grounds of abstracta. The second question will always be empirical, however.

engineers to understand how ideas flow throughout a population in order to better deploy what resources have been allotted to propagation.

2.4: World to Thought

The most relevant world-to-thought causal relations for conceptual engineering are those that affect the motivation of the engineer and the target audience of an engineering project. We have the conceptual or linguistics problems we have because there is some friction between our representational devices, desires, and the world. Accordingly, the perceived state of the world will influence what projects are seen as worth doing and worth buying into.

Looking first at the engineer, contingent facts about the world will affect the conceptual engineer's motivation because contingent facts will affect the desiderata and/or salience of desiderata for conceptual engineering. To illustrate, consider the 2006 revision of PLANET that excluded Pluto from planethood (IAU 2006). In 2005, Eris, an object with a very similar size and orbit to Pluto, was discovered by astronomers. This confirmed the growing suspicion that there might be dozens or even hundreds of Pluto-sized objects beyond Neptune, thereby spurring on the redefinition of PLANET (Brown 2006; Chang 2022). Even if we grant the IAU's definition correctly captures the joints of planethood, the revision was driven by the discovery of Pluto-like objects beyond Neptune in the Kuiper Belt. If the solar system had formed differently and Pluto was the only sizable object beyond Neptune, then the necessary motivation to engineer PLANET would not have arisen at the time it did in the community it did with the urgency it did. Friction between the world and astronomers' thought caused the concept to be revised.

The importance of world-to-thought causal forces goes beyond determining the motivation of the engineers.[8] It also affects the motivation of the subjects of propagation efforts. For an engineered linguistic or cognitive entity to spread, people need to buy into the change and spread it (whether this buy-in is conscious or unconscious). Many of us know that "technically" tomatoes are fruit and that "technically" eggplants are berries. Nonetheless we still use the words in a non-technical way because the costs of changing how we speak and think do not outweigh the benefits of lining up with experts' technical usage (see Abbott 1997). After all, very few of us want eggplants in our berry parfaits. If, in contrast, using "berry" and "fruit" in the folk way caused problems or was seen as uncouth, many of us would quickly change our tune.

Conceptual engineers must therefore think about how the world affects people's motivations. Like the other causal interactions discussed here, if the wrong forces are in place – if facts external to people just do not motivate a project – propagation will face significant, and possibly fatal, headwinds. In contrast, changes that are seen as trendy, necessary, or useful will instead take significantly less work to propagate because people will be motivated to adopt and spread the changes. Some of this will be outside of the conceptual engineers' control – we cannot control what astronomers find in the icy reaches of the Kuiper belt. Nonetheless, many things are in our control; marketing campaigns are a world-to-thought mechanism that use world-based stimuli to change perceptions of products. Therefore, conceptual engineering will be more likely to achieve their projects' purpose if they know what relevant worldly conditions currently exist, how such conditions are perceived, and what world-based messaging will motivate people to take up a designed cognitive or linguistic entity. If a conceptual engineer does not examine the relevant empirical data – whether via sociology, psychology, market research, astronomy, or other empirical methods – the conceptual engineer has a much higher risk of all the work they put into designing

[8] World-to-thought causation goes beyond desiderata and motivation. Abstracta relevant to conceptual engineering play an indirect part in world-to-thought causation. While spooky entities such as Fregean concepts or other Platonic entities do not have a causal effect on our mental states – otherwise they would not be spooky – their grounds often do. These grounds can then play a role in propagation.

the perfect conceptual or linguistic device turning out to be a Quixotic project that no one else ever cares about.

Section 3: From Label to World and Back Again

In mapping some of the causal forces that factor into successful propagation, I have argued that the causal complexity of propagation means empirical research will make successful propagation far more likely. This section explores what this means in practice. Section 3.1 uses semantic externalist conceptual engineering to illustrate that even conceptual engineering projects with abstract targets must muddy their hands with the causal forces described above if they hope to succeed at propagating. This is because while the targets of conceptual engineering are often abstract, the grounds of such targets are often not. Section 3.2 argues the lessons here require a holistic approach to conceptual engineering. Conceptual engineers should not focus on a single device to engineer, and instead should focus on engineering broader clusters of causally interwoven labels, thoughts, and worldly changes.

3.1: An Illustration

Imagine we as conceptual engineers understand the target of conceptual engineering to be the externalist semantic values of words because we want to change the truth conditions of sentences for epistemic reasons (see Isaac, Koch, and Nefdt 2022, 15). Much has been written about what sort of metasemantic moves and changes need to happen in order to revise meaning on an externalist picture (Cappelen 2018; Pollock 2019; Ball 2020a; Sterken 2020; Koch 2021). What is true across the board, however, is that for a meaning to exist on externalist accounts, the

word needs to first be baptized or anchored in some way and then the word-meaning pair needs to be propagated throughout the linguistic community (see Sterken (2020) for discussion). There is significant difference in what this propagation requires or consists in, and depending on the account it might be that experts need to be in place to be deferred to, causal chains need to be broken, linguistic norms have to be revised, or patterns of use have to change.

In each of these cases, changing semantic meaning requires changing thought in some way. Collective linguistic intentions, norms, patterns of use, and other grounds of meaning ultimately emerge because of the beliefs, desires, and other mental states of individuals. Causal chains are perpetuated because people know about the causal chains and are motivated to use them. Someone who has never encountered the causal chain associated with "Caesar" cannot speak with the intention to take part in it, nor will they want to if they think Caesar was just some guy from Rome. A similar story is true for linguistic norms. Someone cannot take part in a norm they have not encountered it, and if they dislike the norm, they will avoid participating in it (see Thomasson 2021; Nimtz 2021).

We can, as conceptual engineers, coin, baptize, or anchor as many new meanings for words as we want. However, it is a necessary (but not sufficient) condition for meaning change that people to pick up on our lexical innovation. We need to think about how and why people might want to pick up the new meaning or resist the change. Accordingly, successful propagation requires harnessing label-to-thought causation to make sure the label is snappy and not misleading and world-to-thought causation to make sure people are motivated to change their ways. On a Kripkean causal-historical story (Kripke 1981; Soames 2003), this will involve figuring out what will make people adopt a causal chain going back to the baptism. On a view inspired by Burge where meaning is determined by experts (Burge 1979; Ball 2020b), changing meaning will involve identifying experts, then identifying and targeting what will change their behavior as a group. Both will be aided by understanding various empirical considerations in sociology, psychology, cognitive science, historical linguistics, and marketing.

In fact, the web of causal influences on propagation is far more complicated than presented here. This paper has only described some of the causal forces that would affect the success of propagation and some of the ways in which the relevant causal forces interact. Nonetheless discussion here is meant to illustrate how engineering even something as abstract as externalist semantic value or corollaries like utterances' truth values cannot solely focus on engineering abstracta. In order to successfully propagate a meaning, we need to pay attention to the conditions that succeed in creating the grounds of a new or revised meaning (thought-to-world causation). This requires creating the conditions that will lead to the right sort of effects at the thought level, which requires, among many other things, thinking about label-to-thought and world-to-thought causation. Thus, conceptual engineering projects that have goals as abstract as truth value still profit from paying attention to things as mundane and empirical as human psychology.

3.2: Notion Engineering

The goal of conceptual engineering is to get certain changes in the mind or world to stick to achieve some purpose. And regardless of what the project is trying to make stick, focus will need to be paid attention to the way labels, thoughts, and the world causally interact. This section argues that rather than avoiding this complexity, conceptual engineers should embrace it. I propose conceptual engineers talk about the combined package of label, thought, and worldly change as a *notion* and that conceptual engineers should design notions instead of standalone concepts, meanings, or words. This is because focusing just on labels, thoughts, or the world risks blinding conceptual engineers to the complexity of propagation.

Focusing merely on a label risks not accomplishing much of value. Labels are unpredictable and finnicky; sometimes the inertia of language keeps problematic labels in circulation for centuries, while other times labels are dropped due to changes in slang. At the same time, the way one person interprets a label will not be universal. While changing a label may prove useful, it is

only useful because of effects at the thought and world level. Dropping a hateful or misleading term may improve things, but it improves things because it leads to changes in inference patterns, behavior, or abstracta. Therefore, labels can and should be paid attention to, but only in light of the rest of a notion.

Focusing solely on changing thought risks either creating an unusable belief and/or concept (understood as a token mental state or disposition) or setting oneself up to let an innovation molder in the pages of academic journals. To pick an extreme example of creating an unusable belief or concept by ignoring the world, imagine I engineered a concept whose content is "every planet in the solar system closer to the Sun than Mercury". There is certainly nothing stopping me from engineering such a concept. Such a planet is conceivable, metaphysically possible, and perhaps even epistemically possible. However, since there is no such planet, there is no use for such a concept, except in certain fictional or counterfactual uses, and thus there is little reason for people to adopt it (see Isaac 2021). At the same time, if I came up with an extremely useful new concept but gave it a terrible name, people would struggle to take the concept up in the same way philosophy students can struggle to remember the distinction between "de dicto" and "de re" statements or "pro tanto" vs "pro toto" reasons.

Focusing too much on thoughts or labels also risks unexpected results caused by the world. As Queloz and Bieber (2021) discuss, the preexisting examples we have of engineered or otherwise new concepts are full of unintended consequences, such as the repurposing of Nietzsche's concepts "Übermensch" and "will to power" by antisemitic nationalists. Less drastically, recently the adjective *woke* shifted in just a few years from a term of praise among members of the political left to a derogatory term used by people on the right. Without a deep understanding of the environment in which we are introducing a label and/or mental state, we have little idea of how the term will be received outside of our own head or our own academic bubble. We also would not know if the concept introduces some sort of unintended injustice (Shields 2021). How feasible those things are to predict is an open question (and likely best answered by sociologists and

linguists), but any engineering project hoping for a certain behavior, belief, mental representation, or label to be spread among some population will be better able to propagate what they want to propagate the more know about the external environment in which their concept or label is being propagated.

That said, focusing too much on worldly changes, including altering abstracta or truth values, risks losing sight of whether a change is doable and how to go about that change. It may be that it would be world-changing to have a word-meaning pair with certain truth conditions, but good consequences are not enough to propagate a word-meaning pair. There are the separate and necessary questions of whether people can represent a token instantiation or ground of the necessary abstracta, whether people would want to have such a token instantiation or ground, and what would get people to form a corresponding token instantiation or ground. All three are questions of psychology, and as discussed in Sections 2.1 and 2.2, the answer to the third question will involve the role of the label as a vehicle of the abstracta.

Importantly, designing a notion should be an empirical project. Because the relationships are causal, they are best examined with the empirical tools of psychology, experimental linguistics, history, sociology, etc., and not the armchair. Returning to labels as an illustration, while a well-designed label or name can carry information and help convey ideas, we need to understand people's psychology to understand with any clarity whether a given label will succeed at doing so. By their very nature, informativeness, transparency, and other label-to-though causal forces depend on the mental representations of the individual encountering the label. A label that seems to carry information to one group of conceptual engineers may not read the same way to another group of conceptual engineers, let alone a normal person. Knowing if a label is a good label requires studying the people who we want to use the label, else we are just providing our own idiosyncratic guesses. Because conceptual engineering needs to consider the complex web of causal forces that labels are merely one aspect of, and complex causal interactions are best studied empirically, conceptual engineering should be an empirical project.

Section 4: Conclusion

Conceptual engineering, regardless of the form it takes, is a complicated endeavor. Focusing too much the target or goal of conceptual engineering risks making conceptual engineering look straightforward – even if the means to do so are not necessarily straightforward to identify or wield. In reality, conceptual engineers need to focus on a wide range of considerations and factors when trying to propagate what they want to propagate. To offer a way to conceptualize and understand the different interrelated causal factors at play in conceptual engineering, this paper suggested the (meta)framework of *labels, thoughts, and the world*. The label is the vehicle by which meaning, concepts, and ideas are presented. The thoughts are the token entities grounded in an individual's psychological states. The world is all other stuff external to us. Because words are causally related to thoughts and thoughts are causally related to the world, any intervention at one level will affect and be affected by factors at the other levels. While attempts to understand these causal forces could be made from the armchair, doing so would unnecessarily hinder attempts at propagation. Instead, approaching these forces empirically would give conceptual engineers the best possible chance at successfully achieving their project's purpose.

Compliance with Ethical Standards

The author has no relevant financial or non-financial interests to disclose.

Works Cited

Abbott, Barbara. 1997. 'A Note on the Nature of "Water"'. *Mind* 106 (422): 311–19. https://www.jstor.org/stable/2254507.

Andow, James. 2016. 'Thin, Fine and with Sensitivity: A Metamethodology of Intuitions'. *Review of Philosophy and Psychology* 7 (1): 105–25. https://doi.org/10.1007/s13164-015-0247-2.

———. 2020. 'Fully Experimental Conceptual Engineering'. *Inquiry*, December, 1–27. https://doi.org/10.1080/0020174X.2020.1850339.

Ball, Derek. 2020a. 'Metasemantic Ethics'. *Ratio* 33 (4): 206–19.

———. 2020b. 'Relativism, Metasemantics, and the Future'. *Inquiry: An Interdisciplinary Journal of Philosophy*, 41.

Baus, Cristina, Manuel Carreiras, and Karen Emmorey. 2013. 'When Does Iconicity in Sign Language Matter?' *Language and Cognitive Processes* 28 (3): 261–71. https://doi.org/10.1080/01690965.2011.620374.

Belleri, Delia. 2021. 'On Pluralism and Conceptual Engineering: Introduction and Overview'. *Inquiry* 0 (0): 1–19. https://doi.org/10.1080/0020174X.2021.1983457.

Brown, Mike. 2006. 'Eight Planets'. 2006. https://web.gps.caltech.edu/~mbrown/eightplanets/.

Burge, Tyler. 1979. 'Individualism and the Mental'. *Midwest Studies In Philosophy* 4 (1): 73–121. https://doi.org/10.1111/j.1475-4975.1979.tb00374.x.

Cappelen, Herman. 2018. *Fixing Language*. Oxford University Press. https://doi.org/10.1093/oso/9780198814719.001.0001.

Cassidy, Kimberly Wright, and Michael H. Kelly. 1991. 'Phonological Information for Grammatical Category Assignments'. *Journal of Memory and Language* 30 (3): 348–69. https://doi.org/10.1016/0749-596X(91)90041-H.

———. 2001. 'Children's Use of Phonology to Infer Grammatical Class in Vocabulary Learning'. *Psychonomic Bulletin & Review* 8 (3): 519–23.

Chalmers, David J. 2020. 'What Is Conceptual Engineering and What Should It Be?' *Inquiry*, 1–18.

Chang, Kenneth. 2022. 'Is Pluto a Planet? What's a Planet, Anyway?' *The New York Times*, 18 January 2022, sec. Science. https://www.nytimes.com/interactive/2022/science/is-pluto-a-planet.html.

Dingemanse, Mark. 2012. 'Advances in the Cross-linguistic Study of Ideophones'. *Language and Linguistics Compass* 6 (10): 654–72.

———. 2018. 'Redrawing the Margins of Language: Lessons from Research on Ideophones'. *Glossa: A Journal of General Linguistics* 3 (1).

Dingemanse, Mark, Damián E Blasi, Gary Lupyan, Morten H Christiansen, and Padraic Monaghan. 2015. 'Arbitrariness, Iconicity, and Systematicity in Language'. *Trends in Cognitive Sciences* 19 (10): 603–15.
Fischer, Eugen. 2020. 'Conceptual Control: On the Feasibility of Conceptual Engineering'. *Inquiry*, 1–29.
Griffiths, Paul, E., and Edouard Machery. 2008. 'Innateness, Canalization, and "Biologicizing the Mind"'. *Philosophical Psychology* 21 (3): 397–414.
Griffiths, Paul, E., Edouard Machery, and Stefan Linquist. 2009. 'The Vernacular Concept of Innateness'. *Mind & Language* 24 (5): 605–30.
Haslanger, Sally. 2000. 'Gender and Race: (What) Are They? (What) Do We Want Them To Be?' *Noûs* 34 (1): 31–55. https://doi.org/10.1111/0029-4624.00201.
Horvath, Joachim, and Steffen Koch. 2021. 'Experimental Philosophy and the Method of Cases'. *Philosophy Compass* 16 (1): e12716. https://doi.org/10.1111/phc3.12716.
IAU. 2006. 'Definition of a Planet in the Solar System'. Resolution B5. Paris: International Astronomical Union. https://www.iau.org/static/resolutions/Resolution_GA26-5-6.pdf.
Isaac, Manuel Gustavo. 2021. 'Post-Truth Conceptual Engineering'. *Inquiry* 0 (0): 1–16. https://doi.org/10.1080/0020174X.2021.1887758.
Isaac, Manuel Gustavo, Steffen Koch, and Ryan Nefdt. 2022. 'Conceptual Engineering: A Road Map to Practice'. *Philosophy Compass*.
Koch, Steffen. 2019. 'Carnapian Explications, Experimental Philosophy, and Fruitful Concepts'. *Inquiry: An Interdisciplinary Journal of Philosophy* 62 (6): 700–717.
———. 2021. 'The Externalist Challenge to Conceptual Engineering'. *Synthese* 198: 327–48. https://doi.org/10.1007/s11229-018-02007-6.
Koslow, Allison. 2022. 'Meaning Change and Changing Meaning'. *Synthese* 200 (2): 1–26.
Kripke, Saul. 1981. *Naming and Necessity*. Oxford, UK ; Cambridge, USA: Blackwell Publishers.
Landes, Ethan. forthcoming. 'How Language Teaches and Misleads: "Coronavirus" and "Social Distancing" as Case Studies'. In *New Perspectives on Conceptual Engineering*, edited by Manuel Gustavo Isaac, Steffen Koch, and Kevin Scharp.
Landes, Ethan, and Kevin Reuter. ms. 'Conceptual Revision in Action'.
Libben, Gary. 1998. 'Semantic Transparency in the Processing of Compounds: Consequences for Representation, Processing, and Impairment'. *Brain and Language* 61 (1): 30–44. https://doi.org/10.1006/brln.1997.1876.
Libben, Gary, Martha Gibson, Yeo Bom Yoon, and Dominiek Sandra. 2003. 'Compound Fracture: The Role of Semantic Transparency and Morphological Headedness'. *Brain and Language*, Brain and Language Special Issue, 84 (1): 50–64. https://doi.org/10.1016/S0093-934X(02)00520-5.
Lidz, Jeffrey, Aaron Steven White, and Rebecca Baier. 2017. 'The Role of Incremental Parsing in Syntactically Conditioned Word Learning'. *Cognitive Psychology*. https://doi.org/10.1016/j.cogpsych.2017.06.002.
Linz, Klaudia Grote Erika, and E Grote. 2003. 'The Influence of Sign Language Iconicity on Semantic Conceptualization'. *From Sign to Signing* 3: 23.
Machery, Edouard. 2009. *Doing without Concepts*. Oxford University Press.
———. 2017. *Philosophy within Its Proper Bounds*. First. Oxford, United Kingdom: Oxford University Press.
———. 2021. 'A New Challenge to Conceptual Engineering'. *Inquiry* 0 (0): 1–24. https://doi.org/10.1080/0020174X.2021.1967190.
Maitra, Ishani. 2018. 'New Words for Old Wrongs'. *Episteme* 15 (3): 345–62. https://doi.org/10.1017/epi.2018.27.

Majnemer, Jacklyn, and Gustav Meibauer. 2023. 'Names from Nowhere? Fictitious Country Names in Survey Vignettes Affect Experimental Results'. *International Studies Quarterly* 67 (1): sqac081. https://doi.org/10.1093/isq/sqac081.
Mandelbaum, Eric, and Steven Young. forthcoming. 'The Sound of Slurs: Bad Sounds for Bad Words'. In *Oxford Studies in Experimental Philosophy*.
Manne, Kate. 2018. *Down Girl: The Logic of Misogyny*. New York, NY: Oxford University Press.
Mompean, Jose A., Amandine Fregier, and Javier Valenzuela. 2020. 'Iconicity and Systematicity in Phonaesthemes: A Cross-Linguistic Study'. *Cognitive Linguistics* 31 (3): 515–48. https://doi.org/10.1515/cog-2018-0079.
Monaghan, Padraic, Morten H. Christiansen, and Nick Chater. 2007. 'The Phonological-Distributional Coherence Hypothesis: Cross-Linguistic Evidence in Language Acquisition'. *Cognitive Psychology* 55 (4): 259–305. https://doi.org/10.1016/j.cogpsych.2006.12.001.
Monaghan, Padraic, Richard C. Shillcock, Morten H. Christiansen, and Simon Kirby. 2014. 'How Arbitrary Is Language?' *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1651): 20130299. https://doi.org/10.1098/rstb.2013.0299.
Nado, Jennifer. 2019. 'Conceptual Engineering via Experimental Philosophy'. *Inquiry* 0 (0): 1–21. https://doi.org/10.1080/0020174X.2019.1667870.
———. 2021. 'Conceptual Engineering, Truth, and Efficacy'. *Synthese* 198: 1507–27.
Napolitano, M. Giulia, and Kevin Reuter. 2021. 'What Is a Conspiracy Theory?' *Erkenntnis*, September. https://doi.org/10.1007/s10670-021-00441-6.
Nimtz, Christian. 2021. 'Engineering Concepts by Engineering Social Norms: Solving the Implementation Challenge'. *Inquiry* 0 (0): 1–28. https://doi.org/10.1080/0020174X.2021.1956368.
Ogden, C. K., and I. A. Richards. 1930. *The Meaning Of Meaning*. http://archive.org/details/in.ernet.dli.2015.221615.
Pinder, Mark. 2017. 'Does Experimental Philosophy Have a Role to Play in Carnapian Explication?' *Ratio* 30 (4): 443–61. https://doi.org/10.1111/rati.12164.
———. 2021. 'Conceptual Engineering, Metasemantic Externalism and Speaker-Meaning'. *Mind*.
Plunkett, David, and Timothy Sundell. 2013. 'Disagreement and the Semantics of Normative and Evaluative Terms'. *Philosophers' Imprint* 13.
Pohlhaus, Gaile. 2015. 'Different Voices, Perfect Storms, and Asking Grandma What She Thinks: Situating Experimental Philosophy in Relation to Feminist Philosophy'. *Feminist Philosophy Quarterly* 1 (1). https://doi.org/10.5206/fpq/2015.1.3.
Pollock, Joey. 2019. 'Conceptual Engineering and Semantic Deference'. *Studia Philosophica Estonica*, December, 81–98. https://ojs.utlib.ee/index.php/spe/article/view/16289.
———. 2021. 'Content Internalism and Conceptual Engineering'. *Synthese* 198 (12): 11587–605. https://doi.org/10.1007/s11229-020-02815-9.
Queloz, Matthieu, and Friedemann Bieber. 2021. 'Conceptual Engineering and the Politics of Implementation'. *Pacific Philosophical Quarterly*, no. n/a. https://doi.org/10.1111/papq.12394.
Rabern, Brian. 2015. 'Descriptions Which Have Grown Capital Letters'. *Mind & Language* 30 (3): 292–319. https://doi.org/10.1111/mila.12081.
Reuter, Kevin, and Georg Brun. 2022. 'Empirical Studies on Truth and the Project of Re-Engineering Truth'. *Pacific Philosophical Quarterly* 103 (3): 493–517. https://doi.org/10.1111/papq.12370.
Riggs, Jared. 2019. 'Conceptual Engineers Shouldn't Worry about Semantic Externalism'. *Inquiry*, 1–22.
Saussure, Ferdinand de. 1916. *Course in General Linguistics*. Edited by Wade Baskin, Perry Meisel, and Haun Saussy. New York: Columbia University Press.

Sawyer, Sarah. 2020a. 'Talk and Thought'. In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 379–95. Oxford University Press.
———. 2020b. 'Truth and Objectivity in Conceptual Engineering'. *Inquiry: An Interdisciplinary Journal of Philosophy*.
Scharp, Kevin. 2013. *Replacing Truth*. Oxford, New York: Oxford University Press.
Schupbach, Jonah N. 2017. 'Experimental Explication'. *Philosophy and Phenomenological Research* 94 (3): 672–710. https://doi.org/10.1111/phpr.12207.
Shepherd, Joshua, and James Justus. 2015. 'X-Phi and Carnapian Explication'. *Erkenntnis* 80 (2): 381–402. https://doi.org/10.1007/s10670-014-9648-3.
Shields, Matthew. 2021. 'Conceptual Domination'. *Synthese* 199 (5–6): 15043–67.
Soames, Scott. 2003. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford, New York: Oxford University Press.
St. Pierre, Thomas, and Elizabeth K. Johnson. 2021. 'Looking for Wugs in All the Right Places: Children's Use of Prepositions in Word Learning'. *Cognitive Science* 45 (8): e13028. https://doi.org/10.1111/cogs.13028.
Sterken, Rachel Katharine. 2020. 'Linguistic Interventions and Transformative Communicative Disruption'. In *Conceptual Engineering and Conceptual Ethics*, edited by Herman Cappelen, David Plunkett, and Alexis Burgess. Oxford University Press.
Tanswell, Fenner Stanley. 2018. 'Conceptual Engineering for Mathematical Concepts'. *Inquiry* 61 (8): 881–913. https://doi.org/10.1080/0020174X.2017.1385526.
Thomasson, Amie L. 2020. 'A Pragmatic Method for Normative Conceptual Work'. In *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198801856.003.0021.
———. 2021. 'Conceptual Engineering: When Do We Need It? How Can We Do It?' *Inquiry* 0 (0): 1–26. https://doi.org/10.1080/0020174X.2021.2000118.
Torregrossa, Clotilde. 2022. 'Experimental Aesthetics and Conceptual Engineering'. *Erkenntnis*, May. https://doi.org/10.1007/s10670-022-00568-0.
Waismann, Friedrich. 1968. 'Verifiability'. In *How I See Philosophy*, 39–66. Springer.
Wakil, Samantha. 2021. 'Experimental Explications for Conceptual Engineering'. *Erkenntnis*, June. https://doi.org/10.1007/s10670-021-00413-w.
Waxman, Sandra R., Jeffrey L. Lidz, Irena E. Braun, and Tracy Lavin. 2009. 'Twenty Four-Month-Old Infants' Interpretations of Novel Verbs and Nouns in Dynamic Scenes'. *Cognitive Psychology* 59 (1): 67–95. https://doi.org/10.1016/j.cogpsych.2009.02.001.
Winter, Bodo, Marcus Perlman, Lynn K Perry, and Gary Lupyan. 2017. 'Which Words Are Most Iconic?: Iconicity in English Sensory Words'. *Interaction Studies* 18 (3): 443–64.