Conceptual Engineering is Old News

(forthcoming in *The Philosophical Quarterly*)

Krzysztof Sękowski*,
Ethan Landes^

* University of Warsaw, k.sekowski@uw.edu.pl
^ University of Zurich, ethan.landes@philos.uzh.ch

Abstract

According to a prevailing view, conceptual engineering introduces a revolutionary philosophical methodology, challenging traditional conceptual analysis. However, in our paper, we argue that closer scrutiny reveals not only the falsity but also the inherent ambiguity of this narrative. We explore four interpretations of the "Anti-Novelty Claim", the claim that conceptual engineering is not a new way of doing philosophy. Discussing the Anti-Novelty Claim from the perspective of a text's producer, the text's consumers, and the exegetical potential of the text, we examine each perspective's metaphilosophical implications and demonstrate that taking each perspective requires different methods. Adopting these different methods, we argue that the different interpretations of the Anti-Novelty Claim range from nearly trivially true to unlikely but untested. Importantly, we emphasize that each interpretation offers unique philosophical insights, yet addressing them requires diverse types of evidence, preventing a singular, straightforward answer to whether conceptual engineering is new.

keywords: conceptual engineering, philosophical methodology, metaphilosophy, history of philosophy, textual analysis

Acknowledgements: We would like to thank Amie Thomasson, Niklas Dahl, Twan Stiekel, Lucien Baumgartner, Katarzyna Kuś, Bartosz Maćkiewicz and the anonymous referees for Philosophical Quarterly for their comments on previous versions of this paper. Both authors contributed to this project equally.

1. Introduction

On its face, conceptual engineering is a huge shift in the methods and self-conception of analytic philosophers. Analytic philosophy, so the common story goes, has largely been interested in descriptive projects. Arguments, thought experiments, formalizations, and other tools of analytic philosophy have been used to discover what knowledge, justice, and freedom *are*. Conceptual engineering, the story goes, rejects the status quo by encouraging philosophers

to *design* the notions under debate. Thus, conceptual engineering brings a new, normative approach to philosophy (e.g., Cappelen 2018; Greenough et al. 2023; Nado 2021).

Some authors have rejected this narrative, arguing that conceptual engineering is a new name for an old method. For example, Deutsch argues that introducing new technical language – such as "credence" and "supervenience" – to talk about something theoretically important that we otherwise lack the expressive power to discuss is an uncontroversial and mundane part of philosophy (Deutsch, 2020). Indeed, we have countless historical examples of philosophers and non-philosophers alike successfully designing and spreading novel stipulations, from "quark" by Gell-Mann to the anonymous coining of "epistemology" in 1847 (Oxford English Dictionary 2023). However, conceptual engineers are typically interested in something more specific, or what Sawyer calls *narrow conceptual engineering*: the *revision* of existing conceptual and/or linguistic resources (Sawyer 2020).

Deutsch is sceptical that narrow conceptual engineering is possible, but several authors have recently argued that works predating conceptual engineering can be interpreted as conceptual engineering in this narrower sense (Thomasson 2017, forthcoming; Andow 2020; Jorem 2021; Misak 2022; Sękowski 2022; Westerblad 2022; Koch & Briesen 2023). According to these authors, attempts at revision have a long history in philosophy. Call the thesis that conceptual engineering is old news the **Anti-Novelty Claim:**

> Conceptual Engineering is not a new method of philosophy but merely a rebranding of an old, commonly used method of philosophy.

The Anti-Novelty Claim is a conceptual engineering-specific form of the cliche that there are no new ideas in philosophy. For the reasons Deutsch discusses, the Anti-Novelty Claim is nearly trivially true when conceptual engineering is understood in a broad way. Philosophers have introduced countless novel concepts or terms. Therefore, this paper will adopt a narrow reading of the Anti-Novelty Claim, namely, whether philosophers have been *revising* concepts all along. While we are focusing on whether past philosophers have revised concepts, there are of course, even narrower readings of the Anti-Novelty Claim, depending on what one thinks makes conceptual engineering conceptual engineering. We will instead stick to understanding the Anti-Novelty Claim as about conceptual revision, broadly speaking, to capture as many contemporary accounts as possible. More restrictive projects – e.g., accounts that take ethical desiderata to be a core aspect of conceptual engineering – are welcome, however, to use the conceptual framework developed here.

Even on Sawyer's narrow reading of "conceptual engineering", the general language of the Anti-Novelty Claim, as formulated here, hides different specific versions of the claim. As is argued below, not only do different readings of the Anti-Novelty Claim have different levels of plausibility, but they require different sorts of evidence to substantiate, and their truth or falsity has different metaphilosophical consequences. To appreciate the ambiguity of the Anti-

Novelty Claim, step back and consider the social epistemology of philosophy. While it is tempting to treat papers as standalone works of scholarship, we should remember that works of philosophy (or any academic work) are instruments of social epistemology (Hills 2020; Landes 2023; Rorty 1980). Texts are created by their authors – or *producers* – and are then read by those text's *consumers* (Deutsch 2015; Landes 2023). In philosophy, the relationship between producers and consumers is often quite loose. Consumers may knowingly or unknowingly misinterpret a producer. Consumers may think a producer is wrong but learn something in figuring out why. Consumers may lightly skim a work and gain inspiration from the general gist. At the same time, a producer may not even accurately record their intended justification or argument in a text, either because of poor writing, mistakes during editing, or stylistic concerns.

To illustrate why it is important to distinguish between producers and consumers when discussing the methodological history of philosophy, consider the following exchange between a particular producer of philosophical text and its consumer. Cappelen (2012), in his interpretation of Chalmers's (1996) zombie thought experiment, argues that Chalmers (1996) does not rely on intuitions to establish the thesis that philosophical zombies are conceivable (and thus physicalism about the mind is true). Cappelen accomplishes this, in part, by arguing that Chalmers's argument needs to be understood through an account of conceivability Chalmers develops in a later paper (Chalmers 2002). Responding to Cappelen, Chalmers refutes Cappelen's reading, saying that "the dialectical structure appealed to in the later text is quite different from that of the earlier text (intuitions play a much less focal role)" (2014: 541).

In this exchange, we can see the value of distinguishing different perspectives towards Chalmers's 1996 argument in *The Conscious Mind*. Most obviously, there is the perspective from the *producer*, Chalmers, who states explicitly (in both 1996 and 2014) that intuitions were intended to be part of the justification of the argument. Second, we can take the perspective of the *consumer*, in this case, Cappelen, whose takeaway from the text was that philosophical zombies are not conceivable due to grounds other than intuitions. Notice, however, that there is a third perspective we can take towards works of philosophy, the *text*. We can abstract away from specific readers and ask about how one can or should read the text, asking, for example, whether Cappelen's interpretation is justified or metaphilosophically useful.

When we consider the different stances we can take towards historical works of philosophy, the ambiguity in the Anti-Novelty Claim becomes clear. Is conceptual engineering not a new method because authors have been consciously engaging in it, because readers have had their concepts changed by authors, because those of us in the present can read old texts as proposing revised concepts, or because we *should* read those texts that way? This paper discusses these four different ways in which the Anti-Novelty Claim can be interpreted – one from the perspective of the producer, one from the perspective of the consumer, and two from the point of view of textual interpretation. It also discusses their metaphilosophical significance, and argues the different versions range from nearly trivially true to unlikely but untested. Note, however, that this paper does not try to definitively quantify how new or widespread

conceptual revision is in the history of philosophy as this would require a systematic historical survey of philosophy. Instead, the goal of the paper is to draw attention to the conceptual distinction between the different ways of understanding conceptual engineering's novelty, mapping each version's general plausibility, exploring how it would be investigated further, and discussing why such investigation is metaphilosophically important.

The next section considers *Producer CE*, the claim that philosophers have long intended for their work to revise concepts. We show evidence that proto-conceptual engineering approaches can be found in texts dating back well over a century both inside and outside of the analytic tradition. Section 3 examines and provides arguments in favour of the plausibility of the two versions of *Interpretive CE* – the position that one can (in one version) and should (in another one) interpret texts as revising concepts, regardless of authorial intent. Section 4 considers the fourth interpretation, *Consumer CE*, which holds that readers of philosophy have had their concepts changed by reading philosophical texts. While Consumer CE might be true, it is ultimately an empirical claim about how texts have in fact been interpreted by the readers. Ultimately, we demonstrate that there is no single sense in which conceptual engineering is or is not old news.

2. Producer CE

The first interpretation of the Anti-Novelty Claim is that the authors of works of philosophy in the past intentionally and explicitly tried to revise concepts or argued in favour of such a practice. Philosophers might do it either by providing arguments aimed at a particular concept or by making methodological claims according to which philosophical methodology should be focused on revision or amelioration rather than descriptive analysis. We will refer to that kind of interpretation of the Anti-Novelty Claim as *Producer CE*.

Whether Producer CE is true of a specific author and text turns on historical facts about the intention of the author. In the case of the living authors, we can ask them directly or conduct more systematic surveys on the methodological inclinations of philosophers (see Bourget & Chalmers 2023). In some cases, we do have authors' explicit declarations of their methodological intentions. For example, in Kripke's introduction to *Wittgenstein on Rules and Private Language* (Kripke 1982), Kripke states that he aims to provide a reading of *Philosophical Investigations* (Wittgenstein 1953) with no ambition to capture Wittgenstein's actual intentions. Sometimes, we have clues about how to interpret works by a specific author thanks to their other writings in which their methodological approach is independently presented and argued (e.g., Knobe & Nichols 2007; Strawson 1959). Another source of information with respect to the methodological approach of a particular author is their private diaries or correspondence, as can be found in Potter's study of Wittgenstein (2009). Finally, in some cases, the authors explicitly present and defend a particular interpretation of their past arguments, such as Chalmers's discussion of his zombie thought experiment discussed above (Chalmers 2014).

These are exceptions, however. Instead, in many cases, all we have about a producer's intentions are clues they leave in the texts themselves. Texts are, of course, imperfect guides to intended methods. Stylistic norms, mistakes, or even subtle changes in the meanings of words over time may mask a producer's true methodological intent. Nonetheless, when we look at historical works, we find evidence across the last century and a half of philosophy that philosophers saw themselves as revising concepts.

The most obvious example of an older work defending conceptual revision, discussed at length elsewhere, is Carnap's works on explication. Carnap is commonly described as the first example of a conceptual-engineering-like approach to philosophy (see, e.g., Cappelen 2018; Eklund 2021; Isaac et al. 2022). However, explicit claims according to which philosophy is (or should be) aimed at ameliorating or revising concepts are much older. Arguments in favour of such an ameliorative approach, and additionally, examples of conceptual engineering in practice, appear in works from the very beginning of the 20th century by the members of the Lvov-Warsaw School. For instance, in a paper from 1905, Łukasiewicz argued that since the term *cause* is often used incompatibly in different contexts, philosophy should at least partly *construct* a concept of causation to make it empirically correct, logically coherent, and explanatory:

> In our everyday life (…), we usually are not aware why, that is, due to what features do we call some things and not other things causes. Hence, we always speak of some cause or other, but a cause in general, as an abstract object, is in our life only an empty sound. In science, this term is used to signify some concept; however, in this case, there is another difficulty, for different scientists determine this concept differently. Each of them creates therefore a different abstract object which he signifies with the same word. Thus, whoever wishes to analyze the concept of cause encounters a serious problem. (...)
>
> There is, it seems to me, only one way around this difficulty. We need to accept that there is no pregiven abstract object called cause that we could analyze, and that such an object has to be yet created. And to create, that is, to construct, some abstract object, is to find certain features, consider which can be combined and which removed, and by this means arrive at a whole set of features connected by relations constituting the very object in question. (Łukasiewicz 1905/2022: 8)

In addition to this, a revisionary aspect of the method Łukasiewicz was arguing for is revealed when he states that a possible consequence of his method is a need for changing linguistic practices when it comes to the usage of the term "cause". This feature is important for two reasons: First, it shows that he was aware of the fact that the result of philosophical work might be incompatible with the concept of cause that is used by scientists, philosophers or folk. Second, this feature reveals that Łukasiewicz was arguing not only in favour of a need for constructing a concept for philosophical purposes but also for putting the

results of philosophical work into practice, which in contemporary literature is usually described as the *implementation* part of conceptual engineering. Here is a passage in which this feature of Łukasiewicz's theory is presented:

> By use of the method sketched above, I want to create a non-contradictory, unambiguous, scientific concept of cause consistent with reality. Perhaps this concept will appear as slightly different from usual definitions of a cause found in logic textbooks or works of metaphysics; perhaps it will not even be always consistent with what our everyday speech names as a cause in a more or less shaky and imprecise way. In case of such inconsistency, I will not be able to provide any remedy. It would simply require *breaking the habit* of calling a cause something that does not fall under the concept of cause, just as one needs to break the habit of calling carbonic acid the compound symbolized by CO_2, which is not an acid, but an acid anhydride. (Łukasiewicz, 1905/2022: 11)

Importantly, Łukasiewicz was not the only member of Lvov-Warsaw School who adopted such a conceptual-engineering-like approach. Similar methodological views can be found e.g. in works about definitions by Ajdukiewicz or Czeżowski, who argued that a proper philosophical definition should sometimes be at least partly regulative in order to make philosophical, scientific and everyday language more precise (see, Brożek et al., 2020 for more). Therefore, the members of the Lvov-Warsaw school developed a systematic methodological account that took a revisionary stance towards concepts decades before Carnap's works on explication.

While the Lvov-Warsaw School is considered to be a part of the analytic tradition, we can find examples of explicit formulations of conceptual-engineering-like views – although perhaps not as systematic as e.g., Łukasiewicz's work – in other philosophical traditions. For instance, as Koch and Briesen (2023) show, Nietzsche took a methodological approach according to which we should revise and improve our conceptual repertoire. Another precursor of conceptual engineering from the continental strand is Foucault and his genealogy and archaeology projects (Dutilh Novaes 2020; Queloz 2021). Indeed, cataloguing the examples across the history of philosophy and engaging in the exegesis required to establish the historical limits of Producer CE could take up this entire paper. Luckily, much of this work has been done already. Historical authors whose works have been interpreted as instances of conceptual engineering include John Dewey (Westerblad 2022), Charles Sanders Pierce and C.I. Lewis (Misak 2022), Gottlob Frege (Jorem 2021), and Martin Heidegger (Jorem 2021; Koch & Briesen 2023). The history of revisionary approaches might go even further back, as the Lvov-Warsaw School's methodology towards concepts is speculated to trace back to Brentano and John Stuart Mill (see Brożek forthcoming).

As just demonstrated, establishing the limits of Producer CE does not require revolutionary methods, but metaphilosophically, this is a project worth doing. First, establishing how widespread revisionary approaches are in the history of philosophy and how revisionary approaches differ across time, space,

and traditions will be fruitful for historians of philosophy and sociologists of philosophy. Such an inquiry is important, for instance, for those who want to determine what drove philosophical inquiry in the past and whether some sociological or political facts had an influence on the popularity of expressing ameliorative or revisionary approaches to philosophy. One might determine, for instance, what kind of aims were postulated within a particular philosophical approach and how they were related to the current political, sociological and theoretical background in which that approach was developed.

Second, when it comes to the impact of Producer CE, differences in frameworks and theoretical assumptions between contemporary conceptual engineering and proto-conceptual engineering can be an asset for scholarship. Correctly understanding texts as engaging in normative conceptual projects can save intellectual labour for both contemporary conceptual engineers and philosophers working on historical texts. Analogues to the contemporary problems facing conceptual engineering may have already been thought through and can be translated with varying levels of effort into conceptual engineering's analytic framework. Even if solutions cannot be lifted out directly, engaging with texts in older historical frameworks that are also revising concepts may help conceptual engineers appreciate new ways of approaching contemporary issues. In addition, engaging with proto-conceptual engineering texts may help problematize contemporary conceptual engineering by highlighting problems recognized by historical scholars that contemporary conceptual engineers have not recognized or have not taken seriously. Flipping things around, contemporary conceptual engineering may even help scholarship of historical work, offering ways of thinking about arguments made in historical texts that scholars of those texts had not yet appreciated.

3. Interpretative CE

So far, we have been interested in authorial intent. Now, let us turn to another reading of the Anti-Novelty Claim, according to which even if authors of philosophy do not intentionally take part in conceptual engineering (or a similar enterprise), their arguments, works, or theories might nonetheless be interpreted as if they were. Such an understanding of the Anti-Novelty Claim we will call *Interpretative CE*, as it is – in contrast to Producer CE and Consumer CE (discussed in Section 4) – not focused on either producers' or consumers' actual attitudes towards philosophical texts, but rather on the possibility of interpreting these texts as revising concepts.

A version of Interpretive CE has recently been defended by Andow (2020) and Sękowski (2022). Andow examined 14 thought experiments from recent publications in analytic journals and Sękowski examined Gettier cases. They independently argue that while these cases appear to advance descriptive claims about topics of interest in philosophy, the cases can be read as advancing revisionary or ameliorative claims. In other words, thought experiments and the associated arguments need not be taken to make object-level claims about what the content of the concepts *are*, but instead they can be taken to be making claims about what these concepts *should* be. Thus, regardless of the intent of the author,

many thought experiments can and/or should be understood as a type of conceptual engineering.

Notice, however, the gap between saying a text *can* be read as revising a concept and saying a text *should* be read as revising a concept. Interpretive CE, the claim that it is possible to interpret one's argument or view as a kind of conceptual engineering, consequently has both a descriptive and normative reading. According to the descriptive reading, not only is it merely possible to read a text as proposing revised concepts – which is an extraordinarily low bar – but the reading is internally coherent. That is, the revisionary reading is not contradicted by anything explicitly stated in the text. Call this *Descriptive Interpretive CE*. In contrast, according to the normative reading of Interpretative CE, adopting such an interpretation is additionally beneficial or desired. Note that one could accept the fact that it is possible to coherently interpret a lot of philosophical work from the past as a part of conceptual engineering, but at the same time, state that such an interpretation is not fruitful or leads to bad consequences. Consequently, the view that it is not only possible but also reasonable to interpret a lot of philosophical works as engaging in conceptual engineering (even if this interpretation does not capture their author's intentions) is *Normative Interpretative CE*. In the next two sections, we will discuss Descriptive Interpretive CE and Normative Interpretive CE, respectively.

3.1. Descriptive Interpretative CE

According to a standard methodological picture, the method of cases establishes truths about a scrutinized concept by considering whether a given concept applies in possible situations (Bealer 1998; Cappelen 2012; Deutsch 2015; Machery 2017). Andow (2020) instead argues that while intuitions serve as evidence in the method of cases, they provide evidence not for descriptive claims concerning scrutinized concepts, but rather for normative constraints that we should adopt for them within a relevant theoretical framework. For instance, Andow considers the following example from Robinson (2019), in which Robinson argues that one can be cruel without causing harm:

> *Imagine that Tom is in love with his neighbour. He believes that his neighbour's husband is severely anaphylactic and that peanuts will kill him. Feigning generosity, Tom makes dinner for the couple each week and sprinkles the dish with crushed peanuts. The husband, it so happens, is not actually allergic to peanuts and ends up enjoying the meal even more than he would have without the peanuts. In the end, good effects are ultimately produced. But I think* ***we can still look at Tom and say that his behaviour is cruel****. He intentionally puts peanuts into the meal, hoping to kill the husband and take his wife. But the effects are not 'successful' in the sense that the aims of the motivational component of his act are fulfilled.*

Andow argues that the bolded passage is evidence that the above case is not offering an accurate account of what cruelty is. Instead, Andow argues that

Robinson is arguing that we *should* think about cruelty in a way that does not require harming the object of the cruelty. Our intuitions that the behaviour is cruel are not, on this reading, revealing what cruelty is but revealing what we want cruelty to be.

Notice that one can argue that Andow (or any non-standard revisionary reading of a thought experiment) is engaging in a sneaky linguistic trick. Robinson may have actually had a descriptive argument in mind – that Tom's behaviour is indeed cruel – but our point here is not to establish that Robinson intentionally engaged in conceptual engineering nor that such an interpretation is the only correct interpretation. Our point instead is that the wording used in the text leaves open the possibility to (mis)interpret the discussed argument as one arguing for revising what cruelty is. Thus, we can read Robinson as making an ameliorative point, making Descriptive Interpretive CE true for the text.

To see Andow's point about how cases can reveal normative constraints, it is worth considering a dialectically similar example from outside of philosophy: Piaget's study on a child's concept of velocity as discussed by Kuhn (1964/1977). The example shows a revision of children's conceptual framework, and its structure can be easily adapted to philosophical thought experiments. In a study, children around 5 years old appear to have two criteria for velocity when judging which of two toy cars on a track are faster. The first is what Kuhn calls "perceptual blurriness", or how fast an object appears to move based on how blurry it appears. The second criterion for velocity is whether the car reaches its destination first – regardless of whether the "faster" object starts much earlier or closer to the goal than the "slower" object. While these two criteria obviously conflict, children below a certain age do not notice the tension. However, when exposed to objects that moved very rapidly but did not reach their goal first, many children in the study recognized they *should* primarily consider how blurry a moving object looks, dropping the other criterion. In effect, thanks to considering particular cases, children changed their concept of velocity, restricting it to how fast an object appears to be moving. Robinson, on Andow's reading is making a similar move based on what our intuitions reveal about the normative constraints on cruelty. Like Piaget's demonstration revealed to children that velocity should not depend on finishing first, Robinson's case reveals cruelty should not depend on causing harm, which helps motivate us to revise our concept of cruelty.

According to Andow's picture of the method of cases, our intuitions concerning the verdict of possible cases can reveal normative constraints on concepts that we were previously unaware of. Notice, however, that the method of cases, traditionally construed, can also be seen as revealing normative constraints. On a traditional understanding of the method of cases, thought experiments uncover something about the underlying structure of reality, such as the structure of cruelty. Insofar as we want our language and thought to reflect reality, on this traditional picture, the method of cases reveals how we should talk and think. What Andow, and by extension, we, are claiming in this section is that you can read the method of cases as revealing constraints that are distinctly revisionary, such as moral or political constraints that might motivate revising concepts. Put another way, on the traditional view, whether we succeed

in revealing relevant normative constraints depends on whether we better represent the structure of certain phenomena (e.g. cruelty), while on a revisionary approach, success depends on uncovering other kinds of constraints (e.g. political or ethical).

Andow (2020) and Sękowski (2022, 2023) argue most thought experiments can be interpreted as being instances of conceptual engineering. Here, we want to go a step further to argue that a sizable chunk of philosophical arguments, beyond just the method of cases, can be interpreted as works of conceptual revision. This is because what is on the page underdetermines how a text is interpreted. Works of philosophy are only so long, so philosophers cannot state every assumption made in a text (Deutsch 2015: 123) – assuming an author is even aware of all their text's assumptions. This includes metaphilosophical stances of the author that would otherwise rule out, for example, a conceptual revision-based reading of the account. While this underdetermination makes Producer CE difficult to evaluate, it means that Descriptive Interpretive CE is true for a wide range of texts beyond just the method of cases.

To illustrate how Descriptive Interpretive CE can be applied beyond thought experiments, we will adapt Sękowski's discussion of Gettier to Hume's (1748) discussion of causation. As we will show, there is a possible interpretation of Hume's argument, according to which Hume is arguing for conceptual revision by referring to theoretically relevant expectations towards causation. As Sękowski (2022) argues, in the method of cases, philosophers present some cases, and they propose arguments in favour of verdicts about these cases by appealing to these philosophers' *expectations* towards a given concept. In Gettier's 10 coins case, the verdict that *Smith does not know that e* is justified by an argument in which Gettier appeals to his expectation towards knowledge that to know something, there must be a non-accidental connection between the origin of the belief and a fact that makes the belief true. Following this interpretation, we should say that Gettier is not assuming that our intuitions are a source of evidence for what knowledge is like. On the contrary, Gettier is presenting a case and arguing in favour of a particular verdict to convince us to revise a given concept to capture relevant expectations towards that concept.

At the start of Hume's discussion of causation (1748), he accepts that causation is a necessary sequence of events. This is his starting concept of causation. Hume agrees that his observation of billiard balls hitting each other is accompanied by a kind of intuition with the content that "there is a causal relationship between one billiard ball hitting a second ball and the movement of the second ball". Therefore, he notes that we can have several intuitions according to which there is a causal relationship between individual events. At the same time, in the light of his empiricism, Hume holds that impressions are the only source of ideas or concepts. This assumption is an expression of a certain expectation towards the desired concept or theory of causal relationship according to which there must be an impression of causation which serves as a source of this idea (or concept), which is partially based on the philosophical context in which Hume worked. Applied to the concept of causation, this entails

that if causation is the necessary sequence of events, then we should have impressions of necessary sequences of events.

However, Hume considers a case of two billiard balls hitting each other, and he states that he finds nothing more in his impressions than a picture of events that occur in order. This conclusion is incoherent with the mentioned expectations. In effect, Hume has to reject either the intuition that there is a causal relationship between the two billiard balls or the requirement that impressions serve as a source for our ideas. Hume chooses the first option, and he argues that we need to revise our conceptual framework due to the incompatibility of more general expectations and intuitions concerning the applicability of a given concept (in this case, causation). As a consequence, we should assume that the concept of the cause of individual token events does not contain any modal concepts. Therefore, causation is not a necessary consequence of events.This discussion of Hume is illustrative of the fact that Descriptive CE is true for a wide variety of texts. If one merely cares about what interpretations are coherent and consistent with what is on the page – and not, for example, accurately capturing authorial intent or other aspects of historical veracity – then texts offer significant hermeneutic leeway to be interpreted as, for example, revising concepts.

The first important metaphilosophical outcome of this fact is that the value of reading a text goes beyond gaining true beliefs about authorial intent. We cannot help but bring ourselves into the process of engaging with a text, bringing our own idiosyncrasies, whether our own idiolect (Fischer 2020) or methodological assumptions (Landes 2023). This is, however, as much a feature as a bug. Setting aside the present issue of conceptual engineering's novelty, even if one believes authorial intent is crucial for determining what a text *means*, the interpretive freedom afforded by texts explains why the philosophical value of works of philosophy is not limited by authorial intent, era, methodologies, or assumptions. We do not need to be logical empiricists to learn about the nature of science from Popper, nor do we need to be Platonists to learn about the nature of beauty from Plato (see also Cappelen 2012: 183). Instead, while reading, we can readily convert their observations to our own understanding of the nature of philosophy. This methodological outcome does not hang on Descriptive Interpretive CE's truth. What it does do, however, is open a way of interpreting texts that might be methodologically fruitful if one is already a conceptual engineer. Because of Descriptive Interpretive CE's truth, conceptual engineering, even if absolutely brand new, can readily draw insights from the philosophical canon.

Another crucial metaphilosophical consequence of the general truth of Descriptive Interpretive CE is that it opens the possibility to the two final readings of the Anti-Novelty Claim discussed below. Because we *can* interpret texts as attempts to revise concepts, it becomes an open question whether we *ought* to interpret the texts this way (Normative Interpretive CE) or whether, as a matter of fact, texts *have* revised consumer's concepts (Consumer CE).

3.2. Normative Interpretative CE

Descriptive Interpretative CE states that it is possible to interpret a lot of philosophical texts as instances of conceptual engineering. Now, let us ask whether one *should* read the texts in this way. Of course, if one is striving to capture authorial accuracy (as far as this is possible), one should read a text as engaging in conceptual engineering if that is what the author intends. Nonetheless, even if a reading of a text is incompatible with an author's intentions, adopting it can be justified by other reasons. This is not to say Normative Interpretive CE should be understood as the radical view that *all* arguments should be interpreted as revisionary or that we should always abstract from the context in which the text was created and from the author's intentions. Normative Interpretive CE is instead a weaker claim, according to which many works in the history of philosophy can be interpreted as instances of narrow conceptual engineering and, crucially, that doing so can be beneficial. We will first look at the benefits such a revisionary reading can offer conceptual engineers but then argue that the value of revisionary readings is not limited to conceptual engineering.

Although sometimes we want to capture authorial intent and we have good reason to think that Producer CE is true of that text, that is not the only reason that one should approach a text as forwarding a revisionary argument. As has already been alluded to, non-standard and inaccurate readings can be inspirational. Reading texts as taking a revisionary stance towards concepts gives conceptual engineers a wider range of arguments and established positions on certain topics. Reading texts in an unintended way is still capable of drawing the reader's attention to considerations, reasons, and arguments that the reader might not have considered otherwise by inspiring the reader to consider things by their own lights (Landes 2023). Therefore, adopting a non-standard revisionary reading of a text may serve as a source of inspiration for conceptual engineers' own revisionary projects – perhaps even providing entire arguments.

To put this point more specifically, for conceptual engineers, there are practical and epistemic reasons for interpreting texts as if they were attempts to revise concepts. Successful, intentional conceptual engineering requires four distinct tasks, *identifying* what existing content exists, *evaluating* whether that content is deficient, *designing* new content, and *propagating* the designed content (Isaac et al. 2022). Reading existing arguments *as if* they propose changing conceptual content will provide insights – and thus potentially save time – at both the evaluation and design stages of conceptual engineering. Similarly, extant analyses of concept, regardless of whether they accurately capture an existing concept's content, can serve as inspiration during the design stage, offering ideas or frameworks to think about what improved content would involve. Arguments for or against analyses may help in designing revised content *or* identifying desiderata by which to evaluate existing content. Therefore, treating descriptive works of philosophy as attempts to revise concepts may help conceptual engineers identify desiderata for their projects, potentially saving time and energy in the process.

Let us give an example. Fine (1994) argues against a Kripkean account of essence that reduces essence to necessary properties by drawing attention to the fact that Socrates has the necessary property "being a member of singleton {Socrates}". As Fine argues, in contrast to a Kripkean account of essences, we should not consider this property to be essential to Socrates because there is nothing in the nature of Socrates that is the source of this property. Even though Fine is a metaphilosophical and metaphysical realist, we may nonetheless adopt these arguments to make a revisionary point about the concept of essence. If Kripkeans are correct that the concept of essence is reducible to necessity, then Fine has provided reasons to revise that concept to meet our intuitions towards the *nature* of some entities (e.g. of Socrates). And what is more, as the last thirty years of literature on metaphysics and modal epistemology show, the Finean account of essence is highly theoretically fruitful (Correia 2024), even if it will turn out to be revised in the future by a different notion of essence. Therefore, normative interpretations of otherwise descriptive arguments can reveal possible ways of formulating arguments for conceptual revision, and what is more, if reframing descriptive arguments into revisionary arguments would produce unsound arguments, it can be useful to establish what kind of fallacies may appear in normative arguments.

Even for those not engaged in ameliorative projects, reading works as works of conceptual engineering can still be fruitful. As discussed in Section 3.1, both Andow (2020) and Sękowski (2022) adopt a revisionary interpretation of the method of cases. However, unlike Andow, who focuses on arguing for the possibility and plausibility of such a non-standard reading, Sękowski argues that we should adopt a revisionary understanding of thought experiments to defend the case method against experimental findings (for why it is fruitful to interpret other types of historical philosophical arguments this way, see also Thomasson (2017; forthcoming)). According to the experimental attack, intuitions are sensitive to philosophically irrelevant factors, and since the case method uses intuitions as a source of evidence or justification, the case method is an unreliable method of philosophy (for an overview of these arguments see Deutsch 2015 and Machery 2017).

Sękowski (2022; forthcoming) responds to this literature by arguing that if we adopt a normative approach towards thought experiments in which the content of intuitions are expectations concerning general features of concepts rather than the applications of concepts, there is a non-standard version of an expertise defence available to defend the use of the method of cases. This normative interpretation disconnects the method of cases' value from intuitions' reliability because concept revision is justified by our expectations. These expectations however do not reflect our actual, shared beliefs about certain concepts. They might be e.g. the consequence of the function we implicitly or explicitly assign the concept or other beliefs we have about the domain in which that concept is used. However, our expectations might be seen by others as more or less convincing, and it is a matter of fact that people may have more convincing expectations towards concepts than other people due to the fact that some of us simply know the functions of certain concepts, or contexts in which they are used better. Expectations' relevancy comes then from the fact that the

person who possesses them is familiar with the purpose and past context in which the scrutinized concept was discussed.

So, for instance, expectations towards epistemological concepts of epistemologists are more apt, not because they track reality, but because they might be more convincing in justifying concept revision within an epistemological dispute on the nature of knowledge. This is due to epistemologists' better immersion in the epistemological context in which the concept of knowledge is discussed. People with such immersion simply know that what epistemologists require from the concept of knowledge is that e.g. it has to address sceptical worries and Gettier examples. Similarly, e.g. lawyers' expectations concerning the concept of alibi used in law are a reliable source thanks to their immersion in legal contexts, and lawyers' knowledge of the function of these concepts in several law contexts. Kids' expectations towards a hide-and-seek game is more apt in the context of negotiating rules of these game among kids than their parents, since kids know better what is to be expected from hide-and-seek *in the context of kids' play* than their parents. This is because parents might think about the game in a different way, for instance as a part of a wider process of child development. The revisionary interpretation of the method of cases therefore makes the method immune to sceptical worries due to the discussed role of expectations and immersion in the context in which the concept is discussed. (See Sękowski 2022; forthcoming).

As discussed in this section, there is value in pretending a work is engaging in conceptual revision, regardless of authorial intent. This stance is enabled by the underdetermination of the text (discussed above). The stance also has parallels to the tradition in the philosophy of literature that has its roots in Barthes *Death of the Author* (1967), according to which from the perspective of the reader, the most fruitful way of interpreting a text is to abstract from the historical context of its origin and to be open to adopt a point of view of any potential reader. Of course, the value of reading a text as revisionary is very much dependent on framework and interest. Repurposing descriptive arguments and accounts to revisionary ends will only help conceptual engineers, and defending the case method against the experimental attack might only be attractive to those who buy into the attack's premise (see Deutsch 2015, Horvath 2023). Thus, we do not take our examples here to be exhaustive, but rather two versions of a more general strategy made available by allowing ourselves to be free from what an author intends.

4. Consumer CE

In contrast to Producer CE, which is a claim about authorial intent, Descriptive Interpretive CE, which is a claim about the possible coherent interpretations of a text, and Normative Interpretive CE, which is a claim about the normative value of interpretation, Consumer CE is a descriptive claim about how texts are actually interpreted by consumers. Consumer CE is true about a text relative to consumers when people's concepts have changed because of engaging with the text – regardless of authorial intent. As discussed above, professional philosophers possess many concepts non-philosophers do not. The

development of novel concepts for technical ends means that it is nearly trivial that philosophers have different concepts (see Fischer & Sytsma 2021; Deutsch 2020; Liu 2023). Rare are the non-philosophers who have the concept of (logical) validity, pro toto reasons, or (Fregean) sense. Because engaging in philosophy is clearly capable of propagating novel technical concepts, the question in the narrower sense of conceptual engineering is whether philosophers have *revised* or *refined* concepts because of reading certain works of philosophy.

Depending on whether the consumer understands the producer as trying to change the concept, we can think of any revision as explicit or implicit. A possible example of explicit Consumer CE from outside of philosophy is the change of the concept of "planet" because of the IAU's explicit redefinition in 2006. Implicit conceptual revision may sound more farfetched, but it appears to be the more common type of conceptual change. There is robust evidence that concepts change during childhood (see Carey 2011). When we propagate to children the conceptual content associated with Newtonian mechanics or real numbers, few teachers, parents, or children explicitly call what is happening "conceptual change". Instead, it is understood broadly as learning. Accordingly, conceptual revision during philosophical training may be presented and understood as some other type of learning or development.

Both explicit and implicit Consumer CE turn on whether concepts have been revised while reading the texts, so the rest of the section will explore the available evidence for whether engaging in philosophy revises concepts. This question ultimately turns on descriptive historical, sociological, and perhaps cognitive facts of philosophy's consumption. To answer definitively whether concepts have changed in response to a text, we would need to draw some sort of diachronic comparison of a reader or readers in response to a text. Looking first at qualitative approaches, Brown (2017) looks at the early consumption of Gettier's 1963 paper among epistemologists. The stability Brown finds about the Gettier verdicts among epistemologists can be interpreted as evidence that the concept of knowledge remained stable among epistemologists in response to Gettier (1963). Similarly, approaches from the point of view of the history of ideas can show how concepts change over time in response to philosophers (e.g., Bennett 2007). Notice, however, that any such project will have difficulty distinguishing changes in beliefs in response to reading philosophers from changes in concepts in response to reading philosophers – a challenge faced by all the methods described in this section.

Looking at *qualitative* diachronic approaches, modern corpus analytic methods may also be able to measure concept change. Computational methods allow for the measurement of meaning through measures like semantic similarity, which allow corpus linguistics to quantify the meanings of words based on their co-occurrence with other words in the corpus (see Navigli and Martelli 2019). In theory, corpus analysis can track the impact of a text on philosophical literature or track the semantic change of a given concept in contexts outside philosophy before and after the appearance of an influential philosophical work. For example, if citation tracking reveals citation of a particular text corresponds to changes in semantic similarity between texts

(Hogenbrick 2023), then there would be evidence that the cited paper led to changes in conceptual content.

Nonetheless, non-historical methods can be used to evaluate the plausibility of Consumer CE by testing conceptual content before and after engaging with texts. While we know of no work testing this directly, for some, albeit limited, insight, we can look towards the literature on philosophical expertise (for a recent overview, see Horvath 2023). Empirical work on the expertise defence has looked for differences in responses to thought experiments between folk and expert philosophers. Work on the expertise defence has interpreted such findings as evidence for or against whether philosophers deploy some sort of expert ability to consume and evaluate thought experiments. This is not the only interpretation of such data, however. For authors who take case judgements and conceptual content to be closely linked (e.g., Machery 2017), changes in case judgements may also be caused by changes in conceptual content. On such an interpretation, studies comparing folk and expert case judgements cannot reveal whether engaging with specific texts has changed people's conceptual content, but they can reveal differences in content between the two groups. Accordingly, data from the expertise defence can offer at least some insight into the historical plausibility of Consumer CE.

Comparisons between philosophers' and non-philosophers' verdicts about thought experiments are mixed. Machery (2012) found no difference between folk and philosophers' judgements about Gödel cases. Schindler and Saint-Germier (2023) also found no folk-expert difference in Gödel cases, nor did they find them in Mary or Twin Earth cases. Schindler and Saint-Germier did, however, find relatively large (about 1 point on a 5-point scale) differences between philosophers' and non-philosophers' verdicts on Gettier cases and Fake Barn cases and found much smaller but nonetheless significant differences in verdicts about the Chinese Room. Horvath and Wiegmann (2016) also found differences between folk and expert responses in a modified Fake Barn case but did not find differences in verdicts about a modified Gettier case.

In addition to the limitations already discussed, this work is further limited by its use of cross-sectional designs that test for expertise by comparing one group (philosophers) to another (non-philosophers). If differences are found, these cross-sectional studies struggle, by their very nature, to distinguish whether differences in the two groups are caused by training or whether philosophers were a non-standard group of people from the start (because, for example, people who have the "wrong" intuitions are pushed out of the field). One study has overcome this by testing how responses change over the course of undergraduate education (Maćkiewicz et al. 2023). Responses to ten thought experiments were collected at seven points over the course of undergraduate study. The study notably found a large and sustained shift in responses about Gettier and Fake Barn cases before and after taking an epistemology course during the first semester. The authors take this as evidence that changes in Gettier and Fake Barn responses were driven by *beliefs* about what the correct responses are, but given the stability of conceptual content over time, this sort of long-term change also suggests a shift in concepts of knowledge. This ambiguity further highlights the difficulty anyone trying to establish Consumer

CE faces distinguishing belief-based and concept-based changes (for one solution, see Fischer 2020; Landes & Reuter ms).

Ultimately, however, there is currently little direct evidence for or against Consumer CE. The findings above suggest *something* is happening with Fake Barn cases that might also be happening to Gettier cases, but due to limitations in designs, we cannot tease apart what caused it, nor can we answer whether the differences are driven by concepts or beliefs. Indeed, as discussed above, authors in the expertise defence literature have, for example, argued that while folk and philosophers have the same concepts, philosophers are better at accurately deploying concepts than the folk or are better at reading and evaluating thought experiments (see Horvath 2023). Therefore, establishing the truth or falsity of Consumer CE relative to particular texts or concepts will require more sophisticated qualitative and quantitative research than currently exists.

Ultimately, the truth or falsity of Consumer CE has consequences for the relationship between the philosophising of professional philosophers and non-philosophers. As mentioned, many have argued that any differences between folk and expert judgements in philosophy are a consequence of experts being better placed to answer philosophical questions or engage with thought experiments. However, as argued here, differences in responses could instead be caused by philosophers having distinct conceptual content from non-philosophers. If it turns out that philosophers' concept of knowledge changes because of reading Gettier, perhaps philosophers and non-philosophers mean different things when they ask themselves, "Do I *know* I have two hands?". This may in turn cause miscommunication between folk and experts.

The truth or falsity of Consumer CE also has an impact on the propagation stage of conceptual engineering. Very little is currently understood about how conceptual engineers can bring about a desired conceptual change (Koch 2021; Koslow 2022; Landes ms). If it turns out that consuming and engaging with a work of philosophy is enough to change concepts, then the propagation stage of conceptual engineering may just require having the public read philosophy. Even if the mechanism through which such a text changed concepts was more involved, the discovery of a philosophy text that revised concepts would be invaluable. Conceptual engineering has hitherto struggled to succeed at the propagation stage, and successful examples "in the wild" would offer insight into what works.

That said, due to differences between conceptual engineering frameworks, no single example would be uncontroversial. Sceptics towards the very possibility of conceptual engineering based on metasemantic arguments, such as Deutsch (2020) (and to some extent Cappelen (2018)), will be able to explain away any change as being merely belief-based where more optimistic theorists might see concept change. However, if one accepts frameworks on which conceptual revision is more achievable, examples of successful revision would provide invaluable insight into what processes work for those frameworks. This is an important benefit not only for conceptual engineers interested in implementing revisions, but it would also put pressure on positions

sceptical of conceptual engineers' control of, or impact on, conceptual revision (e.g. Andow 2021; Machery 2021).

5. Conclusion

Conceptual engineering is old news. Conceptual engineering is also not old news. Contemporary conceptual engineering is not the first philosophical framework to push for the revision of concepts, nor is it the first framework within analytic philosophy to do so. However, conceptual engineering offers us a novel framework for approaching old texts with new eyes. Not only can we approach a huge proportion of texts predating contemporary discussions of conceptual engineering with a new, fruitful perspective – that they are placing normative, not descriptive, constraints on the targets of analysis – but in many cases, we *should* now approach historical texts as if they are engaged in conceptual engineering. Whether specific works of philosophy have actually succeeded in revising concepts in the past, however, is a different matter entirely and requires further empirical investigation. Ultimately, however, there is no one way in which conceptual engineering is new or old because there is no single way to think of a text. As we have shown, different methods will be needed to establish the plausibility of each hypothesis about the novelty of conceptual engineering, and each hypothesis' truth or falsity will have different impacts on how we understand conceptual engineering's place in analytic philosophy and the history of philosophy more generally. Nonetheless, by clarifying the ways in which conceptual engineering is novel or not, conceptual engineers can take different, and more fruitful, approaches to understanding the method's history.

6. References

Andow, J. (2020). 'Intuitions about cases as evidence (for how we should think)', *Inquiry*, 66/6: 1036–68. DOI: 10.1080/0020174X.2020.1767199

——. (2021). 'Conceptual Engineering is Extremely Unlikely to Work. So What?', *Inquiry: An Interdisciplinary Journal of Philosophy*, 64/1–2: 212–26. Routledge. DOI: 10.1080/0020174x.2020.1850343

Barthes, R. (1967). 'The Death of the Author', *A. Leavers, Smith & Hill, New York*.

Bealer, G. (1998). 'Intuition and the Autonomy of Philosophy'. DePaul M. & Ramsey W. (eds) *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, pp. 201–40. Rowman & Littlefield: Lanham, Md.

Bennett, M. R. (2007). 'Development of the concept of mind', *Australian & New Zealand Journal of Psychiatry*, 41/12: 943–56. Sage Publications Sage UK: London, England.

Bourget, D., & Chalmers, D. J. (2023). 'Philosophers on philosophy: The 2020 philpapers survey', *Philosophers' Imprint*, 23/1. Michigan Publishing Services.

Brown, J. (2017). 'The Gettier Case and Intuition', *Explaining knowledge: New essays on the Gettier problem*, 191–211. Oxford University Press Oxford.

Brożek, A. (forthcoming). 'Reconstruction of Concepts from Brentano – through Łukasiewicz – to Tarski: Austro-Polish Contributions to the Methodology of Philosophy', *Przegląd Filozoficzny*.
Brożek, A., Będkowski, M., Chybińska, A., Ivanyk, S., & Traczykowski, D. (2020). *Anti-irrationalism: Philosophical Methods in the Lvov-Warsaw School*. Semper Wydawnictwo Naukowe.
Cappelen, H. (2012). *Philosophy without intuitions*. Oxford: Oxford University Press.
——. (2018). *Fixing Language*. Oxford University Press. DOI: 10.1093/oso/9780198814719.001.0001
Carey, S. (2011). *The origin of concepts*. Oxford series in cognitive development, 1. iss. Oxford Univ. paperback. Oxford: Oxford Univ. Press.
Chalmers, D. (2002). 'Does conceivability entail possibility?', *Conceivability and possibility*, 145: 200.
Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford university press.
——. (2014). 'Intuitions in philosophy: a minimal defense', *Philosophical Studies*, 171/3: 535–44.
Correia, F. (2024). 'Non-modal Conceptions of Essence'. *The Routledge Handbook of Essence in Philosophy*, pp. 124–42. Routledge.
Deutsch, M. (2015). *The Myth of the Intuitive: Experimental Philosophy and Philosophical Method*. Cambridge, Massachusetts: The MIT Press, a Bradford Book.
——. (2020). 'Speaker's reference, stipulation, and a dilemma for conceptual engineers', *Philosophical Studies*, 177/12: 3935–57. DOI: 10.1007/s11098-020-01416-z
Dutilh Novaes, C. (2020). 'Carnap meets Foucault: Conceptual engineering and genealogical investigations', *Inquiry*, 1–27. Taylor & Francis.
Eklund, M. (2021). 'Conceptual Engineering in Philosophy'. Khoo J. & Sterken R. (eds) *The Routledge Handbook of Social and Political Philosophy of Language*.
Fine, K. (1994). 'Essence and Modality: The Second Philosophical Perspectives Lecture', *Philosophical Perspectives*, 8: 1–16. [Ridgeview Publishing Company, Wiley]. DOI: 10.2307/2214160
Fischer, E. (2020). 'Conceptual control: On the feasibility of conceptual engineering', *Inquiry*, 1–29. Taylor & Francis.
Fischer, E., & Sytsma, J. (2021). 'Zombie intuitions', *Cognition*, 215: 104807. DOI: 10.1016/j.cognition.2021.104807
Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis*, 23/6: 121–3.
Greenough, P., Gové, J., & Church, I. (2023). 'Conceptual engineering for analytic theology', *Inquiry*, 0/0: 1–34. Routledge. DOI: 10.1080/0020174X.2023.2244007
Hills, A. (2020). 'Moral Testimony: Transmission Versus Propagation', *Philosophy and Phenomenological Research*. DOI: 10.1111/phpr.12595
Hogenbrick, H. (2023). *Each Book Its Own Babel: Conceptual Unity and Disunity in Early Modern Natural Philosophy*. PhD Thesis, University of Groningen.

Horvath, J. (2023). 'Intuitions in Experimental Philosophy'. Bauer A. M. & Kornmesser S. (eds) *The Compact Compendium of Experimental Philosophy*, pp. 71–100. De Gruyter.

Horvath, J., & Wiegmann, A. (2016). 'Intuitive expertise and intuitions about knowledge', *Philosophical Studies*, 173/10: 2701–26. DOI: 10.1007/s11098-016-0627-1

Hume, D. (1748). *An Enquiry Concerning Human Understanding*., Project Gutenburg eBook.

Isaac, M. G., Koch, S., & Nefdt, R. (2022). 'Conceptual engineering: A road map to practice', *Philosophy Compass*.

Jorem, S. (2021). 'Conceptual Engineering and the Implementation Problem', *Inquiry: An Interdisciplinary Journal of Philosophy*, 64/1–2: 186–211. Routledge. DOI: 10.1080/0020174x.2020.1809514

Knobe, J., & Nichols, S. (2007). 'An Experimental Philosophy Manifesto'. Knobe J. & Nichols S. (eds) *Experimental Philosophy*, pp. 3–14. Oxford University Press.

Koch, S. (2021). 'The externalist challenge to conceptual engineering', *Synthese*, 198/1: 327–48.

Koch, S., & Briesen, J. (2023). 'Conceptual infrastructure and conceptual engineering'. Pinnix A., Volmar A., Esposito F., & Binder N. (eds) *Rethinking Infrastructure Across the Humanities*, pp. 75–86. Transcript.

Koslow, A. (2022). 'Meaning Change and Changing Meaning', *Synthese*, 200/2: 1–26. Springer Verlag.

Kripke, S. A. (1982). *Wittgenstein on rules and private language: An elementary exposition*. Harvard University Press.

Kuhn, T. (1977). 'A Function for Thought Experiments'. *The Essential Tension: Selected Studies in Scientific Tradition and Change*, pp. 240–65. University of Chicago Press: Chicago and London.

Landes, E. (2023). 'Philosophical producers, philosophical consumers, and the metaphilosophical value of original texts', *Philosophical Studies*. DOI: 10.1007/s11098-022-01900-8

Landes, E. (ms). 'Conceptual Engineering Should be Empirical', *PhilPapers*.

Landes, E., & Reuter, K. (ms). 'Conceptual Revision in Action', *PhilPapers*.

Liu, M. (2023). 'X-Phi and the challenge from ad hoc concepts', *Synthese*, 201/5: 166. DOI: 10.1007/s11229-023-04149-8

Łukasiewicz, J. (2022). 'Analysis and Construction of the Concept of Cause'. Jadacki J. J. & Swiderski E. M. (eds) *The concept of causality in the Lvov-Warsaw school: the legacy of Jan Lukasiewicz*, Poznań studies in the philosophy of the sciences and the humanities, pp. 3–64. Brill: Leiden Boston (Mass.).

Machery, E. (2012). 'Expertise and Intuitions about Reference', *THEORIA. An International Journal for Theory, History and Foundations of Science*, 27/1: 18.

——. (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. Routledge.

Maćkiewicz, B., Kuś, K., & Hensel, W. M. (2023). 'The influence of philosophical training on the evaluation of philosophical cases: a controlled longitudinal study', *Synthese*, 202/4: 113.
Misak, C. (2022). 'A new history and underpinning for conceptual engineering', *Inquiry*, 0/0: 1–16. Routledge.
Nado, J. (2021). 'Conceptual engineering, truth, and efficacy', *Synthese*, 198: 1507–27.
Navigli, R., & Martelli, F. (2019). 'An overview of word and sense similarity', *Natural Language Engineering*, 25/6: 693–714.
Oxford English Dictionary. (2023). 'Epistemology (n.)'. Oxford University Press. Retrieved from Oxford English Dictionary.
Potter, M. D. (2009). *Wittgenstein's notes on logic*. Oxford: Oxford University press.
Queloz, M. (2021). *The practical origins of ideas: Genealogy as conceptual reverse-engineering*. Oxford University Press.
Robinson, J. (2019). 'On being cruel to a chair', *Analysis*, 79/1: 83–91. DOI: 10.1093/analys/any006
Rorty, R. (1980). 'Pragmatism, Relativism, and Irrationalism', *Proceedings and Addresses of the American Philosophical Association*, 53/6: 717–38. American Philosophical Association.
Sawyer, S. (2020). 'Truth and Objectivity in Conceptual Engineering', *Inquiry: An Interdisciplinary Journal of Philosophy*.
Schindler, S., & Saint-Germier, P. (2023). 'Philosophical Expertise Put to the Test', *Australasian Journal of Philosophy*, 101/3: 592–608. Routledge.
Sękowski, K. (2022). 'Concept Revision, Concept Application and the Role of Intuitions in Gettier Cases', *Episteme*, 1–19.
——. (2023). 'Reference the untouchable. On the limits of revising concepts using the method of cases', *Synthese*, 201/1: 4.
——. (manuscript). 'Overcoming Modal Skepticism via Conceptual Engineering',.
Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.
——. (1963). 'The Philosophy of Rudolf Carnap'. Schilpp P. (ed.) *Carnap's views on constructed systems versus natural languages in analytic philosophy*, pp. 502–18. La Salle: Open Court.
Thomasson, A. L. (2017). 'Metaphysical disputes and metalinguistic negotiation', *Analytic Philosophy*, 58/1: 1–28.
——. (forthcoming). *Rethinking Metaphysics*. Oxford University Press.
Westerblad, O. (2022). 'Deweyan conceptual engineering: reconstruction, concepts, and philosophical inquiry', *Inquiry*, 0/0: 1–24. Routledge.
Wittgenstein, L. (1953). *Philosophical Investigations*. (G. E. Anscombe, Tran.). Basil Blackwell Oxford.