

# Conceptual engineering and conceptual extension in science

Sandy C. Boucher

To cite this article: Sandy C. Boucher (2022): Conceptual engineering and conceptual extension in science, *Inquiry*, DOI: [10.1080/0020174X.2021.2020159](https://doi.org/10.1080/0020174X.2021.2020159)

To link to this article: <https://doi.org/10.1080/0020174X.2021.2020159>



Published online: 24 Jan 2022.



Submit your article to this journal [↗](#)



Article views: 50



View related articles [↗](#)



View Crossmark data [↗](#)



# Conceptual engineering and conceptual extension in science

Sandy C. Boucher 

School of Humanities, Arts and Social Sciences, University of New England, Armidale, Australia

## ABSTRACT

I argue that the Conceptual Ethics and Conceptual Engineering framework, in its pragmatist version as recently defended by Thomasson (2017, 2020), provides a means of articulating and defending the conventionalist interpretation of projects of conceptual extension (e.g. the extended mind, the extended phenotype) in biology and psychology. This promises to be illuminating in both directions: it helps to make sense of, and provides an explicit methodology for, pragmatic conceptual extension in science, while offering further evidence for the value and fruitfulness of the conceptual ethics/engineering framework itself, in particular with respect to conceptual change within science, which has thus-far received little attention in the literature on conceptual ethics/engineering.

**ARTICLE HISTORY** Received 27 April 2021; Accepted 24 September 2021

**KEYWORDS** Conceptual ethics; conceptual engineering; extended mind; extended phenotype; realism; pragmatism

## 1. Introduction

One notable form of theoretical and conceptual innovation in biology and psychology in recent decades has been *conceptual extension* – the broadening or extension of a concept to include more phenomena than it had previously been thought to by commonsense or received theory. So we have had proposals for the extended phenotype, the extended organism, extended inheritance, the extended replicator, and the extended mind, to name five. While there has been considerable discussion about each particular proposed extension (especially the extended phenotype and the extended mind), there has been less attention paid to conceptual extension in science *as such*. This is especially true of the second order questions, i.e. how are we to understand these disputes, in particular, should

**CONTACT** Sandy C. Boucher  [aboucher@une.edu.au](mailto:aboucher@une.edu.au)  School of Humanities, Arts and Social Sciences, University of New England, Armidale, NSW 2351, Australia

© 2022 Informa UK Limited, trading as Taylor & Francis Group

we take a realist or conventionalist attitude towards them? I define realism here as the view that there is an objective fact of the matter about whether or not minds, phenotypes etc. are extended, such that at most one of the traditional view or the extended view is objectively correct.<sup>1</sup> Conventionalism is the denial of this claim.

Dawkins (1982) explicitly adopted a conventionalist view of the extended phenotype issue, insisting that there is no fact of the matter about whether phenotypes are extended. This conventionalist metathesis, I will argue, is also a plausible (though by no means uncontroversial) interpretation of the other extended concept disputes. However, this interpretation faces the serious objection that it makes conceptual choice in science seem arbitrary, subjective, and unresponsive to empirical facts. In this paper, I attempt to respond to this worry by appealing to recent work in the emerging field of conceptual ethics and conceptual engineering. This promises to be illuminating in both directions: it helps to make sense of, and provides an explicit methodology for, conceptual extension in science, while offering further evidence for the value and fruitfulness of the Conceptual Ethics/Engineering framework itself.<sup>2</sup>

## 2. Conceptual extension in science

In this section, I shall summarise four examples of conceptual extension in science, focusing on the meta-level question of realism vs. conventionalism about the proposed reconceptualisations.

### 2.1. *The extended phenotype*

Dawkins famously argued (1982) that the phenotype of an organism – its physical or behavioural characteristics, as distinct from its genotype – extends beyond the physical boundary of the organism as conventionally

---

<sup>1</sup>This definition assumes that pluralistic realism is not an option here. As an anonymous referee pointed out, versions of pluralistic realism have been defended with respect to some issues in the philosophy of biology. However plausible it is in these other cases it does not seem plausible with respect to the conceptual extensions I shall discuss, and has not, as far as I am aware, actually been defended by anybody in this context.

<sup>2</sup>While the conceptual ethics/engineering framework plausibly applies to other types of conceptual revision in science, I focus on conceptual extension in particular for two reasons. The first is that conceptual extensions have been a specific and particularly prominent form of conceptual revision in biology and psychology in recent decades, with the extensions influencing each other in various ways. The second is that it is with respect to conceptual extension, and not (or not to the same extent) other forms of conceptual revision, that the conventionalist/pragmatist position has been explicitly presented and defended. How to make sense of this meta-level view is what this paper is about. Thank you to an anonymous referee for urging me to clarify this point.

defined. The phenotype associated with a gene includes *all* the selectively relevant effects of that gene, even when those effects reach beyond the boundary of the individual organism that carries the gene. His most famous example is the beaver's dam. The dam the beaver builds is caused by the genes of beavers, and has the function of raising the probability that those genes will be represented in future beaver generations. The dam, just like the beaver's tail or teeth, is a physical, downstream effect on the world the beaver's genes have that are used to leverage those genes into the next generation. There is no theoretically motivated basis for distinguishing the dam from the tail or teeth in this regard, and treating the tail and teeth, but not the dam, as part of the phenotype.

Such is Dawkins' first-order claim about phenotypes. But what I want to focus on here is his metathesis about this claim. It is a conventionalist metathesis. Dawkins states explicitly that, in his view, there is *no fact of the matter* about whether organisms possess extended phenotypes. 'The vision of life that I advocate ... is not provably more correct than the orthodox view ... They are equally correct' (1982, 1). 'What I am advocating is not a ... hypothesis which ... can be judged by its predictions. What I am advocating is a point of view, a way of looking at familiar facts and ideas, and a way of asking new questions about them' (1982). This viewpoint 'has made me see animals and their behaviour differently, and I think I understand them better for it. [It] does not constitute a testable hypothesis in itself, but it so far changes the way we see animals and plants that it may cause us to think of testable hypotheses that we would otherwise never have dreamed of' (1982, 1–2). '[I]t is not a factual position I am advocating, but a way of seeing facts ...' (1982, vi). The gene's eye view is 'not demonstrably more correct' than the orthodox view, but it 'is more elegant and parsimonious' (1982, 232). '[T]he biologist should try both ways of thinking, and choose the one he or she prefers' (1982, 7). He quotes Bonner: 'I do not propose to say anything new or original in these lectures. But I am a great believer in saying familiar, well-known things backwards and inside out, hoping that from some new vantage point the old facts will take on a deeper significance' (1982, 7).

How might we understand this claim? The first thing to note is that in a disagreement over whether A counts as an F, two things could be happening. The disputants could be agreeing on what it is to be an F (on the definition of F-ness) and disagreeing about whether A possesses the relevant properties. Or they could be agreeing on the properties of A, but disagreeing over what it is to be an F. Dawkins suggests this dispute is of latter sort. He notes that he and his opponents do not disagree

about the empirical properties of the beaver's dam; Dawkins cites no empirical facts about the beaver's dam that were previously unknown to biologists, or that orthodox theorists would have any reason to deny. What he argues for is a revised definition of 'phenotype' – essentially broadening the definition from 'observable, physical, selectively relevant gene effects manifested within the organism housing the genes' to just 'observable, physical, selectively relevant gene effects'. And his arguments for the revised definition are largely pragmatic/heuristic. Dawkins claims that thinking of beaver dams as a part of the phenotype of beaver genes has some nice consequences: It indicates connections which had been hidden, it allows us to treat like cases alike, it simplifies and unifies our classificatory scheme, it distinguishes what is important in evolution (that genes lever themselves into the next generation by means of phenotypic effects on the world) from what is unimportant (the distinction between gene effects within and beyond an organism's boundary) and so on. Thus while thinking of it as an extended phenotype isn't *mandatory* – the old and the new view are empirically and factually equivalent, Dawkins suggests<sup>3</sup> – it is *fruitful*. This is suggestive of the conventionalist interpretation of conceptual extension (CICE).

The alternative, realist, construal of this debate would treat it as a factual dispute about the world, such that at most one of the two conflicting positions can be correct. The disputants are disagreeing about whether organisms really do have extended phenotypes as a matter of biological fact: if they do Dawkins is right, if they don't the traditional view is right. Insofar as the dispute is about the correct definition of the term or concept 'phenotype', the standard of correctness for the competing proposals is the structure of the world itself: we should opt for the definition which, as far as we can tell, makes for the closest mapping between our terms/concepts and mind-independent reality; that on which our terms/concepts carve nature at its joints.

---

<sup>3</sup>It may be objected that this illegitimately presupposes that the empirical facts exhaust all of the facts. Could it not be the case that the two sides in the dispute agree on the empirical facts of the case, but disagree about the *metaphysical* facts? After all, it might be said, this is the characteristic situation in metaphysical disputes – to give an example, both sides agree about all the empirical facts concerning the chair and the table, but one side holds that they compose a third object, and the other side does not; that is, they disagree about the metaphysical facts of the case. Whatever the general merit of this way of thinking about metaphysics, it is clearly not an acceptable way of characterising the kind of disputes we're looking at. Dawkins is not arguing that there is some deep, empirically inaccessible, metaphysical fact about phenotypes that makes extended phenotype theory right and the traditional view wrong. He is, as I've said, arguing for what he takes to be a more fruitful way of thinking about certain biological systems – an approach that is not to be regarded as 'true' or 'false' at all, and certainly not *metaphysically* true or false. Deep metaphysical reality is the farthest thing from his mind (and the minds of the other people involved in the debate).

As we've seen, this isn't Dawkins' view about the extended phenotype issue, but of course he might not be right about this meta-level question. My aim here is not to defend Dawkins' conventionalist interpretation, but rather to try to understand it, and to appeal to recent work in conceptual ethics/engineering to offer one possible way in which it *might be* defended.

## 2.2. Extended inheritance

Some theorists, including defenders of the Developmental Systems Theory (DST) approach, have argued that genes are not the only thing an organism inherits. Organisms inherit a number of other resources, which can be put to use in the organism's life-cycle, such as nest-site preferences, bird songs, dominance ranks, migration routes, feeding methods, and gut endosymbionts (Gray 1992, 180).

A realist interpretation of this question is possible (and is presumably favoured by at least some of the disputants), but, again, the conventionalist interpretation is naturally suggested by the arguments that have been offered for the shift in perspective.

DSTers point to the theoretical *benefits* of reinterpreting phenomena as extended inheritance. As Godfrey-Smith notes:

In order to categorise some of these cases as real extragenetic inheritance, DST has to engage in some conceptual battles over what counts as inheritance, and over possible mainstream redescrptions of these in standard genetic terms ... At least part of DST's defence of its preferred way of categorising cases of inheritance must be the claim that *if* the DST categorisation ... is used, new patterns will emerge and new insights will result ... Here as elsewhere, a categorisation is to be judged, at least in part, by its empirical fruits. (2001, 286)

As with the extended phenotype, instances of extragenetic inheritance can be redescrbed from the perspective of the orthodox view of inheritance. Any empirical evidence that may be presented to support the notion of extragenetic inheritance can seemingly be interpreted in terms of the orthodox framework. Orthodox theorists agree that the things DSTers point to exist, and agree with the DSTers with respect to the empirical facts regarding them. In particular, they agree that things like bird songs, nest-site preference, feeding methods, and symbiotic microorganisms, exist, and are passed on from one generation to the next. They just disagree about whether there is value in *calling* these things instances of extended inheritance. I pointed out earlier that Dawkins and his opponents agree on all the facts about the beaver's dam, but that Dawkins defends the usefulness of a broader definition

of ‘phenotype’ on which the dam can count as part of the phenotype of the beaver’s genes. Similarly, DSTers and their opponents seem to agree on all the facts about nest sites, endosymbionts, and the like, but DSTers are defending the usefulness of a broader definition of ‘inheritance’ on which these things can count as inheritance.

Thus, as with the extended phenotype issue, the extended inheritance issue is arguably not an empirical matter, it’s a question of which definition of ‘inheritance’ is the most illuminating, fruitful, and helpful. That is, it concerns whether the orthodox view of inheritance, or the DST view, provides the best perspective on the facts, not which one is factually correct. Such, at least, is the CICE interpretation.<sup>4</sup>

### 2.3. *The extended mind*

Clark and Chalmers (1998) famously argued that the mind extends out into the environment. Artefacts that assist cognitive processes count as part of the mind, such that the distinction between what is in the head and what is outside the head is superficial and unimportant.

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain’t (all) in the head! (1998)

The brain, and the pen and paper being used, form a ‘coupled system’: a unified cognitive system in its own right. If I have a phone number written down, or stored in my phone, then *I know it* just as well as if I had remembered it (in the traditional sense). There is no principled reason to distinguish between these two cases. Both are cases of knowledge. No-one memorises phone numbers anymore: we store them in our phones. Chalmers and Clark would say this is just *a different way of remembering them*. My phone is (literally) *part of my mind*.

The extended mind thesis seems closely analogous to the extended phenotype thesis, and the analogy has often been noted. The two sides in the dispute agree on the properties of the relevant cognitive artefacts, but disagree on the question of what it is to be part of the mind, or to be a cognitive process. Defenders of the extended mind are arguing for a

---

<sup>4</sup>CICE can be thought of as a kind of instrumentalism. It asserts that there is no fact of the matter about whether certain phenomena really are instances of extended inheritance, or whether genes and organisms really have extended phenotypes, but there may be pragmatic reasons to treat them *as if* they are/do. It would be interesting to compare this view to other, more familiar examples of instrumentalism in the philosophy of science, but that would take us too far afield.

broader construal of 'mind' or 'cognitive process' that does not imply such things must be inside the head, but such arguments, according to CICE, are pragmatic, as this is not a question of objective fact. Just as there is no fact of the matter about whether a beaver dam counts as part of the phenotype, there is no fact of the matter about whether pen and paper counts as part of the mind, on the CICE view. On one account of 'mind' it does, on another account it does not, and these accounts do not differ with respect to the facts. They are different ways of looking at the same facts. One may be more *useful* than the other in various ways. But that's all we can say.

It is important not to confuse this assessment of the debate with a superficially similar position on which there can fail to be a fact of the matter about extended cognition. Richard Heersmink (2016) argues that the extent to which artefacts should be seen as 'part of' the cognitive systems of the agent depends on how 'integrated' artefact and agent are, and this, in turn, depends on a range of factors. Artefacts that score sufficiently highly on sufficiently many of these criteria count as part of the cognitive system of the agent. Those that score low on enough criteria do not. Others will be intermediate: neither clearly part of, nor clearly not part of, the cognitive system. There will be no fact of the matter about whether these are part of the cognitive system.

This is a similar result to that produced by CICE, but a very different path to it. On the CICE view, if we adopt one conception of mind, certain items will count as part of mind; if we adopt a different conception they will not; and neither of these conceptions is objectively correct. This is not a matter of there being a set of criteria (that we all may agree on), with the items in question satisfying some but not others. On Heersmink's view, we are all meant to agree, as it were, that given the items in question satisfy only some criteria, they are intermediate cases with indeterminate status. Dawkins is not saying there is a set of criteria for being part of phenotype, and beaver dams satisfy only some of these, such that they are intermediate between being part of the phenotype and not being part of the phenotype. He's saying rather that there are two views that disagree on whether they are part of the phenotype, and there's no fact of matter about which is correct.

Heersmink's view is a version of the extended mind that replaces the necessary and sufficient conditions approach with what he calls a multi-dimensional approach. The conventionalist considers the disagreement between the traditional view and the extended view, and says, at the meta-level, that there's no fact about who is right. Heersmink is

defending, in other words, a first-order version of extended mind theory – it just follows from his view that being part of the mind is a matter of degree rather than being all-or-nothing, and that some items are such that there's no fact of matter concerning whether they are part of mind. The conventionalist view on the other hand is a second order pluralism concerning the first order views.<sup>5</sup>

#### **2.4. The extended replicator**

In response to DST, Sterelny, Smith, and Dickison (1996) developed Extended Replicator Theory (ERT), an attempt to incorporate a number of the insights of DST into a modified version of the gene's eye view. ERTers defend the replicator-interactor distinction (which DSTers reject), and insist against DST on the fundamental importance of replicators, understood as 'elements of the developmental matrix that have been selected to produce [parent/offspring] similarities' (Sterelny and Griffiths 1999, 109). But they concede that genes are not the only things that count as replicators on this definition (thus they propose replacing the name 'selfish gene theory' with 'selfish replicator theory'). Many other developmental resources emerge as replicators, on this view.<sup>6</sup>

As with the previous debates, the debate over ERT largely focused on pragmatic/heuristic issues, such that CICE seems a natural fit. Just as, according to Dawkins, it is arteficial to distinguish between phenotypic effects within, and outside, an organism's boundary, Sterelny et al. suggest it is arteficial to only treat genes as replicators, and exclude elements of the environment, such as burrows and nests, which perform many of the same functions in an organism's life-cycle. The proposed recategorisation is useful, they argue, because it allows us to treat like cases alike, to ignore distinctions that are of no theoretical import, and to focus attention on issues and problems that may be able to be formulated, but would be harder to notice – would be less 'in your face' – from other perspectives (1996, 395).

---

<sup>5</sup>Compare Mishler and Brandon (1987) on the ontology of species. There is a range of criteria for being an individual. Species probably score highly on some, but less highly on others. Different weightings of the criteria are possible. So species may be neither clearly individuals, nor clearly not individuals. This contrasts with the conventionalist view, according to which the view that species are individuals, and the view that they are natural kinds, are maximally adequate and equally correct – there is no fact of the matter. In other work I draw a similar distinction between realist-vagueness and conventionalist-pluralism with respect to the units of selection and the existence of interactors (Boucher 2020).

<sup>6</sup>Dawkins agrees in principle with this shift, but thinks that in practice fewer items are likely to satisfy the fairly stringent criteria for counting as a genuine replicator than Sterelny *et al* suppose. Certainly not nests and burrows, for instance (Dawkins 2004).

Sterelny et al., in making the case for the surprising claim that rabbit burrows should count as replicators, do not seem to advance any novel empirical claims about burrows or burrowers; it is plausible that ERTers and DSTers agree on all the empirical facts about the burrower/burrow systems under consideration. The issue between them arguably concerns the most useful perspective on those facts, and the most useful way of describing them. The ERT perspective, according to Sterelny et al, is fruitful but not mandatory: the relevant phenomena can be accounted for, albeit less perspicuously (in their view), from the perspective of DST, or orthodox selfish gene theory. This is strongly suggestive of the CICE view.

I have discussed four examples of proposed conceptual extensions in biology and psychology.<sup>7</sup> The connections between the different extensions have frequently been noted. The extended phenotype has often been used to argue for the extended mind, for instance (see Schulz 2013).<sup>8</sup> But the general applicability of the conventionalist-pragmatist interpretation Dawkins gave to the extended phenotype thesis tends to be overlooked or flatly denied. Schulz (2013, 246), for instance, claims that conventionalism is true of the extended phenotype but not of the extended mind.<sup>9</sup> We have seen however that CICE also fits quite naturally with the other proposed extensions, given the pragmatic/heuristic character of the main arguments offered for them (whether or not the conventionalist metathesis would be accepted by their advocates). Certainly, CICE is not the only possible interpretation here: the realist metathesis is a possible view. My aim has not been to offer a full defence of CICE with respect to any, let alone all, of the extensions.<sup>10</sup> I want rather to place it on the table as a view worth taking seriously, and that may have something to recommend it.<sup>11</sup>

---

<sup>7</sup>I do not have the space to discuss the 'extended organism' thesis of Turner (2000), but my sense is that much the same points can be made concerning the applicability of CICE to this view, which clearly has much in common with the extended phenotype thesis. (Turner discusses the differences between the two, and the reasons why, in his view, the extended organism thesis (which is centred on extended *physiology*) is preferable to the extended phenotype thesis (which is centred on extended *gene* action and influence) in his 2004).

<sup>8</sup>Schulz himself resists this argument.

<sup>9</sup>... in the [extended mind] case, this is a mind and language-independent fact about the world; in the [extended phenotype] case, this is a mind and language-dependent fact about what we find useful in theorising about the world' (Schulz 2013, 246).

<sup>10</sup>And needless to say one may endorse CICE with respect to some of the extensions I've discussed, while rejecting it for others.

<sup>11</sup>It is also worth noting that CICE is not a *sui generis* anomaly, but has much in common with other popular pluralist and pragmatist views in biology and the philosophy of biology, such as model pluralism about the units of selection (Sterelny and Kitcher 1988; Kitcher, Sterelny, and Waters 1990; Godfrey-Smith and Kerr 2002), and conventionalism about the ontology of species (Gould 2002; Kitcher 1984; see Boucher 2017 for discussion).

But some may think it has very little to recommend it. Does it not make conceptual choice in science arbitrary and unconstrained by the world? Does it not mean we can adopt whatever concepts and theories we like, so long as they are useful for us? Does it not, in short, make questions about the natural world and the way it is structured depend on features of *us* in a way many would find objectionable? In the remainder of the paper, I try to defuse these worries by appealing to recent work in the budding field of conceptual ethics and conceptual engineering. I argue that the pragmatist interpretation of conceptual engineering (PICE) recently defended by Thomasson (2017, 2020) offers a promising way of articulating and making sense of CICE, and, above all, of responding to the worries about the arbitrariness and lack of responsiveness to worldly facts that CICE might seem to invite. It is not my aim to defend CICE, or PICE, but rather to argue that CICE *needs* PICE, if it is to avoid these worries.

### 3. The pragmatic interpretation of conceptual engineering

Conceptual Ethics/Engineering is an understanding of and approach to conceptual work that is arguably not new in philosophy, but was first explicitly articulated in its contemporary form by Burgess and Plunkett (2013), and has become increasingly popular within philosophy (see the major collection of papers Burgess, Cappelen, and Plunkett 2020). It is in the first instance a metaphilosophical conception, but is in principle applicable to any kind of conceptual work (including, I am hoping to show, in science). As the ‘ethics’ in the name suggests, the focus is on *normative* conceptual projects (in contrast to the primarily descriptive projects of traditional conceptual analysis): questions concerning which concepts, or representational devices more generally, we should have, which we should revise, which we should critique and which we should discard.

In this section, I shall summarise Thomasson’s pragmatic interpretation of conceptual ethics/engineering (PICE). Her focus is on the conceptual engineering of metaphysical concepts, i.e. conceptual engineering as a way of doing metaphysics, but we will see that what she says applies quite generally to conceptual engineering in philosophy and science (which is, anyway, what she intends).

Thomasson notes that on the conceptual ethics/engineering approach, metaphysics is not, or should not be,<sup>12</sup> about ‘discovering worldly facts’,

---

<sup>12</sup>This could be a descriptive claim about what metaphysicians actually do, and have done in the past, or a normative claim about what they ought to be doing, or both (Thomasson 2020, 435–436).

but rather involves descriptive and normative conceptual work (Thomasson 2020, 435). It is, or should be, concerned with which concepts, and which intensions and extensions of concepts, we ought to adopt, not (at least, not directly) with what the world is like. There are however two quite different approaches to conceptual ethics/engineering in metaphysics. On the first, ‘metaphysical’ approach, the correctness of our conceptual choices is dictated largely or entirely by the extent to which they equip us with concepts that correspond to the mind-independent metaphysical structure of the world (2020, 437–438). The aim is to arrive at concepts that carve nature at its joints (e.g. Sider 2009). On the second, deflationary approach, our conceptual choices are not answerable primarily to the metaphysical structure of the world, but rather to the various interests and purposes we have in developing, revising, or discarding certain concepts. Thus this is an emphatically pragmatic, and interest- and value-relative approach to conceptual ethics/engineering, congenial to the kind of Carnapian metametaphysical deflationism/pragmatism that Thomasson advocates, and which I shall discuss in Section 5.

Thomasson’s aim is to offer a pragmatic method for conceptual ethics/engineering, based on an analysis of the functions of concepts (2020, 440–441). Much as pragmatist anti-representationalism in the philosophy of language (associated with philosophers such as Huw Price and Simon Blackburn) proceeds by analysing the functions of regions of discourse, i.e. what that discourse enables us to *do* in a practical sense (in contrast to analysing the discourse’s semantics in terms of representational properties such as truth and reference), PICE recommends an approach to normative conceptual work that begins by identifying the function(s) of the relevant concept – what possessing and deploying the concept enables us to do that we couldn’t do, or do as well, without it.<sup>13</sup> On the basis of such an analysis, we may proceed to ‘engineer’ the concept: perhaps replacing our current variant of the concept with a superior variant that we judge will better satisfy the concept’s function(s) than our current variant; perhaps leaving the concept as it is; or perhaps even discarding the concept, should we deem that the function(s) it is serving is not one that we judge to be morally or intellectually legitimate, for

---

<sup>13</sup>There are questions that may be (and have been) raised about this appeal to function. Some have objected to the very idea that concepts have functions (Cappelen 2018). Thomasson concedes that while it may be implausible to suppose that many concepts have functions that are conferred on them intentionally, the friend of PICE can perhaps appeal to standard accounts of biofunctions to offer a more realistic picture of the way in which concepts can acquire functions. These are important questions for PICE but we need not go into them in detail as it is not my aim to offer a full-fledged defence of PICE, but rather to argue that something like PICE appears to be needed to make sense of CICE.

instance the pernicious function(s) served by traditional racial concepts (2020, 454).<sup>14</sup>

As well as suggesting a methodology for conceptual ethics/engineering – one that is congenial to a pragmatist/deflationist view of the relevant subject matter – this appeal to function, Thomasson argues, gives us a way of responding to a general worry about conceptual revision: that in changing the intention or extension of a concept we are in effect *replacing*, rather than revising or extending, our concepts, and are thereby ‘changing the subject’. While there may be different ways of individuating concepts, one seemingly legitimate way of doing so is by reference to the concept’s function: our grounds for thinking that it is one and the same concept that is persisting through shifts in meaning and/or reference is that there is continuity of function (2020, 443).<sup>15</sup>

Thomasson suggests that even on the deflationary, pragmatic view she favours, it will be the case that our conceptual choices will be constrained by worldly factors, and so will not be entirely subjective or arbitrary. This is most obviously true, she suggests, for concepts that feature in natural-scientific laws and theories, for it is their function to contribute to our predictive and explanatory purposes by carving nature at the joints, and picking out natural kinds (2020, 450).

But it is not only the engineering of scientific concepts that must be sensitive to the structure of the mind-independent world:

Even where the function of a concept is not predictive/explanatory ... the pragmatic approach can nonetheless allow that our choices in engineering the concept are not merely arbitrary or subjective. For such conceptual choices also must be responsive to worldly factors ... In conceptual engineering no less than civil engineering, the question of which design (of concept or bridge) will best fulfil the relevant function, given the requirements, does not leave room for a merely ‘arbitrary’ or power-driven answer, and must be addressed while being sensitive to a variety of worldly factors. (2020, 452–453)

Conceptual choices, she points out, are also constrained by worldly factors in that they need to take account of the concept’s connections to other concepts, beliefs, practices, etc., that constrain any modifications we may make to the concept (2020, 453).

---

<sup>14</sup>This procedure – conceptual revision on the basis of a persisting function or set of functions – Thomasson refers to as conceptual *engineering*. But we may also devise new functions for our concepts; she calls this conceptual *ethics*. (This is one way, but not the only way, these terms have been distinguished (Burgess, Cappelen, and Plunkett 2020)). We shall be concerned in this paper with the former only, so shall talk simply of conceptual engineering from now on.

<sup>15</sup>See Prinz (2018) for a more detailed argument for the view that individuating concepts via their function furnishes an adequate reply to the changing-the-subject, or ‘discontinuity’ objection.

Take her marriage example. While there is a clear sense in which the definition of marriage is to a significant extent ‘up to us’, and there is not an objective fact of the matter about whether same-sex couples may be married the way there is a fact of the matter about whether, say, whales are mammals, it does not follow that the definition of marriage is arbitrary, subjective, or insensitive to worldly facts. No-one would, or should, regard the same-sex marriage debate as a trivial verbal dispute – ‘well if you define marriage one way same-sex couples can be married, if you define it another way they can’t, so pick your favoured definition and cease arguing about words’.<sup>16</sup> If we focus rather on the question of which variant of the marriage concept best fulfils the intended function or aim of the marriage concept, the importance of worldly factors becomes obvious, as the answer to this question manifestly depends on a variety of empirical facts about human relationships, institutions, norms and practices, and is far removed from the free and unconstrained adoption of a certain choice of words or meanings.

## 4. Conceptual extension in science as pragmatic conceptual engineering

### 4.1. *The metaphysical vs. the pragmatic approach*

It is presumably uncontroversial that the debates around the proposed conceptual extensions I have focussed on can be understood as examples of conceptual engineering.<sup>17</sup> But, if we apply Thomasson’s distinction between the metaphysical and pragmatic/deflationary approaches to conceptual engineering, we can see there are two quite different ways of construing what is going on in these debates. According to the metaphysical approach, the aim of such conceptual work is to identify and recommend concepts or interpretations of concepts that correspond to the mind-independent structure of the natural world. This is the realist interpretation of conceptual extension I mentioned in section 2. According to this view the

---

<sup>16</sup>Dawkins fell into precisely this error with some comments he made in 2015 about trans people, for which he rightly came under fire: ‘Is trans woman a woman? Purely semantic. If you define by chromosomes, no. If by self-identification, yes. I call her “she” out of courtesy’ (Flood 2021). He was right that this is in a sense a question of definition, but quite wrong to think that this implies the issue of gender is just a trivial matter of ‘semantics’. It is very important, many think (me included), that trans people be thought of as belonging to the gender with which they identify. Whether this is because this gender concept better ‘carves at the joints’ than the traditional concept that ties gender to biological sex, or whether this too invites treatment in pragmatic terms, is a further question I will not delve into. For influential discussions of the ethics/engineering of the concept of gender see Haslanger (2000, 2012).

<sup>17</sup>Cappelen (2018, 10–11) and Simion (2018, 917–918) each briefly mentions the extended mind as an example of conceptual engineering, but neither develops this thought in any detail.

standard by which we assess our phenotype concept, for example, is solely the issue of whether it corresponds to the mind-independent facts about the phenotypes of organisms. In short, the extended phenotype concept ought to be adopted just in case organisms really do have extended phenotypes. Clearly, CICE is sharply at odds with this construal of the question. CICE has it that the debates about conceptual extensions conform (or *should* conform) rather to PICE, according to which our conceptual choices are not answerable primarily to the metaphysical structure of the world, but rather to the various interests and purposes we have in developing, revising, or discarding certain concepts.<sup>18</sup> The question, according to CICE, is not ‘which concept is objectively correct?’, in the sense of being true to the objective facts, but rather ‘which concept is most useful?’, where ‘most useful’ is to be cashed out in terms of which concept best serves the function(s) that concepts of that type are intended or required to serve. Debates about conceptual extension in science, on this view, are indeed closely analogous to the debate about revising the definition of marriage. While the functions of the relevant concepts are clearly quite different – facilitating scientific goals in the one case, facilitating social, ethical and political goals in the other – the projects and methodology are otherwise very similar. We are considering whether to broaden a concept to include within its extension objects, phenomena or relationships not included in the traditional concept. If we have a compelling reason to do so, it is because the new, broader variant of the concept is, we judge, better able to achieve the goals and purposes the concept is/was intended to achieve, and there is nothing other than a misguided attachment to tradition and orthodoxy preventing us from doing so. In neither case do we believe that there is a fact out there in the world determining which definition of ‘marriage’ or ‘phenotype’, is correct; but neither do we think that one is free to adopt whatever variant of the concept one likes, that ‘anything goes’.

Below I will discuss what such a functional analysis might look like for the conceptual extensions in science. For now, I note the potential benefit of this approach for making CICE attractive and defensible. For CICE to be plausible it must steer a middle path between two positions. PICE shows us how this can be done. Conceptual revision or negotiation on this view is about which concepts (or meanings of concepts) are useful to us given our interests and purposes, not about how the world is – not, as

---

<sup>18</sup>In fact, as we’ve seen, Thomasson appears to recommend, for scientific concepts, something quite close to the metaphysical, joint-carving approach, although preferring the pragmatic approach for non-scientific concepts. I discuss this below.

Thomasson says, about discovering worldly facts. But neither is it simply a matter of adopting whatever concepts (or meanings of concepts) we like, or engaging in trivial disputes about definitions. To put it another way, it shows how a pragmatist, pluralist or deflationist position about certain metaphysical or scientific disputes can avoid the charge of subjectivity, arbitrariness, or triviality. This makes it particularly well-suited to articulating and defending CICE.

After all, on the face of it, CICE may seem vulnerable to the charge that it has the implication the debates are uninteresting or are trivial verbal disputes. And there is a sense in which they *are* verbal disputes about definitions, on this view: is the beaver tail part of the beaver phenotype? Well if we define 'phenotype' one way the answer is yes, if we define it another way, the answer is no. This is a disagreement about how to define our phenotype concept, so in that respect it is a kind of verbal dispute. The disputants agree about the facts of the case and are arguing about the meanings of words and concepts. But this is only the very beginning of wisdom on the matter. On the conceptual engineering view such disputes about how to define concepts, which versions of concepts to adopt etc., are much less trivial than traditional verbal disputes. To label a dispute as a verbal dispute is nearly always a way of saying that we should stop engaging in it (see Chalmers 2011 for some exceptions). To label something a conceptual engineering dispute is emphatically not to say this. Conceptual engineering disputes (e.g. should we revise a concept and if so how? What version of a concept should we adopt? Should we retain the concept at all? etc.) are not *mere* verbal disputes. The latter are almost by definition non-substantive and uninteresting, while the former are meant to be substantive, interesting and important, even if they are not, or not directly, about the world *per se*.

Thus, Chalmers and Clark note that

in seeing cognition as extended one is not merely making a terminological decision; it makes a significant difference to the methodology of scientific investigation. In effect, explanatory methods that might once have been thought appropriate only for the analysis of 'inner' processes are now being adapted for the study of the outer, and there is promise that our understanding of cognition will become richer for it. (1998, 13)

Notice that they are here rejecting the shallow-verbal-dispute interpretation (not just a 'terminological decision') without, at least in this passage, committing themselves to the view that the standard of correctness for the extended mind view is worldly facts about whether minds

'really are' extended: they stress rather the pragmatic and theoretical *benefits* of the extended concept (new, fruitful explanatory methods will result, our understanding of cognition will become richer; compare Godfrey-Smith on extended inheritance needing to be judged on its empirical fruits). I am not suggesting they would necessarily endorse CICE, but it is notable that the way they present the issue here is quite congenial to it.<sup>19</sup>

#### 4.2. Joint-carving functions?

As we've seen, Thomasson suggests that even on the deflationist view she recommends, our conceptual choices do need to be constrained by the world and its structure. PICE 'enables us to give due respect to the idea that the world is structured into natural kinds to which our concepts should be responsive, and to the idea that some concepts seem 'objectively' better than others' (2020, 450). Some concepts serve a joint-carving function, such as terms for natural properties. 'These are those that will figure in laws and in our natural-scientific theories' (2020). In other words, we have here the idea that scientific concepts, unlike the sort of philosophical concepts that are her primary focus, have an explanatory-predictive joint-carving function of picking out natural kinds and corresponding to objective structure in the world.

With respect to the conceptual extensions, as I have argued, even on the conventionalist interpretation proposals for conceptual revision must be constrained by and responsive to the structure of the world, empirical factors, etc. Not just any set of objects and properties can be legitimately included in the extension of 'phenotype', 'mind', etc., if these concepts are to play their allotted roles in science, and these constraints will partly be set by the nature of the mind-independent world. On the other hand, talk of joint-carving, and natural kinds, does not fit well with CICE. If the extended phenotype concept better carves at the joints than the traditional phenotype concept, it would seem that extended phenotype theory is more objectively correct than the traditional view, i.e. better fits the facts, and this is precisely what Dawkins rejects. The only possible conclusion is that according to CICE,

---

<sup>19</sup>They appear to reject CICE, and endorse the realist interpretation, when they suggest that the broader concept of belief they favour 'picks out something more akin to a natural kind' (1998, 14). On the other hand Clark comes quite close to endorsing CICE for the extended mind in his (2011), where he notes that the extended phenotype thesis is a 'perspective' that is 'not compulsory', but whose 'virtues lie ... in the different ways of seeing familiar phenomena ...' it engenders (218), and suggests that we understand the extended mind thesis in the same way.

these concepts do not have joint-carving functions, and do not have the function of picking out natural kinds.<sup>20</sup>

This is not to deny that some scientific concepts may carve at the joints and pick out natural kinds (or at least aim to). CICE need not apply across the board, to yield a general conventionalism about all scientific concepts and taxonomies, which, while a possible view, is widely thought to be implausible. But if Dawkins is right there are some scientific concepts that do not have a joint-carving function, and do not pick out natural kinds.

Thus, on the CICE view, Thomasson's dichotomy of natural-scientific concepts that have joint-carving functions and non-scientific concepts that do not needs to be modified. If CICE is correct, some natural-scientific concepts may have joint-carving functions, but others do not. Of course, we may wonder why this is the case, which scientific concepts have joint-carving functions and which don't, and how we can tell the difference, but these are questions we can set aside for now: all that CICE needs to be committed to is that at least some scientific concepts do not have joint-carving functions, and this is consistent with the claim that some scientific concepts may have joint-carving functions, while also being consistent with the claim that none do.

One may wonder, indeed, whether the notion of normative conceptual revision on the basis of the identification of joint-carving functions is really in the spirit of PICE. After all, our concepts would, in the joint-carving model, seemingly be chosen and engineered primarily on the basis of how the world is, not on pragmatic grounds of human interests and utility, and that would seem to return us to the 'heavyweight metaphysical' construal of conceptual engineering, at least with respect to the sort of natural-scientific concepts and theories for which she thinks

---

<sup>20</sup>Some may think that only concepts with joint-carving functions can figure in predictions and explanations. Thus the defender of CICE either needs to deny that figuring in predictions and explanations is among the main functions of the relevant concepts, or insist that concepts can contribute to predictions and explanations despite not being joint-carving. One way in which these options could be combined is as follows. One could deny that the concepts figure in predictions. This may indeed follow from Dawkins' claim that the orthodox and extended phenotype concepts are empirically, i.e. predictively equivalent. (At the very least this equivalence indicates that the extended concept cannot be defended on the grounds of its possessing greater predictive power than the orthodox concept.) One may though retain the idea that the concepts may figure in explanations (despite not being joint-carving), so long as explanation and explanatory power are interpreted in pragmatic terms a la Van Fraassen (1980). Perhaps the extended view and the orthodox view are empirically equivalent, but the former is superior when it comes to the 'theoretical', pragmatic virtues of simplicity, elegance, explanatory power and so on. It is a short step from this to the claim that the phenotype concept can have explanatory functions without having joint-carving functions. Explanation understood in the pragmatic sense would seem not to require joint-carving concepts. It seems to me that this is a defensible position. But if it is thought not to be, that is a problem for defenders of CICE, not a problem for the argument of my paper. Thank you to an anonymous referee for pressing me on this point.

the joint-carving model is appropriate. Thomasson could be accused here of wanting to have her cake and eat it, or of balking at embracing the full implications of her pragmatist approach, and letting the metaphysical approach in at the back-door.<sup>21</sup> In any case, even if we grant that conceptual engineering in science guided by joint-carving functions has a legitimate place in the pragmatist model, what matters for our purposes is that the conceptual extensions, according to CICE, are not examples of this.

We also saw however that according to Thomasson, even concepts whose function is not joint-carving and do not aim to pick out natural kinds, cannot be adopted or revised just as we like, in a manner that is insensitive to empirical considerations, or the structure of the world. With philosophical, ethical, social and mathematical concepts that do not have joint-carving as their function, there are nonetheless a whole host of objective, worldly factors that constrain our adoption or revision of the relevant concepts, and make some such adoptions or revisions superior (in the appropriately pragmatic sense) to others. On the CICE view, it is natural to assimilate conceptual extensions of the sort we have discussed to projects of conceptual engineering for concepts that do not have joint-carving functions, and if Thomasson is correct, this need not force us into the view that conceptual extension is unconstrained by empirical facts about the world. Just as determining which variant of the marriage concept best fulfils the intended function or aim of the marriage concept requires attending to worldly facts about human relationships, institutions, norms and practices, determining which variant of the phenotype, mind, or organism concept best fulfils the intended function of the phenotype, mind, or organism concept requires attending to a variety of worldly facts in the realms of biology and psychology.<sup>22</sup> These facts fill the pages of Dawkins (1982), Clark (2011), and Turner (2000), to give just a few examples.

---

<sup>21</sup>She would claim she is not doing this, as she is careful to distinguish between the *empirical* considerations guiding conceptual work in this domain (whether certain concepts assist us in prediction and explanation) which the pragmatist can accept, from the *metaphysical* considerations about structure that the pragmatist rejects (2020, 450–451). I find this distinction somewhat obscure however. If concepts have joint-carving functions, surely their function is to track objective reality and map onto really existing natural kind structure, whether we think of the latter, and our access to it, as ‘empirical’ or ‘metaphysical’.

<sup>22</sup>As Plunkett notes, ‘Suppose one argued ... that an important part of communication among biologists involves metalinguistic negotiation ... Would that mean that there aren’t facts about animals and their behavior to investigate, and that all biological argument is just about normative issues about word and concept choices? Clearly not.’ (Plunkett, quoted in Thomasson 2020, 438). This also supports the claim, that defenders of CICE should endorse, that from the fact that disputes over conceptual extension are not about discovering worldly facts, it doesn’t follow that nothing in biology or psychology is about discovering worldly facts.

So, in summary, we have here a principled way for the defender of CICE to uphold a conventionalist/pragmatist position, while resisting the charge that her position makes the relevant conceptual choices subjective, arbitrary, or insensitive to the world, or makes debates concerning such choices trivial verbal disputes. Of course, one may not accept Thomasson's analysis; I do not mean to suggest her account is uncontroversially correct, and it is not my aim to offer a full defence of it here, though I find it persuasive, if only in the limited sense of being agreeable to those already sympathetic, as I am, to the broadly Carnapian and deflationist metametaphysical approach she favours (see below). My aim here has simply been to show that something like Thomasson's account will be required if CICE is to be at all plausible. CICE, in other words, *needs* PICE.

### ***4.3. The methodology of pragmatic conceptual extension***

We have seen that on the PICE approach, the methodology for normative conceptual work begins with identifying the function(s) of the relevant concept, then proceeds to retain, revise, extend, replace, or discard the concept or a particular version of it on the basis of a judgment of the extent to which the concept, or a version of it, fulfils these functions. In the case of conceptual extension in science, we would begin by identifying the function(s) of the concept 'phenotype', 'inheritance', etc.; then consider whether extended versions of these concepts would fulfil these function(s) better than the orthodox versions; if so, we should adopt the extended versions; if not, not.

Arguments that have actually been offered for the various extensions have not explicitly taken this form, but my suggestion is that this is in fact the implicit methodology at work. In other words, PICE arguably formalises, or makes explicit, what is actually going on in debates about conceptual extension in science.

But is it even plausible to suppose scientific concepts have functions? Brigandt (2010) has argued persuasively that not only does this make sense, but it is only on the supposition that they do have functions that one can represent as rational revisions in the meaning and reference of such concepts. It is well known that the meaning of scientific concepts (what Brigandt calls their inferential role) can change over time, and across paradigm shifts, as scientific theories and beliefs change. This has given rise to incommensurability concerns (e.g. Feyerabend 1975; Kuhn 1996), to which realists have responded by emphasising continuity (or partial continuity) of reference that may accompany shifts in meaning

or changes in beliefs about the referent (Field 1973; Fine 1975; Devitt 1979). But, firstly, this overlooks the fact that sometimes the reference of scientific terms or concepts can also shift (the conceptual extensions I am focussing on are clearly an example of this), and, secondly, this doesn't address the normative question of when such shifts in meaning count as *rational*. Thus, Brigandt argues, we need to invoke a third semantic dimension of scientific terms/concepts, their epistemic goals, which can be identified with functions in our sense.<sup>23</sup> These goals are determined by the way a scientific community, rather than any particular scientist, uses a term/concept: 'Within an overall field of science, there may be variation regarding the epistemic goal(s) pursued by a term's use, and thereby this component of a concept can also be subject to change in the course of history' (2010, 23). Change of meaning or reference of a scientific term/concept counts as rational relative to, and only relative to, the epistemic goal/function of the term/concept.<sup>24</sup>

Once significant novel empirical beliefs become available, the inferential role may change (possibly leading to a revision of the concept's definition), and this change in inferential role is rational provided that the revised inferential role (of the modified variant of the concept) meets the epistemic goal to a higher degree than the previous inferential role (of the original variant of the concept). (2010, 24)<sup>25</sup>

So consider the arguments for the conceptual extensions, citing as they do the purported benefits of the new variants of the concepts. On Brigandt's view, it is not just that these arguments will be more compelling, and better grounded, if they are supported by appeal to the goal(s) or function(s) of the concept; he makes the stronger claim that this is the *only way* such arguments can have any force, and the only way the recommended conceptual change can be regarded as rational. Appeal to function is not just advisable, in other words, it is mandatory. Whether or not he is correct in this claim, he has made a strong case at

---

<sup>23</sup>He does not claim that 'every scientific concept can be assigned a unique epistemic goal', but we can assign such goals to 'those central concepts (at least in biology) that underwent conceptual change, such that this semantic change can be explained in these terms' (Brigandt 2010, 23). It is not clear why he thinks we can only be confident that those concepts that underwent conceptual change have unique epistemic goals.

<sup>24</sup>This assumes sameness of epistemic goal. But Brigandt acknowledges that sometimes epistemic goals themselves can change, just as Thomasson notes that sometimes the function of a concept can change. For the conceptual extensions, I am assuming continuity of epistemic goal or function, which seems reasonable given the character of the disputes.

<sup>25</sup>Here and elsewhere Brigandt presents the meets-its-epistemic-goals condition as a sufficient condition on the rationality of meaning or reference change, but I think it is evident from the rest of his discussion that he intends it to be a necessary condition as well.

least that an effective, perhaps the most effective, means of rationally revising the meaning and/or reference of a scientific term/concept is via consideration of how best to satisfy the epistemic goals it is taken to have.<sup>26</sup>

How then can we identify the function/epistemic goal(s) of the phenotype concept, the inheritance concept, etc., by reference to which extensions of these concepts (entailing shifts in their meaning/inferential role and reference), according to their proponents, count as rational? In what follows I will be assuming continuity of function/epistemic goal for the relevant concepts, and thus sameness of concept between the orthodox and the extended conceptions. It may not matter enormously whether the extended concepts are construed as new interpretations of the same concept, or as entirely new concepts. And Thomasson points out that it is implausible to suppose there will always be a fact of the matter about such questions anyway (2020, 442; Brigandt makes a similar point, 2010, 25). But we saw that there are advantages to the new-interpretation-of-the-same-concept construal grounded in the functional individuation of concepts, inasmuch as it furnishes a response to the ‘changing the subject’ objection to projects of conceptual engineering. This gives us a *motivation* for treating the conceptual extensions as proposing new variants of old concepts, rather than recommending entirely new concepts. The *evidence* for this, such as it is, is the observation that the arguments for the extensions, as I understand them, are claiming that the extended concepts are broadly speaking doing what the traditional concept was doing (or trying to do), only better, rather than doing entirely new things.

There are (at least) two levels at which we may identify the functions of scientific concepts. At the most general level, we may link the function of scientific concepts to the most general aims of science, i.e. their function is to facilitate scientific explanation, prediction, classification, and the construction of true and/or empirically adequate theories and models. But scientific concepts will also have more specific functions, corresponding to different ways of serving the overarching function: some, perhaps, the function of carving at joints and picking out natural kinds; others,

---

<sup>26</sup>Brigandt calls these *epistemic* goals, but are not the functions of the relevant concepts according to CICE *pragmatic*? In fact there is no conflict here. Even if the functions of ‘phenotype’, etc., according to CICE, are broadly pragmatic in that they consist in something other than joint-carving and are not directly related to delivering objective truth, the concepts also have the function, as we shall see, of contributing indirectly to the most general aims of science, which are largely epistemic. Their specific functions/goals may be pragmatic, in other words, but their more general functions/goals are largely epistemic.

non-joint-carving, more ‘pragmatic’ functions (e.g. facilitating unity and simplicity in our theories and taxonomies).<sup>27</sup>

Consider an analogy. The most general function of our cognitive systems (on an evolutionary interpretation) is to contribute to our inclusive fitness. But there are different ways of achieving this. One way is by tracking the truth, so some cognitive systems will have truth-tracking as their more specific function. But others may contribute to fitness in other, non-truth-tracking ways – perhaps moral or religious beliefs are like this (Griffiths and Wilkins 2015; Boucher 2021). So some cognitive systems, or beliefs, contribute to fitness by tracking the truth; others contribute to fitness by doing other things; similarly some scientific concepts (may) contribute to the aims of science by carving at the joints and tracking objective facts and properties; but others may contribute to these aims in more pragmatic ways.

Clues as to what these pragmatic functions might be can be found by examining the arguments that have been offered for the extensions that highlight their purported benefits. Clearly inferences from the purported benefits of a particular interpretation, or revision, of a concept, to the function of the concept, need to be treated with care. From the fact that one of the benefits of a particular conceptual choice with respect to some concept *C* is *X*, it doesn’t automatically follow that one of the functions of *C* is the promotion of *X*. The benefit might be secondary or fortuitous and unrelated to the original function(s) of *C*. This distinction between functions and fortuitous benefits is a familiar one. For a benefit of *X* to count as a, or the, function of *X*, that benefit needs to be causally implicated in the origin and/or maintenance of *X*.

But as long as we treat such inferences with the caution they require, it seems to me not illegitimate to claim that we may glean some insight into the function of a concept by examining purported benefits of revisions with respect to it. This will be defeasible evidence for the function. At the very least we may assert with some confidence that in most cases arguments in favour of a particular conceptual revision that point to purported benefits of the revision will make more sense, and be more

---

<sup>27</sup>An anonymous referee pointed out that some may regard joint-carving as one of the fundamental aims of science, not a subsidiary one. I don’t need to enter into that debate, as I am happy to accept, for the sake of argument, that joint-carving may be one of the fundamental aims of science and scientific concepts. This doesn’t affect my argument. CICE can accept, as I noted, that some scientific concepts may have joint-carving functions. I have suggested that these functions serve more fundamental aims (that don’t include joint-carving as such), but some may deny this, insisting that the joint-carving functions are fundamental. Even if that is true, it must still be the case, according to CICE, that some scientific concepts do not have joint-carving as their function, either in the fundamental or subsidiary sense. Even if joint-carving is *one* fundamental aim of science, it need not be the only, or even the main, one.

compelling, if they are underwritten by facts about the function of the concept, in such a way as to connect the purported benefits more or less directly to that function. (As we have seen Brigandt argues that appeal to function/epistemic goal(s) is in fact indispensable for conceptual change in science.) So the argument that we should adopt the extended phenotype concept because doing so will simplify and unify our taxonomy of the traits of organisms, will be more compelling if the function of the phenotype concept includes simplifying and unifying our taxonomy of the traits of organisms. The idea behind articulating CICE in terms of PICE is that the pragmatic benefits of the conceptual extensions will be underwritten by, and justified in terms of, pragmatic functions of the relevant concepts (just as the claim that the benefit of some proposed conceptual extension is that it better carves at the joints than the traditional concept will be stronger if underwritten by the claim that the relevant concept has joint-carving as its primary or sole function).

We have seen that proponents of the extended concepts have claimed that they:

- Suggest new questions and fruitful lines of inquiry; e.g. they cause us to ‘think of testable hypotheses that we would otherwise never have dreamed of’ (Dawkins 1982).
- Allow us to develop explanatory methods that may enrich our understanding of the relevant phenomena (Clark and Chalmers 1998), helping us to uncover new patterns, and generate new insights (see Godfrey-Smith 2001).
- Simplify and unify our classificatory schemes (Sterelny, Smith, and Dickison 1996).
- Focus our attention on important issues and problems, while encouraging us to ignore distinctions that are of no theoretical importance (Dawkins 1982; Sterelny, Smith, and Dickison 1996; Clark and Chalmers 1998).

While no doubt some or all of these could be interpreted in the light of the realist/metaphysical version of conceptual engineering – i.e. the extended concepts have these benefits precisely because they better cut at the joints than the traditional concepts (i.e. because minds, phenotypes etc. *really are* extended<sup>28</sup>) – this interpretation seems far from

---

<sup>28</sup>Joint-carving functions and pragmatic functions do not necessarily exclude one another: some scientific concepts may have both kinds of functions, indeed those with joint-carving functions are very likely to possess pragmatic functions as well. My claim is simply that the pragmatic functions do not appear to require joint-carving functions; they can exist without them.

mandatory, and it would appear that these benefits, all of which have a pragmatic and heuristic flavour, may apply on the CICE construal.

Thus we have defeasible evidence that these are, or are being thought of as, among the functions of the relevant concepts. We can construct an argument for the conceptual extensions that exemplifies the methodology recommended by PICE.

P1. The functions of the phenotype/inheritance/mind/replicator concept are to contribute to the most general aims of science, by, perhaps among other things:

- Suggesting questions and lines of inquiry.
- Allowing us to develop explanatory methods that may enrich our understanding of the relevant phenomena.
- Simplifying and unifying our classificatory schemes.
- Focusing our attention on important issues and problems, while encouraging us to ignore distinctions that are of no theoretical importance.

P2. These functions will be better fulfilled by the extended variant of the phenotype/inheritance/ mind/replicator concept than by the traditional variant.

C. Therefore, we should adopt the extended variant of the phenotype/inheritance/mind/replicator concept.

Clearly, the functions being identified here are still quite general, applying as they do to all of the relevant concepts. They don't tell us what the function(s) of the phenotype concept in particular, or the inheritance concept in particular, are. A full account of the function(s) of these concepts, and thus a complete PICE-style argument for each conceptual extension, would certainly have to identify these more specific functions or aims, and show how the proposed extensions also satisfy these functions or aims to a greater extent than do the traditional concepts. But we have been able to abstract from these specific functions due to the similarities in the purported benefits of the extensions. Thus we have been able to construct a more abstract argument schema, applying equally to each case. This is, according to PICE, the general form that arguments for conceptual extensions in science take. But the details remain to be filled in, for example in making the case for Premise 2, for each proposed extension.<sup>29</sup>

---

<sup>29</sup>Thomasson allows for the possibility that two or more different concepts, or variants of concepts, may satisfy the relevant function(s) equally well (2020). Dawkins may appear to hold this view about the traditional and extended phenotype concepts, when he states that the views are 'equally correct',

## 5. Concluding remarks: Carnapian deflationism in metaphysics and science

PICE, for Thomasson, derives much of its motivation from a prior commitment to metametaphysical deflationism, so I shall say a word about this, and its applicability to CICE, in closing. A vigorous debate in analytic metaphysics and metametaphysics has emerged in recent years concerning metametaphysical deflationism, the supporters of which frequently invoke Carnap (1950) as a major inspiration.<sup>30</sup> The position is not a scepticism about metaphysics which accepts that metaphysical questions have answers, and that there is a fact of the matter about which metaphysical positions are correct, but asserts that we are incapable of answering metaphysical questions (due to our epistemic limitations, or for some other reason).<sup>31</sup> Rather, it claims that there is no fact of the matter about (some or all) metaphysical questions – there is something wrong with the questions themselves, at least when conceived as questions with determinate, objectively correct answers.<sup>32</sup> In other words, there is nothing substantive – nothing factual – at issue in metaphysical debates (some are perhaps verbal disputes), and, depending on the version of the view defended, metaphysical propositions may even fail to have a truth value.<sup>33</sup> This anti-realist position is often, in the spirit of Carnap, combined with a pragmatism which accepts that it may nonetheless be pragmatically useful to adopt certain metaphysical conceptual frameworks or ‘ways of speaking’. (As several writers have noted, although Carnap is a key influence on this view, modern deflationists

---

and the ‘biologist should try both ways of thinking, and choose the one he or she prefers’ (1982, 7). But we have seen that he thinks there are pragmatic grounds for preferring the extended concept, and I am interpreting this as the claim that it satisfies the functions of the phenotype concept better than the traditional concept. But one could, in principle, hold that the two variants of the concept are equally good ‘all the way down’, i.e. not only does neither ‘carve at the joints’ better than the other, but neither satisfies the pragmatic functions of the phenotype concept better than the other.

<sup>30</sup>See the volume of papers, Chalmers, Manley, and Wasserman (2009), especially the papers by Bennett, Chalmers, Eklund, Hawthorne, Hirsch, Price, Sider, Soames, Thomasson, and Yablo.

<sup>31</sup>As an example of this sceptical view see McGinn (1993), who argues that metaphysical questions are real questions with real, objectively correct answers, but that the human mind/brain is incapable, due to its inherent limitations, of answering metaphysical questions. This is a realism about metaphysical facts, combined with a scepticism about metaphysical knowledge (in humans). Similarly, Bennett (2009) defends an ‘epistemicist’ position with respect to some metaphysical disputes, according to which they are not just verbal disputes, and there may be a fact of the matter about which position is right, but given our evidential situation, there is little justification for taking one position rather than another.

<sup>32</sup>This distinction between scepticism and deflationism is drawn by several theorists, including Sider (2009) and Yablo (2009).

<sup>33</sup>See Bennett (2009) for a discussion of different forms of what she calls ‘dismissivism’ about metaphysical questions.

generally eschew the verificationism which informed his version of the thesis; see esp. Hawthorne 2009.)

The pluralist/pragmatist position of which CICE is an instance has much in common, on the face of it, with Carnapian deflationism. It holds that there is no fact of the matter about the issues in question (whether phenotypes, minds, inheritance etc. are in fact extended), and thus there is nothing substantive or factual at issue in the debates. This anti-realism is combined with a pragmatism according to which it may be useful, for various purposes, to adopt certain concepts or variants of concepts. Put in Carnap's terms, we can adopt the 'extended phenotype framework' on pragmatic grounds, which is quite distinct from *believing* that extended phenotypes *exist*. Just as, for Carnap, we may acknowledge the usefulness of adopting the 'number framework', without treating this as amounting to the belief that numbers exist, we may also acknowledge the usefulness of adopting a certain way of thinking and talking about organisms and their traits, a 'vision of life ... a point of view, a way of looking at familiar facts and ideas ...' (Dawkins 1982), without treating this as amounting to the belief that organisms have extended phenotypes.

Of course the Carnapian position as standardly formulated is specifically concerned with ontology in the traditional metaphysical sense – e.g. whether numbers, properties, possible worlds, or propositions exist – whereas CICE is a view about disputes within science (and the philosophy of science) concerning the existence and nature of the sort of entities and properties studied by the natural sciences (and we have seen that even Carnapians like Thomasson sometimes balk at applying the deflationary position to debates within natural science). There is no question that the motivation for the Carnapian position in analytic metametaphysics is largely a kind of worry specifically about the legitimacy of the kinds of questions debated within philosophical ontology – the existence questions I mentioned above about numbers, properties etc., as well as questions about material composition, colocation, and so on – questions to which empirical evidence can seemingly contribute very little if anything, and are far removed from the empirical-scientific contexts in which phenotypes, inheritance and so on are posited. But given the structural similarities between the Carnapian view and CICE, I see no reason not to treat the latter as at least belonging to the same pluralist and pragmatist family of views as the Carnapian position in metametaphysics.

And indeed there are precedents for thinking of pluralist/pragmatist views in biology along the lines of conventionalist views in metaphysics. Sterelny and Kitcher, in their influential (1988), likened their pluralist gene selectionist view to the conventionalism about the geometry of space and time associated with figures like Reichenbach, while Kitcher *et al* contrasted the pluralist/pragmatist view on the units of selection they favour with a more ‘metaphysical’ approach, which would attempt to identify *the* unit(s) of selection operative in each selective episode: ‘[A]sking about the real unit of selection is an exercise in muddled metaphysics’ (Kitcher, Sterelny, and Waters 1990, 159; see also Kitcher 2008).

If that is correct, we ought to be in a position to apply insights from the much discussed and elaborated Carnapian tradition to help us understand conceptual extension according to CICE. As we’ve seen, for CICE to be plausible, it needs to steer a middle path between the metaphysical interpretation of the conceptual extensions as the fashioning of concepts that map onto the objective structure of the natural world, and the ‘trivial verbal dispute’ interpretation, on which the conceptual choices are a mere matter of deciding to use or define words or concepts in one way rather than another. It is this middle path that the Carnapian tradition has sought to explore and make attractive and viable with respect to metaphysics, with Thomasson’s work on the pragmatic methodology for conceptual engineering being just one example. To the extent that views such as CICE share much in common with the Carnapian position, there is no reason why ideas and insights from the latter might not be applicable, perhaps in modified form, to normative conceptual work in science and the philosophy of science (see Boucher 2012 for further discussion).

### Disclosure statement

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

### ORCID

Sandy C. Boucher  <http://orcid.org/0000-0002-0575-7497>

### References

- Bennett, K. 2009. “Composition, Colocation, and Metaontology.” In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by D. Chalmers, D. Manley, and R. Wasserman, 38–76. Oxford: Clarendon Press.

- Boucher, S. C. 2012. "Empiricism, Metaphysical Stances and the Philosophy of Biology." PhD diss., University of Melbourne.
- Boucher, S. C. 2017. "Gould on Species, Metaphysics, and Macroevolution: A Critical Appraisal." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 62: 25–34.
- Boucher, S. C. 2020. "Pluralism, Realism and the Units of Selection." *South African Journal of Philosophy* 39 (1): 47–62.
- Boucher, S. C. 2021. "Evolutionary Debunking Arguments, Commonsense and Scepticism." *Synthese* 198 (12): 11217–11239.
- Brigandt, I. 2010. "The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation." *Synthese* 177: 19–40.
- Burgess, A., H. Cappelen, and D. Plunkett, eds. 2020. *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Burgess, A., and D. Plunkett. 2013. "Conceptual Ethics I and II." *Philosophy Compass* 8 (12): 1091–1110.
- Cappelen, H. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Carnap, R. 1950. "Empiricism, Semantics and Ontology." *Revue Internationale de Philosophie* 11: 20–40.
- Chalmers, D. 2011. "Verbal Disputes." *The Philosophical Review* 120 (4): 515–566.
- Chalmers, D., D. Manley, and R. Wasserman, eds. 2009. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Clarendon Press.
- Clark, A. 2011. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Clark, A., and D. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
- Dawkins, R. 1982. *The Extended Phenotype*. Oxford: Oxford University Press.
- Dawkins, R. 2004. "Extended Phenotype – But Not Too Extended. A Reply to Laland, Turner and Jablonka." *Biology and Philosophy* 19: 377–396.
- Devitt, M. 1979. "Against Incommensurability." *Australasian Journal of Philosophy* 57: 29–50.
- Feyerabend, P. 1975. *Against Method*. London: New Left Books.
- Field, H. 1973. "Theory Change and the Indeterminacy of Reference." *Journal of Philosophy* 70 (14): 462–481.
- Fine, A. 1975. "How to Compare Theories: Reference and Change." *Noûs* 9: 17–32.
- Flood, A. 2021. "Richard Dawkins Loses 'Humanist of the Year' Title Over Trans Comments." *The Guardian*. Accessed April 21, 2021. <https://www.theguardian.com/books/2021/apr/20/richard-dawkins-loses-humanist-of-the-year-trans-comments>.
- Godfrey-Smith, P. 2001. "On the Status and Explanatory Structure of Developmental Systems Theory." In *Cycles of Contingency: Developmental Systems and Evolution*, edited by S. Oyama, P. Griffiths, and R. Gray, 283–297. Cambridge: MIT Press.
- Godfrey-Smith, P., and B. Kerr. 2002. "Individualist and Multi-Level Perspectives on Selection in Structured Populations." *Biology and Philosophy* 17 (4): 477–517.
- Gould, S. J. 2002. *The Structure of Evolutionary Theory*. Cambridge: Harvard University Press.

- Gray, R. D. 1992. "Death of the Gene: Developmental Systems Strike Back." In *Trees of Life*, edited by P. Griffiths, 165–209. Dordrecht: Kluwer Academic Publishers.
- Griffiths, P., and J. Wilkins. 2015. "Crossing the Milvian Bridge: When Do Evolutionary Explanations of Belief Debunk Belief?" In *Darwin in the 21st Century: Nature, Humanity and God*, edited by P. Sloan, 201–231. Notre Dame, IN: Notre Dame University Press.
- Haslanger, S. 2000. "Gender, Race: (What) Are They? (What) Do We Want Them To Be?" *Noûs* 34 (1): 31–55.
- Haslanger, S. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Hawthorne, J. 2009. "Superficialism in Ontology." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by D. Chalmers, D. Manley, and R. Wasserman, 213–230. Oxford: Clarendon Press.
- Heersmink, R. 2016. "The Cognitive Integration of Scientific Instruments: Information, Situated Cognition, and Scientific Practice." *Phenomenology and the Cognitive Sciences* 15: 517–537.
- Kitcher, P. 1984. "Against the Monism of the Moment: A Reply to Elliott Sober." *Philosophy of Science* 51: 616–630.
- Kitcher, P. 2008. "Carnap and the Caterpillar." *Philosophical Topics* 36 (1): 111–127.
- Kitcher, P., K. Sterelny, and C. K. Waters. 1990. "The Illusory Riches of Sober's Monism." *Journal of Philosophy* 87 (3): 158–161.
- Kuhn, T. 1996. *The Structure of Scientific Revolutions*. Chicago, IL: The University of Chicago Press.
- McGinn, C. 1993. *Problems in Philosophy: The Limits of Inquiry*. Oxford: Blackwell.
- Mishler, B., and R. Brandon. 1987. "Individuality, Pluralism and the Phylogenetic Species Concept." *Biology and Philosophy* 2: 106–123.
- Prinzing, M. 2018. "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection." *Inquiry* 61 (8): 854–880.
- Schulz, A. W. 2013. "Overextension: The Extended Mind and Arguments from Evolutionary Biology." *European Journal for Philosophy of Science* 3: 241–255.
- Sider, T. 2009. "Ontological Realism." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by D. Chalmers, D. Manley, and R. Wasserman, 384–423. Oxford: Clarendon Press.
- Simion, M. 2018. "The 'Should' in Conceptual Engineering." *Inquiry* 61 (8): 914–928.
- Sterelny, K., and P. Griffiths. 1999. *Sex and Death: An Introduction to Philosophy of Biology*. Chicago, IL: The University of Chicago Press.
- Sterelny, K., and P. Kitcher. 1988. "The Return of the Gene." *Journal of Philosophy* 85 (7): 335–358.
- Sterelny, K., K. C. Smith, and M. Dickison. 1996. "The Extended Replicator." *Biology and Philosophy* 11 (3): 377–403.
- Thomasson, A. L. 2017. "Metaphysics and Conceptual Negotiation." *Philosophical Issues* 27: 364–382.
- Thomasson, A. L. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 435–456. Oxford: Oxford University Press.

- Turner, J. S. 2000. *The Extended Organism: The Physiology of Animal-Built Structures*. Cambridge: Harvard University Press.
- Turner, J. S. 2004. "Extended Phenotypes and Extended Organisms." *Biology and Philosophy* 19: 327–352.
- Van Fraassen, B. C. 1980. *The Scientific Image*. Oxford: Clarendon Press.
- Yablo, S. 2009. "Must Existence-Questions have Answers?" In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by D. Chalmers, D. Manley, and R. Wasserman, 507–525. Oxford: Clarendon Press.