

# Does Metaphysical Deflationism Require Global Expressivism?

A metasemantic methodology for the neo-Carnapian

Ed Willems (she/her)

PhD

University of York

Philosophy

September 2022



## Abstract

Neo-Carnapian deflationism is an approach to metaphysics that is currently enjoying a surge in popularity. The movement behind it takes inspiration from Carnap's (1950) paper *Empiricism, Semantics and Ontology*, and looks to deflate metaphysical questions and disputes, showing them to be insubstantial, since they turn on questions of language rather than metaphysics. I hold that this is all to the good. However, I argue that neo-Carnapianism is problematised by specific metasemantic assumptions, specifically the assumptions of representationalism. If we hold a representationalist view of language, then neo-Carnapianism fails. Over the course of this thesis, I will show how representationalism causes problems for the neo-Carnapian account, and argue for an alternative metasemantic approach – global expressivism – that we can substitute to save it. Global expressivism is itself a difficult position to make clear, so I will spend a significant amount of time demonstrating its core principles, before showing how to combine it with neo-Carnapianism, providing a deflationary approach to metasemantics to complement our deflationary approach to metaphysics.

Contents:

Abstract	p3
Contents	p4
Declaration	p5
Introduction: Deflationism all the way up	p6
1: Carnap's metaphysical deflationism	p13
2: Characterising global expressivism	p40
3: Hirsch's intensional deflationism	p68
4: Thomasson's easy ontology	p88
5: Refining the deflationary theory	p110
6: Price's global expressivism	p127
7: Expressivist explanatory strategies	p149
8: Conclusion	p167
Bibliography	p184

Declaration:

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other university. All sources are acknowledged as references.

A reduced version of Chapter 4 below was presented at the Open Sessions, part of the Joint Session of the Aristotelian Society in 2020. No other material herein has been presented or published elsewhere.

## Introduction: Deflationism all the way up

There is a familiar joke that gets told when philosophers get together. It runs as follows: A conversation will be going on that concerns some ordinary sort of thing, Ks. This conversation may even be completely mundane and non-philosophical. At some point it will hit a lull, at which juncture some participant will quip something along the lines of “Well, if we were speaking *properly*, we wouldn’t be talking about Ks at all.” After some polite chuckling – we’ve all heard the joke before – the conversation about Ks usually resumes.

The joke itself is something close to an in-joke, relying on knowledge that the quipper holds a particular restrictive ontology that does not countenance Ks. In fact, it is a mild form of self-deprecation: The quipper knows that their ontology is restrictive in ways that are not palatable to their audience, and floats the idea of stridently insisting that the audience adopt their conception of proper speaking for comic effect.

The metaphysical deflationist (and here I include myself) wants to sap any fun out of this joke by proposing to take it far more seriously and flat-footedly than is intended. After all, systematic linguistic failure is no laughing matter. If the quipper’s ontology excludes Ks, are they not wrong to be speaking ordinary English, which, we presume, countenances Ks? Speaking of Ks in ordinary English would involve speaking falsehoods, by their lights, and we have a discursive obligation not to assert falsehoods. Wouldn’t silence, then, or extremely careful phrasing be the more responsible tack? Alternatively, if the quipper really is at home with returning to ordinary English, complete with a commitment to Ks, this must surely be because it is, in some sense, fine to speak of Ks so long as we are speaking ordinary English, although not otherwise.

At this point, the imaginary quipper has to try for a way to have their cake and eat it too – try, that is, to have their ontology be *the correct* ontology, to the exclusion of all others, and also to have ordinary English sentences positing Ks be permissible in ordinary English (though obviously not ontologically correct). This is fairly simply done by explaining that the K-committing sentences of ordinary English are *accounted* true in ordinary English, but are not *actually* true; ordinary English is an “as-if” language, that speaks as if the world was a certain way (i.e.: containing Ks), when it isn’t really. We entertain these languages, and are

allowed to make use of them, because they are helpful for getting things done in ordinary contexts, but they don't tell us how the world really is.

However, the deflationist isn't yet done being a fun-sponge. How, they ask, are the "as-if" languages to be distinguished from the *real* languages? What we have been shown so far is that it is possible to take up a stance or set of rules (ordinary English) from which we account sentences such as "There exist Ks" true, and another (the quipper language) from which we account them false. We haven't yet been shown which standpoint we should back for ascertaining which of these evaluations is correct, or how we could arrive at such a standpoint. It doesn't seem that we could use either of the two standpoints yet set up to evaluate such, since each will trivially endorse its own ontological commitments, leaving us again with the task of choosing which one we prefer.

At this point the imagined K-denier has to switch gears somewhat. Here it looks as though they have to give an answer that explains how *real* languages come to "latch onto" the world, according to which the quipper language latches onto the world in all the correct ways, but fails to latch onto any Ks (because there aren't any), and according to which the parts of ordinary English that look as though they *should* be in the business of latching onto the world are actually set up so that they produce false positives (i.e.: endorsing Ks when there are none). In particular, they have to give an explanation of the *real* sort of truth that their preferred language benefits from that distinguishes it from the "as-if" truth of ordinary English.

What we have here is a demand for a substantive metasemantics – an explanation of how the semantic values of a language's expressions get set. This requires a substantive account of semantic relations, such that a language could systematically fail to instantiate these relations, and its expressions thereby fail to have real semantic values. If the quipper can offer such a picture, then they can explain how ordinary English, though undoubtedly useful, actually produces only "as-if" sentences, and why they can view their own ontological framework as telling how the world truly is.

The story above is meant to illustrate the operative link between metaphysics and metasemantics. It presents the thesis in a nutshell: If the quipper is allowed such a substantive metasemantic picture, then they can draw a principled distinction between languages that represent how the world really is, and those that do not. In that case, metaphysical deflationism would be a non-starter; taking semantic relations to be substantial is inimical to taking metaphysical disputes to be insubstantial. In the thesis to

follow, I aim to demonstrate this point by examples. If metaphysical deflationism is going to work, I will argue, it needs to come packaged with a deflationary metasemantics.

The metaphysical deflationism I am interested in pursuing is neo-Carnapian. Metaphysical deflationism can be summarised as the view that the idea of a standard of metaphysical correctness as applied to languages – such that some languages accurately represent the metaphysical structure of the world and some do not – is misconceived. The neo-Carnapian approach to deflationism takes the views put forward by Rudolf Carnap in the seminal paper *Empiricism, Semantics and Ontology* (hereafter “ESO”), and attempts to rehabilitate and build upon them for a modern audience, to obtain a view of language that makes substantive metaphysical inquiry impossible.

In ESO, Carnap distinguishes two sorts of questions that philosophers might ask: Internal questions, which are questions asked using the normal rules of language, enshrined within a particular “linguistic framework” (ESO p22), and external questions, which eschew these rules in an attempt to ask what the world is like independently from how such frameworks hold it to be (*ibid.*). Carnap holds that external questions, while they seem compelling, cannot be made sense of, for the simple reason that they are posed from “outside” the rules of any particular language, so that they lack meaning. Insofar as external questions have a function, it is to call attention to the possibility of alternative linguistic frameworks, which are to be compared to our own and potentially adopted on pragmatic grounds – the latter specifically being because there can be strictly no question of metaphysical validation for linguistic frameworks.

Carnap holds that metaphysical questions such as “Are there numbers?” are best understood as external questions, i.e.: at best, as pointing to a possible new way of speaking that doesn’t involve the contested entity (ESO p25). Full blown metaphysical debates are, therefore, beyond our ability to make sense of, as the disputants, in taking themselves to be utilising an independent, metaphysical concept of reality, are attempting to conduct the debate outside any linguistic framework. Metaphysicians, according to Carnap, therefore make a particular sort of category mistake, namely that of attempting to argue for a pragmatic choice of language on non-pragmatic (i.e.: metaphysical) grounds.

The legacy of the Carnap of ESO has always been in dispute, and remains so today. The chief reasons for this are unquestionably the rebuttals published by W.V.O Quine (Quine 1951a and b). Briefly put, Quine’s rejection of the analytic/synthetic distinction appears to deprive the Carnapian of the means to draw distinctions between linguistic frameworks (I



will discuss Quine's arguments in detail in Ch.1). It appears to be for this reason that Carnapian thought has remained an undercurrent in analytic philosophy for so long, instead of breaking out into greater prominence.

However, over the last decade or so the neo-Carnapian approach has gained in popularity and garnered much discussion, thanks in part to the publication of two influential collections: *Metametaphysics* (eds. Chalmers & Wasserman 2009) and *Ontology After Carnap* (eds. Blatti & LaPointe 2016). The philosophical reason for the resurgence is perhaps best explained by Amie Thomasson's "quagmire" argument (Thomasson 2015, pp14-17): Analytic metaphysics, as it stands in the present day, is bogged down in metaphysical disputes that are at once both irreconcilable and un-decidable. That these disputes are irreconcilable stems from the fact that the respective disputants are convinced that there are substantial metaphysical stakes to the dispute; that they are un-resolvable is due to the widespread adoption of a Quinean desideratum for ontological theories, namely that they contain only those entities that are deemed indispensable by our current best scientific theory. As it turns out, there are many plausible ways to eliminate inessential entities from the picture, resulting in a "morass" of metaphysical theories that each lay claim to containing only those entities (and properties and relations) that are attested by science (*ibid.*).

For the purposes of the discussion presented in the thesis, I want to take it that the neo-Carnapian standpoint is independently desirable *if it can be made to work*. Indeed, most of the thesis will be devoted, in one way or another, to *making* it work. The reason for this is that the discussion centres around the particular issue of the intersection of metasemantics and metametaphysics, specifically posing the question: What are the metasemantic requirements for a metaphysical deflationism? Hence I will not spend time arguing further for the need for a neo-Carnapian deflationism "from cold", as it were. However, in Ch. 1 I will spend considerable time defending the position, or at least making room for it, against Quine's famous critique, in order at least to clear some ground on which to build. Likewise, in chapters 3 and 4, I will myself thoroughly critique different versions of neo-Carnapianism, and attempt to refine them based on this discussion. Hence, although I take for granted a large part of the motivation for neo-Carnapianism, the position will not go unchallenged.

My contention, which will chiefly be argued in chapters 3 and 4, is that the most prominent approaches to a neo-Carnapian deflationism do not do enough to distance themselves from

representationalism, the view, outlined above, that there are substantial representational relations that hold between language and its subject-matter in the world. Ch.3 will deal with Eli Hirsch's intensional deflationism, and Ch.4 with Amie Thomasson's easy ontology, or existence deflationism. For both authors, I'll argue that relevant criticisms of their work (from John Hawthorne and Andrew Brenner respectively) reveal the representationalism embedded in their theories, and that this representationalism allows us to re-engineer substantive metaphysical conclusions, undercutting the authors' metaphysical deflationism. The theories can be refined to meet the objections, but only by removing the representationalism.<sup>1</sup> It follows that there are indeed metasemantic requirements for a deflationary approach to metaphysics, as certain metasemantic views, here representationalism, make such an approach impossible.

I'll pitch global expressivism as the alternative to representationalism. Global expressivism is the view that the link between any given linguistic expression and its subject-matter resists being cashed out in substantive, relational terms. In this, global expressivism parallels metaphysical deflationism, in that it denies that there is any substantive metaphysical story to be told about semantic properties and relations – denying that there is any substantive notion of truth to be used in analysing what it takes for a proposition *p* to be true, for example. Global expressivism is a notoriously difficult view, and one that defies straightforward summary; a significant part of the thesis to follow will be devoted to articulating it clearly, and demonstrating its methods. Hence in Ch.2 I will address the problem of articulating a global expressivism, again by looking at a critique, this time from Matthew Simpson, that states that no such view can be made sense of. This time the response is in the form of a refined "credo" for the global expressivist, which I hope will help anchor an understanding of the sort of global expressivism we want – a form of metasemantic deflationism – going forward. The reason for placing this discussion early on in the thesis is to give the reader a frame of reference for representationalism and global expressivism, in order to make sense of the moves away from representationalism made in chapters 3 and 4, although those moves will not yet constitute a fully expressivist metasemantics for the neo-Carnapian.

In Ch.5 I will address the issue of the compatibility of the neo-Carnapian approaches examined so far, with a view to formulating a neo-Carnapian approach that avoids the

---

<sup>1</sup> Strictly speaking, only Hirsch's theory actually incorporates a form of representationalism; Thomasson attempts to remain neutral on questions of metasemantics, although I argue that this is a mistake, for the reasons given here.

problems highlighted. The key issue with compatibility is quantifier variance; Hirsch holds that different metaphysical languages make use of different quantifiers, since they recognize different entities, while Thomasson's easy ontology is explicitly constructed such that it rules out this idea. Here, I argue for two points: First, that the issue is non-substantive in the same sense that both theories view metaphysical disputes as non-substantive, and second, that quantifier variance is actually not a necessary corollary of Hirsch's intensional deflationism, the core of his view, which is otherwise compatible with easy ontology. Hence the best course, on grounds of utility, is to drop quantifier variance, and adopt a dual view that makes use of both existence deflationism and intensional deflationism; this is the neo-Carnapianism that we want.

If Ch.2 provided enough of a technical sketch of global expressivism to be going on with, chapters 6 and 7 will demonstrate by example how to assemble the machine. Ch.6 presents Huw Price's global expressivism as a way of understanding the sorts of concerns that motivate the expressivist, and the overall worldview that comes with such a theory, while Ch.7 gives worked out examples of expressivist ways of cashing different linguistic concepts, as a way of further illustrating the view. Finally, Ch.8 concludes by presenting expressivist readings of the concepts of intensionality and existence, and examining how these solve the problems outlined earlier for Hirsch and Thomasson.

A key issue throughout the thesis will, of course, be that of deflationism generally: What does it mean to deflate a particular metaphysical dispute, as opposed to, say, adopting an anti-realist metaphysical stance? Both neo-Carnapianism and global expressivism are deflationary theories – the former concerning metaphysics, and the latter concerning metasemantics. I will discuss what exactly we mean by deflationism in each context, since it can appear obscure, particularly as regards expressivism. One source of confusion here is the potential for misconstruing the deflationist's aims as being closer to those of the non-cognitivist's than they in fact are. While both contend that a particular discourse resists substantial metaphysical characterisation, the deflationist, in the sense used here, does *not* hold that that discourse is therefore incapable of yielding truths, having terms that refer to real entities etc. in the way that the non-cognitivist does. The key here is that the deflationist abandons the assumption that a substantial metaphysical characterisation of the discourse is necessary for it to yield these real results. In the thesis I refer to this as a "modern" deflationism, to contrast it with earlier deflationary efforts, particularly verificationist ones, that yielded localized non-cognitivism. The defining feature of modern deflationism, then, is that it allows us to speak real truths using the deflated vocabulary,

even while it places an embargo on certain sorts of metaphysical reasoning using that vocabulary; I will discuss this further in Ch.2.

The thesis to follow will look to demonstrate by examples the incompatibility of metaphysical deflationism with representationalist assumptions. It is difficult to generalize about this connection, since the theories involved are both abstract and varied, and the assumptions difficult to spell out, hence the approach of dealing with particular prominent cases. The gist, though, is that to harbour a representationalist view of language is to build in metaphysics at the metasemantic level. Within such a picture, a thoroughgoing metaphysical deflationism could never really get going, since our representationalist assumptions give us a foundation for posing substantive metaphysical questions, and expecting substantive answers. The solution is to adopt a deflationary metasemantics, global expressivism, that doesn't bear substantive metaphysical commitments – essentially, taking our deflationism sufficiently seriously to adopt it at every level.

In looking to provide an overall view that combines neo-Carnapianism with global expressivism, I will at several points draw distinctions between different theories, and subsequently deflate these distinctions. Disputes as diverse as the Quine/Carnap dispute, the deflationism/primitivism dispute, and the dispute over quantifier variance already mentioned will be subject to this analysis. Although such repetition may appear simply to stem from a lack of imagination, it actually reflects the core point that I aim to articulate in the thesis: If we are to be successful deflationists, we have to be deflationists “all the way up”.

## Ch.1: Carnap's metaphysical deflationism

Metaphysical deflationism is an important philosophical position that is undergoing a revival. Key to the discussion is the influence of Carnap on modern approaches to metaphysics. In the 1950s, Carnap attempted to show that the metaphysical debates of his own day suffered from a systematic confusion that rendered them unresolvable. However, Carnap's disagreement with Quine led to Carnap's deflationism being dismissed in most quarters as reliant on the analytic/synthetic distinction, which Quine dissolved. As a result of Quine's apparent victory, there emerged a dominant brand of Quinean metaphysics that persists to this day, which pursues metaphysical questions under the stated aim of creating the metaphysically-best whole scientific theory.

The present revival of metaphysical deflationism comes largely from the reapplication of Carnap's ideas to these refined metaphysical approaches. That is, it is finding that Carnap's deflationary arguments, suitably adjusted and interpreted, are just as fittingly applied to the neo-Quinean approaches as they were to their original targets that has prompted the new wave of Neo-Carnapians to attempt to deflate modern metaphysical disagreement.

There are several reasons besides sympathy for the Carnapian position to want to do this, the foremost among those cited being the proliferation of metaphysical theories that the neo-Quinean approach allows, with no clear metric for deciding which theories are better than which others. However, the focus on this thesis is not on arguing for metaphysical deflationism, but rather on examining the requirements for a metaphysical deflationism that works.

Taking for granted that such a theory is worth pursuing, I argue that the current proponents of the neo-Carnapian approach underestimate the metasemantic requirements of such a theory. The neo-Carnapian approach is often presented as a curbing of the excesses of metaphysics, particularly in the multiplication of metaphysical theories, and therefore as something that philosophical moderates ought to be on board with. However, I argue that it is incompatible with commonly held assumptions about semantics, namely that semantic meaning is derived from the representational properties of language, in a way that can be specified by philosophical examination. Incompatibility with this assumption forces the neo-Carnapian to adopt an anti-representationalism in the form of

global expressivism, a fairly radical metasemantic stance. This is something that is not clearly or explicitly acknowledged by modern neo-Carnapians.

An understanding of Carnap is a good starting point for understanding the neo-Carnapian movement. That movement and its interlocutors focus almost exclusively on Carnap's views as expressed in one paper: *Empiricism, Semantics and Ontology* (hereafter "ESO"). This represents Carnap's deflationism in a fairly comprehensive way, hence its being able to serve as the foundation stone for the deflationary movement. I will examine the deflationism presented in ESO below, before turning to questions of critique, namely asking whether the deflationism presented is hampered by a commitment to verificationism (I conclude that it is not), and whether the arguments presented by Quine are sufficient to dismiss it (again, I conclude that they are not, although this time more tentatively). I will then briefly examine the semantic and metasemantic views embedded in the arguments in and surrounding ESO, pointing to how they push for a Carnapian deflationism as anti-representationalist.

### 1: Empiricism, Semantics and Ontology

Carnap's seminal paper is concerned with articulating the reasons for taking language at face value, so to speak. ESO is an attempt by Carnap to combat the reductionist tendencies of his fellow empiricists, as he sees them. Carnap argues that all that it is to be committed to the existence of Fs is to accept a language form that supports quantification, in that language, over Fs – his deflationism is thus committed just as much to dissolving views that place stringent requirements on meaningful language, and thus result in error theories, as it is with dismantling substantive metaphysical theories about what the real structure of the world is like.

Carnap's main approach to this task centers around the concept of the linguistic framework. A linguistic framework is a system of rules for the use of terms of a certain type, for instance mathematical terms. These rules govern the correct usage and licensed inferences involving these terms. Carnap also includes reference to a concept of reality grounded within a framework (ESO p22), but makes it clear that this concept is not meant to be a metaphysical one. This will become clearer in what follows, but for now it is perhaps best to think of 'real' as used internally to, say, the framework of everyday things

as the antonym of the everyday expression ‘fake’, rather than the metaphysical expression ‘irreal’, or similar.

The main use to which Carnap puts the concept of a framework is in drawing a distinction between uses of language that are internal to the framework, and those uses that are (intended to be) external to it. By “internal”, Carnap means simply that they use the rules of the framework and its concepts (Carnap’s primary focus is on questions, because metaphysical debates can be construed as attempts to provide different answers to metaphysical questions – see below). For instance, if we were to ask the question “Are there any prime numbers?” internally to the framework of mathematical entities, we would have to use the rules of the framework to work out the answer. According to the rules of mathematics, we can demonstrate quite trivially that there are prime numbers, so the question, construed internally, is trivially dealt with (*ibid.*). This is a mark of internal questions generally: They tend to have trivial answers, and further, when they are an internal reading of questions typically asked by metaphysicians, they tend to be in the affirmative, because our having concepts with which to frame the question is usually the result of having a framework for dealing with entities of the relevant type.

External language use means, for Carnap, attempting to use the terms of a framework without recourse to the rules of the framework itself. It might appear that this use is misconceived, since the rules or correct usage of a term play at least some role in determining its meaning, regardless of one’s metasemantic leanings. This is accurate as concerns one type of attempted external use, namely the “factual” use<sup>2</sup>. This, according to Carnap, is where the metaphysical concept of reality resides; questions about whether Fs *really* exist, regardless of what the framework of Fs says, cannot be made sense of, because the word ‘really’, as used here, is not attached to the rules of any framework, and therefore is not attached to a meaning. Likewise, the concept ‘F’ has no meaning externally to the framework that contains rules for how it is to be used, so we can’t get around the problem simply by dropping “really” and asking simply “Are there Fs?” The question has no answer because it is underspecified – that is, without specifying a framework, it isn’t attached to any conditions the world has to meet in order for it to be answered either affirmatively or negatively.

---

<sup>2</sup> I follow Eklund (2013) in distinguishing between the types of external questions as “factual” and “pragmatic”, acknowledging that these are slightly rough terms. There is no immediate need to apply them to internal questions, as all internal questions appear to be factual in this sense.

This is how Carnap deflates metaphysical debates, namely by showing that the questions being discussed are not framed so as to be cogent, and are thus unanswerable. If the questions being asked by metaphysicians were intended to be internal, they would be answerable, but Carnap maintains that this is not how metaphysicians intend them to be taken; questions such as “Are there numbers?”, which are debated in metaphysics, are so trivial when taken as internal that they don’t warrant asking, at least if one has mastered the framework in which they would be asked.

The demand to be charitable to a speaker is complicated when one takes the options for what they *could* mean by a question to be either utterly trivial such that anyone with mastery of the language could answer it, or thoroughly problematic to the point of meaninglessness. Either seems to mean taking them to be subject to some sort of confusion – the former a linguistic confusion, as the most likely explanation seems to be that they do not understand how the language works, and the latter a metalinguistic one, as they seem to be trying to ask a question without tying it to a language, or at the very least a system of concepts. It is in trying to find the question of substance that is closest to what the metaphysician intends that we find the third category of language relative to frameworks: External–pragmatic sentences.

The idea behind this category is that the metaphysician, by interrogating a linguistic framework as it actually stands, and proposing a theory that seems to contravene our ordinary uses of framework concepts, is actually proposing a new, alternative framework. The question “Are there Fs?” is thus not to be read as the trivially answered version internal to the framework of Fs, but rather as the question whether a framework that served the same pragmatic function but *without* the inclusion of the concept ‘F’ would be better suited to our purposes. Read in this way, the core value of metaphysical inquiry is not in discerning the real nature of the world, and whether language maps onto its structure, but in proposing alternative language forms with different ontologies<sup>3</sup>, and discerning their pragmatic value.

It is obviously up for debate whether this is the kernel of value that drives the pursuit of metaphysics – that is, whether the function of Carnap’s “conceptual engineering” sits too far away from the way metaphysics is actually done to represent its main function. A different but related question regards metaphysicians’ self-conception. To generalize, the

---

<sup>3</sup> Here following the Quinean usage for “ontology”, meaning whatever our language(-framework) ultimately quantifies over. Note that Carnap objected to this usage as misleading (ESO p32, fn2).



model for metaphysics post-Quine *does* largely take itself to be iterating on theories based on criteria other than direct correspondence to reality. However, this idea of correspondence hasn't dropped out entirely, as the neo-Quinean view (in particular Sider 2011) expands upon Quine's notion of theory success as *evidential* support for the theory, to re-engineer a view of theories as capable of being more or less metaphysically correct – essentially attempting to reintroduce a sense of quasi-correspondence of theories to reality post-deflationism.

There is one important aspect of deflationism to discuss before moving on, which I want to focus on as a common misconception concerning the modern sort of deflationism that I want to support. That is the idea of deflationism as rejecting not just the debates but also the language forms that it deflates. It is fairly common to characterise the deflation of a debate as the deflationist's attempting to show the discussion to have no real bite because there is something wrong with the languages being used by both disputants, e.g.: because they are non-representational. A good example is found in the verificationist deflation of ethical debates. The classic position taken by the verificationist towards debates over whether a given act, *x*, is morally right or wrong is that each side's assertion really only amounts to something like exclaiming "X!", and then either cheering or booing respectively. The verificationist claim is that there is no real dispute here because neither side is making any claim capable of being true or false – that the debate can be discarded because the particular corner of language being used, ethical language, is non-cognitive, and is therefore not suitable for sustaining debates. The method for deflating the debate is to claim that neither side is actually saying anything, and so no resolution can be found.

I believe that the common association of modern, neo-Carnapian deflationism with anti-realism springs from a background association with this sort of "deflation". However, this is a misstep. The fact is that the classic verificationist construal is actually much more destructive, and somewhat cruder, than the Carnapian approach. As evidenced in the characterisation above, Carnap does not want to decry any linguistic framework as non-cognitive, or otherwise incapable of supporting a debate. ESO as a paper is effectively an attempt to work out the ontological upshot of applying Carnap's Principle of Tolerance, the principle that any language that is sufficiently systematic and consistent can be used as a framework to support linguistic meaning, without fear that its concepts are in some way bankrupt, or that its terms fail to represent the world in a way that is important to the language's being meaningful (Carnap 1937 pp51-2). With the Principle of Tolerance in hand, it is impossible to adopt the sort of destructive attitude of the earlier verificationism

by dismissing languages for failing to meet some standard of meaningfulness other than systematicity and consistency.

This is the point that I want to emphasize as a recurring theme of this thesis: That the well-rounded, modern deflationism we are here interested in preserving is not in any way that makes sense a form of anti-realism; that in fact it is in the business of showing that different languages or systems of concepts are perfectly in order and usable as they are, provided only that they meet the minimum requirements for counting *as* a system of concepts, not of showing that they are non-cognitive, non-representational, or meaningless. To lump deflationism in with a dismissive attitude towards languages – deflation in the sense of explaining away the entire language form as somehow meaningless or insubstantial – is actually to conflate it with a reductionism about linguistic meaning to which modern deflationism finds itself directly opposed. Such a reductionism takes the form of deriving a criterion for meaningfulness for language generally, and then holding to account any area of language that fails to live up to it.

Verificationism as applied above is a good example. The verificationist reduction of meaning to observable phenomena will be discussed below in section 3.1, but this is what is being used to support the verificationist's dismissal of, here, ethical language as meaningless. In a post-Quinean context, there is the potential for a structurally similar dismissal arising from the naturalisation of meaning. For instance, a causal theory of meaning such as informational semantics requires tokenings of expressions to stand in lawlike relations with things in the speaker's environment in order for those expressions to be meaningful. For the reductive naturalist, non-natural entities such as moral values cannot stand in such (natural) causal relations, hence terms referring to them do not satisfy the criteria for being meaningful.

The Carnapian metaphysical deflationist actually finds it impossible to adopt such an approach to unpacking criteria for meaningfulness, since tolerance requires that, for any such set of criteria, the deflationist will be committed to the meaningfulness of some language form that fails to meet it. I will explore this point throughout the thesis, but for now it is important to flag the central issue of tolerance, that the sort of deflationism we are interested in is not dismissive of the language forms being used in debates subject to deflation, but rather reaffirms them, and their usefulness.

## 2: *Two Dogmas*, and the Analytic/Synthetic Distinction

Quine's rejection of the analytic/synthetic distinction is famous, but its upshot for deflationism is less well understood. Evidence of this can be seen in the fact that the rejection of *Carnapian* deflationism that it apparently entails has since heralded a renaissance for metaphysics under a neo-Quinean banner, despite the fact that the Quinean framework in which the rejection is situated looks at least as hostile to metaphysics as was Carnap's. To see why the two deflationary thinkers were at odds despite these similarities in their views, I will briefly recount the arguments put forward by Quine in his paper *Two Dogmas of Empiricism* against the cogency of the analytic/synthetic distinction, before looking, in later sections, at ways Carnapians have responded and could respond to the problem.

The eponymous two dogmas that Quine attributes to the empiricists of his time are the idea of the analytic/synthetic distinction and the idea that meaning is reducible to an observation language referring only to sense data (Quine 1951a, p20). These two points are not unrelated, specifically because the reductionism of the latter dogma, derived primarily from verificationism, could provide a way to articulate and support the distinction in the former (*ibid.* p35). I will explain Quine's arguments against the analytic/synthetic distinction first, before looking at ways one might try to salvage Carnap's deflationary project in the following sections. The adoption of the reductionist second dogma will be among the responses considered, although it is not the most attractive; other responses will be ways of rehabilitating the analytic/synthetic distinction, or of doing deflationary work without drawing upon it directly. I'll then articulate my own response, which looks to deflate the debate at the metalinguistic level, placing it on a pragmatic footing, and making room for the Carnapian approach.

The main thrust of Quine's critique is that the term 'Analytic' cannot be satisfactorily defined. If we can't define 'Analytic', then we can't state in general terms where the line should fall between analytic truths and synthetic truths. Quine chiefly pursues the line of trying to define an analytic sentence as a sentence that can be turned into a logical truth by replacing certain of the expressions it contains with synonyms (*ibid.* p23). The thought here is that we can capture the class of truths-by-definition, such as "All bachelors are unmarried", as the class of truths that can be transformed, by replacement of synonyms, into sentences that are true in virtue of their logical form, e.g.: "All unmarried men are unmarried." Assuming it is permissible to talk of sentences that are true in virtue of their

logical form, we would thus have a way of outlining the class of analyticities, and thus distinguishing them from those sentences that are merely synthetically true.

The problem here, according to Quine, is that the concept of synonymy is just as mysterious as that of analyticity. Consider that the most intuitive way of unpacking the fact that 'Bachelor' and 'Unmarried man' are synonyms is to say that "All bachelors are unmarried men" or "x is a bachelor iff x is an unmarried man" is analytically true (*ibid.* p28). This is to trace a very tight circle from analyticity to synonymy and back, and does not provide the grounding for the concept of analyticity that we were looking for.

In pursuing other criteria for two expressions to be synonymous, Quine considers whether the notion of their being substitutable, one for the other, *salva veritate* will serve as an elaboration of the concept, although he quickly finds that it will not. The reasoning here is as follows: Mere extensional substitution *s.v.* will not suffice to define synonymy, since in a world in which all bachelors happen to be over six feet tall, the expressions 'Bachelor' and 'Man who is over six feet tall' will happen to be substitutable *s.v.*, even though they are obviously not synonymous. The sort of synonymy we are interested in, synonymy in virtue of sameness of meaning, therefore has a modal component; to define it, we need to be able to say not just which expressions happen to be substitutable, but which would be substitutable in any possible world, i.e.: substitutable necessarily. But, Quine argues (albeit briefly, *ibid.* p29), we could not make sense of the modal operator 'Necessarily', used in this way, unless we already had a functioning notion of analyticity for the language in which it occurs. This is because it looks to apply to all and only analytic sentences, meaning we could only offer its rules of use either using the concept of analyticity, or some completely different concept that itself constituted a definition of analyticity. Substitutability thus seems to be, at best, a red herring, and with no other likely avenues it looks as though analyticity is doomed to go undefined.

The significance of Quine's rejection of the analytic/synthetic distinction for deflationism is in its bearing on Carnap's internal/external distinction. For Carnap, the analyticities attached to an expression, particularly the analytic connection of the expression with its definition, situate it within a certain linguistic framework, which builds those analyticities into the rules for the expression's correct usage. E.g.: x's being extended in space follows analytically from its being a physical object, as enshrined in the rules of the physical object framework, with this analytic connection constituting one of the rules of use for the expression 'physical object', and therefore part of its meaning, within that framework.

It is tempting at this point to blur the two distinctions, analytic/synthetic and internal/external, together, but we must be careful not to do so. For instance, given a physical object framework as described above, the question “Are there physical objects that are not extended in space?” looks like an external question. This is because an affirmative answer is impossible according to the analytic rules of the framework, hence it is unlikely that anyone who understood these rules would intend it as an internal question, posed within the framework. It thus looks like an attempt to make room for discussion of a physical object framework that doesn’t adhere to the rule given above, but an attempt that could easily slip into a misguided metaphysical investigation, i.e.: an external question. However, while this link between the externality of the question and the analyticity of the answer holds in this case, it would be a step too far to generalize this and say that external questions are just those that have analytic answers, so would not be posed internally to the framework.

This is because there are questions that, by Carnap’s lights, ought to be taken as external that have non-analytic answers, such as the more general question “Are there any physical objects?” It is not an analytic truth of the physical object framework that there are physical objects; we can, in accordance with its rules, coherently describe a world in which there are none. In fact, we can assert, without logical inconsistency, that the actual world is one such “empty” world. The triviality with which this claim would be falsified is what marks it as likely being an external one, which no one would seriously assert using the rules of the physical object framework. However, this triviality is empirical in nature, not conceptual. That is, it is trivial because the question is so easily answered if we accept the rules of the framework, given the corresponding empirical procedures for arriving at an answer, not because it is answerable by definition alone. Having an analytic answer is thus not a necessary condition on being an external question, since questions exist that are external in virtue of having trivial but non-analytic answers.

Nor is being answerable analytically a sufficient condition for being an external question; think of obscure or complex mathematical questions whose answers have to be worked out using the analytic rules of mathematics, or questions asked in the classroom to check that students have understood definitions correctly, all of which are correctly construed as internal. In brief, the threshold at which a question goes from internal to external is not the threshold at which its answer goes from synthetic to analytic, but the point where it becomes so trivial that it is implausible to interpret the speaker as adhering to ordinary patterns of usage. The distinction thus has an important interpretive component.

The analytic/synthetic distinction is not identical to the internal/external distinction, but the two are linked. The link relies on the conception of the analyticities attached to an expression as forming its meaning, as distinct from statements of synthetic facts in which the expression occurs, which do not form part of its meaning. This conception of meaning is the basis of Carnapian deflationism. As we have seen above, Carnap dismisses certain interpretations of questions as meaningless pseudo-questions because they are not attached to a meaning, i.e.: it is not plausible to interpret them as holding the analyticities commonly associated with the expressions of which they are comprised. Dismissing the analytic/synthetic distinction means we can no longer draw this distinction between uses of an expression that accord with its meaning and uses that are detached from that, or any meaning (i.e.: internal and external uses respectively). Hence we would no longer have categorical grounds on which to deflate metaphysical questions by showing them to be pseudo-questions.

Likewise, in a Hirschian vein, allowing the analytic/synthetic distinction allows the Carnapian to dismiss metaphysical debates as pseudo-debates because the disputants argue past one another, since they associate different analyticities with pertinent expressions such that neither side can be disproved by the other, because they do not mean the same things by the words they use. Dismissing the analytic/synthetic distinction means we lose the ability to say categorically when two people mean the same thing by a given expression, and when they do not, since we can no longer distinguish between those (analytic) truths that form part of an expression's meaning, and those (synthetic) truths that don't. Hence we can no longer draw a line between factual disputes that debate synthetic facts and pseudo-debates where the difference of opinion is brought about by anterior differences in usage.

Quine's rejection of the analytic/synthetic distinction thus looks to undermine Carnap's attempts to deflate metaphysics by removing the means for specifying when expressions are being used with the same or different meanings. On one level, this seems quite straightforward; it makes sense that Quine, skeptical of the notion of meaning as a property of language, would attack specifically the semantic element of Carnap's semantic deflationism. However, it will be important to look more closely at exactly how this plays out.

### 3: Rewiring Deflationism

In the face of Quine's criticism, there are three responses the Carnapian can make to try to salvage their deflationism. The first is attempting to rehabilitate the analytic/synthetic distinction, and thereby the conception of meaning doing the deflationary work in Carnap's theory. The second is attempting to set up a similar deflationism *without* relying on the distinction – essentially, showing that the distinction wasn't what was really doing the work, and that other concepts can achieve the semantic deflation of metaphysics. The third approach is not necessarily exclusive with the first two, as it involves recontextualizing the Quine/Carnap debate itself by deflating it. Hence it doesn't weigh in on one side of the debate so much as it changes what is at stake in each potential outcome. I will deal with the third option in the following section, and the first two in this section.

#### 3.1: Rehabilitating analyticity

One way for a theory to support a concept of analyticity would be for it to embrace verificationism, specifically the reductionist attitude of verificationism towards meaning. Quine considers this in his paper; verificationism is the second of the titular two dogmas held by the empiricists that were Quine's contemporaries. Broadly speaking, verificationism reduces the meaning of a declarative sentence to the conditions under which it can be confirmed or disconfirmed empirically. Given this reduction, we can execute all the deflationary moves required of a Carnapian theory. For one thing, we can now leverage the verificationist conception of meaning to provide an account of *sameness* of meaning as required for Quine's definition of analyticity via synonymy. The definition of sameness of meaning in question revolves around sameness of confirmation conditions; two sentences have the same meaning iff both sentences are confirmed and disconfirmed in all the same circumstances. From here we can derive synonymy of sub-sentential expressions by looking at what they contribute to the sentences. That is, two sub-sentential expressions will be synonymous iff they are substitutable without changing the confirmation conditions of the sentences in which they occur. The advantage of this model of synonymy is that it does not rely on an undefined concept of analyticity or conceptual truth, which would make it circular, and that it does not bake in the modal concepts necessary to make substitutability *salva veritate* sufficiently powerful to do the job.

There are other advantages to reducing meaning to conditions of confirmation that the deflationist can reap. For one thing, we can now state when disputants are arguing past one another by pointing to the fact that their respective worldviews are empirically identical. Where this is the case, as it generally is in metaphysical debates, the verificationist says that there is no real dispute – it is not simply that the debate is unresolvable because of a dearth of available evidence, but rather that the different sides, when they assert their respective views, are actually asserting the same thing. Likewise, we can directly recreate the external/internal distinction by representing as internal those claims that come attached to conditions of confirmation or disconfirmation, and as external those that do not. Internal *questions* are then those that permit internal answers, and external ones are those that don't. This models Carnap's deflationism by representing internal questions and claims as being attached to meanings, in the sense of rules for applying them in discourse (which here amounts to trying to find out how to confirm them), and external ones as lacking this attachment, and hence being discursively unusable.

It is worth mentioning again, as above, that this approach to deflationism is actually incompatible with what Carnap has in mind in ESO. It is up for debate exactly how deep Carnap's verificationism ran, but it seems clear that at the point where he came to formalize the Principle of Tolerance, he rejected the reductionist excesses associated with the verificationist position, as Quine himself notes (Quine 1951a, p38). Nevertheless, other reductive approaches to meaning have continued in what approaches a Carnapian vein. Most notably, Eli Hirsch (2011) constructs a theory that is structurally similar to the verificationist model above, but that uses possible states of affairs and truth-conditions in the place of conditions of confirmation. I will discuss this approach, and the tension between Hirsch's tolerant, deflationary stance and the representationalism used to achieve it in Ch.3.

Quine himself casts doubt on the ability of verificationism to achieve its reductive aims (Quine 1951a, pp37-8). But aside from anything else, we can see that the reductionism required to make a verificationist deflationism work will not get us the sort of deflationism we want, because the requirement to fit a criterion of meaningfulness is too inflexible, which makes it too destructive as regards the language forms we want to preserve. Verificationism, unlike Hirsch's intensional deflationism, is notoriously prone not just to declaring that philosophical disputants are speaking past each other, but that they are each failing to produce any meaningful sentences at all. A theory that reduces linguistic meaning



to confirmation conditions will therefore be in unacceptable tension with the Principle of Tolerance, since it provides a criterion for meaning that excludes entire linguistic frameworks (most infamously ethical discourse, in the case of hard-line verificationism) based not on the coherence of their constitutive rules, but on an external criterion of verifiability.

I want to look at one more attempt to reinstate the analytic/synthetic distinction in the face of Quine's critique, before moving on, because it will be of some relevance later. This counterargument to *Two Dogmas* comes from Grice and Strawson (1956), who argue that the concept of analyticity is demonstrably usable in a systematic way. Grice and Strawson posit that the consistency with which philosophers apply the concepts 'analytic' and 'synthetic' not just to canonical examples, but to entirely novel cases, is evidence in favour of their being meaningful. In the authors' words:

...In general, if a pair of contrasting expressions are habitually and generally used in application to the same cases, *where these cases do not form a closed list*, this is a sufficient condition for saying that there are *kinds* of cases to which the expressions apply; and nothing more is needed for them to mark a distinction.

*Ibid.* p143 (emphasis in original)

Further, the authors note that these are not just bespoke, philosophical concepts, but that they see at least implicit use outside of philosophy, by Quine's own lights. This is because the ordinary notion of synonymy as sameness of meaning, according to the aforementioned arguments in *Two Dogmas*, is based on a conceptually prior notion of analyticity (*ibid.* p145). It seems to follow that ordinary language-users who are untrained in the philosophical version of the concepts nevertheless display an implicit understanding of them in their ability to apply the derived concept of synonymy, and other concepts in the "analyticity family", such as self-contradictoriness, with a high level of consistency to a non-closed list of cases.

There are two ways that this argument could be understood, which I want to distinguish briefly. One is as the argument that the concepts 'analytic' and 'synthetic' are, by virtue of the consistency of their usage, clearly "getting at something" – that the most likely explanation for the distinction's being systematically applicable is that it picks up on a real distinction. This version of the argument plays particularly well against Quine, because it uses the Quinean notion of evidence, this being that a theory's success at prediction, usefulness etc. is *evidence* in favour of the theory. However, such an application of the

idea of evidence is unlikely to sit well with the sort of deflationism we want to end up with. This is because the practice of counting theory success as evidence that the theory captures something about the world, and is the more correct for it, is all too easily ported over to more explicitly metaphysical disputes, where it would negate any prospective deflationist gains. This is in fact what has happened in the intervening years between *Two Dogmas* and the present; neo-Quineans, using Quine's concept of evidence, count competing metaphysical theories not just as being more or less useful, but as being more or less accurate representations of how the world is, depending on these theories' success against various relevant criteria. To make the claim that our ability to sort propositions into the camps 'analytic' and 'synthetic' is indicative not just of the usefulness of these concepts, but of the fact that these represent two metaphysically real categories into which statements naturally cleave, although it may well be the closer interpretation of what Grice and Strawson had in mind, is obviously too substantively metaphysical for our deflationism.

The other reading of the argument is more compatible with our aims – I will sketch it briefly here, to incorporate it more fully later on. The second option is to advance a language that includes the distinct categories *analytic* and *synthetic* purely on the basis of their usability, and their usefulness. The difference from the first reading is that this reading doesn't take the further step of claiming that these properties of their usage point to the categories' accurately picking up on something deep about the structure of the world; if anyone could adequately make sense of such a concept of "accuracy" as is used here (*cf.* Warren 2016), our deflationist wouldn't be the one to do it. Instead, we should adopt a pragmatist approach, where the use-value of a language form is the only factor in its being acceptable or not. This is in line with the ever-present Principle of Tolerance, which advocates treating any framework of concepts that is sufficiently systematic and consistent as usable, without considerations of accuracy to the world being a factor, at the most basic level, in assessing the concepts being introduced – for how could such considerations end up being anything other than metaphysical? The only way to make the Grice-Strawson argument work for the Carnapian in the long run, then, is to pull our punches, so to speak, and refrain from positing the manifest applicability of the concepts 'analytic' and 'synthetic' as showing anything beyond their usefulness.

### 3.2: Making do without analyticity

The other option for deflationism, besides attempting to reestablish analyticity and related concepts, is attempting to find a way to achieve deflationary results using different concepts left untouched by Quine's critique. A convincing way of doing this is taken from Price (2009) and expanded upon by Thomasson (2015). Price proposes treating the internal/external distinction of ESO in terms of the use/mention distinction:

In my view, it is helpful to frame Carnap's point in terms of the use–mention distinction. Legitimate *uses* of the terms such as “number” and “material object” are necessarily internal, for it is conformity (more or less) to the rules of the framework in question that constitutes use. But as internal questions, as Carnap notes, these questions could not have the significance that traditional metaphysics takes them to have. Metaphysics tries to locate them somewhere else, but thereby commits a use–mention fallacy. The only legitimate external questions simply *mention* the terms in question.

Price

2009a, p324 (emphasis in original)

So, on this view, an internal question is just a question that accords with the accepted rules of use of the terms occurring in it. An external-factual question, then, is a question that eschews these rules of use, and puts nothing in their place. Its status as a pseudo-question is thus earned by its not using the terms in a recognized way – if we were being less cautious, we might say: Without attaching them to a meaning. And, finally, an external-pragmatic interpretation, as Thomasson (2015, p40) notes, is a charitable way of reading a question as opening discussion of a particular non-standard use of a term, which is to be understood as *mentioning* the term in question, since its *use* is what is up for discussion.

Thomasson expands slightly on Price's use/mention formula, showing how it can achieve deflationary results on its own, without appealing to specific frameworks or the analyticities associated with them. Thomasson emphasises the idea that terms such as ‘number’ and ‘material object’ have an accepted usage, and that the metaphysician's usage is non-standard. Not only is the burden of argument thus placed on the metaphysician to show their usage to be better, the nature of the debate is changed from a metaphysical to a linguistic discussion, hence the debate is deflated. This is essentially an application of the ideal of Carnapian Tolerance to ordinary language, just as Carnap himself does in ESO; ordinary language is assumed to be in order as is, and its metaphysical commitments are

taken seriously in their own right, not as beholden to separate criteria of metaphysical realness.

Insofar as the Price/Thomasson version builds in an appeal to frameworks of concepts in any guise, it is only to insist that the metaphysician must be taking themselves to speak *some* language in order for their assertions to be meaningful, even if it is a bespoke language they have created. Says Thomasson:

If we ask a general existence question such as “are there numbers?”, “are there properties?”, “are there propositions?”, using those terms *in the only sense they have – using the rules by which they are introduced into the language*, the answer is a straightforward, easy ‘yes’. If we are spoiling for a debate (if it is to be meaningful), we must undertake it on other territory: regarding whether we should use these terms, along with their customary rules of use, at all.

*Ibid.* p41 (emphasis in original)

One might be suspicious that there is some vestige of analyticity lurking in the Price/Thomasson account in the notion of accepted rules of application for our terms. Put another way, the approach exhorts the metaphysician to pin their terms to a particular language and its rules of use, but for the Quinean, having rejected the analytic/synthetic distinction, what exactly counts as coming under the rules of use for a language (as opposed to the state of the world as expressed by the language’s users) is indeterminate. However, while this may well be a worry, it touches on a stress point in the Quinean account, namely that the concept of analyticity seems to be embedded in many of the concepts we use to analyse language.

The Quinean worry is that having determinate rules of use for an expression creates analyticities. If ‘Bachelor’ is defined as applying to all and only unmarried men, then it becomes analytic that x being a bachelor follows from x being an unmarried man, because there is no way to understand the expressions involved where this is not the case. Since, for the Quinean, there are no unrevisable analyticities, it follows that no rules of use are set in stone; nothing is a rule of use of an expression *as opposed to* a contingent fact about its subject-matter, because a hard distinction doesn’t exist between these two states (e.g.: it isn’t a rule of use that ‘Bachelor’ only applies to men, because it may be revealed to be pragmatically useful to use ‘bachelor’ to refer to some women). The problem that this creates is that the Quinean can no longer distinguish between e.g.: standard and non-standard usage of a word, in the way the Price/Thomasson approach does, because they

hold there to be no fact of the matter about whether the usage accords with the word's rules of use, since these can no longer be delineated.

I don't want to get into this point too deeply here, but it seems that the solution is to point out that even the Quinean presumably wants there to be *some* workable concept of accepted usage, such that we can distinguish, even if in broad strokes, when a word is being used in the accepted way, as opposed to a non-standard way. To so thoroughly eliminate any concept of rules of use that it makes the Price/Thomasson approach unworkable would have the knock-on effect that the concept *Expression* itself would lose coherence, because there would no longer be a distinction between using one expression and using another (I take it that merely uttering a particular string of phonemes is neither sufficient for one to be using a given expression, because of the existence of homophones, nor necessary, because of differences in pronunciation – accepted usage has to feature somewhere). We presumably don't want to go quite this far in eliminating the notion of accepted usage, but even a sketchy version of the notion will do the job the Price/Thomasson approach needs it to do.

This far from concludes the discussion, but, again, I feel that this will best be resolved in the following section, where I will examine further the comparison between Quine's and Carnap's abilities to analyse language within their own respective frameworks.

#### 4: Reconciling Quine and Carnap

Let's review where we find ourselves in the debate. Quine objects to Carnap's deflationism because it relies heavily on a notion of linguistic frameworks that seems to have analyticity built in. On Quine's model, there is no distinction between analytic and synthetic sentences. Carnap disagrees, holding that there are sentences that are squarely analytic, and others that are synthetic. But this is about as far as the disagreement goes. In actuality, we find remarkable agreement regarding the particulars.

For instance, Quine's epistemological pragmatism is a key motivator in his rejecting the distinction. Quine's pragmatism states that any part of a scientific theory is up for revision in the event of contradictory evidence. That is, if we encounter a "recalcitrant experience" – something our theory predicted should be one way but is in fact another – we can choose to revise any part of the theory to bring the experience into line. For example, if we encounter a red thing our theory tells us should be green, we are at liberty to rewire either

the part of our whole theory that concerns which chemical structures produce which colours, the theory of the organization of the colours themselves (e.g.: which colours count as qualitatively identical with which others), or anything else, up to and including logical rules, e.g.: how the conditional works, which governed our conditional inference to the thing's being green in the first place. The only guidelines here are pragmatic: What produces the best overall theory, practically speaking?

This view pushes us to abandon the notion of analytic truths as irrevocable, unequivocal, and eternal. Analytic truths are no longer unconditionally guaranteed under Quine's epistemological holism. The only thing guaranteeing their truth is their usefulness, and that is only contingent on scientific discovery, which may render them obsolete. They are up for revision, and thus not true by necessity. Science tests not individual hypotheses but whole theories, hence the whole theory is up for revision, and no individual posits are intrinsically immune; there is no coherent distinction to be drawn between revising the synthetic statements that make up a view about how the world is, and the analytic statements that comprise the definitions and rules embedded in our language.

We might expect a conflict here, given that Carnap adheres to an analytic/synthetic distinction, but Carnap himself doesn't regard analytic truths as irrevocable. Carnap follows Duhem and Poincaré in accepting the same sort of confirmational holism as Quine (Carnap 1937, p318). The Principle of Tolerance alone shows us this much, as it entails accepting any sufficiently systematic, expressively and empirically viable language, regardless of which posits it takes to be analytic, and whether they appear to contradict those that we ourselves take to be analytic at present, in our language. There is in Carnap, after the Principle of Tolerance is established, simply no room for the concept of analyticities as unrevisable in the sense Quine objects to, or so it appears. So what's going on – why do Carnap and Quine fail to see eye to eye?

We can start to see the how Carnap and Quine actually do differ if we look at the dispute in the following way. Take the case of revising a purported analytic sentence, S. For Quine, it is simply that: We previously took S to be true by definition, that is, true solely in virtue of the words used and S's logical form. Now, because of some recalcitrant experience, it has become expedient to revise some component that went into S's being true. As above, there are no limits on what gets revised, only that the end result be the most pragmatic version of the language available. At the end of the process, we'll still be talking about the same sentence, S, but its truth-conditions (and perhaps its truth-value) will have changed.

Carnap paints the picture slightly differently: We currently speak a language  $L_1$  in which  $S$  is true analytically, that is, by virtue of the rules of  $L_1$ . At some point matters call for a potential revision, and we have before us a proposed language,  $L_2$ , in which  $S$  is no longer analytically true – that is, either  $S$  is false, or  $S$  is true such that its truth isn't guaranteed by the rules of  $L_2$ , but needs investigating empirically.  $L_2$  is a different language from  $L_1$  precisely because it has different rules; it would have to, in order for  $S$  to go from being analytically true to not, since  $S$ 's being analytic was guaranteed by the rules of the language. As above, we have before us a pragmatic decision, but here it involves choosing between *two* languages, namely, deciding whether we ought to speak  $L_1$  or  $L_2$ .

The difference here brings out the sense in which Carnap does regard analytic sentences as unalterable, and unrevisable: Sentences which by the rules of  $L_1$  are analytic will always be true *when spoken in  $L_1$* . We can imagine a world in which a sentence identical with  $S$  in all its particulars (syntax, phonemes, and so on) is *not* analytic, but in order for this to be the case, we would have to change the rules of  $L_1$ . This last step would make  $L_1$  into a different language. It also follows that the non-analytic sentence we are left with – call it  $S_2$  – would not be numerically identical with  $S$ , since the two are each part of a different language. There is, then, a sense in which Carnap's analytic sentences are incapable of being false; since  $S$  can only be stated in  $L_1$ , and will always be true according to  $L_1$ 's rules, there is no situation in which we can state  $S$  and have it come out false. The best we can do is change the rules, and then we have a different sentence in a different language. Note, though, that this is a very thin sense of unrevisability compared to what is being objected to, and does not actually contradict Quine's position that analytic claims can be rejected on pragmatic grounds, and a new language constructed in which they are not analytic.

The distinction between this and Quine's view is that Quine views it as possible to change the analyticities embedded in a language and still have it be the same language, and therefore to change the analytic status of a sentence without making it into a different sentence. Obviously the Quinean wouldn't describe this in terms of analyticities or analytic status – what we would say in the Quinean vernacular is that languages can endure through (incremental) changes to *any* of their licensed sentences, regardless of centrality, and that sentences can endure through (incremental) changes of any sort in the language in which they are situated.

We see this interpretation of the difference in views borne out in a footnote to ESO:

Quine does not acknowledge the [internal/external] distinction which I emphasize above, because according to his general conception there are no sharp boundary lines between logical and factual truth, between questions of meaning and questions of fact, *between the acceptance of a language structure and the acceptance of an assertion formulated in the language.*

ESO p32, fn2 (emphasis mine)

Here Carnap links acceptance of the analytic rules of a language with the acceptance of a language structure. If accepting a set of analyticities means accepting a language structure, then rejecting it in favour of a different set of analyticities means rejecting one language structure in favour of another. Carnap thus draws a distinction between disputes where parties speak the same language but hold different positions, and disputes where parties simply speak at each other in different languages that are similar enough to cause confusion, and draws it along lines cashed out in terms of the analytic/synthetic distinction (as described in section 2 above). Since Quine denies this distinction, he cannot draw this line here, and must provide some other criterion for sameness of language, and likewise sameness of sentence, if he is to hold that these concepts hold up (I will discuss in the following section whether the Quinean can attempt to draw these boundaries).

The next move will not necessarily surprise Carnapians. The debate, articulated in this way, is shaping up to fit the classic template of a deflatable metaphysical dispute, albeit at the metalinguistic level. But disputes about the identity conditions of languages and their sentences are as much metaphysical disputes as those about the identity conditions of tables and their parts, and therefore just as primed for deflation. One way to put this is to say that Quine and Carnap want to apply the sortal concepts 'Language' and 'Sentence' as they occur in the metalanguage in different ways, i.e.: as attached to different identity conditions for their respective entities. At this point, the thing to do is to pull out the Principle of Tolerance and apply it even-handedly to get two (meta)languages we can assess for pragmatic usefulness.

Note that this procedure applies whether we favour a Carnapian or a Quinean approach. Carnap would regard Quine's rejection of the analytic/synthetic distinction as the proposal of a different language-form, since it modifies what are, in the Carnapian metalanguage, analyticities, but Quine's pragmatism means that he too has to regard the proposed change in light of its pragmatic value. This is somewhat at odds with the tone of *Two Dogmas*, which seems to take issue with the analytic/synthetic distinction on conceptual grounds.



The key point for us here is that it would be illegitimate for Quine, either on his understanding of language or Carnap's, to lodge such a conceptual complaint. For consistency, *Two Dogmas* ought to be read as claiming, albeit in a roundabout way, that a metalanguage that does not contain the concept of analyticity has more pragmatic utility than one that does, or else that the inability to define analyticity non-circularly entails, in some way, a lack of such utility.<sup>4</sup>

Both seem quite far from being demonstrably the case, at least given the arguments in *Two Dogmas*. In a later paper (Quine 1951b), Quine comes somewhat closer, explaining that the problem is that Carnap has nowhere principled to draw the line between analytic truths and synthetic facts. Although the paper in question is fairly commonly dismissed as failing to appreciate Carnap's actual position, this much at least seems to follow from the arguments in *Two Dogmas*. This is the proper place to leverage the Grice/Strawson argument given above; the inability to find a satisfactory definition for the concept doesn't invalidate its utility value if we can demonstrate that it is used consistently and to good effect. That our language isn't up to the task of providing a definition doesn't impugn the concept's utility, hence it is not really open to the Quinean to attempt to avert this sort of deflationary take on the debate on those grounds – indeed, strictly speaking we should view *Two Dogmas* as attempting just such an approach.

##### 5: Executing the deflation

First, we should clearly define the difference between the proposed usages. Carnap's view is analogous to the essentialist in an ordinary metaphysical debate about identity (or perhaps persistence) conditions. Carnap holds that the analytic rules associated with a language are essential to the language, in the strictly minimal sense that they are the parts of the language that, if removed or replaced, mean that it ceases to be the same language as before. To reiterate, these analyticities are not held by Carnap to be irreplaceable as such – it's just that if we do change them, we are no longer speaking the same language, since we've changed the rules. We can see this approach in ESO, where Carnap says that

---

<sup>4</sup> Note also that whichever metalanguage we opt for – Quinean or Carnapian – we end up with a similarly deflationary approach to metaphysics, since Quine is in theory no more the friend to “hard” metaphysics than Carnap. I think that this actually demonstrates nicely the sort of equivalence that, for instance, Hirsch struggles to articulate in terms of intensional equivalence (see Ch.3); both frameworks have the same take on metaphysical debates, just articulated very differently, because of their different uses of the relevant sortal concepts.

“The acceptance of a framework of new entities is represented in the language by introduction of new forms of expressions to be used according to a new set of rules.” (ESO, p30)

It’s worth noting as well, with an eye on the following section, that this view cues us up to give a set of identity conditions for sentences with at least some of the nuance these conditions would require; “two” sentences can be seen to be numerically identical iff (or at the very least, only if) they are syntactically identical and are constructed in the same language. If the concept of the numerical identity of sentences is mysterious, we can find a useful contrast with other sorts of relations between sentences, e.g.: sentences being intensionally identical (having the same truth-conditions), which follows from numerical identity but does not entail it, and sentences being phonetically identical, which neither follows from nor entails numerical identity. This points to the relation of numerical identity for sentences as fairly straightforwardly usable, even if it may seem obscure what, metaphysically, it amounts to. Viewed through the lens of Carnapian Tolerance, the fact that it seems obscure e.g.: how one would break this notion of identity down into naturalistic terms shouldn’t discourage us from adopting it if it has a fairly consistent and accepted application (*cf.* Price 1997, pp258-261).

The picture is not so straightforward for the Quinean framework, although that is not in itself a problem. Quine brings up the idea of individuation conditions for languages in discussing “eternal sentences”. In the majority of his work, Quine makes use of the concept of an eternal sentence as the main vehicles of truth, as a replacement for the concept of a proposition. This is because the truth-value of an eternal sentence is invariant under most conditions. E.g.: The sentence “John believes that p” will cease to be true if John later ceases to believe that p, whereas the eternal sentence “At time t, John believes that p” will, if ever true, continue to be true regardless of the time, place etc. of its utterance (Quine 1960, p208).

The difficulty here, as Quine notes, is that the one factor with which eternal sentences *do* vary is language. That is, one could have two languages in which phonetically the same sentence appears, but with different truth-conditions in each. Hence we have to package our eternal sentence with a variable, or as Quine puts it, “a tacit subscript on the word “true,” or “eternal,” specifying the language.” (Quine 1969, p142) This is a problem because Quine does not have a satisfactory set of individuation conditions for languages, and is rather sceptical of the possibility for such, since he suspects that such a way of

individuating languages will have to invoke semantics. In the end, he bites the bullet, and sketches a picture on which the vehicles of truth are “eternal sentence utterance events” – individual utterances, by a particular speaker, at a particular place and time, and in a particular language, of sentences that are, in that language and at that time, eternal sentences.

This is a subtle move by Quine. The language parameter is still present, so languages still need to be individuated, but the parameter is now attached to the individual utterance event, not the eternal sentence itself. Hence Quine dodges the problem of having to individuate languages semantically as eternal frameworks of rules or meanings, and can instead attempt to cash them out in behaviouristic terms, namely as the speaker’s “total present speech dispositions”, i.e.: their dispositions, at the time of the utterance of the eternal sentence, to respond to specific stimuli with specific utterances (*ibid.* p144).

We can give the beginnings of a critique of Quine’s proposal for language individuation to match his critique of Carnap’s, remembering that the point here is not to establish one over the other, but to look at the relevant difficulties accumulated by each that might figure in the pragmatic decision between them. We might first find fault with Quine’s reliance on dispositions. Quine emphasizes the need to distinguish between different senses that one individual could attach to phonetically the same sentence if they spoke two distinct languages (had two distinct sets of linguistic dispositions) that could each potentially license phonetically identical utterances. Quine’s framework seems, potentially, to solve for this limit case, which he himself notes is extremely contrived, but this is not all that we need the language parameter for.

Applications to other cases are revealing. For example, the model should be able to tell us when we have two sequential utterances of the same (eternal) sentence, as opposed to different ones. At a glance, the answer is that the sentences, assuming they are phonetically identical, are different when the speaker’s speech dispositions have changed, since these were what determined the language parameter. However, we get unintuitive results if we apply this completely flat-footedly; the phrase “total present speech dispositions” (*ibid.*) seems to imply that *any* dispositional change, no matter how small, will result in a change in language, which seems at odds both with intuition and with Quine’s attitude to languages as fairly robust regarding changes in their component inferences, as sketched above in §4. Indeed, we might think it almost impossible to state the same truth

twice, as even retaining a memory of saying it a first time will have *some* effect, however miniscule, on the circumstances in which we are disposed to assert it again.

However, if we allow for some small dispositional changes without a corresponding change in language, we have essentially fallen into the same old problem of having to specify a boundary point between those core dispositions that are of the essence of a given language, and those that are not, without a principled place to draw the line. This is exactly analogous to Quine's complaint that Carnap has no principled place to draw a line around the set of truths considered to be analytic (and therefore essential) for a given language (Quine 1951b again), only dealing with dispositions instead of sentences. It thus makes a compelling second horn to the dilemma.

Another way to put the problem is this: We can produce a theory on which to count a sentence as being spoken in the same language as another (possibly numerically identical) sentence is to note that both are uttered by speakers with all the same speech dispositions, if we just individuate dispositions in an optimally fine-grained way – that is, not counting *as a speech disposition* anything too minute to trigger a change in language if it were abandoned, nor discounting anything sufficiently large that abandoning it would obviously bring about such a change. But this just makes the matter one of finding what counts as a speech disposition, instead of which speech dispositions are core to a language. Aside from anything else, this question about dispositions appears no less obscure than the question of which truths count as analytic for a given language L, and which don't. Indeed, given the Grice/Strawson argument, we might even be able to answer that question consistently, at least on a case-by-case basis, where this one just appears mysterious.

So the criteria for individuating languages that Quine provides look to be at least as prone to arbitrariness as Carnap's. But what is the upshot here for the Carnapian? Even if we have successfully deflated the debate, i.e.: shown it to amount to a pragmatic choice between two conceptual frameworks, as I think we have, this doesn't show that the Carnapian way is the correct one. Indeed, it would be wrong to label either version the *correct* approach, as opposed to merely the more workable or useful.

One way to argue for siding with Carnap would be to point to the utility value of the concepts that a Carnapian approach allows us for dealing with metaphysical questions on a semantic basis. Quine obviously rejects the use of explicitly semantic concepts such as 'Meaning', 'Intension' and 'Synonym'. The Carnapian, on the other hand, is able to put these to use in demonstrating metaphysical disputes to be linguistic in nature. For instance,

Eli Hirsch puts forward an “intensional deflationism” in a Carnapian spirit. Hirsch points out that disputes about e.g.: mereology actually consist in proposing different hyper-intensional structures for acceptable sentences (Hirsch 2016, p111). For example, while the mereological nihilist might hold the sentence “There is a table” to be false in all circumstances, they will probably allow the truth of some sentence that is intensionally equivalent (i.e.: It is true and false in all the same circumstances), such as “There are physical simples arranged tablewise.” Hence such mereological disputes come down not to matters of fact, with disputants arguing for different ways the world could be (in the sense of backing intensionally different sentences), but instead down to a dispute over preferred language, with each side recommending a different hyperintensional structure for the disputed intension.

Semantic concepts of the sort that Quine would find suspicious are obviously deeply embedded in such a theory. This means that the strict Quinean will not have access to this sort of reasoning, and will thus not be able to deflate metaphysical debates by using, for instance, the concept of intensionality to point out how disputants appear to be speaking with each other when they take themselves to be disagreeing. In effect, intensionality, as one of the concepts in the “analyticity family” (per Grice and Strawson), is useful enough to earn its keep, and remain part of our vocabulary for talking about language.

There are also concepts that wear their relation to analyticity less boldly on their sleeves, but that nevertheless stand and fall with the “analyticity family”. As we saw above, the concept ‘Language’, used as a sortal, is, as usually construed, highly dependent on semantic individuation. Hence Quine does not have access to the more everyday deflationary observations (i.e.: ones less dependent on obviously semantic concepts such as intensionality), such as that the utterances of speaker x and speaker y are incomparable in virtue of their speaking different (but phonetically similar) languages. Given the problems above for individuating languages in the absence of a concept of analyticity, there also exists, for Quine, the problem of individuating a host of other linguistic items whose identity or non-identity in different instances depend on whether the same or different languages are being used in those instances. For example, it is plausible to think that two words are not numerically identical – are not the same word – when they occur in different languages, irrespective of their phonetic similarity (although assuming different usage patterns in the different languages).

Absent an account of language differentiation, Quine does not have the resources to differentiate words in this way, meaning we are denied the use of the sortal 'Word' for deflationary purposes (consider how many times in Philosophy we hear the phrase "It depends what you mean by 'x'" – meaning roughly: It depends what the accepted usage patterns for 'x' are in your language). Other examples of such analyticity-dependent expressions are 'Sentence', and even 'Expression' (when is an expression not the same as another expression? Answer: When it is attached to different proprieties of usage – but Quine has no access to this notion, and cannot easily re-engineer it using dispositions). Quine himself uses these terms freely. However, as I hope to have shown above, it doesn't appear that they are as easily eliminable from the account as Quine had hoped.

#### 6: Going forward: A workable system

Looking ahead to the thesis to follow, the Carnapianism I will end up developing will not be entirely like that of other deflationists to have come along since Carnap. The key factor is in the metasemantics that supports the deflationary stance from behind the scenes. Hence despite our making room for a Carnapian theory, we do not yet have a full picture of what a workable version of such a theory would look like. One way to explain this is as follows: What the non-verificationist responses surveyed in section 3 have in common is that they fail to appreciate everything that verificationism was doing for the deflationist, albeit at a cost. Some things are accounted for, for instance, the fact that verificationist reductionism provides a basis (conditions of confirmation) on which disputants can be said to be saying the same things, despite employing apparently conflicting linguistic devices, which Hirsch adapts to use the intensional profiles of the speakers' respective utterances to show their *intensional* equivalence.

What goes largely unaccounted for in the newer theories, however, is the fact that verificationism gave the deflationist a *non-representationalist* basis for their metasemantics. To cash this out a little, the verificationist doesn't require for e.g.: a sentence to be meaningful that it bear a representational relation to the state of affairs it is about, where this relation can be cashed out in a non-trivial way. Instead, verificationism only requires that the sentence have empirical conditions under which it can be confirmed or disconfirmed in order for it to be meaningful. This allows the verificationist to be neutral as regards the world's metaphysical structure, where the representationalist is inevitably drawn into doing substantive metaphysics.

To illustrate briefly, consider that taking reference, for example, to consist in a substantive relation, *R*, of words to things already entails an ontology of things eligible to be referents. Hence we could conceive of a perfectly useful language with what looked like referring terms, but where these failed to refer to anything, i.e.: because there was nothing in the world to which they stood in *R*. This obviously violates Tolerance, but also leaves no room for metaphysical deflationism, since we can always ask whether a language actually shows the world as it *really* is, i.e.: whether its terms *really* stand in *R* to anything, regardless of appearances. Likewise, as I argue in Ch.3, we get the same problem at the propositional level if we take propositions to delineate sets of states of affairs within a structured field of possible states of affairs – this structure is just as metaphysical as the more obvious ontological structure in the previous example, and leads to more obviously metaphysical commitments besides.

To put it simply, we were right to ditch verificationism as an unreasonable underpinning for our deflationism, but in doing so, we don't quite seem to have appreciated (going by the available neo-Carnapian literature) exactly what it was doing for us in the first place. Hence, while they offer an excellent blueprint for the most part, the most prominent neo-Carnapian positions currently on the market harbour at least some strain of representationalism, and therefore of metaphysical realism of the "heavy-duty" sort that undercuts their deflationism. The only response available to the committed Carnapian is to find some way to put these theories on a thoroughly non-representationalist footing, as I will endeavour to do, by articulating a global expressivism that supports the neo-Carnapian agenda.

## Ch.2: Characterising global expressivism

In order to approach the task of stripping representationalism out of our neo-Carnapian deflationism, we need to get some idea of where we're going. Representationalism and its antithesis, global expressivism, are not immediately unambiguous positions. Articulating what exactly the global expressivism that is our goal consists in is therefore a fairly significant task in itself, and one that we will have to discharge, at least partly, before seeking to put it to use.

Matthew Simpson (2020b) argues that this cannot be done – there is no coherent way of outlining an expressivism that is truly global. Over the course of his paper, Simpson presents several candidate definitions for global expressivism, finding each one wanting in one way or another. The upshot of Simpson's argument is that it is unclear exactly what global expressivism is supposed to be. According to Simpson, insofar as self-professed global expressivists have a coherent position at all, they are in fact *local* expressivists; it is unclear in what sense their view is supposed to be global, since no plausible global version exists that they could be said to hold.

I will take one of Simpson's early draft definitions and lightly rework it to be sensitive to an important distinction between different types of explanations of meaning, namely semantic and metasemantic explanations. The resulting position, I argue, is what global expressivism really amounts to. I'll then go on to look at the key motivation for such a position, arguing, for good measure, that this actually hedges out *local* expressivism, by the same token that Simpson rebuts the global version. Hence global expressivism is not only a viable position, it is the *only* viable expressivism on the market in the modern day. The position yielded by the discussion in this chapter will enable us to go on to examine the complete metasemantic picture the neo-Carnapian needs for their deflationary programme to succeed.

### 1: Simpson's argument

Simpson characterises expressivism as a suspicion of representational concepts as applied in the analysis of the meaning of language (*ibid.* p143). As Simpson notes, though, here's a



problem with specifying exactly what form this suspicion takes. The most straightforward interpretation would be to have the expressivist treating the discourses they deem expressive as non-representational, and therefore non-cognitive. However, modern expressivists, local and global, don't subscribe to non-cognitivism about expressive discourses. For one thing, a non-cognitivism like that of A.J. Ayer (1952) wouldn't globalize well; we would have to hold that all of language, including the language used to state the theory, was non-truth-apt. Hence if we want to understand local and global expressivism (hereafter "localism" and "globalism", following Simpson's usage) as different versions of the same view, just with different generality, then we can't treat them as denying representation to whatever portion of language they deem to be expressive.

The modern localist gets around non-cognitivism by adopting semantic deflationism for the discourses they deem to be expressive. For instance, the modern ethical expressivist doesn't deny that ethical utterances can be true. What they do instead is offer a different understanding of truth as it is used to characterise ethical utterances – different, that is, from how the localist themselves takes it to apply to language more generally. The sort of truth that ethics gets is a "thin" or "deflated" version, as opposed to the "robust" truth of straightforwardly representational language (*cf.* Blackburn 2015, p843). The same goes for other semantic properties and relations such as reference; expressive discourses get the deflated version, while the rest of language gets the robust version. Hence localism is saved from non-cognitivism by deflationism. The version of expressivism that this yields also globalizes better than non-cognitivism, since we now at least have a vocabulary in which to state the theory, even if it does turn out that it is only true in a thin sense – so are all other truths.

As Simpson rightly notes (Simpson 2020b, pp144-6), the thin/robust distinction bought by deflationism doesn't by itself clear things up particularly well; there is still the question of what makes representation robust, as opposed to thin. Simpson here leverages a key insight (fleshed out in Simpson 2018 & 2020a), which he calls "explanationism".

Explanationism treats the expressivist not as making a claim about the features of language *per se*, but about how the meaning and use of language is best to be explained – what form the best explanation of a particular expression or sentence would take. Simpson proposes to understand treating a term's being robustly representational as holding that the *best* explanation of the term's meaning involves citing the subject matter of the term itself (Simpson 2020b, p150). For example, a robustly representational explanation of the meaning of the word 'tree' might incorporate the phrase "'Tree' is true of...", and then go

on to identify things to which it is appropriate to apply the term by their properties – one way of homing in on a term’s subject-matter (Simpson 2020a for more discussion of this notion of subject matter – I do not intend to dispute the notion here).

So we have a picture of the globalist as denying that any linguistic expressions are ever robustly representational – meaning that there are no expressions whose meaning is best explained by pointing to their intended subject matter. Simpson renders this as the principle (GSM), standing for “Globalism: Subject-Matter”:

(GSM) There is no term *t* such that the best explanation of *t*’s meaning and use cites *t*’s subject matter.

Simpson 2020b, p150

We also get a localist version:

(LSM) For some but not all terms *t*, the best explanation of the meaning and use of *t* doesn’t require mentioning *t*’s subject matter.

*Ibid.*

The problem Simpson notes with (GSM), and the reason he moves on to consider alternative definitions of globalism, is fairly straightforward: There are terms that are supposed to track objects within the environment, as part of their core function. For instance, the correct usage of the word ‘tree’ has to be sensitive, in some sense, to the presence of actual trees within the environment (or at least has to be sensitive to actual trees in some way). If it were not, then it wouldn’t be the word for *tree*. Hence to give an explanation of the term that left out the fact that it tracks trees would be to elide an important fact about its meaning – what gives it the distinct meaning that it has – which a better explanation would incorporate. It follows that (GSM) is false, because the best explanation of environment-tracking terms such as ‘tree’ necessarily involves citing their subject matter.

In the following sections, I will argue that this analysis, in fact, is not sensitive to an explanatory distinction – the distinction between semantics and metasemantics – that is key to globalism. Hence the real problem with (GSM) is that it is underspecified, in a way that actually takes little adjustment to the principle itself to counter. Once we have made this adjustment, we will have a viable globalism.

## 2: The semantic/metasemantic distinction

Is it true that the expressivist has to attest to some sort of tracking relation, or miss out an important fact about meaning for environment-tracking terms? I do not intend to argue the point; let us grant that an explanation of meaning is deficient if it leaves out an account of the relation in which environment-tracking terms stand to their subject-matter, and that a better explanation would include such an account, in some terms. Although Simpson puts it broadly (too broadly, as I will argue) there is a convincing idea in here.

If we are to understand exactly what the importance of subject-matter means for the expressivist, we have to be clear on what sort of explanation of meaning we are talking about. A myriad of explanatory strategies could answer to the phrase “explanation of meaning”. The question, then, is: Should we read the globalist’s claim, (GSM), as restricted to a certain *type* of explanation of meaning, or as ranging over all types? I’ll argue that we should understand the globalist as making a claim purely about *metasemantic* explanations of meaning, as opposed to other types of explanations, such as semantic explanations – terms that I’ll define in due course. It is unclear whether Simpson uses this restricted reading (I do not think so), but regardless, this reading seems to yield a coherent globalism, where the alternative broader reading does not.

In discussing an earlier draft of the core expressivist credo, Simpson states that the expressivist’s characteristic explanatory strategy has two parts. The first is metasemantic, and involves offering an account of how the meaning of the term in question arises – what makes it mean what it does? The second is pragmatic, and involves explaining the characteristic utility of the term – what does it allow us to do that we couldn’t do without it (*ibid.* p147-8)? Simpson stipulates explicitly that an explanation of a term *t*’s meaning and use, in the context of this candidate credo, means this sort of two-part, metasemantic/pragmatic explanation. Importantly, he goes on to discount semantic (as opposed to metasemantic) explanations from consideration:

[W]hile other projects – for instance explaining the difference between declaratives, imperatives and interrogatives – would ordinarily count as explaining meaning, or explaining language, I don’t mean such things here.

*ibid.* p149

It is unclear whether Simpson still has the same sort of explanation in mind by the time he comes to use the phrase “explanation of meaning and use” in (GSM). For consistency, we

should assume so, and that Simpson thus takes (GSM) to be placing a restriction only on two-part metasemantic/pragmatic explanations of meaning. However, it is not entirely obvious that this restrained reading of (GSM) suffers from the implausibility Simpson attributes to it. Recall that Simpson criticizes (GSM) because it forces us to leave out mention of tracking relations that are essential to the meanings of the terms being analysed. Simpson's argument is that this would hamstring the globalist's ability to offer explanations of the meanings of environment-tracking terms, because these explanations couldn't acknowledge the importance of the subject-matter of these terms, which so clearly plays an important role in determining their meanings. But if (GSM) only places a restriction on the explanatory techniques suitable to metasemantic/pragmatic explanation, and not *any* explanation of meaning whatsoever, it remains open to the globalist to offer an account of tracking relations in other terms, for instance, in semantic (as opposed to metasemantic) ones. This would allow the globalist to acknowledge the importance of subject-matter for determining the meanings of environment-tracking expressions after all – they would just have to do it in other terms than metasemantic/pragmatic ones.

Indeed, we can start to see why one might be motivated to take up such an expressivism, namely because one thinks that citing an expression's subject-matter as an explanatory technique is not suitable to the particular sort of explanation the expressivist is concerned with, though it may be suitable for others. What problems, exactly, subject-matter involving explanations might cause will be the topic of the following section. For now, let us note that the globalist who holds (GSM) can still offer explanations of meaning in semantic terms – that is, making use of semantic concepts – that cite the particular expression's subject-matter. For instance, it is still open to them to explain that the expression 'Tree' has the meaning it does because it is true of trees ("truth-of" being the semantic concept doing the explanatory work for us here).

To be clear, I am taking semantic explanations here to be explanations of meaning that use semantic concepts, such as reference or truth. This is different from – and more precise than – the idea of semantic explanation that Simpson gestures towards. On Simpson's view, given above, an explanation of meaning is deemed semantic or metasemantic based on what the goal of the explanation is – whether it is for explaining the divisions between different sorts of sentences or expressions ("explaining the difference between declaratives, imperatives and interrogatives"), or whether it aims to explain what it is that makes an expression mean what it does. I propose to distinguish these different categories of explanation instead based on the concepts they use; semantic explanations of meaning

use semantic concepts and cite semantic properties and relations (truth, reference), while metasemantic explanations look to settle what exactly these properties and relations consist in.

Note that this does not track a difference in generality; on this definition, one can give a metasemantic explanation of the meaning of a particular expression, as opposed to using metasemantics only to cash out whole semantic concepts, although it can certainly be used for that too. The metasemantic explanation of a particular expression will simply eschew semantic vocabulary and concepts, instead cashing them out in other terms, for that particular case. For instance, a metasemantic account of why the expression 'Horse' means what it does might attempt to cash out the semantic concept of satisfying a predicate by explaining the lawlike causal relation that exists between the occurrence of essential horse-properties in a speaker's environment and the tokening of the expression 'Horse'. The point is that this is a metasemantic explanation of a specific expression, because although it analyses the meaning of a particular linguistic item, it cashes out a semantic concept into other terms in order to do so.<sup>5</sup>

As part of their semantic deflationism, the globalist allows, and is even committed to, the possibility of offering at least one sort of subject-matter-involving explanation; "'Tree' is true of trees", is true by globalist lights. What the globalist denies is that we can give a rendering that entirely removes any reliance on semantic concepts. So we can have "'Tree' is satisfied by trees", or even "'Tree' tracks trees", provided we acknowledge that behind the scenes, our understanding this sort of explanation is going to depend, ineluctably, on our prior understanding of how to apply the expression 'Tree'. Hence our understanding this sort of explanation at all depends on our possessing a semantic competence (to apply the expression 'Tree') that itself goes unexplained.

(GSM), then, rather than prohibiting *any* sort of account that makes mention of subject-matter, or even any account that posits a tracking relation broadly construed, should instead prohibit constitutive accounts of what this sort of relation consists in, where these accounts cite the subject-matter of the expression – in other words, subject-matter-involving *metasemantic* explanations.

---

<sup>5</sup> On this definition, primitivism about semantic properties involves a metasemantic quietism, because it holds that semantic properties are unanalysable. I discuss primitivism, and whether it can be distinguished adequately from deflationism, in §6 below.

Simpson's argument for the implausibility of (GSM) thus looks to rely on an ambiguity about how widely we take (GSM) to commit the globalist – whether they are placing a restriction on all explanations of meaning, or just those of a particular metasemantic/pragmatic sort. Taken the broad way, Simpson's argument holds; the globalist who denied that there could be *any* useful subject-matter-involving explanations of meaning would be implausibly overcommitted. But this is not what the globalist says, given their own deflationary approach to semantics – “‘Tree’ tracks trees” comes out true, if only in some “thin” sense, so there is no obvious reason the modern expressivist would want to gainsay its utility in accounting for at least part of the meaning of ‘Tree’. Let us, then, go with the reading of (GSM) that yields the more viable theory, this being the narrower one: Specifically *metasemantic* explanations that cite the subject-matter of the term being explained are prohibited, but others are allowed.

It is potentially the use of the concept of “tracking relations” (*ibid.* p145) that was leading us astray here. The phrase itself conjures images of lawlike causal/counterfactual relations between a speaker's usage and their words' respective subject-matter, thus priming us for an account of this sort (e.g.: Fodor 1998, pp73-5). But just because the expressivist prohibits accounts that involve this sort of tracking doesn't mean that they can't give *any* explanation that acknowledges that terms track their subject-matter. Note again that features such as reference and truth, construed simply and semantically, are tracking relations too, and are apt to be more extensionally accurate than any metasemantic relation we care to construct, as the following section will show. That such an explanation (“‘Tree’ is true of trees”, where we are prohibited from making constitutive, non-semantic claims about what “truth-of” consists in) might appear trite or not the sort we had in mind to account for the importance of subject-matter is a different criticism, and one that is more easily disputed, than Simpson's claim that the globalist *simply cannot* account for the role of subject-matter with the explanatory resources they leave themselves.

Let us clarify (GSM), then, to reflect the more precise, more plausible reading:

(GSM)\* There is no term *t* such that the best *metasemantic* explanation of *t*'s meaning and use cites *t*'s subject-matter.

We can start to see, in this formulation, where the globalist sympathy for pragmatism comes from; when we rule out representationalist metasemantic explanatory techniques, pragmatist explanatory techniques are among the more attractive of the tools that remain at our disposal (*cf.* Price 2011, p232). As will be discussed at length in Ch.7, the pragmatic

portion of the explanation serves to supplement the semantic portion, which, as noted, can appear somewhat bare or uninformative. As we will see in following chapters, the trick to successful expressivist linguistic explanation lies in deriving, from the bare semantics of an expression, an account of its broader linguistic function, and thereby of its utility to the language user.

### 3: Deflationism and rule-following

(GSM)\* is a very specific philosophical view – it has a very niche job, namely placing a restriction on an explanation of the meaning of an expression such that it cannot both mention the subject-matter of the expression and be given in entirely non-semantic terms (per the understanding of metasemantics I have advocated above). However, given our starting point in Simpson’s candidate definitions, it has not yet been important to ask *why* one might hold such a methodological position as regards explanations of meaning. In this section, I will explain what I take to be the core motivation for expressivism – the phenomenon of semantic “placement problems” – and will look at specific symptoms that seem to extend to all of language, necessitating a global form of expressivism, taking the form of a global semantic deflationism.

The core motivation for (GSM)\* is anti-reductionism. The idea here is that the semantic properties of an expression are not the sort of thing that can be reduced to properties that are more naturalistically acceptable, i.e.: that are deemed basic or legitimate by natural science.<sup>6</sup> Huw Price points to the reductionism to which this is opposed as a “proto theory” of language: The relations in play that determine a sentence’s truth conditions, for instance, must be relations of a naturalistically respectable sort (hence so too must the relata), otherwise it will turn out that the sentence doesn’t really *have* truth conditions (Price 2011, p4). Analogously, what makes a referring expression mean what it does is that it stands in a relation of reference to a given entity, where both the entity and the relation must be of the sort compassed by naturalism. Calling this a “proto-theory” signifies that it embodies assumptions about how language ought to work, at which we arrive easily and unwittingly – recall the ease with which we lapsed into a representationalist theory with

---

<sup>6</sup> While other reduction bases may be possible, I take it that naturalism quite obviously main player here, and also that the arguments given against naturalistic reduction will work analogously for other reduction bases.

the “Fodorian imagery” above, which obscured the possibility of a global expressivism earlier.

Price’s point in drawing attention to the “placement problems” that affect such apparently non-natural things as ethical properties, abstract objects and the like is to speak against a particular interpretation of the duties of the naturalist. The point he wishes to make is that the naturalist’s key aim should not be to cash every sort of discourse whatsoever in terms of an ultimate reduction base, namely that domain of entities recognized by current natural science; in fact to do so is to contravene the naturalistic method by assuming this role for such entities against evidence from experience (e.g.: of moral agents) that this is not all there is.

Placement problems form a part of this sort of evidence; philosophical contortions such as non-naturalism and, indeed, local expressivism are responses to the various placement problems. Local expressivism about ethics, for example, is an attempt to explain how we should think of the meaning properties of ethical language as being determined, and what status these meaning properties have, given that the things that seem to be referred to by ethical discourse don’t appear to be natural, and resist clean reduction to natural properties.

The placement problem that concerns us shows in the difficulty one has in trying to make claims about what, exactly, linguistic tracking relations amount to. Recall that we can easily, even trivially, cite the relation between a term and its subject matter in terms of the semantic features truth and reference, thereby rendering a semantic explanation of the link to subject-matter, but – the globalist at least holds – this sort of relation resists characterisation in other terms. This makes some intuitive sense, since semantic vocabulary is, in a manner of speaking, specialized to this task in a way that other vocabulary is not.

Recall that for an explanation of meaning to be a metasemantic one, in the sense of (GSM)\*, it will have to offer a characterisation of the relation between the expression and its subject matter in non-semantic terms. The globalist’s prohibition, then, is against explanations that run afoul of the placement problem by attempting the impossible – a constitutive account of the relation between a term and its subject-matter in non-semantic terms. The solution is deflationism: Continue to use the semantic concepts in play, including for meaning-explanatory purposes, and don’t worry that they don’t reduce to naturalistic (or other) concepts, as this doesn’t compromise their utility.



For symptoms of the placement problem that exists here, we can look to Paul Boghossian (1989) on rule-following. Boghossian notes that, in attempting to delineate the correct rule for use of the word 'Horse' in naturalistic terms, we will have to account for exceptions to avoid building in false-positives. For instance, a simple attempt at such a set of conditions might hold 'Horse' to be used correctly iff there are horses present in the speaker's environment at the time. But this criterion would have it that certain attempts that we would deem inaccurate come out as being correct, e.g.: attempting to use the term 'horse' of a cow in a field otherwise full of horses. The speaker's use satisfies the conditions that there be horses present in the environment, but the word is still not being used correctly, hence those conditions do not properly account for the word's content.

The aim here is to capture the correct use-properties for the expression 'Horse' by specifying, in non-semantic terms, the relation that uses of the expression have to bear to its subject matter in order to count as being used correctly. As Boghossian notes, this means, as a minimum, specifying in non-semantic terms a situation in which the expression is indisputably being used correctly (*ibid.* p538). If we can do this, we stand a chance of generalizing which features of the situation are and aren't essential to the expression's correct use.

However, outlining this set of "optimality conditions" looks to be an infinite job of work. We are here looking to specify a set of features such that a speaker, S, using the expression 'Horse' in a situation that meets our specifications could not possibly be using it incorrectly. However, it appears that no such set of features we care to specify could be sufficient to produce such a result, simply because the vectors for error are too many and too various. For instance, as Boghossian notes, it is not just environmental features we need to worry about, but also the speaker's background beliefs. Even if we could create a watertight set of environmental conditions, there is still the possibility that S will harbour some background belief that will dispose them to use 'Horse' in respect of non-horses (*ibid.* pp539-40). The belief that the owner of the field only owns horses, for instance, could dispose S to use 'Horse' of a partially glimpsed cow. In order to specify our optimality conditions, we have to find a way of excluding all these possible sources of error originating in our background beliefs, but the sheer variety and availability of such disruptive beliefs (Boghossian (*ibid.*) suggests belief in Papal infallibility as an example) speaks against the possibility of doing so.

It might be objected that we are here neglecting an obvious route out of this problem for the naturalist about semantic properties, namely that of taking up a causal theory. Such a theory would, broadly speaking, state that the conditions on S's correctly using 'Horse' come down to S's tokenings of 'Horse' being somehow causally related to actual horses in S's environment. If successful, this would eliminate in one fell swoop the myriad of environmental causes of error that threatened to bypass our earlier attempts at optimality conditions.

Indeed, a causal theory would appear to be more efficient at weeding out these causes of error, but it does come with problems of its own. It is not enough simply that the tokening of the term 'Horse' be caused by some actual horse – it must also be caused in an appropriate way. S could hardly be counted as using the expression correctly if they uttered it in response to a horse accidentally jogging an electrode in S's brain, causing them to utter the word compulsively.<sup>7</sup> Clearly we need to specify which causal chains, originating in actual horses and ending with S's utterance of 'Horse', are and aren't appropriate to confer correctness upon S's use (this is to establish the rudiments of information semantics – *cf.* Fodor 1998 pp12-15). But now it appears we are back at square one, since the wide variety of error-producing causal chains appears no less than the variety of error-producing beliefs in the non-causal theory above.

What the reasoning here suggests is that we cannot, using only non-semantic vocabulary, home in on the conditions under which correct application of an expression is guaranteed by the properties it is supposed to denote. That is to say, there is no extensionally adequate account of an expression's meaning in terms of its subject-matter that is not given in semantic terms. By contrast, we can (and the expressivist does) rely on semantic terms to create extensionally adequate explanations of the meaning of expressions, e.g.: the simple formulation "The expression 'Horse' is true of horses." This is the motivation for the deflationary approach of (GSM)\*, namely that, put simply, attempting to unpack such clauses into non-semantic terms, such as the naturalistic terms with which the representationalist is chiefly concerned, while retaining the extensional adequacy of the explanation, appears to commit us to an impossible task. Note again that this means that the globalist is happy to talk in semantic terms – indeed, a willingness to admit that 'Horse' is true of horses, 'Tree' is true of trees etc. is core to their account. What the globalist

---

<sup>7</sup> This effectively turns the argument into a semantic form of the ubiquitous problem noted by Edmund Gettier (1963) for defining knowledge in relational terms.

objects to is the idea that these semantic properties and relations can be adequately unpacked into other terms, such as naturalistic ones, or that we need to be able to provide an account of semantic properties and relations in a metaphysical register in order that such properties be accounted real.

#### 4: Williams and Price

This helps us to answer Simpson's linked point, that globalists themselves admit that there are tracking relations that obtain between our use of expressions and their subject-matter. If this is true, then (GSM) is not an adequate reading of globalism – but we will need further investigation to see whether this affects (GSM)\*. Simpson's evidence for globalists' accepting subject-matter-involving explanations of meaning comes from Michael Williams (Simpson 2020b, p151). Williams uses a reliable discriminative reporting disposition or "RDRD" as part of an explanation of the meaning of the expression 'red', as follows:

To master 'red' in its reporting use, the speaker must have a reliable discriminative reporting disposition (RDRD): a disposition, given appropriate motivation and conditions, to report 'x is red' only in the presence of a red thing in his field of vision.

Williams 2013, p140

Since this explanation of the meaning of the expression 'red' cites its subject-matter (red things), it violates (GSM). Simpson also cites Huw Price as agreeing that the best explanation of this sort of expression is to be given in terms of "e-representation" – a concept Price uses to refer to tracking relations that are specifiable in non-semantic terms (Simpson 2020b, p152). If (GSM) were correct, then Williams and Price would therefore both be local, not global expressivists, since they allow that explanations of meaning that count as representationalist on (GSM) are sometimes the best such explanations.

The context of the passages in which these apparent endorsements occur is somewhat difficult to parse. Nevertheless, I will try to shed some light, and show that they are in fact in line with (GSM)\*, making use of its more nuanced take on semantics vs metasemantics, and the more limited globalist claim.

Williams' use of an RDRD comes in the context of giving an explanation of meaning in terms of use – an "EMU" – for the expression 'Red'. EMUs, as Williams notes, can vary significantly in character – they can be more deflationary or less deflationary (Williams uses the term "minimalist") depending on the resources they use in offering an explanation, in

particular whether they are ontologically conservative. Williams admits that the EMU for 'Red' is not ontologically conservative (it makes mention of "red things"), but denies that this implies any "representationalist backsliding" (Williams 2013, p141). Williams' idea of representationalism is therefore clearly not the one outlawed by the original (GSM), of citing subject-matter in an explanation of meaning (*any* explanation of meaning), if an EMU can do so without being representationalist.

What, then, could Williams' idea of representationalism be? What prevents the RDRD from being representationalist, by Williams' lights, is that it does not attempt a reduction of any semantic concepts, such as reference; it outlines an external condition on one's mastering an expression, but doesn't look to reduce that expression's semantic properties to causal terms (*ibid.*). Representationalism, for Williams, thus necessarily involves an attempt at semantic reduction. In other words, representationalism doesn't mean a willingness to countenance *any* explanation of meaning that uses subject-matter, but one that attempts to draw a metasemantic relation between an expression and its subject-matter, in terms of which its semantic properties are to be explained.

Williams acknowledges that the EMU for 'Red' is not strictly deflationary ("minimalist"), but is "minimalist in spirit" (*ibid.*). It is unclear exactly what is meant by this. However, it is possible to give a deflationary reading of the RDRD criterion, and thus bring the EMU method of analysis into line with our developing global expressivism. Christine Tiefensee notes that EMUs for ethical terms such as 'Good' can incorporate RDRDs, yet remain palatable to the ethical expressivist. On Tiefensee's reading, the requirement that one have a reliable faculty for tracking morally good things comes down to being a reliable user of the term 'Good' – reliably using 'Good' only when there are good things around (Tiefensee 2016, p2451). This is something the deflationist can get on board with, since it effectively mirrors the explanations of meaning we already saw above. That is, the worst we could say about it is that it is trivial; if one is to know how to interpret the RDRD as stated, one must already know how to apply the term 'Good', since the term is used in the explanation of the RDRD. Again, though, this is unproblematic for the deflationist, as we haven't tried to cash out the operative semantic relation (of satisfying a predicate) in non-semantic terms, instead simply putting the relation to use. The fact that the RDRD is ontologically committing (i.e.: it relies on the existence of morally good things) is likewise not a concern, since our modern deflationist is happy to attest the truth-aptness of assertions that use the deflated vocabulary.

The globalist can apply this reasoning, which works for the expressivist about ethical language, to whatever area of language. So, despite some misleading phrasing, Williams' RDRD for 'Red', above, can be interpreted as requiring only that one use the predicate 'Red' correctly. Granted this renders that part of the EMU trivial, but for the globalist this is as far as we can get in such an explanation before coming under the threat of optimality conditions with infinite exceptions clauses as detailed in the previous section.<sup>8</sup>

One point making itself felt here is that positive expressivist explanations focus little on language entry transitions, which RDRDs for referring terms are supposed to capture, and more on language exit transitions. This becomes obvious when we reflect on Simpson's two-part metasemantic/pragmatic sketch of expressivist explanations; the expressivist can give only a relatively sparse account of the conditions for the correct use of an expression, due to their pessimism about the possibility of cashing these use-conditions in substantive terms, instead focussing their efforts on an account of what we can do with the term once we have established its usage. This supports the idea that pragmatism, for all that it may be a key draw for globalists, is a corollary of the expressivist's core view, which has primarily to do with an embargo on certain sorts of meaning explanations. Speaking broadly, the pragmatist's emphasis on discerning meaning by looking at what we do with language may be necessitated by the restrictions placed on explanations of language-entry transitions, but the restrictions themselves are what primarily determine expressivist thought.

We can be slightly briefer in addressing worries about compatibility with Price's e-representation. E-representation is supposed to denote the relations that obtain between linguistic expressions and the natural objects or properties in the environment that they represent. E-representation is discussed using the language of natural science, typically being given in terms of causal covariation between the tokening of expressions and features of the speaker's environment. Hence e-representation is "naturalistically respectable" from start to finish.

Simpson is correct that this is the basic tool in the representationalist's toolbox. And, further, he is correct that Price endorses e-representation, even for linguistic explanatory

---

<sup>8</sup> Throughout the thesis to follow, I will describe such disquotational formulations as "trivial". This is not meant in a technical sense – in particular, it is not meant to signify that the formulation is *logically* trivial, strictly speaking. What is intended is just what I say above: The formulation is obvious to anyone who already has already mastered the expression being explained, hence it does not contribute directly to explaining the expression's meaning, e.g.: by unpacking it. Rather, in the expressivist explanations in which such formulations occur, their function is to make the explanation adequate to the expression's actual correct usage, and nothing more.

purposes – “Price doesn’t introduce the category of e-representation to argue that it’s empty!” (Simpson 2020b, p145). But, crucially, Price holds that the purpose of e-representation is not to provide reductive metasemantic explanations, i.e.: cashing out semantic concepts (Price 2015, pp146-7). For Price, semantic concepts exist and are to be dealt with purely in terms of i-representation, e-representation’s internal counterpart – in other words, in terms of their functions within particular language games or discourses. Hence if we want to explain what it means to hold a certain statement to be true, for instance, the best we can do is to look at the sorts of inferences (or more broadly, language game “moves”) that it licenses. And if we want a general account of truth as a property, the solution is to look at what features, if any, are common between applications of the concept in different language games, again in terms internal to the functioning of the language.

By contrast, the linguistic-explanatory function of e-representation lies in explaining linguistic behaviour, e.g.: Which environmental properties is it the purpose of certain vocabulary items to track? Note that different areas of language will permit this sort of explanation to different extents, given that, again, e-representation is limited to the objects, properties and relations of a naturalistic sort (so it will have limited utility for explaining the language of ethics, for example). Note also that, in describing the intended function of items of vocabulary in this way, we are not saying anything direct about their semantic features, and so not falling back into representationalism:

In particular, it is open to us to take the view that at least by the time we get to language, there is no useful external notion of a semantic kind—in other words, no useful, general, notion of relations that words and sentences bear to the external world that we might usefully identify with truth and reference.

*Ibid.*

This apparent countenancing of representational relations, and naturalistically delineable ones in particular, in the face of the globalist’s anti-representationalist stance is a large part of what makes the position seem obscure or contradictory. What we can see, upon distinguishing metasemantic explanation as a distinct category, is that this position is possible because the globalist is committed to placing restrictions on that category specifically, and not, for instance, on undermining the concept of representation *per se*. The globalist’s semantic deflationism is a thoroughly “modern” deflationism, in that it aims

at preserving the semantic viability of its target discourse, even while deflating questions about the metaphysical nature of its objects.

#### 5: Why semantic minimalism?

The version of globalism we have set up under (GSM)\* is a global semantic minimalism. That is to say it revolves around the idea that semantic properties do not admit of any substantive explanation. This rules out representationalism, on the understanding presented above, because representationalism is based on the assumption that the nature of semantic properties can be cashed out into other terms (see the following section). As Simpson notes, and as I will illustrate in later chapters, this leaves the globalist with a specific range of tools for metasemantic explanation, specifically based around functional accounts of the roles of different linguistic concepts, and the practical structures that instantiate them.

Global semantic minimalism, though, is not the only option on the table for the globalist. In articulating his global expressivist programme, Stephen Barker (2007) voices a strong opposition to semantic minimalism. Instead, Barker's programme pursues a speculative account of the cognitive structures that underpin linguistic acts that eschews talk of semantics. Indeed, Barker is at pains to distinguish his global expressivism from a theory of meaning *per se*, instead positing it as a theory of language agency (*ibid.* §11). In thus restricting his theory from offering an account of the underlying nature of semantic properties, Barker's intentions are broadly in line with (GSM)\*. However, his critique of semantic minimalism would leave little room for the sort of expressivism to be put to use in aiding the Carnapian in later chapters – indeed, Barker's theory is incompatible with metaphysical deflationism. I will speak briefly against Barker's critique, and attempt to clear some room for the more orthodox semantic minimalist expressivism. In sketching a defence for the globalist as semantic minimalist, I will briefly demonstrate the sort of functional linguistic explanation favoured by that approach.

Barker's critique is brief, but raises many questions for the semantic minimalist expressivist. The critique centres around the objection that the semantic minimalist is unable to provide grounds for distinguishing between various distinct linguistic and conceptual categories. The foremost problem is that there appears at first glance to be no way, given the resources the minimalist has to hand, to distinguish between truth-apt

declarative sentences, the vehicles for assertions, and non-truth-apt sentences, such as imperatives, the vehicles for orders. Consider that the representationalist could, in theory, point to a substantive truth relation that assertions are able to hold, and that imperatives are not, presumably in virtue of their structure; the minimalist can point to no such distinction.

The core problem here is that, while the minimalist *can* outline different use-features of imperatives and declaratives that correlate with their being truth-apt and not, these features themselves stand in need of an explanation. For example, we could say that the difference hinges on the ability to embed sentences within logical compounds, such as placing them as the antecedents of conditionals; declaratives are thusly embeddable, and therefore are truth-apt, while imperatives aren't, so we have our difference. But this seems to assume an account of *why* imperatives should not be embeddable – assuming that language-world relational accounts are off the table, what is the difference in virtue of which declaratives and imperatives differ in embeddability? That the difference is an arbitrary feature of language does not seem like a satisfactory answer (we do not seem to find non-truth apt assertions or truth-apt imperatives in other languages, for instance), but it is not clear how else the minimalist can respond besides invoking truth-aptness itself, and falling into circularity (*ibid.* p30-32).

Our adherent of (GSM)\* can attempt to invoke a functional explanation at the level of linguistic practice, cashing the distinction in those terms. This is the second part of the metasemantic/pragmatic explanatory strategy discussed previously. For example, we might characterise the practices surrounding the issuing of imperatives in terms of language game “moves”. On such a characterisation, a speaker A issuing an imperative “Leave!” to a hearer B places on B certain responsibilities: Either B has to acquiesce and leave as ordered, *or* B has to provide reasons why they are allowed to ignore the injunction and stay put. These are B's only options – simply ignoring the order amounts to refusing to play the game at all, and is not itself a valid move, so to speak. If B is to play the game, then, they are issued with certain obligations.

Where the imperative language game places obligations on the hearer, the declarative language game instead grants the hearer permissions. Specifically, and cribbing from Robert Brandom's inferentialist machinery (Brandom 1994, pp168-170), when a speaker A asserts a sentence S, A thereby grants a hearer B permission to use S as a premise in reasoning. A also undertakes an obligation themselves, namely to respond to objections to



S by offering a justification (we might also think that, just by playing the game, A adopts an obligation to abandon their commitment to S when presented with sufficient countermanding evidence). The distinction between declaratives and imperatives is thus cast in functional terms: When uttering declaratives, we issue permissions, and adopt obligations ourselves, whereas in uttering imperatives, we issue obligations that are recognised by others who take themselves to be playing the same language game.

We can therefore explain the difference in embeddability, accordingly, in functional terms. Declaratives exist in the first place because they allow us to track the commitments of other language users; they embed because this allows us to demonstrate the other commitments speakers have, and ought to have – what Brandom calls “deontic scorekeeping” (*ibid.* p182) – by inferences placing their original commitments in logical compounds. By contrast, a speaker does not undertake a commitment in the same way by uttering an imperative. Rather, they simply attempt to use the imperative language game to produce a desired outcome. Hence it is not *useful* to embed imperatives in the same way as declaratives, because they do not play the same role in the inferential game of deontic scorekeeping. This difference in purpose explains why, although we can imagine a language being arranged such that imperatives are embeddable, our language is not arranged so.

All is not clear just yet, though, for the language-focused expressivist. As Barker notes, imperatives do appear to be able to stand in some inferential relations. For instance, uttering the sequence “Leave! You are physically incapable of leaving.” seems to involve some kind of tension or incompatibility. A speaker uttering such a sequence might well be censured in much the same way as if they had uttered an illegitimate inferential sequence (Barker 2007, p40). This seems to imply that imperatives do carry inferential commitments in some sense, even if they do not do so in the same way as declaratives. The challenge here is, again, to give some account of this difference in minimalist terms. If it turns out that imperatives can bear inferential commitments, then we cannot use commitment-bearing, and its associated permissions and obligations, to distinguish them from declaratives. So what is left that can do the job?

An immediate temptation is to explain that imperatives may require certain inferential commitments in order to be legitimately issued, but declaratives *express* commitments, where imperatives do not. This allows for there to exist a correlation between both sorts of sentences and anterior commitments, but explains why they differ in character. However, while this makes some sense given a colloquial understanding of “express”, we will struggle

once again to articulate this difference in minimalist terms. For instance, Barker notes that the minimalist about truth does not seem to be able to distinguish between uttering a sentence, S, that expresses a mental state, M, and uttering the sentence “I have M,” which reports on M, rather than expressing it (*ibid.* p31). Since both sentences are correlates of M, and both look to grant a hearer permission to use the fact of the original speaker’s having M in reasoning, there seems to be no gap left between the two for the minimalist to lever open, and thus no room for an account of expressing as a distinct relation between discursive commitments and speech acts in minimalist terms.

I think that this temptation to cash the difference in terms of expressing is to be resisted. I share Barker’s pessimism about pursuing this line, particularly in light of (GSM)\*; it seems as though a substantive account of expressing as a relation between a sentence and a commitment, if one could be procured, would end up looking quite similar to the sort of subject-matter-involving relations that are outlawed by the take on expressivism being developed here. The solution I want to put forward, then, involves abandoning the notion of expressing for theoretical purposes. While it is helpful to point to a folk notion of expressing a judgement or desire, for instance, to help home in on the function of a linguistic item within a language game, we don’t need to construct a formula to link these mental states with those linguistic items. What we should instead do is analyse the language game being used on its own terms – what are the rules being followed, and, at a higher level of explanation, why these rules instead of others? Again, this leans on the pragmatic side of the expressivist toolbox referenced above to render the sorts of answers we want, within the constraints set out; in this I mainly follow the Pricean conception of global expressivism, which sees expressivist explanations neither as consisting in pointing out “word-world” relations, nor as explaining pre-doxastic mental machinery, but as pointing to intra-linguistic structures and relations (see Ch.6 for detailed discussion).

To demonstrate this method, we can return again to our toy example and add a little depth. One important point to take into account is just how participants in the imperative language game are supposed to contest the obligations placed on them by other participants. This must perforce involve some sort of reasoning. Typically, though, this does not involve the target imperative itself, but a transformation of the imperative into a sentence using deontic language – “B ought to leave,” for example. Deontic language thus gives an assertoric “face” to imperatives, so that they can be used in reasoning; deontic assertions, as assertions, can be used as premises for arguments (and antecedents for conditionals) by parties to the discussion. This is how we track whether A’s commitments

accord with what is required by the proprieties of the imperative language game in order for A to issue a legitimate command. We can thus clarify the language game of issuing imperatives by noting that the only options available to the hearer, B, who is being issued with an imperative, are either to carry out the issued order or to postpone doing so pending evaluation of the order's grounds *via the assertoric language game*.

This last part bears emphasising. We now have it that language users are able to discuss the commitments borne by imperatives as encoded within assertions using deontic language, but not *qua* moves in the imperative language game. This means that, while we can call on the assertoric language game to settle disputes about what constitutes a proper move in the imperative language game, the only proper response *within* the imperative language game to being issued with an imperative is to carry it out.

This may seem an overly pedantic way of individuating language games, especially since the goal here is finding the grounds for drawing these distinctions. Once again, though, the point is to distinguish between imperatives and assertions in terms of function. By more clearly delineating the language games in the example, we can further clarify the distinct functions that assertions and imperatives perform, in practical terms. As outlined, the function of imperatives is both very straightforward and very limited: If A issues a legitimate imperative to B, then B, insofar as they are committed to engaging with the imperative game, has an obligation to enact it. The job of this language game is thus much more specific than we might previously have thought, amounting, in effect, simply to eliciting actions from others, and not to engaging at all with their obligations or permissions.

It is this drastically limited scope that keeps the obligations imparted by the imperative game from being overbearing. Consider that one only has an obligation to obey commands to the extent that one has an obligation to enter into the imperative game in the first place, e.g.: by relations of authority. For example, subordinates being issued commands by a superior officer are under an obligation to enter into the imperative game, and thus carry out the commanded actions, by virtue of the authority relation that holds between them; by contrast, one has no obligation to obey commands from strangers on the street, except insofar as one opts into the imperative game, in the same way that a stranger moving a chess piece and looking at one expectantly doesn't oblige one to take the next move, unless one opts into the chess game.

The job of discussing the obligations that underpin moves in the imperative language game is then outsourced to the (much more developed) assertoric language game, the function of which is to track our discursive commitments generally. We thus distinguish between assertions and imperatives based on the different functions they perform. Imperatives do indeed bear commitments, but their function is not to keep track of such, hence their not being embeddable. Rather, they are used to translate commitments into actions.<sup>9</sup>

The above is only a rough sketch of how the globalist as semantic minimalist might respond to the critique lodged by Barker, but is intended to show that there is at least a line that the minimalist might take to defend their theory. Barker himself sees the critique as reason to reject semantic minimalism, and instead to delve into “the *noumenon* of cognitive structure” to deliver an expressivism that focuses on explaining linguistic acts in terms of the pre-doxastic mental states that produce them (*ibid.* p87). However, although the expressivism he presents is thoroughly worked out, it does not provide the sort of metasemantic account we need to support a neo-Carnapian deflationism. As mentioned above, Barker is at pains to point out that his expressivism, by design, does not offer any sort of theory of meaning, instead offering an account of the means by which human beings produce and interpret linguistic signals. Indeed, Barker notes that his global expressivist programme has metametaphysical consequences that run counter to the deflationist’s agenda. Barker’s focus here is on “real definitions”, which amount to metaphysically explanatory identity claims, such as “The property, F-ness, is the set of F-things.” The deflationist wants to maintain these sorts of claims, but as platitudes, in an attempt to make them “as low-key as possible.” (Barker 2020, p279) By contrast, Barker’s expressivism denies that such real definitions can even be thought, on the grounds that the metasemantic facts needed even to think such claims do not obtain (*ibid.* p280). This clearly conflicts with Carnapian Tolerance, which requires not only that the various metaphysical commitments of different linguistic frameworks *can* be thought, but that they come out true, strictly speaking.

---

<sup>9</sup> Dividing the relevant functions in this way means, conversely, that we have to take it that deontic assertions do not, *qua* assertions, impose obligations on the hearer, since if they did, we would again lose the functional difference between assertions and imperatives. I do not think that this is implausible. We might take it that A, in using a deontic assertion to place an obligation on B, is using it as a stand-in for an imperative. For instance, A might utter the assertion “You should leave” as a way of eliciting action from B, instead of using the imperative “Leave!”. On the current picture, this is something like a “soft” move; A is leaving B to fill in the gaps, analogous to a chess player who announces “Mate in four”, and allows their opponent to capitulate rather than play out a foregone conclusion.

What the neo-Carnapian requires, as we will see more clearly in the following chapters, is a semantic methodology, in the sense of an approach to examining semantic properties themselves. In order to support the metametaphysical consequences of their commitment to Tolerance, the neo-Carnapian needs to be able to explain, for example, how apparently conflicting metaphysical statements in different linguistic frameworks can simultaneously be true; the requirement placed on their supporting metasemantic theory is to outline a sense for “true” that makes this possible, and likewise for other semantic properties. Where Barker takes the limits of expressivist explanation to be drawn at linguistic acts such as assertion and expression, the neo-Carnapian needs a theory that will outline an understanding of semantic properties themselves that is compatible with their deflationism. Barker is explicit in stating that not only does his theory not aim at such a project, it actually precludes it. I believe, as I will attempt to bring out in the following chapters, that a semantic-minimalist expressivism provides the only semantic methodology that will give the neo-Carnapian what they need, by providing an understanding of semantic properties that is consistent with their deflationism.

#### 5: Representationalism reviewed

The argument in section 3 is not quite the standard way of pointing out placement problems – that largely takes the form of gesturing to particular properties of individual areas of discourse, such as ethical properties, that appear to resist placement within the natural order. What I hope to have done above is to provide a demonstration of the character of the placement problem that afflicts semantic properties specifically. This placement problem is not new to expressivists (see Price 2011, p191), but in articulating it in detail I have attempted to shed some light on the representationalism to which global expressivism is opposed. This is especially important now that I have proposed to understand global expressivism, as (GSM)\*, as primarily a rejection of representationalist explanation; if, with Simpson, we are still confused about what exactly representationalism is supposed to amount to, then we will likewise still be in the dark about what the expressivist intends by rejecting it.

According to the version of globalism outlined above, representationalism is *not* simply the thesis that language has a representational character, or that some linguistic expressions are capable of representing. Such a thesis could hardly be denied on a global front. I take representationalism, rather, to be the general philosophical commitment to the possibility

of explaining meaning by providing some sort of metasemantic account of the relation between an expression and its subject matter. Paradigm examples are causal theories of reference, in which the role of metasemantic relation is played by a causal chain of an appropriate type, of which the referring expression and the referent are part (Kripke 1980, Millikan 1984), and correspondence theories of truth, in which it is played by the relation of structural similarity between a true sentence and the arrangement of things in the world (Wittgenstein 1922). What these theories have in common is that they assign to semantic properties a metaphysical nature, which the theories seek to unpack. This is what makes such theories metasemantic accounts of the relevant semantic properties, in my terminology; they attempt to go beyond the semantic properties, to translate them into non-semantic terms (typically naturalistic ones).

This means that representationalism is necessarily a reductionist theory, in a broad sense – it looks to reduce the representational semantic features of linguistic expressions to non-semantic terms. This is, in a sense, where the problem of optimality conditions presented above gets its hooks in: If one wishes to present a causal theory of reference, for instance, one needs to identify which sorts of causal chains are appropriate for supporting (correct or successful) reference, and which are not. As noted, this sort of specification looks to be extremely difficult.

It is worth commenting on the link between representationalism and reductionism. For one thing, there are several types of reductionist theories out there, in the sense in which we are interested. It would not be especially interesting for the globalist to deny, for instance, a naturalistic reductionism that states that for each semantic property/relation, there is one basic physical relation that constitutes it. Such a theory appears implausible on the face of it; even the Siderian idea that semantic concepts carve the world at its logical joints allows that semantic properties are realisable by many possible physical relations (Sider 2011, §3.2).

It is possible to take the above argument from Boghossian as responding to this implausibly strong reductionist claim, but this is not its intended target. Instead, the argument, as I have used it here, is intended to target a weaker form of reductionism, and to do so in a specific way. The target is a representationalism built on the idea that each occurrence of a semantic property is instantiated by *some* natural property – i.e.: that semantic properties are multiply-realizable by natural properties. This thesis looks extremely plausible, and indeed the globalist does not actually want to deny it outright. What the argument above is

intended to show, rather, is that no informative explanation of the underlying nature of semantic properties can be retrieved from this theory, because the set of natural properties that are capable of instantiating the semantic property in question could never be specified, even in general terms, provided we limit ourselves to non-semantic vocabulary.

What the argument is doing here is not targeting the theory itself – the idea that there is some relation of supervenience or grounding, for example, that holds between semantic properties and natural ones – but the idea that this theory, or this relation, can support an informative account of the meaning of semantic terms, and therefore that it can furnish the blueprint for a metasemantic account of other expressions by demonstrating the nature of their relation to the physical. Note also that the account works regardless of our candidate reduction-base; I have addressed the possibility of a naturalistic reduction here because it is by far the foremost contender among representationalist accounts, but the analogous arguments could be produced for e.g.: a mentalistic account of semantic properties.

#### 6: Deflationism vs. primitivism

There are other arguments that pick up on what are, essentially, symptoms of the placement problem, but I will leave further discussion of these for Ch.6. For now, I want to begin addressing a different concern about the optimality conditions argument: Granted that it does point to a placement problem for semantic properties, does it in fact motivate a global expressivism, or could it motivate a distinct position in the form of a semantic primitivism?

The problem is that global expressivism is not the only anti-reductionist take on semantic properties we could offer. One option is to take it that semantic properties are metaphysically basic – that their metaphysical nature cannot be unpacked or analysed by reduction, but that they nevertheless do *have* such a nature, in contrast to the expressivist's deflationary claim that they do not. The argument presented above looks to speak at least as much in favour of this position as that of the expressivist. Given that both are strongly anti-reductionist, an argument against reduction doesn't immediately seem to favour one over the other.

In fact, I think that there are two ways of understanding primitivism, as regards the current discussion. On one reading, the primitivist view is equivalent with that of the expressivist, and there is no conflict; what motivates one will naturally motivate the other, because they are batting for the same team. On the other reading, what distinguishes the primitivist's picture of semantic properties from a deflationary one, such as the expressivist's, is relatively mysterious, and requires further elaboration if it is to be made sense of.

The solution here hinges on how we understand what the primitivist means by the "metaphysical nature" of semantic properties. On one plausible understanding of the expression, the primitivist is committed to the simple fact that, while such properties are un-analysable, we can cite them *as* properties in truth-apt sentences, quantificational contexts etc. On this view, a primitivism about the property of reference might take it that the sentence "The name 'Aristotle' refers to Plato's brightest student" expresses a truth, and that it is likewise true that, as a consequence, the name 'Aristotle' bears at least one real relation, namely reference, to the man, Aristotle. This reading has it that the semantic properties' having a metaphysical nature means, simply, their being metaphysically real, but in a way that resists explanation.

The problem, here, for the primitivist who wants to outline a distinct position, is that the above characterisation is consistent with the expressivism outlined thus far. Recall that a "modern" deflationist, meaning one who rejects non-cognitivism about the deflated discourse, is happy to accept the reality of the entities and properties of that discourse, while rejecting certain analyses of their nature. For this deflationist, being "metaphysically real" means simply being the subject of true existence claims. Hence neither the claim that semantic properties are metaphysically real nor that their nature is un-analysable in metaphysical terms is sufficient to distinguish the expressivist's view of semantic properties from the primitivist's (*cf.* Sosa 2001, §5 for a similar argument<sup>10</sup>).

This is, of course, a deflationary understanding of metaphysical reality (one on which, for instance, the phrase "metaphysically real" is redundant compared to simply "real"). It is open to the primitivist to offer a substantive account of the idea of a metaphysical nature, such as they impute to semantic properties; such an account would distinguish their view

---

<sup>10</sup> Note that Sosa appears to take the difficulty in distinguishing the positions as a problem for the *deflationist* ("minimalist") about truth. I take it that his reasoning is, similarly to Simpson's, that the alleged problem is to do with how the deflationist can conceive of truth as both being a real property *and* lacking a substantive metaphysical nature – the concept of a metaphysical nature (even an un-analysable one) is deemed unproblematic, hence primitivism is the default position. I hope to have spoken against this conception above.



from deflationism, since the deflationist would not assign this substantive nature to the properties in question. However, the problem is that the primitivist is constrained in offering such a substantive account by their criterion of un-analysability. Any account of what it means for semantic properties to have a metaphysical nature, in terms of their intrinsic features, that goes beyond their simply being real properties is necessarily going to go some way to unpacking that nature, which runs counter to primitivism.

The discussion presented here is not conclusive. I hope, though, to have presented the structure of the expressivist defence – or, at least, the defence I think the expressivist has to take, if we are to define expressivism as I have above. In summary: One can opt for a semantic primitivism as a form of anti-reductionism, in response to the argument presented in the sections above. However, either such a primitivism comes down to the same thing as the expressivist’s semantic deflationism, or else, in order to be distinguished, it relies on a conception of metaphysical nature or metaphysical reality that is mysterious, and needs further elaboration.

#### 7: What is *local* expressivism?

The localist is now the one left with a problem, namely how to differentiate the discourses that need the expressivist’s deflationary treatment outlined above from those that do not. As Price has noted (e.g.: Price 2009b, pp119-120), once we have a deflationary plan in place to get truth, reference and so on for our expressive discourses, it looks difficult to resist the pull to treat all of language in this way.

To put it another way, the concerns that motivate treating language as expressive, in particular the placement problem given above, don’t look to be restricted to a particular discourse. Rather, they are features of metasemantic explanation *as such*, and so will apply to any area of language for which we try to give a metasemantic account. The non-cognitivist can fall back on ontological worries to differentiate discourses, supporting localism; it is only the entities referenced by certain discourses that are non-existent or otherwise unreal, so it is only those discourses whose sentences fail to be truth-apt, because they fail to refer. But the modern localist can’t do this, since for them it is true, albeit in a “thin” sense, that such entities really exist. Neither can the localist fall back on this thin/robust distinction to achieve the required distinction between discourses, since clarifying that distinction was one of the key reasons to try to cash out the expressivist’s

core commitments – to cite it as definitional of localism would be for the localist to fall into exactly the obscurity Simpson attributes to the globalist.

The root of the problem is that the localist appears to be in the grip of a category error of a certain sort. We previously understood modern localism as guaranteeing cognitivism for expressive discourses with a deflationary semantics for those discourses – their sentences were true and false, but in some thin sense to be explained non-representationally.

However, we now see that semantic deflationism is not an extra ingredient, a specific sort of semantics that guarantees a specific “thin” sort of truth, reference etc. Instead it is a restriction on the resources available to metasemantic explanation, and further, it is a restriction that applies to metasemantic explanation as such, in virtue of its intrinsic limitations, hence not one that can apply to some discourses and not to others.

The deflationism that guaranteed cognitivism has turned into a trap for the localist, so that they face the following dilemma: On the one hand, they could demonstrate what features non-expressive discourses have that permit them representationalist (i.e.: subject-matter-involving) metasemantic explanations, thus reserving semantic deflationism only for expressive discourses. If the restrictions cited above really are intrinsic to metasemantic explanation as such, this doesn't look to be possible. Alternatively, the localist could outline a sort of deflationary account distinct from the globalist one given here, but which does the job of differentiation of discourses. This is the more promising route. However, the localist's deflationism will have to be radically different from that of the globalist, to the point where it is unclear how this account could proceed. This approach thus falls prey to the same criticisms of obscurity that Simpson levels at a deflationary globalism: It is not clear in what sense they are different versions of the same position, just applied with different generality. Localism, then, not globalism, is the view in need of clarification; globalism itself is clear, plausible and stable.

Conclusion:

I have outlined what I think a viable global expressivism looks like, against Simpson's attempts to show that no such thing is possible, by taking one of the definitions he rejects and making it sensitive to the semantic/metasemantic distinction, on which the globalist relies. This allows us to understand semantic deflationism, and its key role in globalism. The latter hedges out localism, for the same reasons that Simpson uses to take the globalist to

task. With semantic deflationism, the expressivist can have their cake (rejecting representationalist metasemantics) and eat it (continuing to use the deflated vocabulary), but they cannot pick and choose where in language to do this.

More detail is required to flesh out the features of the sort of globalism we are looking for, but the sketch presented above should be enough to be going on with. Globalism construed as semantic deflationism is unquestionably the sort needed to support the neo-Carnapian project. The task of the following chapters will be to demonstrate why exactly the neo-Carnapian needs to adopt a global expressivism, so construed – i.e.: what problems orthodox representationalist assumptions cause for their deflationary approach to metaphysics.

### Ch.3: Hirsch's intensional deflationism

Eli Hirsch holds a prominent position among neo-Carnapian deflationists. Hirsch's position, quantifier variantism, has the potential to be a powerful tool for critiquing and dissolving apparent metaphysical disputes by showing them to be merely verbal, although the position itself is controversial. The pushback against quantifier variantism comes not only from metaphysicians, on which more below, but also from other deflationists, such as Amie Thomasson, who sees Hirsch's idea of multiple senses of the word "exists" as a liability for deflationism, inviting accusations of anti-realism.

I will deal directly with quantifier variance, and whether it is necessary for Hirsch's theory, in Ch.5. For now, I want to argue that Hirsch's metaphysical deflationism appears to rely on representationalist presuppositions. These presuppositions actually undermine the metaphysical deflationism itself, by forcing us to accept that a variety of metaphysical debates are substantive. Since deflationism is incompatible with representationalism, it cannot have a representationalist metasemantics underpinning it. Therefore, because the deflationism is worth preserving, we need to come up with an anti-representationalist, or expressivist, metasemantics for the theory, something I will proceed to do. The argument put forward here is structurally similar to the one to be put forward in the following chapter regarding Thomasson's "easy ontology". However, where Thomasson remains non-committal as regards representationalism, Hirsch seems to endorse a representationalist metasemantics for his theory, as I will attempt to make clear. Hence it may be that the way to salvage the theory is to abandon Hirsch's own version of it.

I will proceed by recounting Hirsch's deflationary approach in section 1. Then, in section 2, I will note a problem area where the approach seems to admit of substantive metaphysical answers to disputes, as raised by John Hawthorne. In section 3 I will examine the concept of the "intensional landscape", essentially homing in on the representationalist roots of the problem. Section 4 will then see an attempt to put the theory on an expressivist footing with a more comprehensive metasemantic picture, thus saving it from making room for substantive metaphysics.<sup>11</sup>

---

<sup>11</sup> A brief note: In this chapter I will use "deflationism" as a shorthand for "metaphysical deflationism". Since the argument forms part of an attempt to show how *metaphysical* deflationism

## 1: Deflation by equivalence

Hirsch is famous for his theory of quantifier variance. This states that we are largely free to pick what our language quantifies over, that is, to pick which quantifiers our language uses (Hirsch 2011, Ch.5). Hirsch's strategy for deflating metaphysical debates uses this idea to posit that the two sides in such a debate are actually speaking past one another. This is clearest when the debate is a straightforward ontological one, e.g.: between the organicist (Van Inwagen 1990) who says that living things are the only composite objects there are, and the speaker of ordinary English, who holds that many non-living composite objects exist. Hirsch explains this debate away by modelling the organicist as using a non-standard quantifier (relative to ordinary usage) that is only satisfied by living things and metaphysical simples. Hence a contested sentence such as "There exist tables", that it seems one side should accept and the other should deny, actually means something different in each language. In ordinary English, the sentence means that there exist tables (in the ordinary English sense of "exist"), but this is not the same meaning the organicist takes themselves to be denying. The organicist-language sentence the organicist takes themselves to be denying, for instance, translates into ordinary English as "Some tables are living things or metaphysical simples", something the ordinary English speaker would concur in denying – that is to say there exists a conciliatory translation scheme that allows us to understand the organicist and the ordinary speaker as speaking the same truths in different languages, hence there is no real disagreement (Hirsch 2016, p106).

Hirsch's main strategy here revolves around a concept of charity. Hirsch holds that metaphysical debates take the following form: There is some controversial sentence C that one side takes to be intensionally equivalent to a claim  $N_1$  and the other side takes to be equivalent to a different claim,  $N_2$ , with neither  $N_1$  nor  $N_2$  being controversial for either disputant (*ibid.*). Since both sides agree on these non-controversial claims, the dispute looks to be over their equivalence with C. The dispute thus turns not on whether the truth-conditions for C are met, but what, in fact, C's truth conditions *are*. One can see this by adopting an attitude of charity; applying the charitable principle that one ought to interpret an interlocutor as speaking the truth wherever possible, it is relatively easy to construe the

---

and *metasemantic* deflationism dovetail, there is much potential for confusion. Hence I will take care to flag up when an issue is specifically metasemantic in character, to try to avoid ambiguity.

other side of the debate as speaking a distinct language in which the equivalence between C and their N-claim holds. The alternative is attributing to them a failure of linguistic expertise by seeing them as staunchly maintaining a more-or-less trivial falsehood in one's own home language.

Once we realise that the two sides are linking the disputed sentence to different sets of truth conditions in this way, it becomes obvious that the way to treat the debate is as a linguistic dispute, not a factual one per se. That is, in terms of charity, because of the systematic character of each side's application of relevant concepts and sentences, it behoves us as charitable hearers to approach them as speaking their own bespoke language in which they are systematically correct, rather than viewing them as systematically speaking falsehoods in another language. Hence the disputants don't disagree on any facts, they just express the relevant facts differently, according to their language.

Of course, the question still remains whether the equivalence held by either disputant holds in ordinary English, and further, whether either disputant actually *needs* to argue that theirs is the view that captures the truth-conditions used in ordinary English (*cf.* Hirsch 2011, Ch.10, §1). Regardless of whether either of these questions is answered in the affirmative, though, charity still upholds the deflationary conclusion; even if English does not in fact license organicism, as seems to be the case, that isn't to say that we couldn't construe a coherent organicist language, and thereby achieve a deflation of the debate. Further, it is open to the organicist to claim that we *should* be speaking their language *instead* of ordinary English – they just can't do so on metaphysical grounds, since their language doesn't make any coarse-grained claims that an interlocutor would deny (I will address the notion of "coarse grained" or "unstructured" facts in more detail in the following section).

There are several principles at work here. One is quantifier variance, the idea that these differences between languages are grounded in what it is each language means by the quantifiers, and which entities it holds to fall under them. Quantifier variance itself has come under fire from both metaphysicians and neo-Carnapian deflationists (e.g.: Sider 2011, Ch.9, §6, and Thomasson 2015, Ch.1, §3). I will discuss these criticisms in a later chapter, but I want to focus on another way of dissecting Hirsch's view here, this being an analysis of the truth-conditional semantics underpinning the above deflationary move.

Hirsch's semantics is truth-based rather than reference-based, in that it proceeds by the roughly "top-down" method of looking at the uses of sentences deemed correct by a language (Hirsch 2011, pp238-40). Individual expressions can then be analysed in terms of the sentences in which they can be used in order to home in on their individual use-properties. This is what prompts Hirsch to cite quantifier variance as the difference between alternative metaphysically committing languages, as the quantifier expressions have different inferential characteristics in each (*ibid.* p188). Hirsch puts the doctrine in truth-conditional terms as follows:

...the basic idea of quantifier variance can be nicely formulated by saying that the same (unstructured) facts can be expressed using different concepts of "the existence of a thing," that statements involving different kinds of quantifiers can be equally true by virtue of the same (unstructured) facts in the world.

*Ibid.* pp78-9

The method Hirsch follows effectively comes with a principle of charity built in, as it requires analysing the sentences of a language in the language's own terms; the possibility of their being mistaken in some sense never really comes into it. This top-down method stands in contrast to the "bottom-up" method that would analyse individual expressions in terms of reference or sensitivity to properties, for instance, and then compose sentence meanings out of these. Such a view of language as deriving meaning from representing structures in the world would make deflationism impossible, as it builds in metaphysics at the ground level. However, I will argue that Hirsch's "intensional" deflationism doesn't entirely escape the representationalism that problematises the bottom-up approach, and that we need to go even further to find a complete, thoroughgoing deflationism.

## 2: Intensional structure

There is a notion of "coarse-grained" or "unstructured" facts already at work in the background of Hirsch's deflationary programme as elaborated thus far. Hirsch notes that the equivalences in the above schema that each party draws between the disputed sentence C and the sentence in their own language are truth-conditional equivalences, or intensional equivalences (Hirsch 2016, p114). It follows that applying the schema is basically a way of getting clear on how we want to *individuate* metaphysical assertions. That is, counting locutions that are phonetically identical as expressing the same assertion

will have the two sides genuinely disagreeing on the truth of the controversial sentence C, or in the example, the truth of “There are tables”. But individuating assertions by their intensional features (namely their truth conditions) shows us that the disputants are not really disagreeing, because they are not making the same assertion – “There are tables” expresses one intension for the organicist and another for the common-sense ontologist.

It will be helpful to have a definition of an intension. I take an intension to be, simply, a set of states of affairs. A single intension thus constitutes the truth conditions of any given unambiguous sentence; the sentence is true if and only if one of the states of affairs in its intension actually obtains. The sentence itself can be said to express the intension in question. Ambiguity at the level of sentences is thus ambiguity between different intensions the sentence could be expressing – between different sets of states of affairs that would make the sentence true. The usefulness of the concept of a proposition resides largely in its being a shorthand for talk of intensions, with “proposition” itself being a functional synonym for “intension”. Part of our job later will be to provide an expressivist theory that retains the usefulness of talk of intensions – I regard this as supporting also the usefulness of the concept of a proposition, since the two are inter-defined.

Hirsch calls the states of affairs that make up intensions “unstructured facts” (Hirsch 2011, p78). The intensions expressed by sentences have to be understood in terms of unstructured facts, for Hirsch, because different languages will potentially differ in the hyperintensional structure they attribute to them. To illustrate: The organicist can assent to the common-sense ontologist’s sentence “There are tables” in the sense that they (the organicist) acknowledge that the truth-conditions attributed to that sentence by their interlocutor are fulfilled. This holds true even in the organicist’s language – the organicist would simply express those truth conditions differently, i.e.: without quantifying over tables. So we have an instance of intensional equivalence between two statements in different languages that nevertheless differ with regards to their quantificational structure, quantificational structure being an example of hyperintensional structure. In brief, we won’t get the deflationary results we want if we individuate assertions by quantificational, predicative, referential, or other hyperintensional structures, because we won’t thereby be able to generate N-claims in each language that speakers of the other language can assent to.

The role of intensional equivalence and unstructured facts in Hirsch’s theory can be explored further, and potentially problematised. In order for us to reason via Hirsch’s



method that the dispute between two competing metaphysical theories is not actually to be settled by whichever one better reflects reality, since in fact they both do so equally well, we first need to establish a theoretical language for each theory in which the claims they make can be mapped onto sets states of affairs, and second, to establish that these two theoretical languages are intensionally equivalent in all respects, via this intensional mapping. In other words, for any sentence that might fill the role of  $N_1$  in language  $L_1$ , we need to be able to find a corresponding sentence  $N_2$  in  $L_2$  that has the same intensional profile.

This has a number of important corollaries. First, it establishes that if either theory is intensionally adequate – if all the sentences that come out true according to its theoretical language express true propositions – then both theories are. This is important for obvious reasons, namely that we don't want to end up wrongly deflating genuine debates about matters of fact. But it also establishes that each language has equal expressive power by establishing that any unstructured fact expressible in one language is also expressible in another, a necessary requirement for wholesale intensional translation.

John Hawthorne makes the point that “certain ontologies multiply possibilities in ways that are resisted by other ontologies.” (Hawthorne 2009, p221) What this means is that there can be pairs of theoretical languages for which we cannot construct all-encompassing intensional translations on Hirsch's model, because their ontological commitments affect *which possible states of affairs* they regard as real, or more precisely, as really distinct from one another. Hence Hirsch's strategy of bypassing differences in hyperintensional structure between languages by focusing on intensional structure will not always work, because intensional and hyperintensional structure can be interlaced.

Hawthorne gives an example of how this is supposed to work:

Let us begin by contrasting someone who embraces a unique fusion axiom in mereology with someone who thinks that many distinct things can fuse the same set of objects. Imagine, then, that one philosopher has an ontology of particles and lumps of matter, while another believes that the same set of particles can compose a lump and a distinct statue at the same time. [...] [The second philosopher] can then, for example, imagine one world where Ben is the last statue made, Bill the penultimate, and infinitely many before them, and another where Bill is the last statue made but where there is duplication of the first world when it comes to particles and lumps. Thus one of the philosophers has a wider range of haecceitistic possibilities in his vision of things and this makes for a wider range of envisaged possible worlds.

The point Hawthorne is making with this example is that there seem to be genuine intensional stakes to the debate depicted here; either the statue lover is making what Hawthorne calls an “intensional advance” over the lump theorist, asserting as part of their language mastery a true intension that the lump theorist would wrongly deny, or they are “intensionally regressing”, asserting a false intension that the lump theorist would be correct in denying. In either case, it looks like a substantive resolution to the debate is, in some sense, available, even if as a matter of fact our means of enquiry are not up to the task of discovering it (*cf.* Ecklund 2006, §V). Since one of either the lump theorist or the statue lover is wrong in their view on whether the world with Bill and the world with Ben represent genuinely distinct possible states of affairs, the debate is substantive in a way that not only defeats Hirschian efforts to dissolve it, but that actually tracks the dispute over metaphysical commitments.

One thing we might be suspicious of is the way in which Hawthorne heads off translation between the two languages in the example. The fact that the Bill-world and the Ben-world are identical with respect to all the properties the lump theorist is capable of expressing is a core premise, but it also doesn’t track most real-world applications of the Hirschian schema. For instance, it is taken as a given that most mereological nihilists can understand the conditions under which speakers of ordinary English apply the referring expression ‘table’, since they are at least able to feign speaking ordinary English themselves, and further that they could construct in a nihilist language sentences that are intensionally equivalent to sentences in ordinary English that (falsely, according to the nihilists) appear to refer to tables. This results in the well-known formulations that replace occurrences of ‘table’ with ‘particles arranged tablewise’. Provided the application of expressions such as ‘tablewise’ is sufficiently regimented and systematic, it appears as though the mereological commitments of nearly all languages can be shown to be intensionally equivalent in the way Hirsch wants by using such formulations as a fallback.

We would have to go quite far to find two fully functioning languages whose expressive resources were so different that we couldn’t translate any given sentence in one into an intensionally-equivalent sentence in the other – indeed, Hawthorne’s own example doesn’t seem all that plausible, because the statue lover’s concept of statuehood would have to have it attributed almost completely arbitrarily, so that it doesn’t depend on any of the features that the lump theorist has the expressive resources to talk about. That is to say,

for any property F that the statue lover could attribute to the last statue that makes it Bill instead of Ben (or lack of which makes it Ben instead of Bill), it looks as though the lump theorist could come up with a translation that distinguishes between the two possibilities based on whether the lump in question possesses F. If no “F-wise” rephrasing is possible, then it appears that there is no feature in virtue of which Bill is Bill and not Ben, and the distinction is arbitrary.

Nevertheless, truly incompatible languages are in fact conceivable. Cian Dorr proposes an “astronomically impoverished” language in which one is only able to refer to and quantify over entities within a certain distance of the Earth’s surface (Dorr 2005, pp237-8).

Astronomically Impoverished English would genuinely lack expressive power compared to ordinary English, because there are sentences in ordinary English that cannot be expressed in AIE, and further, for which an intensionally equivalent translation can’t be found via “tablewise” rephrasing, because the state of the universe outside our solar system (to pick a boundary for the domain of AIE) is underdetermined by the state of affairs inside the solar system – there are possible distinct states of affairs in the universe at large that do not entail differences in the state of the solar system, i.e.: the domain that AIE is limited to talking about. Therefore we have a case of a disagreement between languages that fits Hawthorne’s structure, where one language lacks the expressive resources to frame intensionally equivalent translations of sentences in the other language.

It is tempting to read off this example a victory for ordinary English, on the grounds that it has greater expressive power than AIE. This is the lesson Matti Eklund takes from the case (Eklund 2013, p241-2), namely that ordinary English has the ability to express facts that AIE can’t. As part of a critique of the Hirschian project, Eklund extends this conclusion to encompass the Platonism/nominalism debate, stating that Platonism claims a straightforward victory on factual grounds (as opposed to being the more pragmatic choice, for instance) because “only a Platonist language is adequate for expressing all the world there is, [therefore] the Platonist objectively wins.” (*ibid.*) In Hawthorne’s terms, the Platonist and the ordinary English speaker make an intensional advance over the nominalist and the AIE speaker respectively.

I think there are fairly straightforward reasons that Hirsch need not be concerned by Eklund’s version of the criticism. This isn’t to say that the Hawthorne-style criticism is of no concern at all – quite the contrary, as I will discuss below – it is just that this is not the version we need worry about. For one thing, Eklund elides the alternate possibility that the

Platonist and the ordinary English speaker are intensionally regressing. In the Platonism/nominalism example, this amounts to taking the Platonist's side, no questions asked; if we hold that the Platonist can express more of the world, we already have to accept the Platonist's view of what there is of the world to express, in the sense of having adopted their language and its ontology. From within the nominalist language, the claim that the Platonist has the resources to quantify over more entities than the nominalist makes no sense, since we would have to state that the Platonist could quantify over things that don't fall under the (nominalist's version of the) unrestricted quantifier, a nonsense. We can state, in the nominalist's language, that the Platonist *claims* to be able to do this, but this simply means stating that the Platonist makes a nonsense claim about their quantificational abilities. Claiming an objective victory for either side requires taking an ontological stance *prior* to declaring victory, according to which the victory can count as "objective", and this simply means that we aren't taking Hirsch's deflationist method sufficiently seriously.

The AIE case is somewhat harder to dissolve, if only because of the strong pull towards ordinary English as against AIE itself. What the deflationist needs is *not* to show that we can't make the decision between the two languages, but rather to show that that decision is not made on the basis of which language is metaphysically correct. The deflationist *can* thus opt for ordinary English on the basis of its being more expressively powerful, where this is cashed differently from Eklund's ontologically committing version. Again, stating that one language captures more of what there is requires an antecedent notion of "what there is", and one that stands independent of the ontological commitments of our language (a requirement that looks close to a straightforward contradiction). The deflationist should instead unpack the idea of having usefully greater<sup>12</sup> expressive power in terms of being able to express what we want to, what helps our pragmatic aims, and what allows us to frame the best theories according to non-metaphysically committing theoretical virtues such as parsimony, simplicity, explanatory power and the rest. The deflationist diagnosis is that greater expressive power in this sense pulls us strongly towards ordinary English, from

---

<sup>12</sup> I qualify with "usefully" here because we can conceive of languages with greater expressive power, in the sense of being able to frame sentences that can't be paraphrased in another language, where this does not immediately make these more powerful languages preferable. For instance, a language that was ontologically committed to ghosts would fail to exemplify the virtues mentioned here as against the ghost-impoorished ordinary English. The fact that we could make statements about ghosts in "ghostly" English that resisted intensional translation into ordinary English wouldn't make it a preferable language to ordinary English, although this is not straightforwardly because it bears faulty ontological commitments – rather it is because it fails to exemplify the non-metaphysical virtues of a good language (or at least exemplify them as well as English does).

which we then make the judgement that AIE is expressively impoverished in Eklund's ontological sense.

### 3: Addressing the problem – The intensional landscape

We can articulate a version of the problem that Hawthorne states that looks more plausible, and potentially therefore more threatening to a Hirschian deflationism. This version focusses not on the possibility of claiming a straightforward victory for either side in the metaphysical dispute, but rather on establishing that, although it may not easily be resolvable, it is substantive – that the metasemantics that Hirsch uses to attempt to deflate it actually entails that there is a factual disagreement between disputants in at least some metaphysical debates. This is arguably what Hawthorne had in mind (*cf.* Hawthorne 2009 p224), although he doesn't pursue it to its fullest realisation.

Hawthorne has correctly identified that different languages recognise different possibilities. This threatens to make translation between languages impossible in some cases, since in so translating, we want to identify sentences in one language that are true and false in all the same circumstances as corresponding sentences in the other language. The problem extends further than simply being unable to deflate debates by finding translation schemes, though. The real problem is that Hirsch's intensional semantics, which is integral to his deflationism, requires us to be realists about the "intensional landscape" (my phrase) - about the structure of possibility space, meaning which possible states of affairs are and aren't numerically distinct, and therefore what intensions there actually *are* for sentences to express. If this realism is correct, then it looks as though we can read off which theoretical language is objectively correct simply by looking at whether it acknowledges (only) real possibilities, and expresses (only) real intensions. To use Hawthorne's example, if Bill's being the most recent statue and Ben's being the most recent statue really are distinct possibilities even with the world being otherwise qualitatively identical, then the statue lover's language and its accompanying ontology is the correct one. If they are not really distinct possibilities, and the distinction is merely imagined, then the lump theorist's language is correct, and the statue lover is wrong.

In brief, the problem is that a language's structure at the hyperintensional level, which is the level at which we do ontology, interacts with its structure at the intensional level. If we agree that a language can have a correct or incorrect intensional structure, judged

according to a single real intensional landscape, then we can draw conclusions about which hyperintensional structure, and therefore which ontology, is correct (or at least which is closer to reality, i.e.: which allows us to express more real intensions and fewer unreal ones). Attempts at an intension-based deflationism would therefore be self-defeating.

The problem, then, centres on the idea of the real intensional landscape – a single real set of possible states of affairs that different languages attempt to key into, and succeed to different degrees. Hirsch’s talk of “unstructured facts” seems not to account for the fact that, intuitively, talk of different possibilities assumes some set of features in terms of which these possibilities are to be distinguished – that if two possible states of affairs really are numerically distinct, they must be qualitatively different in some way. The issue is that to know whether those states of affairs really are distinct, something we need to know to execute a Hirschean deflationism, we first need to know whether those features that distinguish between them are themselves real, and so we fall back to metaphysics.

What might be our options for getting rid of a set intentional landscape? Possibly the most straightforward response would be to be anti-realists about intensional talk – to reject it as somehow empty or defective, in keeping with our anti-metaphysical approach. If there are no intensions, then there is no intensional landscape from which to infer substantive ontological facts. But this response is too strong; the concept of intensionality is, arguably, responsible for underpinning our concept of meaning. More than this, it has a significant use value, especially for the deflationist, which we should want to preserve if at all possible. If we took statements about intensions to be somehow systematically flawed, so that they had no bearing on reality, then we would be incapable of saying, for instance, when two sentences have the same meaning, because we couldn’t say that they were true and false in all the same circumstances. In making a truth-conditional semantics impossible in this way, such a move would actually make a Hirschean deflationism impossible, defeating the point of making the move in the first place; while such a move might allow us to dissolve metaphysical debates, it would entail a metametaphysical anti-realism, so is not a move for the deflationist.

An alternative, although perhaps a slightly far-fetched one, would be to find some other footing on which the different theoretical languages *are* equivalent, since intensionality seems to have failed us. The verificationist movement, for instance, took empirical equivalence as its cornerstone for a deflationism on a similar format. This empirical equivalence of metaphysical theories is what brought Carnap, at one time a verificationist,

to suspect they were really “saying the same thing”, and talking past each other when they seemed to be at odds. Notably, there is no explicit mention of a reliance on empirical equivalence in Carnap’s paper *Empiricism, Semantics and Ontology*, the *locus classicus* for the deflationism that inspires modern day neo-Carnapians. Carnap had reportedly moved away from hard-line verificationism at the point of writing the paper, giving up on the idea of reducing linguistic meaning to a basic ground of evidential states, and had already formulated his Principle of Tolerance (Richardson 2007, p296; see Carnap 1937, pp51-2), in many ways the ancestor of Hirsch’s appeals to charity. Nevertheless, Hawthorne seems to feel that Hirsch, a neo-Carnapian, is somewhere relying on verificationist principles to make his own deflationism appear attractive (Hawthorne 2009, pp213-15), hence the failure of the theory when the reliance on empirical equivalence is removed in favour of more palatable intensional machinery.

I do not think that Hirsch lapses into verificationism, nor that he harbours secret verificationist tendencies. As mentioned, it is the later, tolerant Carnap from whom Hirsch takes his cues. In fact, as I’ll argue below, the underlying problem for Hirsch is representationalism. Hirsch relies on a vision of sentences as picking out areas on a particular intensional landscape, meaning a field of possibilities with a structure that disparate languages attempt to key into. The problem with basing any deflationism on a representational picture like this is that we can always attempt to settle the debate to be deflated in a substantive way by asking how things are with what is being represented, here the intensional landscape – i.e.: by asking whether possibilities, derived from the contested ontological facts, are really distinct.

There is a picture emerging of what we will need from our deflationism if we are to overcome Hawthorne-style problems. This engenders three desiderata for a reworked Hirschean deflationism: First, and most straightforwardly, it will need to cover the bases already covered by the Hirschean theory, namely deflating metaphysical disputes by drawing attention to the different languages spoken by participants, in virtue of which their apparently conflicting claims really speak past each other; second, it needs to retain the utility of intensionality and adjacent concepts (propositions, truth conditions) – we want to end up with a picture that allows us to talk about sameness and difference of meaning, rather than adopting some form of anti-realism about these things; and third, we want to dissolve Hawthorne-style problems, showing, specifically, that there is no substantiality in the intensional landscape to be mined for metaphysical debates. All this will revolve around stripping out the representationalism on which Hirsch’s view, on reflection, hinges.

#### 4: An expressivist footing

We thus turn again to expressivism, this time to provide an anti-representationalist account of the concept of intensionality. Recall that expressivism holds that we shouldn't seek to cash out semantic concepts in metaphysical or representational terms. For the concept of reference, for example, this means we shouldn't look for an explanation that hinges on establishing a substantive relation between the proper usage of referring terms to the things being referred to. Such an explanation will never be informative; either it will fail to outline the correct relation, or it will be correct, but trivial, of the form "The term 't' refers to t" – something that fails to render new information to anyone who has already mastered the term 't'. What creates the placement problems we saw in Ch.2 is, by the expressivist's lights, the attempt to render an account of the relation between the expression and its subject-matter as a substantive relation.

Analogously, the form of explanation the expressivist wants to avoid for the concept of truth has it relating sentences to states of affairs. This is the explanation built into Hirsch's theory as it stands at the moment; sentences pick out sets of states of affairs, the landscape of possible states of affairs, and therefore possible intensions, being set independently of language. As we have seen in Hawthorne-style cases, this representationalist approach to truth is incompatible with a thoroughgoing metaphysical deflationism, since we can use it as a springboard to derive substantive metaphysical resolutions to the debates we want to deflate. Thus, in the following, I'll look to present an alternative, expressivist version.

For this to work, we need to get rid of the idea that there is a single structured "intensional landscape" that each language's sentences map onto. But we needn't go as far as the anti-realist about intensionality to do so. Instead, we should proceed analogously to Hirsch on ontological structure; just as each language harbours a quantifier with a different structure, so too does it harbour a differently structured intensional landscape all its own.

Proceeding in this way is actually supported by Hirsch's own charitable reasoning. To treat the statue lover, for instance, as mistaken in their assertion that differences in terms of statues entail different states of affairs would be to attribute to them a systematic falsehood (assuming their concept "Statue" is systematically applicable); better to treat them as speaking a different language in some respect, e.g.: using quantifiers differently



regarding states of affairs, or attaching a different meaning to the sortal “state of affairs” (which might amount to much the same thing – see discussion in Ch.5).

This relativisation means that there is no metaphysical debate possible, at this stage, between the speakers of each language. This is because, each language having its own intensional landscape, it is not possible for the disputants to point to claims that they would affirm, and that their interlocutors would deny; at the present stage, it is impossible to assert a relation of intensional equivalence between different sentences in the different languages. Such an assertion must be made against the background of a shared intensional landscape to be meaningful, and there is as yet no agreed upon intensional landscape that the two languages share. This is not yet a complete deflationist picture; to preserve the utility of intensional concepts, as we want to do, it will be necessary to detail how the two languages can be made directly comparable via the introduction of a metalanguage with its own intensional landscape. I’ll discuss this below. Note, for now, that it is an important premise that the absence or impossibility of such a metalanguage does not entail that the metaphysical disagreement really is substantive – rather, it makes this sort of disagreement impossible.

This characterises the initial stage of the debate, before the adoption of a metalanguage for use comparing the two object languages. This next stage adds a complication, but also enables the new deflationism to reflect the use value of intensional concepts for addressing sameness and difference of meaning – our other desideratum, besides metaphysical deflationism, for a successful theory. It does this without reopening us to intensional disagreements in a substantive way, although it will have to allow for those disagreements in some guise.

The introduction of a metalanguage for intensional comparison between object language sentences involves formulating an intensional landscape onto which the intensional commitments of the respective object languages can be mapped. This is somewhat analogous to working out a translation scheme between two disparate dialects in a more widely spoken version of the same natural language, to facilitate communication between provincial speakers. This means the project ought to be responsive primarily to the behaviours and the structures of the inference-practices of the different speakers, although the meanings of dialect sentences do not necessarily reduce to these factors. The important difference from the original Hirschian version is treating this as an effort of construction, namely of constructing a way of mapping one language onto another, rather

than an effort of derivation, i.e.: an attempt to key into what possibilities those languages represent by their sentences.

What is happening here is a reframing of the concept of intensional equivalence as a metalinguistic tool, rather than a concept denoting a representational feature, itself undergirded by metaphysical properties. Such a treatment of intensional equivalence is analogous to the expressivist treatment of other metasemantic concepts, such as truth and reference. That is, it proceeds by analysing them in terms of the features of their use, specifically as they are used for metalinguistic analysis. Hence the expressivist account should of necessity leave in the use value of the concept of intensional equivalence, because it is in terms of its usefulness that it is to be defined.

Taken this way, it is possible to see the concept of intensional equivalence as a tool for expressing the intersubstitutability of sentences. Substitutability between languages is a special case, so we can start by looking at substitutability of sentences within the same language. It is fairly easy to see that we could define intensional equivalence here as two sentences' being correctly assertible in all and only the same circumstances. The role of the language's intensional landscape is thrown into relief here as the structure that underpins and allows us to understand the phrase "all the same circumstances" – this requires a picture of which circumstances there *are*, in the sense of which states of affairs are numerically identical and non-identical. The expressivist, then, sees a language's intensional landscape and intensional equivalence fundamentally as tools for expressing patterns of substitutability (*cf.* Köhler 2018 for a similar picture)<sup>13</sup>.

The idea of "correct assertability" of a sentence used above is obviously a way to appeal to the truth-conditions of a sentence. But the expressivist's reliance on truth conditions to explain intensional equivalence shouldn't be construed as a lapse back into metaphysics; the expressivist offers the same type of account of truth as we are offering here for intensionality, namely analysing it as a concept with a metasemantic function, not as a

---

<sup>13</sup> Köhler's paper offers an expressivist metasemantics based around substitutivity to support a truth-conditional semantics, much as I want to do here. However, Köhler takes the expressivist's metasemantic deflationism too far, inferring that the expressivist shouldn't draw *any* metaphysical conclusions from metasemantics on pain of falling back into representationalism, and that they should therefore be fictionalists about semantic properties and relations, such as truth and reference. To my mind, this is a response to a lingering motivation of the "old style" expressivism of Ayer. The point is that Köhler misses, in my view, the take-home message of deflationism in any domain: That we shouldn't look for substantive analyses of the relevant properties (here truth, reference and so on), *but* that we should nevertheless go on being realists about them (*cf.* Price 2013, pp177-8).

metaphysical property. The function most often pointed to by expressivists when explaining truth is one of compression and substitutability (E.g.: Williams 2013, p135). If I say “Everything Tom says about fly fishing is true”, I use the concept of truth to compress an assertion of all the things Tom might tell you about fly fishing. Essentially, I hand you an “inference ticket”, allowing you to add my credibility to Tom’s regarding any of the class of statements mentioned, without my having to enumerate them, or even share all the same information as Tom. This is a good example of the sort of metasemantically deflationary account the expressivist favours for semantic concepts see Ch.7 for a deeper analysis).

In the present discussion, the notion of truth-conditions, or more specifically sameness of truth-conditions, plays a similar substitutional role, as does the concept of intensional equivalence, which amounts to the same thing. As mentioned, within a single language, the comparison necessary to ascertain whether two sentences are intensionally equivalent can be carried out against the backdrop of that language’s intensional landscape. The same is true for intensional equivalence *between* languages; an assertion of intensional equivalence here is just an assertion that, according to the intensional landscape that we have constructed for comparing the two languages, the two sentences are correctly assertible in all and only the same circumstances – that they are, in that sense, intersubstitutable. Where Hirsch implicitly justified this intersubstitutability with the help of a metaphysical relation, namely a representationalist conception of truth, here we have analysed it solely in metalinguistic terms. This non-representationalist conception of intensional equivalence, then, can go on being useful in the same way as always, namely as a tool for assessing which sentences are inter-substitutable – provided, that is, we have a relatively stable metalanguage with an agreed upon intensional landscape.

How do we arrive at such a metalanguage? The metalanguage works as a device for translation, so the structure of its intensional landscape must mirror those of both object languages as far as possible, allowing us to map between them. The formation and potential refinement of the metalanguage itself is, like any other language, sensitive to pragmatic pressures.

Hirsch cites as a guiding principle the idea that his schema portrays conflicting languages as each talking about the same world, in terms of unstructured facts, saying:

...I am inclined to agree with Putnam that, once we’ve accepted quantifier variance, there is no point in trying to hold onto language-shaped facts that are in the world independent of language. However, we can retain the notion of an unstructured fact. I think this is indeed

our most basic notion of “reality,” “the world,” “the way it is,” and this notion can remain invariant through any changes in our concept of “the things that exist.”

Hirsch 2011, p79

Within Hirsch’s theory, as we have seen, this idea is intrinsically representationalist. The expressivist will thus have to refine it into the normative pragmatic principle that, as far as possible, sentences from disparate languages that the metalanguage treats as intensionally equivalent should entail the same behaviours, be receptive to the same stimuli, and so on, as one another. Where this is not the case, this will be a pragmatic reason to restructure the metalanguage, and thus its intensional landscape, the better to realise this goal.

For instance, consider the metalanguage used to map  $L_1$ , a language that countenances tables, onto  $L_2$ , the mereological nihilist’s language. The intensional landscape of the metalanguage will outline a set of possible circumstances that it treats as identical to both the circumstances in which the  $L_1$  sentence “There is a table at  $x$ ” is true, and the circumstances in which the  $L_2$  sentence “There is a set of simples arranged tablewise at  $x$ ” is true. If we observe that the two statements appear to license different or exclusive behaviours or inferences – for instance, the  $L_1$  speaker will eat off the thing at  $x$ , but the  $L_2$  speaker won’t eat off the collection of simples at  $x$  – then this gives us a reason to reconsider our intensional mapping. We can imagine this process happening in any language; we have a practical interest in having intensionally equivalent sentences in the *same* language license the same behaviours, and not exclusive behaviours, hence we have a reason to restructure if we come across these anomalies. The only difference here is that our metalanguage has the purpose of bridging a gap, as between two metaphysical “dialects”.

The point here is that while we can ask how we arrive at a metalanguage for translation, we cannot ask how we arrive at the *correct* metalanguage, with the correct intensional landscape. Just as each object language trivially affirms its own intensional landscape, so too does each metalanguage. The question whether one metalanguage or another more accurately depicts the structure of possibility space is thus an external question in Carnap’s sense, that is, a question intended to be posed outside of a framework of concepts within which it can be made sense of. In this way, it is analogous to the earlier comparison between the two object languages in the absence of a metalanguage to frame the comparison. Put slightly differently, the idea that some metalanguage might be correct, and that this might be a criterion for validating it over other metalanguages that attempt

the same job of translation, would be a regression to the representationalist idea of a language's intensional landscape attempting to mirror the structure of possibility space as it stands independent of the language itself. As this idea would entail the failure of metaphysical deflationism, we have already jettisoned it.

We can now see how the expressivist picture hits our bases for a successful theory. First, it provides a metaphysical deflationism: There are no substantive answers to metaphysical questions of the sort that Hirsch was principally interested in deflating. This is because these debates can only be framed by coordinating the languages of the participants with the aid of a metalanguage, such that, in the core cases at least, metaphysical debates are dissolved by intensional equivalence, exactly as in Hirsch's version.

Second, our theory supports the usefulness of intensionality and adjacent concepts, such as "Proposition" and "Truth-conditions". We have not fallen into the anti-realist trap of denying that any sentences can ever be intensionally equivalent, but nor have we based intensional equivalence upon representational concepts, such as the idea of mirroring the pre-given structure of possibility space. Rather, we have simply noted that these concepts are part of the metalanguage, which is perfectly in order as is. The correctness of claims using these intensional concepts depends on the intensional structure built into the metalanguage, but this is no different in principle from the way in which any concept depends for its use on the conceptual structure of its language.

And finally, our theory dissolves Hawthorne-style problems of intensional non-equivalence between languages. This is because, again, to compare the intensions of two languages as a step towards resolving a metaphysical dispute, an intermediary metalanguage must be constructed with its own intensional landscape. If the two first-order languages are as described in the Hawthorne case, then no metalanguage will be a satisfactory intermediary for both parties – any intensional landscape we construct for the purposes of comparison will side with either one first-order language or the other. But this only shows that the two first-order languages are incompatible, not that there is a substantive answer as to which (if either) is correct; the two first-order languages are intensionally incommensurable, so to speak, but the question of which metalanguage we should use for translating between them (and thus which first-order language gets validated) is an external question, and thus has no substantive answer. We are at liberty to pick one or other of the translation schemes on offer on a pragmatic basis, but there is no further intensional framework by

which to validate it, so we must be clear that this is precisely a pragmatic decision, and not a matter of metaphysical correctness, nor of intensional-structural accuracy.

## Conclusion

I hope to have shown how underpinning Hirsch's theory with an expressivist metasemantics heads off Hawthorne-style objections that look to recover some substantivity for metaphysical debates. In a sense, what I have been doing is taking Hirsch's top-down method of analysis and applying it through and through, stripping out assumptions that treat language as representational; the theory is bottom-up in at least one sense if it postulates a structured field of states of affairs – the single intensional landscape – that sentences use for truth conditions. What Hawthorne-style cases do is essentially to bridge between intensional and hyperintensional structure, allowing a representationalism about states of affairs to bleed into realism about ontological structure. In contrast, the kernel of expressivism was present already in Hirsch's top-down method; it just needed pursuing further.

It's worth noting, very briefly, that the metalanguage-introduction process described above has the potential to iterate. That is, just as we looked for a solution to the Hawthorne-style metaphysical debate by trying for a metalanguage to serve as a translation scheme, and got two potential, exclusive candidate metalanguages, hence not solving the debate, in the same way we could look for a further metalanguage to compare the intensional landscapes of these two translation schemes. While pursuing this to a higher level might produce new complications, I don't believe that it would change the nature of the standoff significantly; we would still find two ways of mapping the intensional landscape of metalanguage  $M_1$  onto that of metalanguage  $M_2$ , one of which confirms  $M_1$ , and one of which confirms  $M_2$ , just like at the lower level. Adopting either of these new metalanguages would allow us to frame the debate as substantive, just as before – e.g.: adopting the one that favours  $M_1$  would allow us to say the sentence, true in that metalanguage, " $M_1$  makes an intensional advance over  $M_2$ ." But, just as at the lower level, we can frame no substantive reason for adopting either of these new metalanguages over the other – and so on up the chain, *ad infinitum*.

We can see, here, that the principle that no intensional comparison between languages can be framed without the benefit of an agreed-upon intensional landscape is of a piece with

Carnap's principle that ontological claims can only properly be made "internally" to a linguistic framework. The two are in fact the same principle; a language's intensional landscape just is its ontology of states of affairs, and one makes claims about such when one posits an intensional translation scheme. What we have ultimately been doing, then, is attempting to bring Hirsch back into alignment with Carnap, for the benefit of the deflationary project generally. This has involved pulling Hirsch in the direction of expressivism, itself a metasemantic deflationism, thus uncovering further how these two spheres of deflationism go together.

#### Ch.4: Thomasson's easy ontology

Amie Thomasson's influential deflationary theory, outlined in her 2015 book *Ontology Made Easy*, aims to give the ability to reason ontologically back to ordinary language users. Among the core posits of easy ontology is the idea that simply having mastered the use of the term for a particular type of entity (coupled, occasionally, with some empirical investigation) is sufficient to be able to answer ontological questions concerning that type of entity. Thomasson's argument for this position represents a deflationary approach to ontology: Since ontological questions are resolved simply by competent language use, the only debates that generate ontological implications will be either linguistic or empirical in nature, meaning that there is no distinctly ontological discussion to be had. In this, Thomasson is firmly a part of the growing movement of neo-Carnapians applying a deflationary approach to metaphysics more generally, taking their cues from the Carnap of *Empiricism, Semantics and Ontology*.

Thomasson's position leans heavily on the notion of "application conditions" for the terms of ordinary language. However, her discussion of this notion leaves the broader semantic picture with which she is working largely open and unexplained. The issue is forced by the potential for application conditions to generate a regress, as noted by Andrew Brenner (2019), which necessitates a fleshing out of the easy-ontological picture if it is to be avoided. I will show that there are several ways to go in supplying a semantics for application conditions that meet the requirements of the theory, but that each of these incorporates ideas that are potential deal-breakers for many moderate metaphysicians. With easy ontology as it stands, we know where the empiricism and the ontology fit in, but we still do not have a clear picture of the semantics.

In section 1, I will explain application conditions and the role they are supposed to fulfil in the easy ontological programme. Section 2 will introduce the ambiguity that I want to look at, which hinges on a distinction between "assertability conditions" and "satisfaction conditions", and will try to show that it is not in itself a problem for Thomasson's account. In section 3 I introduce Brenner's regress, and in sections 4 and 5 I look at alternative ways of solving it, along representationalist and global-expressivist lines respectively, and will assess their compatibility with the aims of easy ontology. I conclude that the only way to



avoid a Brenner-style regress is to side with the expressivist. Fleshing out an expressivist take on Thomasson's programme will be the work of later chapters.

#### 1: The role of application conditions

Thomasson introduces the idea of application conditions for terms in her 2007 book *Ordinary Objects* as a solution to the "qua problem" of causal theories of reference. The problem applies to theories that look to ground reference entirely in causal chains established between speakers (or their mental representations or the like) and the objects about which they are speaking. This sort of theory lacks a criterion for when reference should fail when we are trying to coin a new referring expression. Suppose, for example, that I sight what appears to be an island while at sea, and, pointing to it, name it 'Crocker island'. If this supposed island turns out actually to have been a mirage, it is unclear what 'Crocker island' actually refers to; I have definitely undertaken the baptismal ceremony necessary to fix a reference for the name, but nothing seems to match my idea of its referent satisfactorily. Many things have had a hand in causing my speaking; does 'Crocker island' name the sea spray, the optical phenomenon, or perhaps the tip of my finger? A similar problem seems to apply even if there does exist a landmass to which I can point. Again, several things could be said to be the cause of the baptism, e.g.: the present time-slice of the island, the layer of material composing it that faces me when I point, and so on. This is the *qua* problem, namely the problem that reference, on pure causal accounts, is highly indeterminate (Thomasson 2007, pp37-40).

Thomasson originally introduces application conditions to solve this problematic indeterminacy. On her theory, new terms such as 'Crocker island' actually come laden with some conceptual content, in the form of at least one sortal term, e.g.: 'Landmass' (*ibid.*). If, at the point of creating a new referring expression, there is nothing that answers to the sortal term at one end of the relevant causal chain that has myself or my utterance at the other end, then the attempt at reference fails, and I fail to coin a new referring expression. So Crocker island has to be a land mass, and cannot be a trick of the light or the foam on a wave. This likewise solves the problem of indeterminacy for apparently successful attempts at reference – Crocker island is the *landmass* that is involved in causing me to attempt to refer, not the time-slice or island-part. My attempts at reference, then, must come with an intention to refer to a certain sort of thing, although this can be more or less specific, depending on the sortal. What we have then is a "hybrid" theory of reference; reference is

still determined causally, in part, but it also comes with at least some conceptual content necessarily attached.

Thomasson introduces application conditions as a way of understanding how sortals can perform this function. A sortal term has conditions under which it is correct to apply that term. These application conditions determine whether circumstances are appropriate for the purported referring expression being coined to successfully refer. So if the expression 'Crocker island' comes with the sortal 'Landmass' attached, and that sortal has in its application conditions that there be an area of dry ground (perhaps of a certain size), bounded by a region that rises above sea level, then something meeting this description would be required to be within the relevant causal chain in order for my attempt at reference to succeed.

The usefulness of application conditions is not restricted to solving the qua problem for hybrid theories of reference. Thomasson's (2015) book *Ontology Made Easy* revolves around using the concept to reach ontological conclusions relying only on mastery of a language. This follows a simple formula: If Ks are a kind of entity, then the statement "There exists a K" is true iff the application conditions for the term 'K' are fulfilled (Thomasson 2015, p86). Hence disputed ontological questions concerning ordinary objects (for instance, as being composites with parts), as well as institutional terms such as 'Marriage' and 'Company', and abstract objects like numbers, can all be resolved by asking, for the relevant sort of thing, K, "Is it correct to apply the term 'K'?"

We have moved, then, from identifying the thing at the end of a chain of reference to the simpler, more basic use of application conditions as establishing whether there is anything that answers to that term at all. It is important for Thomasson's deflationary approach that this be all that there is in determining whether or not an entity of a given type exists. There is nothing by way of a more fundamental sortal 'object' or 'entity', for instance, that existents must answer to. Nor can application conditions be circular; the application conditions for a term 'K' cannot include that there be Ks (*ibid.* pp96-98). Either of these options would leave room for a specialised metaphysics that could only be done by philosophers, and would not be achievable simply by being acquainted with ordinary language use. In this respect, Thomasson follows the line laid out by the Carnap of *Empiricism, Semantics and Ontology*, but without relying on the concept of a linguistic framework that Carnap introduces there. Thomasson thus sidesteps criticisms to do with quantifier variance, the idea that the existential quantifier can have different meanings

within different linguistic frameworks or domains; for Thomasson, the existential quantifier has a single, invariant meaning, given by the schema:

E: Ks exist iff the application conditions for 'K' are fulfilled.

(Thomasson 2015, p86)

In stipulating (E), Thomasson provides a deflationary account of existence. That is to say, on Thomasson's account, there is no substantive first-order property, such as possessing causal powers, that existence amounts to; once we have established that the application conditions for a sortal, 'K', are satisfied, there is no extra check needed to see if there exists a K. This has important ramifications for our understanding of the quantifier 'Exists'. For one thing, it implies that "bare" quantificational sentences such as "There exist three things" are actually incomplete, since, if we take them to be genuinely bare, they do not specify a sortal whose application conditions would provide truth conditions for the sentence. Thomasson provides an analysis of such sentences on which they make use of "dummy sortals", such as 'Thing' and, potentially 'Object', to stand in for other sortals within the language (*ibid.* p109). The sentence "Some object exists", while apparently a instance of bare quantification, would then actually mean something like "Some sortal is satisfied". If correct, this preserves Thomasson's deflationism – the alternative, that truly bare quantificational sentences in natural language were actually meaningful assertions, regardless of their not relying on sortals, would require a substantial reading of the quantifier. I will further discuss the details of the deflationary quantifier posited by Thomasson in the following chapter.

Another upshot is that the easy ontologist now has access to an extremely wide range of ontological judgements, essentially re-establishing Carnapian tolerance from a different direction. These include judgements derived from so-called "pleonastic" inferences, i.e.: inferences to an explicitly ontological conclusion from premisses that are not explicitly ontologically committing. For instance, we can infer from "The cups and plates are equinumerous" that "The number of cups is the same as the number of plates" is true, and from there that "There is a number (that of the cups and plates)" is true (*ibid.* pp135-6). While Stephen Schiffer, who coined the term for pleonastic inferences, regards such conclusions as reflecting our thought and language more than the world itself, and therefore as being more ontologically "thin" than more hard-won ontological conclusions (Schiffer 2003, p60), Thomasson argues for a simple realism as regards pleonastic entities;

according to the deflated quantifier, there is no difference in principle between the “sort” of existence attributed to a trivially-inferred abstract object, and an empirically-inferred concrete one.

## 2: An ambiguity for application conditions

The above seems to give us enough to assess whether Ks exist on the basis only of having mastered a term ‘K’. However, before moving on to look at other objections, there are some areas of ambiguity to clear up. In order to bring out problems for the theory as detailed so far, it will help get clear on the argument Thomasson puts forward for easy ontology, so that we can examine individual premisses. We can reconstruct this argument as follows:

P1: Terms are associated with application conditions.

P2: Speakers must master the application conditions associated with a term in order to use that term properly, and do master terms this way.

P3: The application conditions for a given term ‘F’ do not include that there be Fs.

P4: For any given referring term ‘F’, an F exists iff the application conditions for ‘F’ are fulfilled.

---

C: There is no further information (such as is sought in metaphysical discourse), over and above that to which ordinary speakers could have access, that determines whether Fs exist. Ordinary speakers’ mastery of language is therefore sufficient for answering existence questions.

P1 and P4 are integral to Thomasson’s system; they introduce application conditions, and hinge upon arguments to do with Thomasson’s “formal” conception of the term ‘exists’ (Thomasson 2015, pp84-86); P4 is just Thomasson’s principle (E). P2 and P3 unpack relevant properties of application conditions; P2 establishes that they are actually usable by speakers, a necessary condition for easy ontology to be possible, and P3 wards off the idea that it is the satisfaction of our ordinary terms *plus something else* – i.e.: something that

might be uniquely accessible to metaphysical discourse, such as satisfying a substantial criterion for existence – that establishes whether something exists.

There exists a tension between P2 and P4. This tension is not unresolvable, but it will serve to direct the discussion to a genuine dilemma for easy ontology and the theory of application conditions. As per Thomasson's theory, application conditions are required to be the conditions under which ordinary language users actually do apply terms in their language (*ibid.* p91), as well as being the conditions under which such terms actually do refer to entities in the world – i.e.: the conditions according to which those entities actually do exist (*ibid.* p93). There are at least some cases in which these sets of conditions can be seen to come apart. One fairly clear case is when using natural kind terms, the precise application conditions of which (in the latter sense of the conditions on the term's successfully referring) may be known to a few experts, but may not be known to a wider language-using public. For instance, I may not be able to enumerate the exact physiological or genetic features that make a platypus a platypus, but I, and most English speakers, could probably pick the platypus out of a line up when asked. It seems that in cases such as this the ordinary language user has access to, if not the actual conditions for applying 'Platypus', then some conditions under which we are *warranted* to apply it. That is, we may not have the detailed, biologically articulated concept, but we seem to have access to some sort of proxy concept that gives us warrant to believe that the stricter concept would also apply.

I want to introduce a distinction here between these two functions of application conditions. I will call the conditions under which a term can be warrantably applied its "assertability conditions", or "ACs" for short. The conditions for a term's actually applying, irrespective of warrant on the part of speakers in attributing it, I will call its "satisfaction conditions" or "SCs". We can link this terminology as used here to its more standard use regarding propositions: The satisfaction conditions, in our sense, of a given term 'F' will be fulfilled iff the proposition "There exists some F" is true, and the assertability conditions of the same term will be fulfilled iff the same proposition is warrantably assertible. It seems uncontroversial that a distinction of roughly this type should exist in at least some areas of language, simply on the basis of how easy it is to come up with examples, using holograms, hallucinations and the like, in which ordinary speakers are *justified* in asserting the existence of some thing, F, where there is in fact no F to be found, or indeed vice versa where we have no such warrant but there is in fact some hidden F.

There are a few things to note once we introduce this distinction into the theory. First, it appears possible that the conditions represented by a term's ACs could also be those that determine its SCs. In other words, a term could truly apply in all and only the cases where it is warrantably applied. For instance, the existence of subjective appearances such as a red quale or a déjà vu might have exactly the same requirements as our having warrant to assert that these things exist.

Another thing worth looking at is how this refines the picture of what is going on above in the case of identifying a platypus. On the most immediately plausible reading, the fact that I can identify a platypus without knowing what strictly, biologically defines it as such is down to my having a functional stand-in, in terms of a set of assertability conditions for the term 'Platypus' as ordinarily used – these assertability conditions here standing in for the stricter, ontologically-committing satisfaction conditions. This is important, because it gives us a “handle” on language that we as ordinary language users can grasp, in whatever sense (ACs), and a “business end” that grounds the truth conditions of what we want to say (SCs). Of more specific importance, though, is the fact that this allows us to do easy ontology, even without knowing satisfaction conditions for our terms. To hold that ACs confer warrant or justification for applying a certain term, 'F', is to hold that they confer this same warrant or justification for believing that the application conditions that make it *objectively* correct to apply 'F', its SCs, are also fulfilled wherever its ACs are fulfilled. Put more simply, just knowing ACs gives us full *justification* for easy answers to ontological questions – which, really, is all we could hope for as language users.

So what does Thomasson mean when she talks about “application conditions”? Are they conditions of assertability or satisfaction? There is definitely some ambiguity in Thomasson's work as regards this question. Although Thomasson does not draw the distinction I have made here, application conditions are required to have these two functions, which are in tension. Having said that, it appears clear to me that Thomasson's argument requires application conditions to be SCs – that is their more fundamental function, and the one that cannot be done away with if we want to draw ontological conclusions easily, as evidenced in P4 in the argument above; in order to draw first-order ontological conclusions by the methods Thomasson employs, we must have a set of conditions embedded in a sortal, K, that are met iff there really are Ks. As ordinary language users it is enough that, in any given situation, we can hold our easy ontological conclusions to be justified. However, from the philosopher's perspective, it is required that we show how these conclusions can hit the mark. If we were to treat 'application

conditions' as referring to ACs rather than SCs, we would not have this assurance, since ACs do not determine the truth or falsity of any ontological facts, merely whether we are justified in asserting them (P4 in the argument would therefore not hold). Easy ontology, then, has to start by explaining how ontological conclusions drawn on the basis of language alone (plus some empirical investigation here and there) can be reliably accurate, and then proceed to explain how, via heuristics or some other mechanism, ordinary language users acquire the knowledge necessary to draw such conclusions.<sup>14</sup>

### 3: Brenner's regress

It looks as though this distinction could be used to solve a regress posed by Andrew Brenner. According to Brenner, a sortal term that has application conditions that are non-trivially fulfilled must make reference in those application conditions to other sortal terms (Brenner 2018, p607). E.g.: the sortal 'Platypus' might include in its application conditions that there be required to exist a mammal with a bill that lays eggs. If mastery of 'Platypus' requires mastery of these conditions, then it seems that it will require functional mastery of the terms 'Mammal', 'Bill' and 'Eggs'. But each of these terms will likewise require mastery of the application conditions of new sortals (such as 'Vertebrate' for 'Mammal'), and so on. The regress has no clear end point, Brenner argues, because of Thomasson's requirements that application conditions be non-circular (as above) (*ibid.* p609). Brenner's main criticism along these lines is that this introduces an infinite backlog of sets of application conditions that we must have learned previously, if we are to acquire a new term. Similarly, if we want to apply a given term in our language, we appear to have an infinite checklist to run through to see if we are correct to do so. Given that we are finite beings, Brenner deems these to be implausible as accounts of what happens when we learn and apply terms in our language (*ibid.* p608).

Thomasson has a reply ready to Brenner's objection, and other critiques that might put pressure on the two main functions of application conditions, as I have done briefly above. This reply is that application conditions need not be stateable by speakers to be usable (Thomasson 2015, p91). If we took this proviso simply and at face value, it would seem to

---

<sup>14</sup> This is the same reason we noted in Ch.2 that a "mentalist" expressivism such as that of Barker (2007) does not work for the neo-Carnapian, i.e.: because the latter needs an account not only of how we come to make ontological assertions, for instance, but also of how those assertions come to be true.

solve the issues discussed; we would not need to use any other sortal terms to understand a given term's application conditions, so Brenner's regress never gets going. By the same token, my simpler point that terms appear usable in the absence of reportable knowledge of their application conditions would be a non-issue. However, this is only the first stage of a reply. In order to flesh out such a response, we need to explain how exactly application conditions are known – what sort of thing knowledge of them amounts to – such that they can be at once un-statable, and at the same time able to underpin and justify ontological judgements.

In the remainder of the chapter, I will put forward two broad responses to this question. The first, put forward in section 4, is broadly representationalist. This will encompass three potential ways of dealing with the regress, namely letting it run, adopting a primitivism about certain sets of application conditions, and adopting a particular sort of reductionism about application conditions generally, to be discussed in the following section. What these responses have in common is that they pitch application conditions, and SCs in particular, as providing a semantic account of our sortal terms. I will evaluate the compatibility of each of these responses with easy ontology, ultimately finding each to be incompatible with the theory and/or its aims. In section 5, I will then go on to detail the alternative, expressivist reading of application conditions. This view does not understand application conditions as a semantic theory, but rather as a theoretical tool for deflation. Despite some tensions, I will argue that the expressivist reading thereby avoids the regress, and remains compatible with easy ontology.

#### 4: The representationalist picture

If we hold that ACs and SCs are different in kind, then we have the beginnings of a solution to the problems above, and, further, of a fuller semantic picture than previously. First, as regards Brenner's regress, note that the problem has been split. The original problem was that application conditions are required to be both known (in some sense to be determined) and used, and simultaneously grounded by terms with other application conditions. Now, we have ACs, which need to be known and used, but not grounded in the same way as SCs, and SCs, which are understood in terms of other sets of SCs, but need not be directly known and used.



It is not hard to see how we could make ACs immune from the regress. Because ACs do not ground the truth or falsity of any ontological facts, but merely whether or not we are warranted in believing them, there is no obvious problem in supposing that ACs are trained rather than explained propositionally in the way Brenner supposes application conditions to be taught. Training someone in a term's usage, in this sense, is not a matter of explaining it to them or conveying information in any direct sense, apart from telling them whether their attempted uses are correct or incorrect. Rather, it is a matter of demonstrating cases where the term applies, and thereby helping the learner to see "how to go on" – that is, how to apply the term in a range of unseen counterfactual cases.<sup>15</sup>

We might think of mastery of ACs as possession of a type of "knowledge-how", as opposed to factual knowledge; correct usage according to ACs could then be trained and adhered to, but not necessarily stated or explained to another person propositionally. The fact that explanation of skill-based knowledge in terms of propositionally known facts would be infinitely long is one of the key arguments that there must exist a type of knowledge distinct from propositional knowledge. Ryle (1945, p6) notes a regress along these lines that bears a striking resemblance to Brenner's: Suppose we want to demonstrate a valid argument to a student of logic, but cannot get them to acknowledge that the argument follows, in spite of our demonstrating the premises and the conclusion several times. If we are bound to characterise the understanding of logic that we have but the student lacks in terms of factual knowledge (knowledge-that), we must conclude that they are missing something like a maxim – that where we have premises like these, this conclusion follows. But this really only introduces another argument whose validity the student has to understand, namely that where these premises are true *and* the new maxim holds, the conclusion of the original argument follows. Further, if we view factual knowledge as the determining factor in the student's understanding, there exist no other ways of trying to make them understand than by supplementing in maxim after maxim.

If this was really what understanding consisted in, Ryle maintains, we would never be able to teach anything of this sort, as this would require relating an infinite number of rules in order to foster the understanding of just one. Similarly, it follows that one could never understand any given argument, as doing so would require the antecedent checking of an infinite number of maxims, which seems an impossible task. The mistake lies in this idea

---

<sup>15</sup> Cf. Thomasson 2015, p90; Thomasson's account here of learning application conditions seems to draw on the Wittgenstein of *Philosophical Investigations*.

that we need to teach the student new propositions; in fact, what we need to do is teach the student a skill, namely how to apply the first maxim. The similarity to Brenner's regress should be clear. Analogous arguments go for any action that involves applying a skill or following a rule: Because seeing how to apply a rule or skill is itself a skill, explaining skill in propositional terms will always lead to regress. Ryle therefore postulates a type of knowledge distinct from propositional knowledge, namely ability knowledge or "knowledge-how".

Analogously with Ryle, I want to propose treating mastery of ACs as a similar type of ability knowledge, as a step towards averting the present regress. On this account, Brenner's regress never gets started for ACs, because we need not look at any other terms to ground a given term's ACs – we are not told a set of conditions that we then internalise and call up again when we need to use them, but instead told which individual uses are correct and incorrect, until we gain an understanding of how to go on applying the term correctly.

Why can we not construe SCs along the same lines, and thereby end the regress closer to home for those as well? One answer is because this fails to explain adequately what is required of the world in order for a term, 'F', to apply in general, or at least to explain it in a way that has language represent the world somehow. This is the distinction I want to draw between the representationalist and the expressivist; the representationalist, understanding the meaning of language as stemming from its representing the world, will want application conditions, in the sense of satisfaction conditions, to concern the world and the arrangement of things therein, and cannot settle for the ability to train someone to detect relevantly similar situations, without supposing that there can be some account given of just how they are relevantly similar – some account, that is, of how the semantic values of the terms get set. Hence the representationalist approach to SCs will be to attempt to use them to specify the subject-matter of the term in question, this being, for sortal terms, the set of properties that they ultimately pick out.

We have part of an answer, then, to why we do not need to be able to state application conditions to be able to use them, namely, because the application conditions we actually make use of (ACs) are not of a form that needs to be stateable, and the ones that determine whether or not our ontological assertions are actually true (SCs) need not be available for direct recall, because of the availability of ACs to stand in for them. There is still a loose end though, namely in how to deal with the regress for SCs. I will put forward a few proposals below, but I do not view any of them as entirely satisfactory options, some

because they do not fit with easy ontology, others for more general reasons. This will help to distinguish between the different versions of this semantics that are on the table.

#### 4.1: Response 1: Ignore the regress

It is open to the easy ontologist at this point to hold their work done, and the problematic aspects of the regress averted. As mentioned, the regress poses a problem regarding how we get to know about application conditions if they require an infinite backlog of information. Although it makes for rather a bare theory, it seems as though the reply above, which shows that the application conditions that we do actually get to know about are not the ones with the infinite backlog, is sufficient to avert the specific worry about teachability, learnability and usability. On this view, then, there remains an infinite regress for satisfaction conditions, but it is, as it were, out of sight and out of mind.

Although this may be an attractive option, it does not appear to be one that is open to the representationalist. As we saw in Ch.2, the representationalist approach has it that meaningful linguistic expressions ought to admit of a metasemantic explanation that picks out their subject-matter, thereby unpacking the semantic relation that holds between this subject-matter and the expression itself. In the case of sortal terms, this semantic relation is the relation of satisfaction; on the representationalist picture, we ought to be able to specify, using SCs, the subject-matter of our sortal terms, i.e.: the properties that the terms pick out, and which satisfy them. If we simply let the regress run, then we have no answer as to what, ultimately, the subject-matter of our sortal terms is – what properties they are ultimately related to via the satisfaction relation.

We can put this point another way, using terms that Thomasson herself provides: If the SCs for a term 'K' are regressive in this way, it becomes unclear what is "required of the world" for Ks to exist. Thomasson responds to this concern when talking about the possibility of circular application conditions (Thomasson 2015, pp96-7). For instance, if I say that one has an xliver iff one has a liver and does not have an xheart, and an xheart iff one has a heart but does not have an xliver, then it is still unclear in what circumstances one would have either an xheart or an xliver (*ibid.*, pp260-2) – in representationalist terms, it is not clear what properties ultimately satisfy the expressions 'xheart' and 'xliver'.

Thomasson's reason for explicitly outlawing application conditions that are circularly defined in this way is precisely to avoid such "bad company" objections, where we appear

to be able to outline application conditions for – and therefore legitimise – such otherwise empty terms. Brenner likewise references this passage as a reason why his regress cannot end in a circularity (Brenner 2019, p608). If circularity of this sort is not permissible, then neither should an infinite regress be. The fact that we know what it takes for there to be Rs in the sense that we can state the application conditions for ‘K’, even though those application conditions are subject to a regress, does not solve the issue; we can state the application conditions for ‘xheart’, but that fact by itself does not prevent its lacking a semantic value.

#### 4.2: Response 2: Find some basic terms

If some terms could be found at which the regress simply ended, then the problem would of course be solved. What remained would then be to explain how these terms got their meanings, and what sort of terms they are. Here, the account of these basic terms’ meanings would be intertwined with an explanation of what form their SCs take (such that they don’t refer to any other terms with SCs of their own), and how they came to have them. What is needed, then, is not just a primitivism about a certain set of terms, but a primitivism that explains how these terms can be meaningful without violating the rules binding SCs.

A promising option here is to distinguish which terms are basic along the lines of which are the most natural, i.e.: those that purportedly refer to the most basic natural kinds or entities. A prominent example of such a project is that put forward by Ted Sider. Sider (2011) outlines the concept of a “joint-carving” ontology, which contains only the most fundamental natural entities, reflecting the real structure of the world by carving it at its logical joints. Sider’s theory is essentially takes the idea of reference magnetism – that some entities are more eligible to be referred to by our linguistic expressions than others – and works it through to its logical conclusion, providing an account of the concept of worldly structure that underpins this magnetism. This sort of primitivist view, then, could provide us with a set of basic sortal terms that are meaningful not because they have a standard set of SCs, but because they carve at the joints. With this foundation, we could then build up the SCs for our other, non-fundamental sortal terms.

It seems quite clear that, however attractive a theory of this sort might be, it is unacceptable to the easy ontologist. Sider’s theory, and those like it, hinge on the idea that

there is a fundamental set of entities to which reference attaches, provided our language use fits the correct causal profile. This entails a reading of the quantifiers on which they have a substantial meaning, i.e.: on which unrestricted quantificational sentences such as “There exists some thing x” state substantial ontological facts. This is why Sider and others speak of different versions of the quantifiers as being more or less joint-carving, and why it follows that there can only be one really accurate set of quantifiers; they represent, one might say, a univocal notion of objecthood. As mentioned above, Thomasson has no truck with this. For one thing, easy ontology is based on the idea that the meaning of existence claims is intrinsically tied to sortal terms. Hence the bare quantificational statement “There exists some thing x” does not state a definite, first-order ontological fact. At best, it is equivalent to the claim “Some application conditions are satisfied”, without specifying *which* application conditions those are, a very different reading from the Siderian version.

Further, if there was a reading on which the quantifiers carved at the joints, then ontological conclusions could not be generated “easily”; while we could assert that the application conditions for our sortal terms were fulfilled, it would be a further, substantial question, and one only answerable by distinctly metaphysical inquiry, whether the entities thus attested satisfied the requirements for objecthood, and whether they could be quantified over. Another way of putting the point is that, on this view, it is possible for ordinary speakers to make widespread, systematic linguistic mistakes with a metaphysical origin, e.g.: using expressions for composite objects, when in fact there are no such objects, because a sufficiently joint-carving quantifier tells us that there exist only physical simples. The possibility of this sort of systematic error obviously runs counter to Thomasson’s easy-ontological project.

#### 4.3: Response 3: Reduce SCs to ACs

A third option that bears talking about is the possibility that SCs for our terms could reduce to ACs. That is, the ultimate terms in which application conditions are to be understood are those of an “observation language”, understood in terms of assertability. Since ACs are knowable, usable etc. we therefore have a grounding for SCs that likewise guarantees their usability. Further, it is clearer on this response how the SCs for terms can be set by stipulation, namely by giving them in terms of observation-language concepts we already know how to use (even if we can’t provide a verbal explanation of the details). Obviously some work is required to explain the exact links between SCs and the ACs in terms of which

they are to be understood; one's being justified in thinking that there is an electron cannot be sufficient for there to be an actual electron, or the concepts of illusion and deception (and thereby the notion of an intersubjectively apprehended world) would go out of the window. There are, however, mechanisms that can be implemented to restore some nuance to this sort of view. One such mechanism is the idea, culled from pragmatism, of an "ideal state of information": The SCs for a sortal 'F' could be phrased such that, in an ideal, maximally informed state, one would be justified on the basis of the basic observation-language concepts in attesting to the existence of an F. This solution still allows us to cash out the states of affairs in which Fs actually exist in terms of warranted assertability, but retains the possibility of our being mistaken about our individual assertion of existence, an outcome we want to preserve (of course, the notion of an ideal state of information might be problematic in itself).

Thomasson does talk about entities whose terms have assertability and satisfaction conditions that cross over (such that they exist iff we are justified in believing that they exist), although she does not consider this for the basis of a semantics of application conditions. In discussing whether there is any stable sortal version of the term 'object', Thomasson mentions the "Spelke object" (Thomasson 2015, p111). This is a set of criteria hypothesised by Elizabeth Spelke (1990) that infants use to identify objects in their environment prior to obtaining a conception of object-permanence. As such, the Spelke-object criteria involve criteria that reliably pick out cohesive lumps of matter – criteria such as boundedness and continuity of movement within a visual field – but are light on complex reidentification criteria (infants in the relevant age bracket tend to favour continuity of movement over shape and colour in reidentifying an object that is passed behind a screen, for instance, which results in many false reidentifications). These criteria are plausibly basic, in that they are satisfied by the activation of "sensory input analysers", which are supposed to be cognitively basic (Thomasson 2015, p111). The point here is that the circumstances in which a Spelke object exists are exactly the circumstances in which one is justified in asserting that one exists; both sets of conditions are met when the relevant sensory input analysers register boundedness, cohesiveness etc. Hence, according to Thomasson (*ibid.*), one could not be mistaken about whether there is indeed a Spelke object.<sup>16</sup>

---

<sup>16</sup> One could contend that Thomasson is going for a stronger position in the relevant passage, namely that the Spelke-object criteria actually *are* the application conditions for a basic version of a straightforward sortal 'Object'. Brenner follows such an interpretation, and therefore rebuts

Thomasson, then, does allow for the possibility of entities of the type described, even though this apparently allows some ideal objects to enter into easy ontology – but then, easy ontology was already very ontologically profligate, and there are some entities that are plausibly ideal. Armed with a basic language of ideal entities then – our observation language – we can begin to construct application conditions for terms that are not clearly sensory in nature, or are not as obviously to be understood in terms of justification, with, at the base, an understanding of our basic terms’ use that arises from ordinary linguistic training. This, at least, is the broad outline of the reductionist picture. This sort of theory has been put forward before, notably by Carnap (1928) and Husserl (1970), both of whom Thomasson notes as precursors to easy ontology (Thomasson 2015, p22 and Thomasson 2019, p297 respectively). However, for both of these theorists, ease of doing ontology is just a happy side effect of an otherwise very heavy-duty theory about meaning that looks to explain how such judgements, among others, get their meaning in the first place. It is worth noting also that, despite their respective authors’ professing their realism, both theories look to edge much closer to idealism than Thomasson wants easy ontology to do. If easy ontology requires an idealist outlook in order to function, then Thomasson looks to bury the lede; the fact that we can make ontological assertions easily looks, as with Husserl and Carnap, like a secondary feature of a much more heavy-duty theory about how those assertions get their meaning in the first place.

Again, I do not think that the above is what Thomasson has in mind. The easy ontological project is pitched to convince an audience of philosophical moderates, who will likely balk at a version that has full-blown idealism as a cost. More than that, this refocuses the theory to be an idealism with an easy-ontological addendum. In the spirit, then, of finding a semantic picture to supplement easy ontology, rather than the other way around, I will look at the possibility of an expressivist semantics for easy ontology in the following

---

Thomasson’s claim that infants cannot be mistaken about the existence of Spelke objects, as that would entail infallibility about the existence of objects *simpliciter*. However, I think that Brenner misdiagnoses what is going on here. First, it would be counter-productive for Thomasson to introduce a basic-object sortal after arguing so strongly that none exists. Second, and more importantly, if the application conditions for ‘Spelke object’ are to be understood in terms of sensory input analysers, Thomasson must be after criteria for an ideal object, in which case it stands to reason that one couldn’t be mistaken about its existing. This last point supports her claim that this doesn’t provide the sort of objecthood that the metaphysician might be looking for at this point. I bring this up because it seems like the point at which both Thomasson and Brenner get closest to the sort of distinction that I want to draw between ACs and SCs. Brenner, however, crosses the streams, assuming that application conditions must always perform both functions, while Thomasson leaves off pursuing that level of semantic detail.

section, taking on board lessons learned above about the specific problems facing the theory.

#### 5: The expressivist picture

As outlined in Ch.2, the global expressivist eschews the possibility of a representationalist explanatory account of the semantic features of linguistic expressions. Instead, the expressivist offers a use-theoretic account comprised of two parts, namely the linguistic structures in terms of which the expression's correct usage is to be understood, and the linguistic function that underpins its utility value – an account of why we have this expression, rather than another, in our language.

The main drive for this sort of account, I posited above, was the occurrence of placement problems for representationalist theories. These occur because representationalism requires us to delineate the entities or properties in the world that the correct usage of our linguistic expressions is sensitive to, in order to then outline the representational relations between these expressions and these world-properties. The problem is that, in attempting to outline these world-properties and representational relations, we end up having to adjust for an endless array of potential exceptional circumstances to make our account of the usage of the expression in question adequate to its actual usage. Hence we end up having to introduce ballooning – and, in theory, infinite – exceptions clauses into our account, curtailing hopes for its completeness. Global expressivism can be seen as a direct response to this sort of worry. The globalist holds a semantic account to be complete if it is extensionally adequate and gives an explanation of the function of the expression in our language. Since disquotational formulations such as “The term ‘Tree’ is true of trees” are trivially extensionally adequate, the expressivist avoids the placement problem by using these instead of attempting other, non-trivial ways of outlining the expression's subject-matter, as the representationalist does.

The new regress shares certain features with placement problems, specifically in that it demands that we provide a full account of the subject-matter of the terms being explained, which ends up looking like an infinite, and therefore impossible, task. However, there are some differences in the format of the problems. The regress demands, for each set of application (satisfaction) conditions, that we be able to state what properties satisfy them, where this leads us on to be required to do the same for the application conditions for the



terms for *those* properties, and so on. The problem here is not to do with engineering extensional adequacy using exceptions clauses, but with finding a point at which to stop unpacking application conditions at some basic set of terms, for basic entities.

Another difference is that the placement problems that chiefly motivate expressivism are problems of vocabulary, paradigmatically the vocabulary of natural science; representationalism typically wants to understand the properties that expressions “latch onto”, as well as the semantic properties and relations by which this occurs, in naturalistic terms. The traditional form of the placement problem is therefore the problem of mapping different vocabularies onto the vocabulary of naturalism. In contrast to this, the regress we are currently dealing with is not concerned with naturalistic reduction, but rather with reduction generally – with the problem thrown up by requiring that every term be mappable onto *some* other set of terms, without ever producing circularity.

The approach the expressivist took to the placement problems of Ch.2 was to stop at the first step, offering a trivial account of the subject-matter of the term in question, and refusing to be drawn into attempts to break it down further. Essentially, the response here is the same; the only outline of the meaning of the sortal term ‘K’ guaranteed to be complete is the trivial formulation “‘K’ applies iff there are Ks.” The expressivist can then supplement in use-theoretic explanations regarding the term’s utility, to sit atop this extensionally accurate base.

Immediately we run into a problem: The formulation above looks to fall into exactly the sort of circularity we were worried about when setting up the regress. That is, we seem to have offered an explanation of the form “The application conditions for ‘K’ are fulfilled iff Ks exist”, which we earlier prohibited, as such explanations look to be unclear as to what exactly the world has to be like in order for there to be Ks. However, there are different reasons we might hold this worry. One, as we saw above, concerns the bad company objection, which attempts to show how meaningless terms can be introduced into the language using easy ontological principles. This is the worry that leads us to the regress, which pushes us to prove that our terms are meaningful by providing them with a determinate set of application conditions; if we can provide such an account for a particular term, we will have shown it to have a substantial meaning, and thus head off concerns about its being “bad company” for easy ontology to keep.

Note that this version of the worry frames the bad company objection as expecting a representationalist solution. Indeed, it is this framing that gets the regress started: The

method by which we have to prove our terms to be non-empty, on this reading, is to provide a complete, substantial account of their subject-matter by unpacking their application conditions, and thereby getting at the properties the terms represent. However, reading with the original worry, this is necessary only insofar as we suspect the term of being bad company. For terms whose content we are assured of, such as terms whose usage we can master, there is no problem, because there are no doubts about their being empty. In other words, the problem with the terms 'xheart' and 'xliver' is not that its explication in terms of application conditions seems to be circular *simpliciter*, but that there appears to be no way of understanding them independently of this circular explanation. For non-dubious sortal terms, whose usage language users can be trained in as in section 4 above, there is no such problem; we assume, as part of the force of the example, that there is no such training that could be offered in the correct usage of 'xheart', because the term is intended to be meaningless. Avoiding the regress for this reading of the problem thus necessitates abandoning the representationalist assumption that our terms generally require it to be possible, in principle, to give a reductive explanation of the sort required by the regress in order for them to be contentful.

The other version of the circularity worry is deflationary. Thomasson initially puts forward the non-circularity requirement as a guard against reiterating a non-deflationary conception of existence using the apparatus of application conditions (Thomasson 2015, pp96-7). The idea here is that the sentence "The application conditions for 'K' are fulfilled iff Ks exist" is a way to re-substantiate the concept of existence by making the application conditions for 'K' dependent on some further constraint – that Ks hold some substantial property, which would then be identified with existence. On the easy ontological approach, the sentence is actually necessarily true, since it is a restatement of principle (E). The point Thomasson is making is precisely this: The sentence "Ks exist" is nothing more than a statement in the object language of the metalanguage sentence "The application conditions for 'K' are fulfilled." Hence including "Ks exist" as an application condition would be uninformative – indeed, we could only think it informative if we attributed a non-deflationary meaning to "exist", which the easy ontologist does not do.

Thomasson therefore does not deny the truth of the sentence "'K' applies iff Ks exist" – indeed, according to easy ontology, it is necessarily true. However she does say of the sentence that "While [it] will always be true, it will not count as an application condition in our terms." (*ibid.*) This stipulation means that the expressivist cannot offer their circular formulation as the application conditions for 'K'. However, with the above, the expressivist

has already rejected the idea that we need to be able to offer (complete, non-circular) application conditions for a term in order for that term to be meaningful. Explanations of meaning, instead, should take the form outlined in Ch.2. The expressivist view on application conditions, then, amounts to a rejection of their utility for complete semantic explanation; we cannot, on grounds of regress, usefully conceive of the meaning of a term as exhaustively underpinned by a set of application conditions.

The expressivist's response, then, looks very similar to the suggested response of section 2.1 above: Ignore the regress, and continue to use the vocabulary anyway. As noted, the reason against adopting this response that motivate the representationalist – the requirement that it be possible to give a metasemantic account relating sortal terms to their subject matter – do not motivate the expressivist, who is at best pessimistic about the possibility of such an account. If this response is not immediately clear, it is because we are limited here by the limited expressivist theory that has so far been set up. I will return to the point in Ch.8 and flesh it out more clearly. The intervening chapters will provide a more in-depth analysis of global expressivism, which will furnish us with the concepts and machinery necessary to address the problem more fully.

Note that the expressivist does not have to reject application conditions wholesale – just the particular purpose to which we were above tempted to put SCs specifically, namely that of cashing out the subject-matter of our sortal terms. There are still uses to which the expressivist can put application conditions. For one thing, application conditions still have a mundane use in reporting and questioning the commitments borne by our use of a sortal term; if I allege that a certain animal *x* is a platypus, you might oppose this on the basis that you have reason to believe that *x* isn't a mammal. Despite anti-representationalist reservations about the ability of application conditions to render a full semantic account of the term, we can still invoke them to explain why this is a sensible response – namely, because the application conditions for 'Platypus' contain those for 'Mammal'.

The above provides a pragmatic (if brief) argument for keeping the concept of application conditions available, as well as an argument that it is a concept already implicit in our ordinary language use. The sort of pragmatic argument here is fully at home in the expressivist approach; that application conditions have this mundane use in ordinary language does not run counter to our reservations about putting them to a reductive semantic-explanatory use. The expressivist, then, can acknowledge the primary function of

application conditions as cashing *specific* necessary and sufficient conditions for correct word usage, and therefore as cashing commitments.

Likewise, we can supply a similarly pragmatic account of the difference between ACs and SCs, for instance, as supporting the idea of mistaken but warranted application, or of deference to experts. The idea of a “no-fault” disagreement, where both sides are licensed to hold their respective views despite their being in conflict, as well as the idea that there is in fact a right answer to be had, also seems to hinge on this concept. Note that these explanations involve only the intra-linguistic moves and structures (deference, licence, correction and the like) attached to the distinction. These structures will be discussed in further detail in Ch.7, where I unpack Huw Price’s expressivist account of truth.

## Conclusion

The above picture shifts the conception of application conditions from being the semantic underpinnings of a term’s usage that fix its subject-matter to being a linguistic tool, defined in terms of their use-properties. This is not to say that application conditions cannot be theoretically useful. Quite apart from examining individual conditions for applying a sortal term, we can make use of the idea of the entire set of application conditions for a sortal term as a theoretical device. Such whole sets of application conditions are somewhat alike to theoretical particles – useful for reasoning and talking about, but difficult or impossible to observe in the wild. We can therefore run a version of easy ontology on an expressivist footing by adopting a minimal conception of the application conditions of a term ‘K’ as whatever conditions must be true of the world in order for ‘K’ to apply – a characteristically “thin” definition that ought to satisfy the expressivist, but should be enough to serve as a foundation for easy ontology. Application conditions then become a theoretical stepping-stone on the way to existence deflationism – not a semantic theory in their own right, nor a substantial component of another semantics, but rather a trivially obtainable theoretical entity.

This is not exactly a quietism about application conditions, but more like a deflationism; the role of application conditions in lending substance to the meaning of our terms is gone, having been replaced by the minimal linguistic structure necessary to guarantee their utility in everyday contexts. From this we can expand to the theoretical construct necessary to

engineer existence deflationism, but without assigning to this minimal theory of application conditions the role of a semantic theory.

Although I have attempted here to resolve the regress problem for Thomasson's theory in a satisfying way on an expressivist footing, the solution is limited by the limited sketch of expressivism we have so far seen. In chapters 6 and 7, I will return to global expressivism, expanding the range of expressivist concepts we have available, and presenting and analysing detailed examples of expressivist semantic explanations. I will then return to the above solution in Ch.8 with the benefit of this deeper understanding of the position, and lay out an expressivist version of Thomasson's easy ontology more clearly.

## Ch.5: Refining the deflationary theory

We have so far seen two distinct ways in which the neo-Carnapian style of theory we are interested in preserving can be approached. In particular, we examined the attempts of Thomasson and Hirsch to provide a deflationary account of metaphysical debate in keeping with Carnap's suggestions in that area. Despite the deflationary outlook they share, and their common origin, Hirsch's and Thomasson's approaches nevertheless appear to be exclusive; Hirsch embraces quantifier variance, while Thomasson rejects it. If such exclusive theories are possible under a neo-Carnapian umbrella, we must pose the question: What sort of deflationary, Carnapian theory do we want?

Continuing in the spirit, from previous chapters, of dissolving apparent theoretical differences, I want to argue that Hirsch's and Thomasson's approaches are actually reconcilable. Specifically, I do not think that the revised versions of the accounts arrived at in Chapters 3 and 4 are at odds in the way they first appear to be. In this chapter, I will attempt to drill down to the real differences between Hirsch and Thomasson's respective theories, and attempt to argue that these do not disqualify them from working in tandem. Section 1 will look at some potential differences between the theories, eventually arriving at quantifier variance (QV) as the key point of contention. Section 2 will pose the question whether QV is necessary for a deflationism of the sort we want, arguing that it is not, and moreover, that the deflationist ought not to recognize the dispute over QV as substantive. Section 3 will examine another potential point of difference, namely on interpreting speakers in metaphysical debates as speaking the truth in their own language, or speaking falsehoods in ordinary English; I argue that both theorists ought to offer the former interpretation. I conclude by arguing for a fusion of Hirsch's and Thomasson's respective theories – one on which QV is dropped as the only (thin) pretext for differentiating the two.

### 1: Initial differences

The primary difference between Hirsch's and Thomasson's theories that we have seen so far is in their basic building blocks: Hirsch undertakes metaphysical deflation in terms of the quantificational commitments of whole languages, while Thomasson works in terms of individual expressions and their application conditions. Note that, on the face of it, this

difference in scope alone does not need to imply that the two theories be strictly exclusive. At this point, it may simply be that Hirsch's deflationism works at one scale or resolution, while Thomasson's works at another, more fine-grained one. With the correct way of "translating" between the two, we could create a unified theory that reflects the commitments of both authors, and achieves the deflationary outcome they are both working towards.

We must note that there is in fact a key difference attested by the authors themselves, namely the issue of quantifier variance. This is the idea that the quantifiers must work differently in different languages, since those languages recognize different entities as existent. Stated more precisely, a quantificational sentence such as "There exist tables" might be true in language  $L_1$ , but false in language  $L_2$ , hence the word 'exist' must be doing different work, with a different meaning, in each language.

Thomasson sidesteps quantifier variance by proposing a reading of the existential quantifier that consists of the formal properties abstracted from the supposedly different quantifiers attested by quantifier variance. Thomasson's formal quantifier is in a sense metalinguistic, in that it contains no application conditions of its own, but instead relies on the fulfillment of the application conditions of some sortal term,  $K$ , as follows:

E:  $K$ s exist iff the application conditions for ' $K$ ' are fulfilled.

Thomasson 2015, p86

We should recall here the refined concept of "application conditions" in use. As per the previous chapter, it is not necessary that application conditions be specifiable. This notion is a reductionist hangover from representationalism, specifically, the idea that an expression's application (read: "satisfaction") conditions need to be explicable in terms of the building blocks of the one correct ontology in order for those conditions to be coherent, with expressions that fail to meet this test being applied semi-arbitrarily. Obviously, Thomasson's easy ontology has no truck with such an application of the idea of ontological correctness. Hence there is no guarantee that our ontology is up to the task of elucidating the application conditions of our expressions, where this is not, of course, to be seen as a damning failure of either our ontology or our language. As argued above, mastering the correct application of a linguistic expression is a skill, and is learned as such. It does not follow from someone's having acquired skill-based knowledge that they can outline the basis on which it is applied as an orderly, finite set of conditions, nor even that this should be possible in principle.

Application conditions are thus to be seen as a theoretical device; we can infer from an expression's being applied systematically that there are conditions under which it is correct and incorrect to apply it. We may even be able to unpack some of these conditions, but we should not expect to be able to create a general, informative principle for its application, nor to be able to unpack the conditions "all the way down". Nevertheless, we can note apparent differences in application conditions, hence their usefulness as a tool for deflation. This conclusion is an upshot of the metasemantic quietism licensed by our putting easy ontology on an expressivist footing, as opposed to the previously assumed representationalist one.

With that groundwork laid, the question that now confronts us is whether Thomasson's formal reading of the quantifier can be made compatible with Hirsch's views about language differentiation and quantifier variance. As a first point, note that merely talking of differences in language is compatible with the formal quantifier; Thomasson's schema still allows us to use the notion of a language to draw a line between e.g.: phonetically similar but semantically different sentences, as a tool for deflation. All we need for this to work is to be able to apply the sortal 'Language' in the appropriate way, that is, with individuation conditions that imply that sentences can have different meanings in virtue of being uttered in different languages. Given our ability to apply the term 'Language' in this way, with apparent consistency and coherence, it does not look too controversial, on the version of easy ontology we are working with presently, to propose that the notion of language can serve as such a deflationary tool within the easy-ontological framework.<sup>17</sup>

I suspect that Thomasson's reason for engineering a theory that doesn't *rely* on the differentiation of languages lies in her attempts to appeal to a broadly Quinean audience. Quine himself was generally wary of the concept of languages so individuated because of his scepticism regarding semantics, and his suspicion that such a concept of languages would necessarily have them individuated at least partly semantically. As addressed in Ch.1, Quine held out hope for a thoroughly non-semantic way of individuating languages to support his concept of sentences, although it is doubtful whether his theory here is really

---

<sup>17</sup> Obviously, this is a rather uniquely philosophical use of the word 'Language'. The everyday usage – noting the difference between English and French, for instance – is obviously not primarily put to work as a deflationary tool, there being very few phonetically identical sentences between these sorts of languages. I do not view this as a problem; whether the philosophical notion of a language is derived from this usage or is a fresh creation of philosophy, I would argue that the present theory validates it on the basis of consistency of application alone (this essentially recaps the earlier discussion of Carnap's own take on languages).



workable. As regards easy ontology, regardless of whether its deflationary machinery relies on language individuation as that of Hirsch does, it looks to follow quite straightforwardly from the theory that such languages do exist – we can judge as much using our mastery of the sortal ‘Language’. The two approaches therefore look to be broadly compatible on the language front; it looks as though we can still make use of this tool for Hirschian intensional deflation while running the sort of quantifier deflation Thomasson uses.

With access to the sortal ‘Language’ established for both theories, we begin to see how we could enact much of Hirsch’s deflationary method within the framework of easy ontology. The next point, then, is to question the *centrality* of quantifier variance to Hirsch’s theoretical framework. In Ch.3 we reconstructed Hirsch’s deflationism without having to rely on quantifier variance to do the work for us. As outlined, it was sufficient to successfully deflate a metaphysical debate that we showed that the participants’ key utterances were intensionally equivalent. Quantifier variance entered the picture because these utterances were apparently oppositely ontologically committing. But there are other responses here than differentiating the meanings of the quantifiers themselves, the most obvious being to differentiate the meaning of the sortal(s) used by each speaker. Taking seriously Hirsch’s top-down approach – the approach that advises us to analyse expressions based on the use-properties of the sentences in which they are used – the difference in the willingness of speakers of languages  $L_1$  and  $L_2$  to attest “There exist tables” could be down to a difference in the meaning attributed to ‘table’ just as easily as in the meaning of ‘exist’.

An initial consideration preventing us from pursuing this reading of the difference in meaning here is the widespread, systematic nature of the disagreement. In the sort of case under discussion, a metaphysical dispute between philosophers, the disagreement concerns not just entities (potentially) falling under the one sortal, but under an extremely wide range of sortals. The organicist, for example, will reject as false *all* sentences asserted by their opponent of the form “There exists some  $x$ ” wherever  $x$  does not answer to either the sortal ‘Physical simple’ or the sortal ‘Living organism’, rejecting, by the lights of the ordinary speaker, the existence of most things. It is understandable that one would infer from this broad pattern of use that the organicist is proposing a language in which the existential quantifier *itself*, not just a handful of sortals, has a different meaning from that in everyday usage – namely a meaning tied to the categories of physical simples and living organisms.

Another way of bringing out the difference in how Hirsch and Thomasson view this dispute is to note that, although they agree that metaphysical disputes are wrongheaded, they disagree as to how. According to Hirsch, there is no factually right or wrong side in any deflatable disagreement over a metaphysical question, since each side is speaking the truth in their own language; for Hirsch, closely following Carnap, the debate is wrongheaded inasmuch as its participants mistake its character *for* that of a factual debate, when it will really only be settled by a practical decision about which language performs better as a tool for scientific inquiry, navigating the world, or whatever other pragmatic tribunal.

For Thomasson on the other hand, it looks at first blush as though there is a fairly clear answer about who is in the right, factually speaking. As usual, it seems the easy ontologist will have to side with the common-sense ontologist, meaning here siding against the organicist. Further, this follows not from the practical decision put in front of us by Hirsch, but for reasons of ontological commitment; reading with Thomasson's formal quantifier, the organicist claim that no tables exist comes out as being that the application conditions for the sortal 'Table' are never fulfilled. Since this is something any speaker of ordinary English can straightforwardly disprove, the easy ontologist holds the organicist to be wrong not on pragmatic but on factual grounds. In this sense Thomasson appears to depart from Carnap in holding the debate to be resolvable factually, with one side being correct in their ontological commitments, where the other side is straightforwardly incorrect. The mistake the disputants make is to think that their language is tracking something deeper, about which ordinary language users are able to be mistaken, and attempting to discuss that, rather than simply looking to ordinary language itself.

As tempting as it is to read this response off Thomasson's easy ontology, I do not think that it is actually the response we ought to take to the case in question, in the spirit of Thomasson's theory. Although Thomasson herself attempts in *Ontology Made Easy* to steer clear of commitment to semantically defined languages as part of the building blocks of her theory, presumably to mollify her neo-Quinean interlocutors, deflation in virtue of language differentiation does still play a part in her deflationism (Thomasson 2007 is, for instance, less reserved about this fact). Hence it is not quite right to read the easy ontologist's formal quantifier as a quantifier *per se* – since, for one thing, the formal quantifier does not come attached to a language. The difference, then, between Hirsch's quantifier variantism and Thomasson's formal quantifier is not quite so clear cut as their providing different deflationary results; the following section will look more closely at whether it is possible to tease the theories apart along these lines.

2: Do we need quantifier variance?

On the easy ontological view, quantifier variance is explained away in terms of application conditions. It is not, for instance, that the organicist and the common-sense ontologist are using different quantifiers, but rather that they are attaching different application conditions to their sortals – the quantifier is used in the same way by each speaker when they state an ontological truth, namely to assert that their personal sortal's application conditions are fulfilled.

The main obstacle to this interpretation is that the difference is systematic and widespread, but that it looks to be specific to quantificational sentences. Given these facts, it looks very much as though the speakers really do harbour different rules for the application of the quantifiers themselves. It might also occur to us to ask what application conditions exactly the organicist is supposed to be attributing to sortals such as 'Table', given that they never quantify over tables, and therefore never hold the sortal's application conditions to be fulfilled.

I think we can answer both worries in one go. The organicist view is that the only real existents are living organisms and physical simples. Rather than alleging a modification of the quantifiers themselves, we can instead interpret this as the introduction of a linguistic rule stating that any sortal used in quantification should have a disjunction containing the application conditions for 'Living organism' and 'Physical simple' appended to its own application conditions. In other words, the organicist will only assent to the existence of a table if the application conditions for 'Table' are met, *and* either the application conditions for 'Physical simple' or 'Living organism' are also met by the same entity.

There are a few things to note about the proposal for interpreting the organicist as introducing this specific linguistic rule. One is that it achieves the same effect as quantifier variance without resorting to varying the meaning of the quantifiers themselves, instead placing the emphasis on differently used sortals. Despite this fact, we still find a use for the concept of a (semantically individuated) language; the organicist is proposing to introduce a rule that would entail widespread changes for our language use, to the extent where we could justifiably say they are speaking a different language altogether, since, for instance, ordinary speakers taking them to be speaking everyday English will have trouble understanding them unless they learn the new rule. Granting that the individuation

conditions for different languages may yet be slightly vague (limited vagueness in application not being damning for a sortal concept), it seems like the notion of a bespoke language for the organicist has some deflationary traction here, despite the fact that we haven't given it its own distinct quantifiers.

Another thing to note is just how closely the proposed interpretation fits the organicist's own conception of what they are doing. The organicist contention is that if something exists, then it must either be a physical simple or a living organism. We can transform this ontological claim by semantic ascent into the purely linguistic claim that for any sortal 'F', if the application conditions for 'F' are met, then the application conditions for either the sortal 'Physical simple' or the sortal 'Living organism' must also be met by the same thing – exactly our formulation above. Tables obviously fail this test, meaning that the existentially committing sentence "Tables exist" (read: "The application conditions for 'Table' and either 'Physical simple' or 'Living organism' are satisfied by the same thing") comes out false, while the sentences "There are cats" and, potentially, "There are spacetime points" come out true. We have, by following the organicist's own intentions, essentially re-engineered the specifically organicist quantifier alleged by Hirsch using the concept of application conditions, and the organicist rule that wherever the quantifier is used, it must be used in conjunction with the stated application conditions. Those application conditions now function as a sort of "suppressed sortal" – "suppressed" because it is not vocalised outright in the sentences in the new language, although its application conditions are in use.

Stated thus, the difference between the quantifier variantist and easy ontological interpretations might seem rather trivial – what, really, is the difference between saying that the metaphysician changes the application conditions embedded in the quantifier itself, and saying that the new application conditions are located in some "suppressed sortal", independent of the quantifier? This question leads us to an interesting point: Neither Thomasson nor Hirsch has the metalinguistic resources to express any categorical difference between their theories. The reason for this is that the dispute over whether such a change in meaning as the one discussed above is "located" either in the quantifier or outside it must perforce assume the sort of representationalism that is incompatible with both sorts of deflationism. This is because both Hirsch and Thomasson (going by the reconstructions in chapters 3 and 4) are committed to the sort of top-down metasemantics that analyses meanings of expressions in terms of the changes they bring to the use-features of sentences in which they are used, as contrasted with a representationalism that

allocates discreet meanings to linguistic expressions based on a purported connection between those expressions and the subject matter that they represent.

What this means is that if we want to differentiate between two semantic explanations of a sentence's meaning – here differentiating between the quantifier variantist explanation of the organicist language's non-standard practices and the suppressed sortal explanation – we have to be able to put the options in terms of different use conditions for sentences. These sentences don't necessarily have to be the sentences being explained, but can instead be other sentences in which the relevant expressions are used. The problem for the current debate, on this schema, is that there is no way to cash the difference between the quantifier variantist and suppressed sortal explanations in sentential terms, since, in effect, both explanations are consistent with the same data; because the linguistic rules proposed by each theory (using the quantifier itself differently and introducing a suppressed sortal respectively) are implemented whenever we construct a quantificational sentence, and both rules have the same effect, namely that of rendering false any sentence that quantifies over anything not falling under the noted sortals, there is no way to parse the difference between the theories in terms of a difference in the conditions under which it is correct to assert a given sentence. In Hirschian terms, the difference between QV and easy ontology is reduced to a mere difference of hyperintensional structure, rather than any more substantive intensional or metasemantic dispute.

I would conjecture that the felt difference between the two theories is a hangover from representationalist thinking. While the metaphor of "locating" meaning within this or that linguistic expression is not disallowed by the present metasemantic approach, since we can still talk of the meaning brought to a sentence by a particular expression, it is still apparent that that metaphor can be misleading, as when it creates the appearance of categorical difference between theories where there is none storable, as has happened here. My suspicion is that the bottom-up approach of representationalism lingers in our predisposition to attempt to assign meanings to individual expressions, even where this makes no difference in terms that our current metasemantic theory can recognise; the metaphor is itself a metaphysical one, in that it depicts meaning as an intrinsic feature of an expression, and not, as we have it on our current framework, an extrinsic relation comprising the effect of the expression on the use-properties of sentences in which it is used. It seems that there is really no question, then, of whether the application conditions introduced attach to the quantifiers, as on QV, or to sortals, as on easy ontology – that is, unless we can show that either produces different results in isolation from the other.

One way to attempt this might be to consider, briefly, the “cookie cutter” metaphor for mereology. In this thought experiment, due originally to Hilary Putnam (see Putnam 1987), mereologists with different views are presented with a universe containing only three things,  $x_1$ ,  $x_2$  and  $x_3$ .<sup>18</sup> When asked how many objects there are in such a universe, the mereologists will give different answers: The mereological nihilist will say three, while the universalist will say seven (three dots, then three groups of two dots each, and finally one group of three dots). Each mereologist divides up the same arrangement of entities differently, as one can use differently shaped cookie cutters to stamp out different shapes from the same cookie dough.

Putnam’s thought experiment is potentially important for us because it appears to hinge only on the bare logical apparatus of quantification, not on the specific application conditions of any sortals; if the different answers are obtained by virtue of different languages being spoken, then there can be no question that these languages must have different quantifiers, because there are no other linguistic components present that could bring about the change in meaning. This puts the matter in terms our deflationary metasemantics can deal with, by noting that, in this case at least, the differences in hyperintensional structure of the sentences under examination are ineliminable without eliminating the quantifiers themselves from those sentences.

We can go some way to solving the issue by looking at how Thomasson might unpack the case under consideration. The question we must immediately ask is whether any disputant is using the quantifier in Thomasson’s deflationary “covering” sense, as quantifying over anything that answers to some sortal. If this were the case, then once again the temptation would be to say that this disputant was correct in their ontological assertion, and the others were wrong. However, this is not quite correct. The fact is that Thomasson’s formal use of the quantifier is generally not meant to cover all *possible* sortals ever – that is, all sortals *in any language* that are capable of being outlined with a well-defined set of application and coapplication conditions.<sup>19</sup> If it were so, then it would be practically unusable, as there is surely an infinity of such sortals possible. It would be doubly useless for counting, because the answer it turned up would almost always be that there were indefinitely many things, because there are indefinitely many such possible sortals whose

---

<sup>18</sup> This is, of course, a loose description, the act of describing the case being potentially problematic in and of itself.

<sup>19</sup> Caveats on “well-defined” as per the above discussion.

application conditions are met. This seems to follow from the relative ease with which we can come up with novel and completely arbitrary sets of application conditions for sortals.

This, though, is not the only way to apply the deflated interpretation of the quantifier. With regards to counting specifically, Thomasson uses the quantifier to cover only the sortals that are part of the language in which the counting is taking place (Thomasson 2007, pp121-5). It is not hard to see why this is: We are interested in modelling the quantificational structure of sentences such as “There exist three things”, which have to be understood as being formulated in a particular language. Thomasson is clear (*ibid.* p119) that quantificational claims, including purportedly “bare” quantificational claims, such as “There exist three things”, need to draw on a stock of sortals or be counted incomplete, in the sense that their application conditions are underspecified. The word “things” here can act as a dummy-sortal, able to be substituted for any of a range of sortals, but that range needs to be delimited (i.e.: as the range of sortals available in the language being spoken) in order for the sentence to have any sort of applicability, something reflected in the highly contextual use of the word in ordinary language.

Therefore, the way to model the different quantificational claims of the mereologists in the cookie-cutter case is to hold each as speaking a different language, and therefore drawing on a different range of sortals. This reveals that the case is, in fact, underspecified from the start, since it provides no sortal(s) that  $x_1$ ,  $x_2$  and  $x_3$  answer to, by virtue of which they meet the criteria for the formal quantifier. On Thomasson’s schema, we must assume that, for the case to be intelligible,  $x_1$ ,  $x_2$  and  $x_3$  satisfy the application conditions of at least one sortal. For simplicity, let us assume that they all satisfy the same sortal ‘F’ (e.g.: ‘Physical simple’). Then the sentence “There exist only three things” in the mereological nihilist language would translate metalinguistically as the claim that there are only three instances of sortals having their application conditions satisfied. Given that the nihilist, we presume, counts only ‘F’ among their language’s sortals and no others, they are using the formal quantifier, and simply drawing from the limited stock of sortals their language provides; there are only three things in the universe that satisfy any of the sortals recognised in their language. The same goes for the universalist, who introduces sortals for (at least) first-order groups, and therefore gains composite objects in their worldview. In both cases, the supposed “bare” quantificational sentences actually hinge on drawing from their language’s stock of sortals.

This leads to the job of the suppressed sortal, as I see it. So far, we only have it that in each ontological language the stock of sortals drawn on by the formal quantifier is idiosyncratically limited. At present, this limitation looks arbitrary – we do not yet have an explanation of why, for instance, the nihilist is exempt, when drawing ontological conclusions, from considering the sortals recognised by the universalist, given that these sortals are available and in working order. The suppressed sortal answers this by explaining that the difference comes about from the linguistic rule proposed by the nihilist that the application conditions for ‘F’ be appended to the application conditions of every other sortal, whenever they are used for quantification. In this way, we don’t need to posit a reason the nihilist removes the other available sortals from consideration; they use as wide a range of sortals as the universalist, it’s just that they add to all of them the criteria for satisfying the suppressed sortal ‘F’. Hence the only things that satisfy any sortal, on the nihilist’s rules, are Fs, since only Fs could satisfy *both* the application conditions for whatever other sortal *and* the application conditions for ‘F’, which are appended to them.

The answer above should help to spell out, finally, the use to which Thomasson’s formal quantifier is to be put. It is tempting to see easy ontology as itself an ontological position, i.e.: a position on what exists and what doesn’t, hence the appeal of the flat-footed easy-ontological answers to the disputes examined above. But easy ontology isn’t designed to weigh in on these disputes with just another position (or by backing an existing position). Instead, its primary function, like all such deflationary theories, is to defuse the debate at the metaontological level by diagnosing its causes. In practice, this can mean validating particular ontological assertions made by speakers, such as when the speaker of ordinary English disagrees with a metaphysician over the existence of some ordinary object or other. However, such validations of “easy” attributions of existence are backed by the competence of the ordinary English speaker regarding the concepts of their own language, not by their adhering to the formal quantifier.

Ordinary English does not contain an inexhaustible supply of sortals, and a workable deflationism should allow that there can be sortals too arbitrary or contrived to be widely useful, such as Hirsch’s famous ‘Incar’ and ‘Outcar’ (Hirsch 1982, p32). The point of the formal quantifier is not simply to extend the reach of our ontology to these entities, but to highlight that the question of whether we should do so is a question of pragmatics – of whether it is useful for our language to include the relevant sortals or not, or equivalently, of whether we want to introduce a particular suppressed sortal to exclude them – not of metaphysical existence. Thomasson is not interested (and as deflationists, nor should we



be) in coming up with a “final” answer to what things exist, or how many. Even with the maximally broad covering use of the quantifier, no such answer would be language-independent in any interesting way, since the main factor influencing it would end up being the number of possible sortals in every possible language, which is almost entirely a linguistic fact. Hence it tells us nothing ontologically interesting. This is because the deflationary version of the quantifier does not serve as the final arbiter of ontological disputes, as if it were an abstraction of the content underlying all quantificational expressions; it is such an abstraction, but the point of it is that there is no such content, only a metalinguistic structure, which, as such, needs to be situated within a given system of sortals in order to be meaningfully applied. The formal quantifier is not *a* quantifier in the sense in which the quantifier variantist speaks of comparing different quantifiers, it is instead the linguistic structure of quantification abstracted so as to make clear the commitments, in terms of sets of application conditions, borne by its particular uses.

What the apparatus of the suppressed sortal shows is that we can model the results of quantifier variantism using Thomasson’s theory of application conditions for sortals, even in cases where it looks like the two theories should come apart. Characteristically for a deflationary move, this does not tell us which of the two theories we ought to adopt, only that the grounds for choosing between them are a lot less substantive than we might have thought. What we have seen is that quantifier variance is not necessary for a neo-Carnapianism – not even one that takes a deflationary stance towards purportedly bare quantificational disputes. I propose, then, that we drop quantifier variance proper, while retaining the chief insights of Hirsch’s intensional deflationism.<sup>20</sup>

### 3: Metaphysical claims: True or false?

In §1 I noted that one way to draw the difference between Hirsch and Thomasson was in how they diagnosed the motives misleading metaphysicians – where they attributed the “wrongheadedness” of metaphysical debates. The idea presented was that Hirsch sees the disputants as each speaking the truth in their own language where Thomasson sees them as making different claims, some incorrect and some (the ones that agree with ordinary usage) correct, and all in the same language. Much as we have begun to dissolve the

---

<sup>20</sup> Note that the deflation in this section, between Thomasson’s and Hirsch’s diagnoses of the dispute, was Hirschean (i.e.: intensionally deflationist) in character, but did not rely on quantifier variance, further illustrating the point that it is not an essential part of Hirsch’s theory.

differences between the two authors in metasemantic terms above, we can also begin to dissolve this difference of diagnosis.

This can be done by looking at the relation of each speaker to ordinary English. Where Hirsch takes the common-sense ontologist to be speaking ordinary English and the organicist to be speaking a language of their own devising, complete with bespoke quantifiers, we previously assumed that Thomasson should take both speakers to be speaking ordinary English, with different degrees of success. The dispute, then, is over whether the organicist is speaking ordinary English, or their own organicist language. I will outline the reasons the easy ontologist might view the organicist as speaking ordinary English, and then speak against these reasons, hoping thereby to show that easy ontology in fact pushes us to agree with Hirsch's intensional deflationism that both sides are speaking the truth in their own language.

There are two different reasons the easy ontologist might feel compelled to view the organicist as speaking ordinary English. First, and most straightforwardly, it may in fact be that the organicist takes themselves to be speaking ordinary English. This is the revisionist route, where the organicist essentially claims that ordinary speakers apply the concepts of their own language, particularly the concept of existence, incorrectly much of the time, but that the correct metaphysical theory will offer the clarity or coherence that this usage is lacking. As we have it above, the easy ontologist makes the exact opposite claim: The organicist is in fact the one who is applying the concepts of ordinary language wrongly.

This case is tricky to untangle, but we can do so if we look to the idea of interpretive charity, an idea that ought to guide the easy ontologist just as much as it guides Hirsch, as discussed in Ch.3. In order to diagnose the linguistic mistakes that cause the debate to invoke metaphysics, we want to treat the organicist as charitably as possible as regards their linguistic competence. However, this is complicated by two conflicting facts about their linguistic usage; on the one hand, they themselves claim to be speaking ordinary English in the sense outlined, while on the other their ontologically committing assertions do not reflect the concepts of ordinary English, as the easy ontologist knows from having mastered those concepts themselves.

So which is it – are they speaking ordinary English badly, or are they mistaken about which language they *are* speaking? As charitable interpreters, we want to take the organicist to be making the lesser of the two mistakes, and therefore the smallest possible lapse in linguistic competence. As I see it, we should do this by taking them to be mistaken when

they describe themselves as speaking ordinary English. This is because, first, this is a single mistake, as opposed to the myriad of errors we would have to attribute to them if we held them actually to be using ordinary English. Further, and more compelling to my eyes, this mistake does not have the character of a first-order lapse in linguistic competence, but of a philosophical error; the organicist is not displaying an inability to apply basic sortal concepts that most ordinary speakers have no trouble mastering, but is instead showing themselves to hold a misconceived philosophical theory about the link that such concepts have to metaphysics. This is an error not of competence but of credence, so taking this to be their misstep seems to be, categorically speaking, the more charitable option.

The second case, slightly more difficult than the first, comes when the organicist takes themselves to be speaking a different, more ontologically precise language, such as “ontologese”. The problem here for the easy ontologist is that they can make no sense of such a language. Ontologese speakers take their language to “carve the world at its logical joints”, or at least to come closer to doing so than ordinary language, but the easy ontologist holds such a notion of joint carving to be confused in the extreme. It follows that, strictly speaking, there is no such thing as ontologese according to easy ontology. To be clear speakers are fully at liberty to construct a language that reflects their metaphysical commitments, and have statements of those metaphysical commitments be truths in that language. However, according to the deflationist, the would-be ontologese speaker’s stated aim of using the quantifiers in a maximally joint-carving way cannot be made sense of, so they can never actually be speaking ontologese as they construe it.

This appears to narrow our options for charitable interpretation; since we can’t take the organicist to be speaking ontologese, our only option seems to be to take them to be speaking ordinary English less than competently. However, this is not the only option on the table. As in the previous case, it is open to us to hold our organicist to be mistaken about the language that they are speaking, specifically by holding them to be speaking a bespoke organicist language, just as before. Note, though, that this language is *not* ontologese (because nothing is), but a language that, as before, incorporates a suppressed sortal into all of its quantificational assertions. This gives us the most charitable reading of the organicist who takes themselves to be speaking ontologese, by once again attributing to them the fewest errors, and of the most understandable kind. The systematic character of their deviations from ordinary English speaks against their simply failing to use the rules of ordinary language correctly, and for their actively adopting a different set of linguistic rules, but the incoherence (by the easy ontologist’s lights) of their assertion that they are

speaking ontologese prevents us from taking them at their word, even while being charitable. The alternative is that they *are* speaking a bespoke language, but that they are mistaken as to its character, with this mistake once again being rooted in philosophical error (i.e.: the concept of joint carving), not in a failure of linguistic competence (*cf.* Hirsch 2011, p212 for a similar point).

As a side note and complement to the above point, it is important to remember that although charity has been our guide here, we should not bend over backwards to be charitable to a deviant speaker. If, for instance, the speaker's deviations from ordinary English usage were apparently random, the reasonable thing to do would be to assume that they in fact lack a mastery of ordinary English. We could, of course, still hold them to be speaking in a language of their own, where this language was arbitrarily constructed to validate their aberrant utterances, but such would be an unwarranted stretching of the sortal 'Language'. The key for the organicist and comparable metaphysicians is that they are proposing to introduce a linguistic rule not present in everyday usage (although they themselves might not cash it directly in this way); we can, with minimal effort, outline the rules their language is following, and so treat it *as* a language, rather than as an *ad hoc* collection of "spot rules". This criterion bears flagging up here because of the arguments above that we need not take the speaker at their word regarding the language they take themselves to be speaking, at least where charitable interpretation trumps such an avowal. That is, we should remember that charity is not the overwhelming consideration, and that we still need to construct a set of linguistic rules the speaker *could* be following if we are to maintain that they are, in fact, speaking any sort of language. Further, this highlights the consideration mentioned in favour of Carnapianism in Ch.1 that the sortal 'Language' itself, as part of the "analyticity family", has significant deflationary utility in its own right.

The point to take away here is that, in the sorts of cases under discussion, both Hirsch and Thomasson ought to take the metaphysician to be speaking their own language. Contrary to the previous assumption, then, there is no difference in how Hirsch and Thomasson should interpret this debate. It should also be noted that whether we take the metaphysician to be speaking their own language or not, both Hirsch's and Thomasson's theories agree on how the metaphysician's utterances are to be treated. If we take them to be speaking ordinary English, then we must regard them as doing so badly by making far reaching if systematic errors. If not, then they are to be seen as proposing a new linguistic rule to be applied to quantificational sentences, the widespread adoption of which will be subject to its utility, not being sensitive to any metaphysical argument.

## Conclusion

I have attempted, in this chapter, to show that Hirsch's and Thomasson's theories actually disagree only on fairly minor points, and that those points of disagreement that are to be found are rather thin, to the point that either side would struggle to find sufficient tools to outline them. Specifically, I argued that the standout difference, quantifier variance, amounts only to a difference in the hyperintensional semantic structure that each theory attributes to quantificational sentences. Since attributing semantic facts based purely on hyperintensional structure, absent any difference in terms of usage, is a preserve of a representationalism that, as argued previously, is incompatible with both theories, I conclude that, contrary to appearances, there is no substantive difference between the authors' respective deflationist theories.

Of course, the word "substantive" here is important; while the two theories may well state the same facts, in the sense that Hirsch attributes to metaphysicians who argue past one another, just like these metaphysicians' respective languages each theory still has a different structure. The standard procedure in such cases is to ascertain which theory-structure has more practical value, and to go with that one. However, in this case, it seems possible, instead, to use both theories in tandem. Once we have established, via mechanisms like the suppressed sortal, that the results of the two theories are compatible, we can start to see how, rather than our rejecting one or the other wholesale, the two can potentially be useful in different ways in different situations.

As noted previously in passing, the two theories work at different levels of resolution; Hirsch's approach is to deal with whole languages (and, as we saw in Ch.3, whole intensional landscapes), whereas Thomasson's is to home in on individual sortals and their application conditions. Of the two, Thomasson's approach offers a more detailed diagnosis of the disputes it is used to deflate, since it is able to explain not only that disputants are using language differently, but also in what respects their usage of the expressions of their languages differ. The theory of application conditions is very versatile, largely because it is possible to rework many metaphysical questions into explicitly ontological questions. Nevertheless, there are metaphysical disputes we might want to deflate where Thomasson's theory does not clearly look to apply.

One example is the above deflation of the debate around quantifier variance in §2; it is not clear how we could cash this in terms of sortals and application conditions, but it nevertheless looks like something we ought to acknowledge. To be clear, this is not to state definitively that the dispute could not be cashed in such terms, but rather that, given the ability to draw on an intensional deflationism, there would be no need. If we regard Hirsch and Thomasson as providing tools or mechanisms for deflation, then, in light of the compatibility of their theories, we can feel free to mix and match which tools we use according to the situation.

In all of this, the one point of incompatibility that does get dropped from the combined deflationism is quantifier variantism. As we have seen, QV is not an integral part of Hirsch's toolkit, which is in reality chiefly built not around quantification, but around intensionality. The key part of Hirsch's deflationism, sans QV, is directly compatible with that of Thomasson. Hence we should drop QV as a liability, and make use instead of the other useful tools and concepts Hirsch offers (such as the notion of "unstructured facts", and his Carnapian language differentiation), alongside Thomasson's existence deflationism.

## Ch.6: Price's global expressivism

Huw Price provides the model for both global expressivism and a worked-out Carnapian deflationism. In discussing expressivism so far, we have had reason to call on Price's reasoning to clarify the exact commitments of both the global expressivist and the Carnapian. It has served us to do this piecemeal, because the overall stance Price takes towards semantics and metaphysics can be somewhat impenetrable, based as it is on much bespoke terminology that is very tersely defined. One of the purposes to the following chapter, then, is to provide a "key" to reading Price, clarifying some of his ideas and argumentative strategies, and critiquing others where they appear to conflict with the aims of the thesis – and thus with Price's own aims, since I take him to be perhaps the only author unambiguously aligned with both sides of the developing neo-Carnapian/expressivist view.

The main aim of the chapter is to flesh out the approach taken by a prominent global expressivist, and one whose views are closely linked to Carnap's. In Ch.2, we saw a fairly sparse characterisation of global expressivism. Here I will attempt to flesh out the approach of a particular global expressivist, to illustrate the independent motivations for expressivism, and the different concerns it is sensitive to as a methodology. In turn, this will give us a framework for understanding the sorts of semantic explanations the expressivist will want to give (Ch.7), and on which the success of the project will hinge (Ch.8). The chapter will be divided into two sections. In Section 1, I will attempt to make clear three key distinctions Price draws, upon which his position is built, as a means to understanding that position. Section 2 will articulate and examine Price's functional pluralism, the structure of his expressivist worldview, and assess its relations to neo-Carnapian metaontology, tying it back into the main concern of the thesis.

### 1: Three Distinctions

Price often presents his views in the form of dichotomies, often paring apart two positions or phenomena that are usually taken to be identical, and exposing both the differences and relations between them. We can establish a baseline of Pricean terminology by examining

some of these distinctions, and help to clarify some common misconceptions by critiquing others, as follows.

### 1.1 Thick and Thin Properties

Price (2009b) explains his expressivist view in terms of “thick” and “thin” properties. I believe this to be a mistake, if only in a strategic sense, as I’ll elaborate below. I’ll first recount briefly Price’s explanation of the thick/thin distinction, and how he uses it to outline expressivism, before detailing the misleading picture the distinction paints – a picture that has historically led to interpretations that are directly at odds with what global expressivists, and Price in particular, actually believe.

Price sets up the thick/thin distinction as tracking several other distinctions as used by the deflationist. One is the distinction between talking in the “theoretical voice” and talking with the folk. The deflationist is perfectly happy to talk with the folk about semantic properties, since the language of “folk semantics” contains predicates such as “True” that seem to be applicable in a consistent, rule-governed way. However, when it comes to theoretical talk of such properties, the expressivist must adopt a “strategic silence” (*ibid.* p258). This silence is necessary because, while the expressivist doesn’t want to provide an account of semantic properties in the theoretical voice on the one hand, neither do they want to use their theoretical voice to deny such properties. Were they to do so, they would fall victim to the old criticism, originally from Boghossian (1990), that since they deny that there are any properties such as truth, they must deny, in the very theoretical tone in which they voice it, that any given statement of expressivism is true. Hence the expressivist, in order to prevent their expressivism eating its own tail, must refrain from giving the sort of analysis that would either affirm or deny the status of semantic properties as thick properties, while still continuing to speak of them in a thin sense.

This captures Price’s conviction that semantic *properties* have no load-bearing role to play in a philosophical theory of language, while allowing that semantic *concepts* are still usable, at least in folk discourse (and, of course, in an explanatory capacity, by noting their utility value). Hence representationalist (i.e.: reductionist) analyses of semantic properties, and therefore analyses of particular areas of discourse that might proceed by using such a view of those properties, are disallowed.



The presentation of semantic properties as “thick” and “thin”, however, bears several problems, mainly by being misleading regarding the expressivist’s take on the semantic properties themselves. One way of making this clear is to note that there is no such property as “thin” truth. One fairly straightforward reason for this is that there is no possible theory, by Price’s own lights, that countenances a thin version of truth, as opposed to a thick version.

The temptation here is to think of the distinction in the terms of the local expressivist: Most areas of language are blessed with thick truth, allowing, as they do, assertions about real things that can be translated into respectable naturalistic language, but there are some frameworks within language for which this is not the case, and that can therefore only aspire to thin truth. On this view, thin truth is an “as-if” notion of truth; the linguistic frameworks of e.g.: ethics, mathematics or probability-theory allow us to speak, according to their rules, *as if* certain sentences were true, but in fact these sentences fail to have a truth-value, because their terms fail to refer, for instance – they are *merely* expressive, and do not really assert anything.

However, Price has spent much time arguing, as noted above in Ch.2, that there is no coherent version of such a theory. That is, there is no stable theory that attributes a thin or deflationary version of truth to some areas of language and a thicker notion to others; any theory, such as local expressivism, that allocates deflationary truth to some areas of language necessarily collapses into semantic deflationism across the board, deflating truth across all of language. This means that there is no coherent worldview on which truth, or other semantic properties, can be either thin or thick. Further, for Price, there can be no such thing as “as-if” truth. The point of global expressivism as a global semantic minimalism is that the minimum sufficient criteria for an assertoric framework’s containing a “thin” concept of truth, or allowing for “thin” truth-aptness (namely, that it is arranged in such a way as to contain the linguistic structures that underpin a concept of truth as a linguistic tool – see Ch.7) are the minimum sufficient criteria for truth-aptness *tout court*. In other words, once we have in place an explanation of truth in terms of its structural role in language, any reductive theory we might want to bolt on is going to be extraneous, since we have already explained what the language’s being truth-apt consists in.

The root of the problem lies in the relation of the semantic properties themselves to their respective philosophical explanations. “Thickness” and “thinness” are not in fact features of semantic properties, but features of the explanations of those properties that are possible.

Indeed a property's being "thick", as Price uses it here, essentially amounts to its being naturalizable, that is, being restatable in terms of the sorts of properties countenanced by the physical sciences. But this is as much a feature of the relation between the relevant linguistic frameworks, and the concepts belonging to them, as a feature of the property itself. The idea that admitting of such a reduction makes the property somehow more substantial or full-blooded – the connotations of the label "thick" – comes from an inclination towards the object naturalist stance, which Price does not share. Unless we privilege naturalism in a way that is inconsistent with Price's own functional pluralism (see below), there is no reason we should call these concepts, much less the properties they denote, "thick" or "thin" based on their relation to the linguistic framework of the natural sciences. By using these labels in the way that he does, Price invites misunderstanding.

The reason Price uses these evaluative terms is to contrast between semantic concepts as they are used in a folk semantic theory, and the same concepts as they might be used in more regimented inquiry. But here, again, we come up against a distinction Price himself should be wary about endorsing. The traditional picture of folk theory vs academic or philosophical theory is that the latter is apt to correct the former, because, speaking roughly, the folk version is set up to be helpful, while not necessarily reflecting reality, where the more rigorous academic version shows the world as it truly is. Again, this view falls to semantic minimalism: The true propositions yielded by the folk theory are *no less true* than those in the rigorous one. Nor are they less "robustly" true, since the concept of truth applicable to each theoretical framework is the same "thin" or deflated one. In other words, for Price, the thin concept of truth is as thick as a concept of truth can ever get. Thick truth is then not so much thick as bulked out with extraneous theoretical padding.

I think the reason Price adopts the "thick/thin" terminology in the first place is a tactical one. Price invokes the distinction, as it is used here, in response to Frank Jackson, whose representationalism Price sees as arising from tenuous semantic assumptions (Price 2011, pp267-8). These assumptions are based on certain observations about the everyday usefulness of language hinging on its representational properties – for instance, that notes with directions are useful *because* they represent the correct route, that the sports section in the newspaper is useful *because* it represents the outcomes of sporting events, and so on. The expressivist, of course, does not deny this usefulness in virtue of representation. Rather, what Jackson does, and Price deems illegitimate, is to infer from these examples of representational properties in their folk sense that we need a substantive philosophical theory of such properties that showcases their explanatory value. Price's usage, then, is

designed to distinguish the folk sense and the theoretical sense to make it clear that Jackson's examples don't warrant this inference; language is clearly representational in some "thin" sense, but we ought not therefore expect to be able to give a theoretical account of its representational features.

Unfortunately, this comes at the cost of a terminological choice that has surely contributed to the misapprehensions that surround global expressivism, making deflationary truth out to be "as-if" truth by associating it with a colloquial, "thin" version of the concept. This conception is arguably what leads to a perceived conflict between expressivism and realism, where none exists. As Camil Golub (2017) notes, the realist is hard-pressed to articulate a version of realism that explicitly conflicts with expressivism, once semantic minimalism is on the table, since for any fact that the realist cares to state in the target discourse, the minimalist can help themselves to the same. This stands even for explanatory facts; for instance, the realist about ethics who holds that genocide is wrong can, as part of their realism, explain that this is the case because genocide has the property of wrongness, that its having this property constitutes a fact, that this fact explains why the sentence "Genocide is wrong" is true, and so on. While previous expressivist theories would deny these explanatory locutions, the semantic minimalist can accede to all of these points (*ibid.* p1403-4).<sup>21</sup>

The interpretation the semantic minimalist gives of these explanations is that they are, effectively, tautological reformulations of the interlinked definitions of 'Fact' and 'Truth', and likewise for other related concepts (*ibid.*). The non-minimalist will likely disagree, preferring an inflationary account of the content of such sentences. However, this does not constitute a preference for realism, but rather for representationalism. The key, in the terms laid out in Ch.2, is that these explanations are not metasemantic in nature, since they don't set out to cash out the semantic concepts involved into non-semantic terms, instead simply putting them to use. It follows, in line with our definition of expressivism in Ch.2, that the expressivist (as semantic minimalist) has nothing to say about them. Hence the realist and the expressivist have no metasemantic grounds for disagreement, because, while representationalism and semantic minimalism are metasemantic theses, realism about the target discourse is not, being instead a first-order ontological thesis.

---

<sup>21</sup> Note that Golub takes ethical discourse as his example, as his focus is on ethical (i.e.: local) expressivism. I take the point to expand readily to cover any target discourse, and hence to apply to a global expressivism.

Likewise, Golub notes that there is no way of articulating the difference in terms that both sides would accept on grounds of the independence of facts of the target discourse from our attitudes (the expressivist endorses the idea that “Genocide is wrong” can be true even at worlds where our counterpart selves judge that it is false – *ibid.* p 1388), or on grounds of the metaphysical structure of the domain of the target discourse itself (*ibid.* 1395-6).<sup>22</sup> The upshot is that on any understanding of realism that does not conflate it with representationalism, the expressivist as semantic minimalist has no problem with being counted a realist.

What, then, should Price have said on the “thick/thin” distinction? Essentially, the answer is that Price should have followed what he has said elsewhere about representationalism: Language, in some of its uses, is truly representational, and is capable of bearing representational semantic properties such as truth and reference in the most robust sense in which these properties are capable of manifesting. However, such properties do not admit of a reductive explanation, specifically a reduction to (object) naturalist terms (see the following section). This is not a result of the metaphysics of the properties themselves so much as it is a feature of our language; the linguistic frameworks (in the least controversial sense possible) of the physical sciences and of semantics simply do not map cleanly one-to-another. Their structures and concepts were set up independently of each other, in response to different influencing factors, and they remain independent systems. This doesn’t compromise their use for talking about the world, hence it doesn’t preclude taking a realist attitude towards their domains – this is semantic minimalism at work. All that this disconnect between frameworks means is simply that certain forms of philosophical explanation are off the table.

## 1.2 Subject-naturalism and object-naturalism

Price’s distinction between subject-naturalism and object-naturalism is a distinction between two views of the role of science and its relation to philosophy. The two positions

---

<sup>22</sup> The latter perhaps merits some discussion: Since the semantic minimalist holds that the analyticities licensed by the target discourse are genuinely true, they are committed to a metaphysics by acceptance of this discourse, much like the Carnapian who accepts pleonastic inferences – see Ch.4. As regards semantic discourse, this commits them to a metaphysics of semantic properties (well-formed sentences about truth-aptness are themselves truth-apt, for instance), although not an inflationary one. The potential difficulty for interpreting what this means, and which semantic claims the expressivist is and isn’t allowed to endorse, was the reason for the attempts at clarification of their position in Ch.2.

do not appear to be exclusive at first glance. Indeed, according to Price, most philosophers who hold themselves to be naturalists will hold some version of both. Nevertheless, Price argues that, first, subject-naturalism is the more fundamental, and second, subject-naturalism actually undermines object-naturalism, or at least challenges the grounds for holding it.

It is perhaps easiest to characterise object-naturalism first, then juxtapose subject-naturalism. If naturalism generally is the idea that where philosophy conflicts with science, science wins out, then object-naturalism is a particular version of how this conflict is to be assessed. Object-naturalism establishes the metaphysical priority of the scientific worldview. For the object naturalist, the sorts of things studied by science, that make it into our scientific ontology, are the only sorts of things that really exist (Price 2011, p185). Object-naturalism is thus committed to a physicalist ontology, because it takes it that the ontology of a unified scientific theory would be put in the terms of the physical sciences. In the most straightforward cases, this means that philosophical theories positing entities that do not fit within the natural order, such as non-naturalism, are, by object naturalist lights, mistaken.

In the other direction, this means that the object naturalist is under pressure to show how certain domains of entities fit within this physicalist framework, since anything that doesn't is deemed unreal. This pressure is particularly felt in apparently non-natural domains on which scientific theorizing nevertheless relies, such as mathematics, modality and semantics. This is the origin of placement problems. These have a semantic root: We use language, such as ethical or semantic language, that appears to be about a particular type of non-natural entity, and have to figure out how to show that this language is *actually* about some region of the natural world, since this is all there really is (*ibid.* p188).

The various anti-realist and non-cognitivist theories about these different regions are common ways out of placement problems that can't be resolved – if we can't make, say, moral talk reduce to talk about the natural world, then we have to take it to be committing some sort of systematic metasemantic mistake. This anti-realist line is less easily taken when the region of language in question is one relied upon for scientific theorizing. For instance, it complicates the object naturalist's view of science as delivering the metaphysically accurate picture of the world if number terms do not refer to anything, since these terms are so entrenched within so many scientific theories.

The most important placement problem for the object naturalist to solve, however, is the one concerning semantic properties themselves. This is because, as noted, the object naturalist viewpoint finds its footing in the metasemantic thesis that only scientific language actually portrays the world as it really is, being the only sort of language that refers to and quantifies directly over real entities, and so on. On this sort of view, even non-scientific terms, such as 'Table', that do refer will do so in virtue of having a usage that is reducible to the language of natural science, the determiner for linguistic success. This metasemantic stance, that linguistic items can only really bear semantic properties if they are part of the language of science or in virtue of reducing to the language of science, is what motivates the object naturalist's criterion that any area of language that is not reducible to scientific language should be rejected as merely a colloquial tool that does not represent the world as it really is. As such, the object naturalist position becomes unstable if we take it that semantic terms themselves do not refer to real properties, as the distinction tracked by this naturalization criterion itself, between those areas of language that are capable of bearing semantic properties and those that are not, ceases to be naturalistically respectable.

Price positions subject-naturalism, by contrast, as the view that human beings are natural creatures, and that "Philosophy needs to begin with what science tells us *about ourselves*." (*ibid.* p186, emphasis in original) Price does not expand directly on this rather opaque definition, but the view can perhaps be better understood by the position it occupies in his argument against the object naturalist outlook. According to Price, subject-naturalism is the more fundamental of the two positions. As such, it actually underpins the object naturalist viewpoint. This is because the object naturalist's metasemantic thesis – that only the naturalistically respectable areas of language possess genuine semantic properties – relies on a substantial reading of those same semantic properties to be viable, but this further thesis of the substantial nature of semantic properties comes under the domain of subject-naturalism. As Price puts it:

Without [a substantial reading of semantic properties], there can be no subsequent issue about the natural "place" of entities such as meanings, causes, values, and the like. Object-naturalism thus rests on substantial theoretical assumptions about what we humans do with language— roughly, the assumption that substantial "word–world" semantic relations are a part of the best scientific account of our use of the relevant terms.

However, these assumptions lie in the domain of subject-naturalism. Moreover, as the conceptual possibility of deflationism already illustrates, they are non-compulsory [...]

Hence my Priority Thesis: given a linguistic conception of the origin of placement problems, subject-naturalism is theoretically prior to object-naturalism, and object-naturalism depends on validation from a subject naturalist perspective.

*Ibid.* pp189-90

It might seem at this point that subject-naturalism is being defined somewhat arbitrarily – there does not seem to be much reason, other than helping the present argument, to allege a category distinction between these two versions of naturalism. Nevertheless, Price’s distinction does bring out a threat of circularity in the object naturalist setup as I will attempt to make clear.

The object naturalist’s core principle states that only languages that are naturalizable really represent reality. However, as noted, this relies for its coherence on the assumption that semantic properties are substantial in nature, meaning here, for the object naturalist at least, that the language of semantic properties itself has to be naturalizable. Since semantic deflationism is a live option, we need some justification for this substantial take on semantic properties. The problem is that we can’t offer this as an argument in object naturalist terms, because that would assume the object-naturalism in support of which we were hoping to use the substantial reading. For instance, we could argue that semantic properties are real because they are required by our best scientific theory, and that they must therefore be reducible to physical terms, because all real properties are. This last step, though, would clearly rely on the object naturalist’s naturalization criterion, and would thus be circular as a defence of object-naturalism.

Shortening the circular route of the argument this much renders it unconvincing, but the point here is to demonstrate the difficulty with supplying an argument for the substantial nature of semantic properties from an object naturalist perspective, since that perspective assumes they have such a substantial nature. The task left to us, then, is to try to analyse the nature of semantic properties from a perspective that bears no such preconceptions. For this, we have to retreat to mere subject-naturalism. In practice, this means treating language and its attendant semantic properties not in terms of reductive metaphysical commitments, but in terms of their role in our lives as subjects, specifically by analysing linguistic practices, both in terms of their structures and the purposes to which we put them. What we find when we adopt this starting point is that we have a similar difficulty in motivating object-naturalism, because these subject naturalist explanations look adequate to explain the features of language that interest us without any input from object-

naturalism, and without recourse to its metaphysical convictions. What we end up with, then, is a deflationary theory of semantic properties, to which any further metaphysical explanation appears superfluous.

### 1.3 I-representation and e-representation

The distinction between i-representation and e-representation is more fine-grained than the subject-naturalism/object-naturalism distinction, in that individual semantic explanations can contain both an i- and e-representational component. Briefly put, e-representation is the familiar concept of representation in virtue of covariance with features of the environment; a fuel gauge in a car is an e-representational device, for instance, in the sense that its only function is to vary in parallel with the amount of fuel in the tank. I-representation, on the other hand, is representation in virtue of the internal role something has within a practice – usually the inferential role of a linguistic expression within a linguistic framework (Price 2013, pp36-8). Note that Price himself does not give an example of a clearly i-representational linguistic expression, although Legg and Giladi (2018, p69) propose the expression ‘2’ as a purely i-representational device, as it seems that, insofar as it represents reality, it does so purely in virtue of its being situated within a network of inferential relations.<sup>23</sup>

As with subject- and object-naturalism, Price’s goal in distinguishing these two types of representation is to draw attention to the way in which they are usually blended together. The specific target here, once again, is representationalism, which can be rendered as the thesis that all representation must eventually reduce to e-representation. Here we can see the echoes of the above distinction; e-representation is (supposed to be) naturalistically explicable, which, according to object-naturalism, speaks to its being the real version of representation. On this view, anything that *only* i-represented would at best be quasi-representational, and anything that was both i- and e-representational would only really represent by virtue of the latter.

---

<sup>23</sup> Part of the problem with finding examples is that it seems that many, or perhaps all i-representational expressions or devices must also be e-representational. Note, however, that if this were true, it would not invalidate globalism, as an expression’s being e-representational is not the same as there being a reductive naturalistic explanation of its semantic properties in terms of word-world relations.



In contrast to this, Price claims that “these are not two competing accounts of a single species of representation but two quite different beasts.” (Price 2013, p38) On Price’s view, each type of representation forms a different part of a semantic explanation. As such, a complete semantic explanation might feature both types of representation, or only one, and still be complete. Most pertinently, this means we can offer a complete account of certain linguistic items in terms of their inferential role within a linguistic practice, without having to point to a direct representational function, in terms of covariance with environmental features (*ibid.* p40).<sup>24</sup>

Price uses the notions of i- and e-representation to underpin his standpoint of functional pluralism. Price notes that our language is home to many functionally distinct assertoric language games. Since i-representation amounts to representation underpinned by the function of an expression within a given practice, and these language games represent distinct linguistic structures within which different functions are called for, this means that we have many different ways of representing the world, or many different forms of i-representation, none of which need necessarily correlate with one another, since their respective language games are structured differently (*ibid.* pp53-4). Price calls this totality of assertoric language games, and the overlapping picture they paint via i-representation, the “i-world” (*ibid.* p55).<sup>25</sup>

In contrast to the i-world stands the e-world. Essentially, the e-world is the picture of the world from the naturalist perspective – the version of the world in terms of which we give the e-representational part of a semantic explanation. The claim that completes Price’s functional pluralism is that the e-world is in fact part of the i-world, namely the part given in terms of the language game(s) of the natural sciences (*ibid.*). That is to say the language of the natural sciences is i-representational, at least in part, and it is in this language that we conceive of e-representation. Hence the difference in kind between the two – the sense in which they are “different beasts” – is that i-representation depicts an expression’s significance to the language game of which it is a part, while e-representation analyses the expression by representing it in terms of a *different* language game (that of the scientific investigation).

---

<sup>24</sup> This point is important for the expressivist’s rescue of the neo-Carnapian deflationary approach; I will discuss it in detail in Ch.8, §2.1.

<sup>25</sup> Price’s “world” terminology is apt to mislead here. As a disclaimer, note that no metaphysical distinction between worlds is intended – see the discussion of discourse pluralism below.

What this means for Price is that the idea of the priority, ontological or methodological, of the object naturalist worldview is debunked at the semantic level; the language of physical science does represent the world, and does so in a different way from other linguistic frameworks, but this difference is just the difference between different ways of i-representing, such as might also be observed between, say, the difference between the ways that the framework of ethics and the framework of mathematics represent the world. If we want to show that the scientific image shows us what the world is like fundamentally, we will have to make the argument on some other grounds than the manner in which scientific language gains its representational character, because this ends up being the same process (of setting up the inferential structures required for i-representation) as any other form of language.

## 2: Price's functional pluralism

Price characterises his overall view as a functional pluralism. This name is meant to capture the idea that there are many different areas of language, each of which has a distinct function, none of which are any less concerned with the real world than any others, in the sense in which the metaphysician is interested. In this section, I will characterise Price's functional pluralism, and attempt to bring out its relation to both Carnap's thought and the neo-Carnapian approach we have seen so far.

### 2.1: Carnap and Quine

Price derives his functional pluralism from the "Carnap thesis". This is the idea set up in Carnap's *Empiricism, Semantics and Ontology* of the distinction between questions asked internally to a linguistic framework, which are supported by the rules and procedures of the framework, and those questions that are asked externally to the framework, which lack this support. In a rough slogan, the Carnap thesis states that there is no theory-independent stance from which to ask metaphysical questions (Price 1997, p250). What this implies is that there is no sense to the question, for any given linguistic framework, whether the entities of that framework really exist, independently of how we would state things "according to the framework".

Neither, therefore, is there any answer to metasemantic questions such as whether the referring terms of the framework really refer, and so on. From this follows the falsity of the

“bifurcation thesis”, which states that we can distinguish, on metasemantic grounds, between, on the one hand, areas of discourse whose sentences are really truth-apt, whose terms really refer and so on, and on the other, areas of discourse that, despite appearances, are non-cognitive, whose sentences are *merely* expressive, etc. – what Price calls “non-factual” discourses (Price 1992, pp393-4). The crucial phrase here is “despite appearances”; what might seem to the thoroughgoing externalist, with their own idea of what constitutes a genuine metasemantic underpinning for an area of discourse, to be mere surface features of language – facts such as assertoric structure, inferential norms and practices, and so on – are, for the Carnapian, the only facts one can legitimately use to ground one’s judgements about the framework’s factuality. If one were to go in for the sort of view that demands a substantive metasemantic account of the meaning of linguistic expressions, then one could make sense of the idea of a framework that contained all of the structures and practices required to be genuinely factual, but nevertheless failed to “hook onto” the world in the correct way, or stand in the correct sorts of substantive relations with things in the world, and therefore fails the final test of factuality. However, since the Carnapian rejects substantive metaphysical questions as lacking meaning, they must also reject this demand to back our semantics with metaphysics, on which the bifurcation thesis rests, as misconceived.

Price laments that Quine, as the most famous voice to speak against Carnapian deflationism, has come to be seen as the saviour of substantive metaphysics (Price, 2009, p321). Indeed, Price holds that a proper reading of Quine understands him as a deflationist ally, and the neo-Quinean approach to inflationary metaphysics, therefore, to be misconceived. Briefly, the neo-Quinean school takes their cues from Quine’s principle of indispensability, which governs ontological commitment, specifically as stating that we ought to take as real all and only those entities that are indispensable for our current best scientific theory (e.g.: Van Inwagen 2009, §4; Sider 2011, p14). This is taken to justify metaphysics as an attempt to understand the ontological commitments of our current best science, and metaphysical debate as the comparison of attempts to model its commitments.

Price argues that this reimagining of the task of metaphysics does not actually sit at odds with Carnap’s deflationism. Indeed, it pulls in the same direction with more force, reconstruing not only metaphysics, but also natural science along pragmatic lines. As Price puts it:

A metaphysician who takes this as a vindication of his position [...] is someone who has not been told the terrible news. Quine himself has sunk the metaphysicians' traditional boat, and left us all, scientists and ontologists, clinging to Neurath's Raft.

(Price 2009. pp326-7)

We might put the point this way: By Price's lights, substantive metaphysics hinges on representationalist assumptions, and representationalism in turn hinges on a substantive criterion for reality, as per §1.2 above. But Quine's own criterion for reality – that real entities be quantified over by our best natural scientific theories – is not substantive but rather pragmatic, namely because Quine determines the best scientific theory on pragmatic grounds (Quine 1951b, p72).

Mary Leng puts the point a slightly different way: Quine's criterion for ontological commitment is not simply scientism, i.e.: adherence to natural science as against other discourses because of taking science to be different in kind, and its methods superior. Instead, Quine emphasizes the continuity between current scientific methods and previous frameworks for theorizing about that world (Leng 2016, pp250-3). This reading of Quine places the emphasis not on his privileging of the scientific worldview *per se*, but on his noting that, whenever we do theorizing of either a scientific or philosophical kind, we necessarily do so from within our inherited worldview. What makes the body of current scientific theory special, then, is that it is the most recent point in a process of inquiry that encompasses scientific and non-scientific methods alike (*cf.* Quine 1957, pp5-6). Indeed, this opens the door to acknowledging other discourses that have survived, in relatively stable condition, this process (Leng focuses on mathematical discourse about sets, and normative discourse):

In particular, it is hard to see how we could exempt from this any stable and considered aspects of our inherited worldview that remain through our processes of refinement. And if this is the case, then our most stable mathematical and normative beliefs do have 'something rather special' about them that allows us to see them as authoritative [...] What is special is their history, and the fact that they have survived in our worldview through our collective best efforts at trying to determine what we ought to believe and how we ought to behave.

Leng 2016, p252

Accordingly, there is a potential for other stable areas of discourse to be legitimised under a Quinean framework, and to sit alongside natural science, undercutting the neo-Quinean

insistence on a relatively austere, monistic metaphysics. Leng thus provides an example of how to draw Quinean methodology closer to the sort of functional pluralism espoused by Price, based on Quine's rejection of a substantive criterion for metaphysical reality in favour of a pragmatic one, and of the potential for Quine's methodology to make room for other stable discourses than that of natural science.

## 2.2: Dimensions of pluralism

The consideration that we cannot evaluate linguistic frameworks using the traditional externalist metasemantic criteria leads to a picture of a plurality of such frameworks, each with an equal claim to viability. The point here is two-fold: First, no single, independent criterion is available for judging the adequacy of different frameworks to representing the world, and thereby sorting them into representational and non-representational camps. And second, because there is no such single metasemantic account of representation available, we cannot offer a reductionist view<sup>26</sup> that consolidates the vocabulary of all the truly representational frameworks under one vocabulary – one set of entities that they are all really “about” (traditionally, of course, this would be the vocabulary of natural science).

Given the fact of placement problems, this creates an interesting picture of disunity. On this view, there are many linguistic frameworks that, provided they are individually coherent, consistent, practical and so on, all seem to describe the world perfectly aptly, but in different ways. While these different modes of description are not incommensurable per se – they do not necessarily provide irreconcilable interpretations in the sense of making inconsistent claims – they do look to resist coordination with one another, in the sense of the formation of determinate rules for relating the expressions of one to the expressions of another, i.e.: placement problems.

This is Price's standpoint of functional pluralism. To help characterise the view, Price invokes yet another dichotomy, this time the distinction between “vertical” pluralism, the camp to which functional pluralism belongs, and “horizontal” pluralism. Horizontal pluralism is the idea that there are multiple different languages or frameworks that perform the same function, where none of these has any claim to being “truer” than any

---

<sup>26</sup> Using the broad understanding of reductionism from Ch.2.

other. Price uses Quine's idea of ontological relativity as an example of such a view (we might also point to Hirsch as another clear example). Quine proposes:

...the possible existence of a range of alternative scientific worldviews, each empirically adequate to more or less the same degree, and none, even in principle, having privileged claim to provide a "truer" description of the world.

Price 1992, p389

This is a horizontal pluralism because the different scientific worldviews on offer are all aimed at performing the same function – they all occupy the same linguistic "level". By contrast, a vertical pluralism is committed to the idea of a multiplicity of functions across different discourses, which, by and large, cannot be collapsed into one another. This view follows fairly readily from the Carnap thesis, as noted above.

It is possible to construct the beginnings of an explanation of the origins of placement problems from what we have so far. Price discusses this origin in one sense, by pointing to the original mistake that gives rise to placement problems. According to Price, this is tied to the notion of semantic descent: We can move easily from talking about a referring term to talking about the object it represents, simply by disquotation (e.g.: moving from talking about the name 'Carnap' or the sentence "Carnap was a philosopher" to talking about the man, Carnap). The issue arises with the assumption that this link between word and object has a representational character. On such a conception, devices of semantic descent, such as the disquotation schema, take on the character of tracking procedures within language, and semantic properties take on the role of unerring tracking relations. These demand to be explained, leading to the substantial conception of semantic properties, and the placement problems we have seen above.<sup>27</sup>

While this points us to the origin of the initial mistake that generates placement problems – the demand for a substantive account of semantic properties – it does not explain the origin of placement problems in the sense of explaining *why* we cannot, for instance, give a naturalistic account of the meaning of moral terms. Given that both relevant linguistic frameworks are in some sense concerned with the same world, why should we not be able

---

<sup>27</sup> I take this to be the point made in Price 2011, pp187-190, although the passage is open to interpretation.

to construct some set of rules to coordinate their respective terms and inferences (and thereby the objects with which they are concerned)?

Price's "world" terminology is not especially helpful here. As noted above, the different "i-worlds" (of which the "e-world" appears to be one) each comprise the totality of facts according to one specific linguistic framework (Price 2013, p54). This definition already assumes the incommensurability of such frameworks, in the sense that it individuates totalities of facts by the different sets of linguistic rules. Whatever meaning we have in mind for the word "world" in asserting that the different areas of language are each concerned with the same world, it doesn't seem to be the same one that Price has in mind for his "i-world/e-world" terminology.

Essentially, what seems to be going on is that the different linguistic frameworks are, in some sense, constructed to be *too different* to coordinate in the way we might want for, say, a reductive naturalism. The problem for the expressivist here is that they are restricted by their theoretical commitments from providing the sort of substantive account of semantic relations in terms of which we might frame this incommensurability – for the expressivist, there is no substantive account available of how different frameworks latch onto the same world in different ways, simply because there is no substantive account available of how they latch onto the world *tout court*. To venture a metaphor, the non-deflationist conception of the different areas of language has them akin to different gears; if we can but find the correct ratio, we can translate work done by one into work done by another. By contrast, for the functional pluralist, the different areas of language are entirely different machines, with different purposes and structures. The reason we can't parse between them is because they do different work, and because they were designed and structured in response to different concerns.

### 2.3: Non-Naturalism

At this point, we might pose the question whether the view put forward by Price comprises a sort of non-naturalism. The functional pluralism outlined above allows for the existence (in the most full-blooded sense) of apparently non-natural entities such as numbers, meanings and so on, as well as allowing for this existence, per the rejection of object-

naturalism<sup>28</sup>, without these entities' being reducible to a physicalist ontology (Price 2011, p136). Nevertheless, Price denies that this commitment makes him a non-naturalist. To see why, we have to take a Carnapian stance; from the standpoint of the Carnap of ESO, while there do exist non-natural objects, the theses of non-naturalism and, contrasting with it, object-naturalism are both external claims, and therefore meaningless. As Price puts it:

The point is that the judgment of unity or plurality [of the natural and non-natural world(s)] could only be made from the framework-independent external stance, which the Carnap thesis disallows. Without that stance, functional pluralism is neither monist nor pluralist, in a primarily ontological sense – for there is no such sense.

*Ibid.*

This does raise some further questions. For instance, while it may seem plausible to interpret the thesis of non-naturalism as a metaphysical claim, and therefore a pseudo-claim by Carnap's lights, it is not clear that this is necessary. The main reason to think that this interpretation is apt is to think of the non-naturalist thesis as a denial of object-naturalism, which is clearly a metaphysical thesis; if the denial of a metaphysical claim is itself a metaphysical claim, then Price would appear to be right that the Carnapian cannot make sense of either object-naturalism or non-naturalism. However, Price specifically argues above, in pushing for "strategic silence", that rejecting a metaphysical claim does *not* necessarily mean stating the opposite claim in the same metaphysical register. For instance, the Darwinist, in denying the creationist's theory that God created the different animal species, does not have to posit that "God did not create the different animal species", thus foreclosing the possibility of atheism if their posit is to avoid suffering reference failure.

Although I argued above that Price's line of passive rejection was a strategic liability for the expressivist, it's difficult to avoid its application here. Read in this minimal fashion, of simply denying the object naturalist's reductive thesis, non-naturalism does seem to be in line with Price's functional pluralism, at least on the face of it; if we simply take the non-naturalist to be committed to there being, for instance, moral facts that do not reduce to natural facts, then Price can hardly object. The curious issue is how to square this with the

---

<sup>28</sup> I take it that the position that the non-naturalist is characteristically opposed to is *object-naturalism*, as the assertion that the objects of study of the physical sciences are ontologically fundamental – just "naturalism" is too broad, in light of the foregoing discussion.



idea of the non-naturalist as making a metaphysical claim for the distinctness of, say, the world of moral facts from the world of natural facts.

The issue here seems to be the same issue that the Carnapian has of how to interpret metaphysical questions. In such cases, as we saw in Ch.1, the Carnapian has three options: 1) Take the speaker to be asking a (likely trivial) internal question, 2) More charitably, take the speaker to be engaged in the speech act of *proposing* a new linguistic framework using the format of a question, or 3) Take the speaker to be attempting to adopt an external perspective from which to ask a genuinely metaphysical question. Due to the nature of the case, the speaker's intentions may not even be determinate as regards any of these three options.

Analogously, I propose that we could interpret the non-naturalist either as making a trivial internal claim (e.g.: the claim from within the framework of ethics that ethical facts, values and principles exist); as proposing that we treat linguistic frameworks for non-natural entities as independent of the framework of physical entities; or as making, by Carnapian lights, a metaphysical pseudo-proposal about such entities, e.g.: that they exist in different worlds, or have different essential characters, where these concepts are taken by the non-naturalist to be metaphysically loaded. Given that the non-naturalist's own intentions in asserting their position may not be determinate between these options, there seems little hope of settling the question whether Price is *really* a non-naturalist. In any case, one suspects that the line Price takes on non-naturalism is part of his attempt to portray his views as intuitively naturalistically respectable, without relying too heavily on his having shifted what this means by having reversed the priority of object-naturalism and subject-naturalism. All that we can really say here is that there is at least one (deflationary) way of understanding non-naturalism on which Price does fall into that camp. I would not be opposed, therefore, to characterising Price as a non-naturalist, provided we are careful to do so from a deflationist standpoint (and disregarding the fact that this makes him at once a naturalist and a non-naturalist, given that those labels have had their meanings substantially altered by this point).

#### 2.4: Existence deflationism

We have seen that Price derives a metaphysical deflationism from a Carnapian starting point – in this sense, he is a neo-Carnapian. However, Price's focus is very much on the

metasemantic side, on the global expressivism discussed. The point to highlight here is that Price's metaphysical deflationism flows from his metasemantic deflationism; the latter, in the guise of anti-representationalism, is responsible for establishing the subject-naturalism that embodies Price's metaphysical deflationism. This would imply that from global expressivism, metaphysical deflationism must follow – a sequence that seems difficult to dispute, if we take Price at his word. In contrast with this, the concern of this thesis is to argue that from a neo-Carnapian deflationism, a global expressivism must follow – that a neo-Carnapian deflationism *requires* global expressivism in order to function. If both theses were accepted, this would present an interesting insight: That these theories present two sides of the same picture, and that they stand and fall together.

This biconditional claim is too significant for the present piece of work to do it justice. Likewise, it will only be possible to flesh out the initial stages of the combined view (Ch.8); further illustration will have to be the topic of future work in this area. However, we can here note where Price's deflationary views dovetail with other neo-Carnapian theories, in particular Thomasson's existence deflationism. Price offers a functional pluralism on which different linguistic frameworks function in ways that are sometimes radically different, despite appearing similar in form. A non-deflationary take on existence would be in tension with such a theory, because it would involve tying the existence of the objects of all linguistic frameworks to the properties or relations of one such framework (e.g.: by stating that existence hinges on physical causal powers), when a functional pluralism has it that such attempts at mapping are what creates placement problems.

Indeed, Thomasson argues that deflationism about semantic properties, such as is involved in Price's rejection of representationalism, actually entails existence deflationism – accepting semantic deflationism is incompatible with rejecting existence deflationism. In Ch.4 we noted the availability of so-called “pleonastic inferences” – inferences to existential claims from premises that are not explicitly existential, such as from “John has two apples” to “The number of apples John has is two”, to “There is (exists) a number.” Thomasson notes that we can perform the same trick using truth claims. For example, from the claim that “n is P” is true, we can infer, by disquotation/semantic descent, that n is P, from which it is a short step of the sort already demonstrated to the claim that there exist Ps (Thomasson 2014, pp195-6).

As Thomasson has it, given the triviality of the moves involved, there is no way to deny the derived existence claim in this sort of case except by denying the original truth claim. Given

obvious truth claims as a starting point (“The sentence “Snow is white” is true”, “The sentence “The number 5 is odd” is true”), one could only be inclined to reject the initial claim on metaphysical grounds, such as suspicion about the nature of its truthmakers (*ibid.*). Such reasons are, however, themselves rejected by the semantic deflationist. It follows that a non-deflationary metaontology, which might give us analogous metaphysical reasons to deny the derived existence claim, is incompatible with such a semantic deflationism, because from accepting semantic deflationism it follows that we must accept the derived existence claims.<sup>29</sup>

What makes Thomasson’s theory complement that of Price is the shift to a structural analysis of language. Thomasson construes the concept of existence in terms of linguistic structure, namely the linguistic structure common to all well-formed existence claims, a move that parallels Price’s own expressivist view of truth. This sort of analysis of existence is compatible with a functional pluralism because it is able to acknowledge the different sorts of procedures present in different linguistic frameworks, as well as the different functions they serve, while espousing a univocal concept of existence as a structure shared between these frameworks. This structure is abstract, essentially amounting to assessing the application conditions of sortal terms for satisfaction. The structure’s abstractness allows for its compatibility with pluralism; we can see the same structure at work whether the sortals we are interested in are part of the physical, mathematical or ethical frameworks. This avoids a picture of existence as a framework-independent property, or a property grounded in one framework that gets “picked up on” by many.

## Conclusion

As we have seen previously, a significant problem facing global expressivism is that it appears obscure to the outsider. Much like the metaphysical deflationist, the globalist seems to want to have their cake and eat it, maintaining that semantic properties are

---

<sup>29</sup> Marschall and Schindler make the point that there are reasons other than a substantive metaphysical account of truth that one might want to deny the purportedly obvious initial sentences, such as rejecting the existence of composite objects to avoid the threat of causal overdetermination – essentially, eliminativism about the relevant domain of objects (Marschall & Schindler 2021, p109). However, this and the other examples the authors offer (*ibid.* fn 8) seem to me to beg the question against the metaontological deflationism being argued for. Interestingly, the authors themselves note the converse: The only watertight argument for *accepting* the initial truth claim as obvious is accepting easy ontology, and making the argument itself circular by begging the question against substantive metaphysics (*ibid.* p108). It is an interesting question how one can choose a principled starting point with such a threat of circularity apparently present on all sides.

insubstantial, yet truly real, that representationalism is false, yet language does represent the world, and so on. I hope in this chapter to have continued clearing up some of the obscurity of the global expressivist's methodology by outlining the concerns motivating Huw Price, a key author in the field, as well as by attempting to clear up potential confusions and, where appropriate, critiquing misleading terminological and theoretical choices.

Not all of Price's choices are necessary ones for the neo-Carnapian deflationist. For example, the neo-Carnapian needn't be a pluralist on the same order as Price. Indeed, with the exception of Price, a vertical pluralism is not a position really occupied in the neo-Carnapian literature. Carnap himself appears to have been both a vertical *and* a horizontal pluralist, but the only other pluralism on offer is Hirsch's quantifier variantism, which should properly be characterised as horizontal pluralism (this trend away from pluralism in general is presumably a result of the backlash against quantifier variantism, as seen in Ch.1). There appears to be nothing in a Carnapian metaontology that specifically motivates us to a horizontal pluralism, although the view might start to seem attractive given the concern of placement problems, and given the responses, both in favour of Carnap and in Thomasson's existential deflationism, to problems with quantifier variantism.

The purpose of this chapter has not been to make recommendations for the sort of global expressivism the Carnapian *must* adopt, but to flesh out a possible form it might take, and to try to emphasize the points of common cause between the two approaches. This will give us a framework, going forwards, for understanding the various sorts of semantic explanations the expressivist will give, with detailed examples given in the following chapter, as well as the sorts of concerns we'll need to address when combining the two methodologies in Ch.8.

## Ch.7: Expressivist explanatory strategies

In this chapter, I will offer examples of positive expressivist explanations for various categories of expression. As outlined previously, the criterion for success for these explanations is that they do not use substantive representational relations to bear explanatory weight – more precisely, as per Ch.2, that they do not cite a substantive relation between the linguistic expressions being explained and the subject matter of those explanations in an explanatory capacity. In outlining these explanations, I hope to give an impression of the versatility of the expressivist approach.

The purpose of this chapter is illustration – although it has been discussed at length in previous chapters, the expressivist approach to semantic explanation may yet seem obscure without many examples to draw on. Although the examples that follow are meant to give a feel for the expressivist approach, they will not be entirely comprehensive; I will save an analysis of the central concepts of Hirsch’s and Thomasson’s programmes – intensionality and existence respectively – for the concluding chapter, where I will deal with them more thoroughly, in the light of the modifications to their respective theories put forward in previous chapters.

Some of the examples below will involve an explanatory tool created by Michael Williams: The “EMU”, or “Explanation of Meaning in terms of Use” (Williams 2010, 2013). Williams’ EMU is designed to offer a specific structure for non-representationalist explanations of the meaning of linguistic concepts. It does this by separating out the different tasks that such an explanation has to undertake. Williams’ EMU gives a metasemantic explanation in three parts. The first, the inferential portion, lists simply the language game moves that can be made using the relevant expressions; the second, the epistemological part details the warrants pertaining to those moves – when one is justified in making them; the third, functional part is perhaps the most recognizable, in that it offers, essentially, the pragmatist take on the concept being explained – what role or utility value it has within our linguistic practices, in virtue of which it is retained. I will elaborate on the EMU further below, as Williams himself offers a good illustration with our first example, the expressivist account of the concept of truth.

## 1: Truth

The minimalist account of truth as an abbreviating or generalizing device is a relatively popular one, and descends from Peter Strawson's theory of truth as a device for noting agreement (Strawson 1948). Truth's role as an abbreviating device is evidenced in statements such as "Everything Jim tells you about biology is true." In such instances, the concept of truth allows the speaker to add their endorsement to particular claims without having to utter the claims themselves. Not only this, the concept allows the speaker in the example to make use of expertise that is not their own, in that they can commit themselves to, if not directly assert, a number of different claims about biology without having to go to the trouble of sorting true claims they want to endorse from false ones they don't – this task effectively being delegated to Jim, who has the expertise to perform it.

Williams himself provides an EMU for truth that captures this function, quoted in full below (with minor formatting changes to clarify section titles):

- Inferential: Excepting sentences that generate paradox, the inference from "Snow is white" to "It is true that snow is white", and vice versa, is always good; the inference from "Grass is green" to "It is true that grass is green", and vice versa, is always good, and so on.
- Epistemological: Such inferences are *primitively* acceptable, (*a priori*). They are "free" moves in the discursive game.
- Functional: The truth predicate is important exclusively as a generalizing device. It enables us to do things we could not otherwise do: Endorse or repudiate claims that we cannot otherwise explicitly state because we do not know what they are ("You can trust John: Anything he tells you will be true") or because there are too many ("Every proposition of the form "p or not-p" will be true").

Williams 2013, pp134-5<sup>30</sup>

What makes the EMU expressivist, by our lights, is that it doesn't attempt to state what truth *consists in* – only the inferential moves that it allows, and the function for which the

---

<sup>30</sup> Note that Williams later (2015) introduces a different, four-part version of the EMU; the only distinction seeming to be the splitting of the "Functional" section into separate sections specific to the word under discussion (here 'True'), and the speech-act function attached to the concept. As nothing major seems to hang on this distinction for present purposes, to keep the version from the earlier discussion of Williams' 2013 chapter.

concept is used. We can put this another way by noting that the account contains no material regarding the assessment of sentences for truth, such as any mention of truth-conditions, instead noting only that the inferential/epistemological requirements for asserting a sentence and taking it to be true are identical. Thus the EMU assumes the competence of language-users to assess sentences for warranted assertability, and to perform the occasional disquotations, but nothing else; no general criterion for a sentence's being true is supplied, only the relevant language game moves.

The EMU's completeness hinges on what we can derive from it. For instance, Williams leaves out an account of exactly how the inferential role of a concept of truth enables its function as a generalizing device; nevertheless, we can show how such a function arises, given only the materials that Williams presents here, thus showing how the EMU hangs together. The functional role of a concept of truth as given above exists thanks to the introduction of a truth predicate that can be applied to clauses that can function as assertoric sentences in their own right. As per the inferential and epistemological roles of truth, if we have sufficient grounds for applying this predicate to such a clause, then we also have sufficient grounds for making the assertion, and vice versa. It is this function that allows us to generalize, since we can make quantificational statements applying the truth predicate. This is worth spelling out, because it demonstrates how the two sides of the explanation are interlinked, namely by showing how the functional side is made possible by the language game moves introduced with the truth predicate.

One thing that does *not* appear to be made possible by the language game moves detailed above is the function attributed to a concept of truth by Price. Although the generalizing function of truth is undoubtedly important for its linguistic role, the more important role Price sees for the concept of truth is as underpinning norms of disagreement. This is the idea that one ought to take exception to other language users' expressing factual commitments that are inconsistent with our own. We might regard this as explaining the "unitary" conception of truth – the idea that there is only one way the world can be, and that there is therefore no such thing as a "no-fault" disagreement. Price's idea here is not just to restate the law of non-contradiction, but to bring in linguistic behaviours surrounding, effectively, its enforcement for the set of beliefs of a linguistic community, as opposed to an individual, and incorporate these behaviours into the explanatory account of truth.

Once we have such a picture, it is not hard to see the utility value of truth in communication. If you believe there is a cliff to the west, and I am unconvinced, a norm telling us that only one of us can be right, and that we ought to resolve which one of us it is could end up saving my life. While the expressivist is happy, much of the time, to use biological evolution as an analogy for the way certain linguistic structures or frameworks out-compete others with less utility, this particular structure of truth as a norm of non-contradiction within a linguistic community could plausibly be a genuine result of natural selection, since it appears so basic to the structure of assertoric discourse, and thereby grants such obvious survival value for language users who adopt it.

The relation of truth to assertion in this explanation is important for Price, and helps to illustrate the importance of the expressivist form of explanation to the current project. Other forms of explanation take truth to be conceptually prior to assertion. For instance, an account might explain that truth is the goal of assertion, at which speakers aim when making particular claims. Further, truth-aptness is used to distinguish genuine assertions from sentences that fail to assert anything, despite having an assertoric form. This latter is especially important, because it is relied upon by both representationalists and those whom Price calls “non-factualists”, namely theorists who want to dismiss a particular region of language, such as ethical or mathematical language, as not genuinely assertoric but *merely* expressive (Price 1988, p21). As argued previously, the local expressivist necessarily falls into this latter camp, along with the early positivists. By contrast, Price’s expressivism has it that assertion is conceptually prior to truth, since it is in terms of assertoric practices that we can make sense of a concept of truth. This gives us a clue as to how the account avoids the all-important placement problems associated with truth, namely by explaining the concept of truth in terms of the structures of linguistic behaviour that it enables, rather than forming an idea of truth as a property, and then having to explain both how our linguistic practices manage to track it, and how it can be conceptually independent of those practices.

On the face of it, while it does appear to account for the function of generalization, Williams’ EMU does not appear directly to enable the above function of “friction”, or establish norms of disagreement. This is in line with Williams’ focus on the function of the truth predicate specifically, rather than on the structural concept of truth more generally (and with his emphasizing that “compact and exhaustive EMUs are the exception, not the rule” (Williams 2015, p166 fn9)). Nevertheless, Price’s account does, in principle, fit within the framework of the EMU.



The Pricean theory requires the following additions to the EMU:

- Inferential: b) For any sentence *S*, one can always freely infer from “ $\sim S$ ” to “The sentence ‘*S*’ is false”, and vice versa.
- Epistemic: b) The inferences in Inferential (b) above are a priori. Further (as a matter of epistemic normativity), if a subject asserts a sentence that is implied to be false by a sentence one holds to be true, one has a defeasible obligation to correct them as appropriate, according to the specific circumstances.
- Functional: b) The above structures work to diminish the possibility of contradictory assertion-pairs, and therefore conflicting actions, within a linguistic community, meaning the community operates more efficiently in general, since people are less likely to work at cross-purposes. Further, they help to marshal the epistemic resources of the community by triggering the procedures used by a given language game for resolving disagreements, such as reviewing assumptions and inferences for error, and gathering empirical data as appropriate. This strengthens the community’s epistemic standing, specifically regarding the “problem cases” that trigger disagreement.

There are a few notes to make about these additions. First, I have expanded on Williams’ own proposal for adding Price’s “friction” view of truth into the mix; Williams (2013, p136, fn3) proposes merely relaxing the functional clause of the EMU to allow for the new function. I do not believe that this is sufficient – one needs, at the very least, to supplement an epistemic norm to motivate initiating the process of challenge and resolution. This is what explains why speakers don’t simply rest content with “no-fault” disagreements, and why they ought to respond when their views are challenged – both intrinsic parts of Price’s account.

The conditions under which the norm of correction can be defeated are, of course, not outlined in the EMU. This is a matter of the procedures for settling disputes being different for different language games; the procedures for improving one’s epistemic standing are not detailed for the same reason. There are, of course, cases where disputants simply have to “agree to disagree”, primarily, it seems, when the setting or circumstances for holding a debate or gathering evidence are too limiting of the sort of detail that resolving the dispute

would require. One does not have a duty to challenge conspiracy theorists on street corners, for instance, because convincing them of their error would take more time and evidence gathering than is feasible given the circumstances. Nevertheless, it seems plausible that, in the right circumstances, and given enough time and resources, one would have a social-epistemic duty to attempt to disabuse fellow epistemic agents of such misconceptions.

The inferential role stated for truth in Inferential (b) is not really new material, but simply makes explicit the moves enabled by the original account that set us up for an account of Pricean “friction” (although it does introduce ‘False’ as a new vocabulary item in the object language, which is then used in subsequent clauses). Once again, the EMU relies on the language user’s already being able to interpret and apply the individual sentences that are to be assessed for truth; just as Williams’ initial version left out an account of what it took for the example sentences to be correctly assertible, so now we leave out an account of what it takes for pairs of sentences to be co-assertible or inconsistent.

## 2: Reference

Reference is the other key semantic property besides truth that appears to be paradigmatically representational. As with truth, the expressivist doesn’t want to deny that reference is a genuine property that linguistic expressions are capable of holding; all the expressivist would wish to deny is that reference can be informatively explained in terms of its representational character.

At first glance, reference appears primarily to have a disquotational function, similarly to truth. Where a concept of truth licenses moves from sentences mentioned by quotation to assertions, and vice versa, the concept of reference allows us to move from mentioning individual expressions – chiefly nouns and noun phrases – to using them. For instance, we may freely move from successful uses of the term ‘Aristotle’ to the sentence “The term ‘Aristotle’ refers to Aristotle”, and vice versa from “The term ‘Aristotle’ refers to Aristotle” to using ‘Aristotle’.

What could be the practical usefulness of such a basic linguistic function? In the discussion of truth, we saw that there was some utility to moving from assertions to statements of the truth of those assertions, namely the utility of being able to generalize over the first-order assertions. Similarly, the concept of reference has a direct practical usage wherein it

encodes or cashes out terms, as in the sentences “‘Claimant’ will hereafter refer to Mr. Smith”, and “When Aquinas mentions ‘The Philosopher’, he is referring to Aristotle”, respectively. From this we can derive a generalizing function for reference, similar to that of truth, where its use value lies in the fact that it bypasses the need to enumerate or even know a relevant group of singular terms, as in “Every spy named in the document is compromised”. Again, the practical advantage granted us by a concept of reference in such cases is that it allows us, in a manner of speaking, to move from mention to use, via disquotation. For instance, in the latter example the relevant grouping is one of names, but the pertinent quality is held by the people to whom those names refer; reference allows us to jump between these without going via an intermediate step of using or even learning the names themselves.

Paul Horwich (1998a) notes some corner cases to watch out for in defining reference diquotationally. The primary pitfall to avoid is the familiar problem of reference failure in virtue of non-existent referents – the term ‘Atlantis’, for instance, does not refer to anything. Horwich thus proposes a quantificational phrase to define reference:

$$(x)(\text{Tokens of } *n* \text{ refer to } x \text{ iff } n = x)$$

*Ibid.*, p119

This phrasing bypasses the problem because there is no  $x$  such that  $x$  is Atlantis, hence we are not committed to ‘Atlantis’ referring to anything.<sup>31</sup>

---

<sup>31</sup> Horwich also presents workarounds for indexicals and attributing reference to foreign languages; I take these to be beside the point for present purposes.

We thus have the materials for a simple EMU for reference, as follows:

- Inferential: With the exception of indexicals, for any  $x$ , tokens of a term 'n' refer to  $x$  iff  $x = n$ .
- Epistemic: The inferences from " $x = n$ " to "The term 'n' refers to  $x$ " and vice versa are a priori.
- Functional: The above enables generalization in cases where specific names are not known to the speaker (e.g.: "Everyone John named is a good friend.").

There are more options on the table for the expressivist that simply treating semantic properties in terms of pure disquotation. For example, Robert Brandom (1984) gives an alternative expressivist treatment of the concept of reference in terms of anaphora. According to Brandom, the core use of the concept of reference is in creating indirect descriptions, such as "The person to whom Jones referred in yesterday's meeting". Brandom argues that such descriptions are anaphoric on previous utterances – the example, for instance, is anaphoric on Jones' utterances in yesterday's meeting, since the subject of the description is determined by being the subject of those utterances. This distinguishes descriptions that use the concept of reference from other relational descriptions, such as "The person Jones upbraided in yesterday's meeting", which are not anaphoric.

The distinction between anaphoric and non-anaphoric descriptions, Brandom argues, comes down to the iterability. Briefly, being an anaphoric dependent is a transitive relation; if  $y$  is anaphoric on  $x$ , and  $z$  is anaphoric on  $y$ , then  $z$  will be anaphoric on  $x$ . Hence we would expect our reference-involving descriptions, if they do involve anaphora, to display this transitivity. The referential operator in the description "The person to whom Jones referred in yesterday's meeting" should, then, be able to iterate while preserving the subject of the original phrase – as indeed it can, giving the description "The person to whom "The person to whom Jones referred in yesterday's meeting" refers", which picks out the subject of the original phrase (*ibid.* p478). By contrast, other relational examples do not allow for iteration. For instance, the description "The person Jones shocked in

yesterday's meeting" becomes (something like) the phrase "The person who was shocked by "The person Jones shocked in yesterday's meeting"". Whether or not we get rid of the smaller-scope quotation marks (whether we treat someone as being shocked by a *person* or by a *description*), it is possible for the subject of the new phrase to be different from the subject of the original, so the description does not iterate in the required way (*ibid.*).

With this analysis of the core usage of 'Refers' in place, Brandom can parse other uses. This includes uses that, on a relational account of reference, might appear to be more basic than the anaphoric use above. For instance, the disquotational sentence "'Leibniz' refers to Leibniz" is unpacked by reworking it into an identity statement that more clearly displays the anaphora at work, namely "The one referred to by the term 'Leibniz' is (=) Leibniz" – an expression of the identity of Leibniz with someone picked out by the indirect description 'The one referred to by the term 'Leibniz'', where this description is anaphoric on some previous use of the term 'Leibniz' (Brandom does not seem to think it matters which use or when – *ibid.* p482). The above identity claim is thus comparatively trivial in the way we would expect of such a claim. Further, since it uses only intra-linguistic relations (i.e.: not relations between expressions and their subject-matter), it is non-representationalist (Brandom is emphatic on this issue – *ibid.* pp487-8; cf Brandom 1994, pp325-7).

It might seem that, thus far, although we have outlined the possibilities for an expressivist account of the intra-linguistic role of reference, we have not yet accounted for its extra-linguistic significance. That is to say we have explained how to use the word 'Refers', with specific attention paid to its disquotational use, but we have not explained why it is, for instance, that we use 'Aristotle' to refer to *Aristotle*, and not someone else – what is it about the term, or its use, that means that it refers as it does?

The expressivist is committed, in one sense, to denying that a satisfying answer can be given. Although such a question appears pressing, it effectively amounts to a request for an account of the underlying nature of reference in terms of a substantive relation between word and referent – in other words, just the sort of account that the expressivist denies is possible. Nevertheless, treading carefully, we can gesture towards an answer that stays within the bounds we have set.

The first step to doing so is to reframe the question as being about reference-attributing *behaviour*. As noted, the question of the link between a term 't' and its referent, t, is one that for which the expressivist cannot provide a substantive answer. However, the expressivist can begin to provide an account of the sorts of factors that prompt language

users to use 't' in respect of t. Stephen Barker proposes understanding reference-attributing behaviour in terms of mental "modules" possessed by speakers. Modules effectively encode concepts in pattern-recognition terms, such that they provide speakers with an immediate grasp of whether a certain sort of thing – the sort encoded in the module – is present in their environment. For example, the visual information that speakers are trained to correlate with the presence of tables in their environment might be encoded in a module: TAB-Module. This module will be either active (when the speaker recognises a table), inactive (when they do not), or undecided (Barker 2015, p12). Modules can be activated directly, by exposure to sensory inputs encoded in the module, or indirectly, by testimony, inference etc.

The main explanatory purpose of such modules is in outlining the mental state being expressed by referential utterances. In a paradigmatic case of recognising a table on the basis of sensory inputs, TAB-Module activates, tokening a mental state,  $\Pi$ .  $\Pi$  is something akin to a state of recognition, but is pre-doxastic, and therefore not a full belief. The idea here is that a referential utterance caused by TAB-Module's activating expresses  $\Pi$  (*ibid.* p14). Here the referential term is the predicate 'Table', which refers to the property of tablehood. Barker uses the system as established so far to explain reference-behaviour for singular terms by modelling these as "files" – essentially bundles of modules, connected so as to allow recognition of individuals. We thus have simultaneously an account of the cognitive architecture underpinning referential behaviour, into which we can begin to slot specific features for specific modules, to account for the usage patterns of different referring expressions, and the beginnings of an account of the mental states expressed by referential utterances.

Barker's modules bear a resemblance to Williams' reliable discriminative reporting dispositions (RDRDs), in that both denote a capacity by which speakers recognise the application of linguistic expressions to be warranted. However, there is an important difference: RDRDs are perfectly extensionally accurate, while modules need not be (*ibid.* p12). The function of RDRDs was to establish that speakers were correctly trained in the use of an expression, namely by specifying that they are only disposed to use the expression when it is correct to use it, and not when it is incorrect. This function requires RDRDs to specify the extension of the expression in question, so that they can specify its correct use (recall that the extension had to be specified in a minimalist way, as discussed in Ch.2). By contrast, the function of mental modules is to underpin recognitional capacities in actual speakers, not to offer a standard against which speakers' correct usage can be

judged.<sup>32</sup> In other words, the point of modules is to explain what is going on in actual speakers' heads – which pre-doxastic mental states they are responding to causally, and which they are expressing – when they undertake to utter referring expressions, and there is no requirement that actual speakers (or their dispositions) be infallibly accurate. To reiterate, we have dropped as illegitimate the question what the character, in substantial terms, of the relation between a referring expression and its referent actually *is* fundamentally; the only question that remains for the expressivist is that of a causal explanation of what speakers are doing when they attempt to refer.<sup>33</sup>

### 3: Modality

Amie Thomasson's recent (2020) book *Norms and Necessity* offers an explicitly expressivist account of modal language. Put briefly, Thomasson's account envisions modal operators such as 'Necessarily' and 'Possibly' (and especially deontic expressions such as 'Shall' and 'Must') are best construed as tools that help in expressing linguistic rules. This is in contrast to the proposal to understand these operators in terms of their ability to represent a realm of distinctly modal facts, which poses difficult metaphysical and epistemic worries. As Thomasson puts it:

[M]odal talk enters into language as a way of expressing rules or norms in the indicative (rather than imperative) mood. On this view, the most basic function of modal terms is to enable the expression of rules and their logical consequences, while making the regulative status of what is said more explicit than it is in the other forms of expression, enabling us to express conditionals connecting them to other rules, and to make explicit our ways of reasoning with them, and enabling permissions as well as requirements to be expressed.

Thomasson 2020, pp62-3

---

<sup>32</sup> Zalabardo (2019) offers an abstract sketch, in pragmatic terms, of a parallel view for the reference relation between predicates and the properties they stand for.

<sup>33</sup> I include this point to sketch briefly how the expressivist might respond regarding questions of reference attribution; for a much more in-depth discussion, see Barker 2007, §§44-47 (although the explanation there does not use the "Module" terminology). Note that, although I pointed out in Ch.2 that Barker's take on expressivism would not yield the results we wanted in support of neo-Carnapianism, the machinery provided therein can still be of use to the expressivist who adopts semantic minimalism, as shown here. A high degree of cross-compatibility of explanatory strategies is a mark of expressivism.

This expressivist reading allows us to make sense of the sentence “Bachelors must be unmarried” (in the instructional sense used to teach non-fluent English speakers, not as a general, contingent quantification over bachelors) as a way of expressing the permission “If x is a bachelor, then one may infer that x is unmarried” or the imperative “Apply ‘Bachelor’ only to unmarried things”. Similarly, this reading of modal language tracks with Williams’ use of “always” in the EMUs above. By stating that a given inference is “always good”, Williams doesn’t mean first to adopt a particular ontology of states of affairs, then to attempt an explanation of x by coordinating it with that ontology (a non-conservative explanation, in a way that makes it a non-expressivist one, as we saw in Ch.3. Rather, Williams is stating that the inference forms part of a semantic rule of the language, something we can express using modal terminology.<sup>34</sup> Thomasson’s proposal thus has a deep significance for the expressivist specifically, as it makes available to them the tools for constructing expressivist explanations.

There are other advantages to modal vocabulary that underpin its usefulness. One key advantage is that it allows us to reason about rules using conditionals, something we can’t do by expressing rules in the imperative mode. For instance, having a way of expressing a rule propositionally allows us to construct sentences such as “If bachelors must be unmarried, then no bachelor can be a spouse.” By contrast, the imperative version “Apply ‘Bachelor’ only to unmarried things” cannot be embedded in the antecedent of the conditional, so cannot be used to express such reasoning (*ibid.* p60). Besides this, a dedicated modal vocabulary serves to make explicit the fact that the stated commitments are rules of the language, rather than simply commitments of the speaker. E.g.: A speaker uttering the sentence “All bachelors are men” could wrongly be taken to be attempting to voice a hypothesis rather than attempting to express a linguistic rule, where the sentence “A bachelor must be a man” less ambiguously expresses the latter (*ibid.* pp83-4). Note that, in linking the rules of use for modal vocabulary to the rules of use for the relevant language (the rules of which it is the job of modal sentences to express), Thomasson implicitly makes modal vocabulary part of the “analyticity family”, specifically because its primary use is in drawing attention to the analytic rules of a given language *as* analytic for that language.

---

<sup>34</sup> I take it that something like being a basic rule of the language is also what is intended by the use of ‘*a priori*’ in these EMUs. Certainly the expressivist will want to give an explanation of a-priority in terms of its functional role, particularly in linguistic explanation.



We thus have the pieces to construct an EMU for, at least, the concept of necessity, as follows:

- Inferential:
  - o a) If “p” is an object-language expression of an actual semantic rule (or a logical consequence of actual semantic rules), then you are entitled to introduce “Necessarily p”, regardless of any subjunctive suppositions (Thomasson 2020, p83).
  - o b) If you have “Necessarily p” as a premise, you may use p as a premise in your reasoning anywhere, under any subjunctive suppositions (*ibid.* p84).
  
- Epistemic: The inference from “p” expressing a semantic rule to the claim “Necessarily p” is a priori.
  
- Functional: Stating that necessarily p expresses p’s status as a rule of the language; ‘Necessary’ and its cognates are thus used for expressing linguistic rules in teaching contexts. Sentences containing ‘Necessary’ also offer a way of encoding imperatives concerning linguistic rules in assertoric sentences, making them available to be embedded in conditionals, allowing conditional reasoning about language rules to be made explicit.

One worry we might have about this explanation of necessity is that it itself relies on implicitly modal concepts such as the one it is meant to explain. For instance, the phrase “you may use p as a premise in your reasoning anywhere” above might well appear to have a modal ring to it, since it seems attractive to cash it out as meaning “in any circumstances”. The case here is part of a recurring theme for expressivist explanations such as the ones presented in this chapter, which have an interesting relation with circularity and tautology. In such cases, the expressivist is committed to maintaining that the circularity is not vicious. Indeed, the apparent circle is just a product of the expressive linguistic tools that are being explained being the same tools that are doing the explaining. This is perhaps clearest in the modal case, since, on Thomasson’s view, modal concepts are characteristically used to encode exactly the sorts of imperatives and permissions that we are here seeking to unpack with respect to modal terminology itself.

The apparent problem is that we seem to be led back to a substantive account of modality, because our use of modal phrases in our EMU for necessity means that that EMU is structured so as to represent, for instance, a field of possible worlds, and to use this representation in an explanatory capacity, i.e.: representing every possible world as a world in which “p” (a necessarily true proposition) is usable as a premise in reasoning. This view on the objection is answered by simple appeal to Thomasson’s main thesis that modal terminology does not have primarily a representational function, but serves to express the rules of a language. The phrase “you may use p as a premise in your reasoning anywhere” is thus simply a way of expressing the rule that p is usable as a premise without having to satisfy any further conditions, much as the phrase “A player can always fold on their turn” is primarily meant to present a permission, equivalent to the rule “Nothing can oblige a player not to fold if it is their turn.” It is not that the word “always” isn’t being used with a modal meaning – it definitely is – but rather that, per the account, this does not entail a reliance for our EMU on its being representational in a substantive sense.

This sort of circle is relatively familiar; we saw something like it in Ch.3. But the circle is not a vicious one – rather, as already pointed out, it is a result of the fact that we are here turning our language’s semantic-explanatory tools on themselves. Thomasson makes a similar point regarding logical consequence: It might appear as though an antecedent understanding of modality is required to understand logical consequence, if, for example, we understand an argument’s being valid in modal terms, i.e.: as there being no possible worlds in which its premises are true and its conclusion false (*ibid.* p125). Since Thomasson’s modal normativism relies on logical consequence, it might appear to be self-refuting, as it seems already to require a substantive account of normativity in order to get off the ground. Thomasson, though, assures us that this is not the case; taking an inferentialist approach to logical consequence, we can see that the modal vocabulary at work in stating logical necessity serves the same function as that which is used to state metaphysical necessity, namely of encoding permission and obligations. Specifically:

A claim that it is a logical necessity that, if the premises are true, then the conclusion is true, on this model, is not a description of some feature of this world or of all possible worlds. Instead, it expresses a norm regarding acceptance and rejection: that one mustn’t accept the premises and reject the conclusion.

*Ibid.* p127

Again, as with Hirsch, we note that this form of deflationism has to be recursive if it is to be truly deflationary – a semantic deflationism articulated in terms that built in metaphysical assumptions would not achieve its goal of rejecting metaphysics, but would only bury the metaphysics beneath a single semantic layer. To find fault with the recursiveness of Thomasson’s theory, therefore, would require begging the question against it, as we would have to assume that Thomasson’s modal normativism was incapable of supporting semantic accounts, such as the modal normativist’s own explanation of modal terms and of logical necessity. Modal normativism simply relies on modal-normativist accounts of modality. Recursive though this may be, as Thomasson puts it, this is “not circularity but thoroughgoing consistency.” (*ibid.* p126)

#### 4: Environment-tracking terms

While semantic properties are defined largely in reference to their role in structuring discourse, particular concepts localized within a particular linguistic framework have to be treated more directly. Terms such as ‘Red’ and ‘Tree’ that refer to or pick out particular properties or things in the environment are a particularly tricky case, because of the expressivist’s commitment to ontologically conservative metasemantic explanations. This was discussed earlier, in Ch.2, so I will only briefly recap the form that expressivist explanations of these sorts of term must take.

An expressivist explanation of an environment-tracking term will focus on the term’s patterns of use, specifically its inferential patterns. As mentioned in Ch.2, following Williams (2013, pp140-1), we can take “x is red” to warrant such assertions as “x is not green”, “x is coloured” and so on. Likewise, we can infer the truth of “x is red” freely from more specific facts such as “x is crimson”, “x reflects light with a wavelength of 700nm”, etc. Insofar as environment-tracking terms such as ‘Red’ have an expressive function, it is in making possible such linguistic moves as these inferences allow – assertion, endorsement, demonstration of evidence and so on. Following Price’s terminology, this portion of the account details the i-representational function of the term.

The other part of the explanation is an explanation of *when* the term gets used, meaning the particular conditions under which it is correct to use it. As discussed in detail in Ch.2, the expressivist here need only specify, in whatever terms, the cases in which the term is applied correctly, even if this amounts only to a disquotational truism such as “‘Red’ is

correctly used of red things.” Williams himself flirts with the line here, stating that “To master ‘Red’ in its reporting use, the speaker must have [...] a disposition, given appropriate motivation and conditions, to report ‘x is red’ only in the presence of a red thing in his field of vision.” (*ibid.*) However, while non-ontologically conservative (because it quantifies over red things), Williams notes that this explanation is non-reductive, and thus is non-representationalist; the disposition to utter ‘Red’ only in correct cases is defined using the term ‘Red’, so barely advances beyond the truistic definition above.

In Price’s terminology, this latter is the e-representational portion of the explanation. While not disallowed outright, there are significant constraints on the sort of content it can contain, according to the tenets of global expressivism as we have outlined them up to this point. Hence there is little room to rest significant theoretical weight on, for instance, the specific relation between utterances of environment-tracking terms and the things or properties they track in the environment, in large part because these properties, and thus this relation, can only be specified in the truistic manner noted above.

## 5: Sortals

Functionally speaking, sortal terms are frequently employed with a similar environment-tracking purpose to expressions like ‘Red’ – as noted previously (Ch.2), it is an important part of the purpose of e.g.: the term ‘Tree’ that its correct use “tracks trees” in the speaker’s environment. As also noted, this tracking function is not a function we can leave out of the account, first if we want the account to be extensionally accurate, but especially if we want to give our account in terms of the linguistic “purpose” of such terms. We might expect, then, that an EMU for sortal terms – at least in a basic descriptive use for tracking objects in the environment – would have an Inferential clause that looks rather like that for ‘Red’, i.e.: that for some sortal ‘F’, one uses ‘F’ correctly only if one demonstrates a reliable reporting disposition to utter ‘F’ only when there are Fs in one’s environment.

Sortals (in their tracking use) plausibly have at least one other function beside the one attributed to ‘Red’, namely of establishing conditions of identity – what Thomasson calls “coapplication conditions” (Thomasson 2015, pp223-6). The idea here is that the sortal to which something answers grounds the conditions under which it is correct to reidentify it as the same thing, e.g.: Whether Theseus’ ship pre-rebuilding is numerically identical with Theseus’ ship post-rebuilding comes down to a question of the coapplication conditions

embedded within the sortal 'Ship', and whether they are satisfied in respect of the ship post-rebuilding. That the Theseus' ship case itself provokes different intuitions on this score is a sign that it is a corner case to which the existing coapplication conditions for 'Ship' (or other sortals with which we might run the paradox) do not unambiguously extend. What is called for is simply a decision on the best way to extend those rules to such cases – this showcases in microcosm the deflationary utility of recognizing the function of coapplication conditions.

Broadly speaking, the practical function of such cases of reidentification seems to be to note when properties are attached to the object that are not otherwise self-evident from observing it in its present state. For instance, properties something has in virtue of its history "carry forward", as licensed by reidentification via a sortal. The customs status of a ship coming into port, for example, depends on the port from which it sailed, so being able to reidentify the ship wishing to enter port B as the same ship that left port A is necessary for determining this status. This sort of property highlights the usefulness of coapplication conditions, and thereby the reason that such conditions are maintained within our language; it would not be possible to perform functions tied to these historical properties, such as regulating trade, as effectively if our language did not contain fully fledged sortals with coapplication conditions.

The above hardly amounts to offering a comprehensive EMU for sortal terms. However, it is unavoidable that the sort of expressivist characterisation of sortal terms in general that we are equipped to give here should be somewhat bare, for two reasons. The first is that much of the heavy lifting in terms of the function of sortal terms within language is done by attaching them to quantifiers, in particular the existential quantifier. The following chapter will formulate an EMU for existence that takes Thomasson's deflated concept of existential quantification and shapes it into an explicitly expressivist account. Since the function of sortals is largely to provide application conditions to make quantification possible, the account there will also serve to further substantiate the role the expressivist sees for sortals in language.

The other reason the expressivist is unable to give much of a general account of the function of sortals beyond what I have detailed here is that such terms might serve a wide variety of functions depending on the language games in which they are used. For instance, terms tracking medium-sized dry goods may have various uses across many different language games (there are all sorts of reasons to want to track chairs, for example), while

“social-ontological” terms such as ‘Government’ or ‘Marriage’ will likely be more closely tied to particular clusters of practices governing the operation of those social structures. What I have done here is to attempt to characterise in functional terms the structures common to sortal terms as such, and thereby some of the reasons we might have sortal terms at all, a task that is inevitably going to result in a rather abstract explanation, analogously to how the expressivist explanations of truth and reference lead to complete metasemantic explanations of those concepts that nevertheless do not tell us how to assess, in particular cases, whether a sentence is true, or whether a term successfully refers.

## Ch.8: Conclusion

At this point, we have most of the argument of the thesis: We have examined several key approaches to neo-Carnapian deflationism and found problems for them. In presenting these problems, we have seen that they are necessarily representationalist in character, i.e.: they hinge on demands for a substantive account of how semantic properties relate linguistic items to their subject-matter (for Thomasson's easy ontology, the problem lay in attempting to leave room for a representationalist metasemantic account, while Hirsch's intensional deflationism required work to remove such an account, which was already built in). In each case, the problem was that the assumed representationalist metasemantics provided a fertile ground to reestablish substantive metaphysical positions (and therefore substantial metaphysical debate) of the sort the neo-Carnapian sought to deflate. Hence, given the problems' origins in representationalism, we have also examined the anti-representationalist global expressivism, which looks to provide a metasemantics that doesn't necessitate the word-world relations that proved problematic for the neo-Carnapian.

What remains to do is to examine the links between these positions; what, exactly, is it that global expressivism does for neo-Carnapian deflationism that solves these representationalist problems? And why, exactly, are neo-Carnapian thinking and representationalism apparently incompatible? The answers to these questions are not immediately obvious. In this chapter, therefore, I will attempt to clarify the need for a connection between neo-Carnapianism and global expressivism, and, to a degree, the nature of the connection itself. I will begin by analysing the core concepts of Hirsch's and Thomasson's positions in expressivist terms, and showing exactly how putting their views on an expressivist footing prevents these specific problems from threatening the re-emergence of substantive metaphysics from within these systems. I will then attempt to draw some general conclusions about why a Carnapian rejection of metaphysics in general seems to be incompatible with representationalist methods and assumptions.

## 1: Hirsch

### 1.1: Intensional deflationism

To recap briefly, Hirsch's model for intensional deflationism involved deflating the claims made by opposing metaphysical theories by showing their truth-conditions to be identical, and therefore demonstrating them to be in agreement. The problem presented by Hawthorne, under the analysis presented in Ch.3, was that Hirsch's method involved mapping the respective claims' truth conditions onto a shared intensional landscape, which basically amounts to a set of commitments concerning which states of affairs are numerically identical and non-identical. If we assume a single correct intensional landscape, as Hirsch seems to do, then we can derive substantive first-order metaphysical truths, and provide metaphysical validation to one side or the other in several debates we would want to deflate. This is because the commitment to a particular intensional landscape over the alternatives is itself a metaphysical commitment, which erodes the deflationism Hirsch sets up. The solution is to be deflationists about which intensional landscape is the "correct" one; since this question itself is metaphysical, the deflationist will regard it as being determinable only by practical considerations, and not by an appeal to metaphysical correctness. It follows that, while we can still derive validation from an intensional landscape, this validation is not metaphysical but pragmatic, as it simply extends from whatever pragmatic justification we marshal for supporting that intensional landscape, instead of an opposing one.

What we discovered in Ch.3 was that although Hirsch's deflationism *utilizes* semantics, it does not *extend to* semantics – what Hirsch offers is, in fact, a substantive theory of truth-conditions, upheld by an assumed correct ontology of states of affairs. This amounts to a representationalism about the semantic concept of intensionality: An assertoric sentence's having truth-conditions is understood as its representing a certain portion of the intensional landscape. Assertoric sentences with the same truth-conditions are thus *de re*-equivalent representations, i.e.: they are truth-conditionally equivalent in virtue of their representing *the same* portion of the intensional landscape.

The key here is the direction of explanation. In Huw Price's phrasing, the above interpretation is representationalist because it uses the concept of truth-conditions in a semantic-explanatory or theoretically load-bearing capacity – a capacity that such a representational concept should not occupy (Price 2011, p193). It is this substantive explanation in terms of representing portions of the intensional landscape that underpins



the substitutability of truth-conditionally equivalent sentences, which itself supports Hirsch's deflationary move. Hence Hirsch's deflationism depends directly on using a representational conception of truth-conditions in a metasemantic-explanatory way; for Hirsch, intensional equivalence *explains why* sentences are inter-substitutable, i.e.: by explaining that they represent the same thing.

The expressivist version, by contrast, reverses the order of explanation: Rather than holding that two equivalent sentences are intersubstitutable in virtue of their representing the same part of the intensional landscape, the expressivist analyses the concept of intensional equivalence, and the truth-conditional machinery that goes along with it, as being a tool for establishing intersubstitutability. This is a subtle shift; the expressivist still uses the concept of intensional equivalence in explaining the inter-substitutability of equivalent sentences. However, the expressivist account of intensionality *itself* is given in terms of its linguistic role – the language game moves that it licenses – rather than in terms of mapping onto a particular ontology of states of affairs.

As we saw in Ch.3, this requires engineering a metalanguage with rules regarding which object-language sentences count as intensionally equivalent. The expressivist version of the account has it that the semantics of this metalanguage is likewise non-representationalist. That is, we should not think of object-languages as attempting to represent the same intensional landscape of states of affairs; therefore, we should not think that finding the "correct" metalanguage will be decisive in the way metaphysicians want. In other words, in choosing between different such metalanguages (that map intensional equivalences differently), we ought to follow Carnap and treat the choice as an "external" issue – which metalanguage, and therefore which structure for an intensional landscape, best suits our practical purposes as language users?

This is what is meant by treating the concept of intensional equivalence as a practical linguistic tool; its purpose is not to pick out features of the world, but to help us get our commitments in order by ruling on which assertions are and which aren't interchangeable. To be sure, we can still derive metaphysical conclusions of a sort from the intensional metalanguage, specifically in Hawthorne-style cases. For example, we can derive a metaphysical conclusion from the fact that the sentences "The statue was destroyed" and "The lump was destroyed" have different truth conditions (according to a certain intensional metalanguage), namely that the statue and the lump are distinct entities. However, these conclusions are not metaphysical validations of particular object-

languages, so much as manifestations of our pragmatic choice between metalanguages. In this they are akin to the ontological conclusions of easy ontology, in being metaphysical truths of a kind acceptable to the deflationist. Recall that the “modern” deflationist regards their job not as rendering meaningless the deflated vocabulary, but explaining it in an acceptable, non-metaphysically committing way. For our modern deflationist, the fact that we *can* resolve metaphysical debates by appeal to an intensional metalanguage (if only we can agree on which metalanguage to use) is a feature, not a bug.

We might ask whether the concept of intensionality has a linguistic function other than at the level of philosophical debate. So far, intensionality, atypically for the linguistic concepts analysed over the course of the thesis, has appeared to have primarily a philosophical role. Moreover, the function we have primarily focused on has been one of deflation, and dissolving rather than enabling philosophical discussion. It would be curious if this were the primary function of the concept – for one thing, Hirsch’s appeal to intensional equivalence to dissolve metaphysical debates has the appearance of applying an independently understood tool of common sense in an abstract metaphilosophical context. The expressivist, then, will want to provide a “ground level” analysis of intensionality in terms of its use in more everyday contexts, even if only as a sketch or a “just-so story” (Blackburn 2013, p75), on which model we can begin to understand the metalinguistic version present in our revised Hirschian account.

It seems that the primary function of a concept of intensional structure within a language is to disambiguate and establish a speaker’s discursive commitments by making explicit which proposition they intend to express. For instance, when a speaker utters an ambiguous sentence A, we can usefully ask whether they intended to express the proposition p or the proposition q, where these bear different commitments. Likewise, we can use the concept of an intension to cash out these commitments directly. Phrases such as “Does that mean...?”, as in “Does that mean that I’ll have to pay a fine?”, play such a role. On the expressivist account, such a locution isn’t primarily concerned with the representational character of a sentence uttered by an interlocutor, as it may appear to be, but rather with establishing the interlocutor’s commitments – their position in the “game of giving and asking for reasons”, to use Brandom’s phrase (Brandom 1994, p141).

Intensional vocabulary, then, appears to be primarily concerned with disambiguating commitments via the mapping of sentences onto propositions. It is natural to view this function as a necessary one for a language of sufficient combinatorial power. Non-

combinatorial language forms, such as the language of Wittgenstein's bricklayers (Wittgenstein 1950, §2), if they are taken to be assertoric, are incapable of ambiguity of this sort because they map one-to-one between well-formed "sentences" and intensional commitments (or at least intensional commitments that the language is capable of expressing). The price of a generative language capable of creating complex assertions from sub-sentential components is the possibility of a one-to-many mapping between the sentence thus constructed and the possible sets of commitments it could incur for the speaker, i.e.: the possibility of ambiguity. The function of intensional language ('Intension', 'Proposition', 'Meaning') is to provide a check on this ambiguity.

## 1.2: Intensionality and Modality

Note that, although we are steadfastly concerned here with language-game "moves", and speakers' commitments therein, this model does bear ontological commitments, namely to an intensional landscape for the relevant language. Essentially, this is to say that if we want to explain intensional language in terms of its mapping sentences onto intensional commitments for speakers, we have to have an ontology of intensional commitments that it is possible for speakers to express by their sentences; in order to disambiguate A by asking whether it expresses p or q, for instance, we have to be assured that p and q are distinct propositions, and this means commitment to an intensional landscape.

The key here is that the expressivist takes a deflationary attitude towards this ontology. As discussed in Ch.3, the structure of the intensional landscape – the ontology of possible propositions, and the ontology of states of affairs in which it is grounded – is tied to an intensional metalanguage in the same way that the Carnapian sees the ontology of, say, mathematics as tied to the linguistic framework of mathematics. This is how our picture of propositions as sets of states of affairs is nevertheless able to remain non-representationalist, and metaphysically deflationary. The utility value, then, of intensional language is *not* that it allows us to clarify which of several metaphysically significant sets of states of affairs an ambiguous sentence is intended to represent, but that it allows us to clarify which of several possible sets of discursive commitments a speaker intends to express, where these could lead to different responses, courses of action etc.

One interesting point to consider is that the metalinguistic account of intensionality offered in Ch.3 actually complements Thomasson's expressivist account of modality. Intensionality

is intrinsically a modal phenomenon; the claim that two sentences, A and B, are intensionally equivalent is equal to the claim that whenever A is true, B is true, and vice versa. Nevertheless, it is not immediately clear how our expressivist accounts of each go together. Consider that Thomasson has it that a modal claim expresses a rule of the *object* language (e.g.: “Bachelors are always unmarried” expresses the rule *If x is a bachelor then x is unmarried*). It is not immediately obvious how this gels with our maintaining that such claims are dependent on an intensional metalanguage.

I will attempt to outline how the theories fit together, although such an explanation will have to be brief, as a thorough account of the topic would involve much detailed work that is beside the point here. To start with, note that Thomasson’s theory does not actually extend to the context in which Hirsch is working; Thomasson notes (Thomasson 2020, p83) that modal claims express the analytic rules of the language, but Hirsch is dealing with the interface between two distinct languages. In particular, the Hawthorne-style problems that gave us reason to engineer the metalinguistic theory hinge on the analytic rules in each language being different. Thomasson’s object-language oriented modality would be insufficient to deflate such debates, as restating the different rules in the object-languages using modal vocabulary would only result in the disputants continuing to speak past each other, as they remain tethered to their own languages.

By contrast, the metalinguistic conception invites the disputants to map their sentences onto an ontology of states of affairs (the intensional landscape) as an intermediary for mapping sentences between object-languages. Thus the intensional-cum-modal sentence “A and B are true and false in all the same circumstances”, where A and B are sentences in the disputants’ respective languages, does express an analytic rule, as per Thomasson’s theory – but here it is a rule of the metalanguage. The metalinguistic analysis of intensionality thus fills a gap in Thomasson’s modal expressivism, namely, in explaining how modal claims can be made that appear to bridge across languages with different analytic rules.

## 2: Thomasson

### 2.1: Deflating the problem

Thomasson’s easy ontology posits a deflationary understanding of the concept of existence. To recap: Thomasson’s stance is that a cogent statement of existence comes tied

to a sortal term. Sortal terms have application conditions, which can be mastered by speakers, under which they correctly apply to the world. The core posit of easy ontology is that, for any sortal term 'K', Ks exist iff the application conditions of 'K' are fulfilled. This has a deflationary effect on ontological debates, primarily because it alleviates the need to outline a substantive criterion for existence that things have to satisfy, over and above fulfilling the application conditions of a sortal (such as being a physical simple, or having causal powers) – an assumption on which a substantial number of ontological debates hang.

The problem we saw for easy ontology in Ch.4 was the potential for an infinite regress, posed by Andrew Brenner. Briefly, Brenner alleged that no sortal term could have finite application conditions, as any given set of application conditions for a sortal term makes mention of other sortal terms, with their own application conditions, and so on.

In Ch.4, we derived two potential readings of this problem. One reading has it as an epistemic problem: Being finite creatures, human beings could never know the application conditions of their terms, nor learn them, impart them, or check that they were fulfilled. We noted that this was solvable by treating the mastery of application conditions as a species of ability knowledge, or knowledge-how, which seems in line with how Thomasson intends them to be taken. This solves the problem of showing how ordinary language users can be justified in asserting ontological claims on the basis of their language mastery, a key aim of Thomasson's programme.

However, as noted in Ch.4, it is not enough simply for ordinary language users to be justified in making ontological claims – easy ontology also requires that the truth of these claims be sensitive to the same considerations to which ordinary language users are sensitive in articulating them. That is to say, application conditions for sortals not only determine when language users are *licensed* to make ontological claims, *those same* application conditions also determine whether the claims are in fact true. Hence Brenner's regress poses another problem, this time a broadly semantic one: It is unclear, for any given sortal, what is "required of the world" for its application conditions to be fulfilled. In other words, the content of the application conditions for a given sortal looks to be indefinite, on the grounds of those conditions being non-finitary.

We can now show more clearly that this second problem only arises given representationalist background assumptions about the requirements for semantic explanation. The requirement that we articulate the application conditions that make it

correct to apply some sortal, 'K', in terms *other* than the simple formulation that mentions semantic relations as such, constitute a requirement that we offer a substantive metasemantic account of what it takes for 'K' to apply. The revised Brenner objection alleges that there is a problematic regress regarding the truth conditions of the sentence "There exists some K" – problematic specifically because it means that those truth conditions are not finitary, and therefore do not carry a determinate content; we cannot specify in what its truth consists.

In Pricean terminology, the requirement lodged by the semantic reading of the objection is that we offer a finitary set of conditions under which the term 'K' successfully e-represents, that does not depend on any terms that do not themselves have a finitary set of conditions for e-representing. This forms the e-representational portion of the account because it is concerned with detailing the extra-linguistic conditions under which it is correct to apply the term, rather than with its linguistic function, or its structural role within a system of inference.

That it is not possible to satisfy this demand for a finitary e-representational explanation is not a problem for the expressivist, because the e-representational part of an expressivist metasemantic account is concerned only with outlining extensionally adequate truth-conditions, which can be done by the disquotational schema: "The sentence "There exists some F" is true iff there exists some F." The regress is only a problem if what we want is a further explanation of what the correctness of applying the sortal 'F' *consists in*, in some substantial sense – what properties in the world the sortal represents. We can now see how the sort of explanation the regress demands amounts to exactly what is prohibited by our expressivist credo of Ch.2, namely an attempt to analyse in substantive terms (i.e.: not the disquotational schema) the relation between an expression and its subject-matter – that subject-matter, here, being the set of properties the sortal is supposed to be sensitive to in its application. Hence the metasemantic version of the Brenner objection seeks to make substantive the e-representational portion of our metasemantics of application conditions, while expressivism holds that this is inessential for a complete metasemantic explanation.

There is one objection one might see here, to which I will attempt to respond as a further illustration of the point. In setting up the idea that we must be clear on what is "required of the world" by a term's application conditions in order for the term to be determinately meaningful, we drew on Thomasson's own example of the terms 'xliver', whose application

conditions are fulfilled iff there is a liver but no xheart, and 'xheart', whose application conditions are fulfilled iff there is a heart but no xliver. The significance of the example for our purposes is that the circularity of these terms' application conditions leaves us in doubt not just about when we, as speakers, would ever apply them, but about what the world would have to be like for either an xheart or an xliver to exist, since we can give no sort of analysis of their application conditions in other terms. But the expressivist is themselves committed to the viability of non-substantive explanatory strategies that themselves look about as informative as the 'xheart'-'xliver' case – for instance, accepting something as trivial as the use-mention schema "'K' applies iff there are Ks" as an explanation of the application conditions of 'K'. If the expressivist thus eases the requirement for the sort of reductionist analysis of terms that Brenner does, are they (perhaps absurdly) committed to the idea that terms such as 'xheart' and 'xliver' are perfectly meaningful, usable etc.?

I don't think so. First, recalling to mind Ch.2, we should note that the above is only one part of the explanation of the meaning of the term 'K', namely the e-representational part, designed to establish the accuracy of the explanation to actual usage. First off, this relies on an understanding of the target term, which presumably is lacking in the case of 'xheart' and 'xliver'. For instance, a language user can understand the disquotational use-mention schema in respect of 'K' if they have mastered the use of 'K' – and if not, they can be trained in its use by other means. By contrast, we presume that there is no independent way to master the use of 'xliver' or 'xheart'.

Further, and more importantly, there is no i-representational portion (we presume) of the xheart/xliver explanation. In fact, this is how the expressivist can clearly account for the comparative meaninglessness of the terms, namely by pointing to their not having a role in any sort of linguistic practice. Had they such a role, they would at least have meaning in virtue of the inferences (or, more broadly, the language game moves) in which they were involved. This sort of meaningfulness would be "purely" expressive, i.e.: it would have an expressive linguistic function, but not primarily a representational one.

I believe that this point is highly illustrative of the role of i-representation in accounting for the meaning of expressions. The concept of such circular application conditions is not entirely fanciful – one might think, for instance, that the application conditions for basic physical terms will end up being inter-defined similarly to those of 'xheart' and 'xliver', such that any one of them can only be understood in reference to the others, and nothing else (axiomatic concepts in maths and logic might follow a similar pattern). A semantic

explanation of these terms will be informative only insofar as it displays their functional role in the language structures of which they are a part, as simply tracing their application conditions will only lead us in circles. Nevertheless, despite the circularity of these e-representational explanations, such concepts are both usable and meaningful in virtue of the surrounding linguistic apparatus in which they play an inferential role.

What makes 'xheart' and 'xliver' problematic, then, is not strictly the circular e-representational account of their application conditions, but the paucity of i-representational structures in which they are situated. Hence the expressivist can explain what is wrong with these terms without having to fall back on representationalist tropes; the expressivist neither has to countenance the absurd position that these terms are perfectly in order as they are, nor agree with the representationalist that "non-substantive" semantic explanations are insufficient.

## 2.2: An EMU for 'Exists'

One interesting upshot of the combined neo-Carnapian program we have been putting together is that existence ends up looking much like a semantic property on the order of truth or reference. At the present stage, we can see that these properties share a lot in common: They are, for instance, all used to bridge between linguistic discussions and metaphysical discussions. As an example, as noted by Price (2011, p189), the property of reference gets used to propose ontological questions about a particular domain of entities by asking whether the language used to range over that domain can successfully refer, i.e.: whether it has anything to refer *to*. By expressivist lights, this is illegitimate – the question whether a certain group of terms successfully refer can only be answered by marshalling the conditions for correctly applying them, which is characteristically done "internally" by simply using the language, not by "external" metaphysical analysis. This is exactly analogous to Thomasson's diagnosis of ontological questions, discussed earlier in detail; the question whether Fs exist really comes down to whether the application criteria for 'F' are satisfied, which is best assessed from within the language by its speakers. If we approach ontological questions from this end then there is no converting them into hard metaphysical problems without assuming in a non-deflationary account of existence, analogously with the cases of truth and reference. Absent this account, the question is as much or more about the relevant sortal term, and whether its application conditions are



fulfilled, as it is about objects in the world, just as a question of successful reference is as much about the referring term as it is about the purported referent.

I have already argued above that Thomasson's deflationary analysis of existence works best against an expressivist backdrop, and we can now start to see why: Easy ontology exposes the "semantic basis" for a concept of existence – the satisfaction of the application conditions of a sortal term – placing it at the start of the inquiry. This is essentially the same move that Price makes in declaring that we ought to start with subject-naturalism, not object-naturalism. Once we have worked out how to understand existence on this quasi-semantic footing, we notice that there is no longer any room for a further metaphysical explanation – explanations of existence as a property of entities, aside from this quasi-semantic one of satisfying a sortal term, will not explain any further what it is to exist, or what is required of an entity for it to properly be said to exist, any more than has already been done by the easy ontologist's deflationism. We can thus create an EMU for 'Exists', which ought to reflect this deflationary stance by capturing all cases of the correct use of the expression, and explaining its linguistic function, without relying on external conditions.

I am not the first person to propose an EMU for 'Exists'. Kyle Mitchell (2014) proposes cashing a deflationary conception of existence using Williams' EMU schema in order to address concerns about quantifier variance (p578) and logical conventionalism (pp586-7) for the deflationist. Mitchell's paper takes as its central subject Thomasson's existence deflationism – notably, though, it was published in 2014, the year before Thomasson's *Ontology Made Easy*, which more fully fleshes out Thomasson's position. Mitchell instead largely draws from the less thoroughly worked-out versions in Thomasson (2007) and (2009) for textual support (Mitchell was likewise presumably unable to make use of Thomasson 2014, which contains a useful functional analysis of 'Exists').

The EMU below, then, takes Mitchell's version and builds upon it by drawing from the subsequent detailed work done by Thomasson. In particular, while a version of Inferential (b), below, is present in Mitchell's EMU (Mitchell 2014, p577), the inferential section offered below is far more comprehensive (at slight risk of redundancy). The EMU below also draws upon the "doch" function of a concept of existence presented in Thomasson (2014).

Our EMU for 'Exists' looks like this:

- Inferential:
  - a) With the exception of intensional contexts (and *pace* Meinong), for *any* predicate F, the inferences from "x is F" to "x exists" and "something exists that is F" are always good.
  - b) The inference from "x exists" to "For some sortal predicate K, x satisfies K" is always good.
  - c) Inferences from sentences that use referring terms to the existence of the entity referred to via intermediary steps are always good; e.g.: From "There are two apples" to "The number of apples is two", and then to the targets "The number two exists" and "There exists a number".
  - d) "Pleonastic" inferences from sentences that omit reference to entities they nevertheless rely on for their truth are always good. E.g.: From "The apples and oranges are equinumerous" to "The number of apples is the same as the number of oranges" to "There is a number".
  
- Epistemic: The inferences presented above are all *a priori*. The epistemic grounds for asserting "x exists" *simpliciter* are therefore parasitic on the epistemic grounds for holding that some sortal term applies to x (as per (b) above); the assertion that x exists is defeasible iff it can be shown that there is no sortal term that applies to x. Existence claims are subject to deference to experts, whose expertise pertains specifically to the relevant sortal.
  
- Functional: As per the Inferential analysis, part (a) above, existence is a necessary condition on having properties. A primary function of existence-discourse *per se* is thus to call into question assertions about entities by asserting that they *lack* existence, by appealing to part (b), and therefore that attributions of properties to those entities are in error. E.g.: "King Arthur ruled Britain" is defeated by "King Arthur did not exist", which is equivalent to the claim that the expression 'King Arthur' does not refer to anything that satisfies some sortal predicate.

The above EMU captures the key points of the usage of 'Exists' that give it the function that it has in our language, and therefore that underpin its meaning. As Thomasson herself notes (2014, p199), the point of the expression 'Exists' consists primarily in its negative usage in a "doching" role, i.e.: in countering particular claims about the properties of entities on the basis that they do not exist to hold those properties (this is especially evident within Philosophy, namely in error theories about various domains). This discursive move, Thomasson posits, is a key reason for having an expression for existence at all, since it allows us to make the claim that some term fails to refer without having to move to a metalanguage that mentions the term – we simply use the term to say, for instance, "King Arthur did not exist" (*ibid.*).

The introduction of sortals and their application conditions in *Ontology Made Easy* shows us that this negative existential move provides a minor inferential shortcut: To disprove that x is F on existential grounds, one does not need to show that x fails to bear *any* properties (although it will entail this, as per part (a)), but only that x fails to answer to a sortal, as, following part (b), such is a necessary condition on its existing. In practice, this usually takes the form of demonstrating that x fails to satisfy the specific sortal the interlocutor has in mind, e.g.: That 'King Arthur' doesn't refer to any *human being*. Even if the interlocutor can subsequently prove that something we refer to with the name 'King Arthur' satisfies the application conditions of a different sortal, e.g.: 'Fictional character', they will have proved the existence of something other than originally intended, and not rebutted the existential point.

We now have an account, then, of the concept of existence as it is used linguistically that explains both the use to which it is put and how language users are able to make use of it (i.e.: how they are able to learn to apply sortals). This account is deflationary in the way required by easy ontology, hence it can be used to generate genuine ontological conclusions, but it does not encounter either the regress or circularity problems that previously threatened the easy-ontological account.

### 3: Why incompatibility with representationalism?

We have seen, in each of the main objections to neo-Carnapian deflationism that we have examined, that the root cause of the issue came down to representationalist assumptions embedded in either the theory itself, or the objection. In Ch.3, we saw that Hirsch's

intensional deflationism was hamstrung by its own assumption of a set ontology of states of affairs, required to establish a representationalism about intensionality. In Ch.4, we saw that attempts to fashion application conditions for sortal terms into a representationalist semantics undermined the easy ontology that those sortal terms looked to establish. Indeed, as far back as Ch.2, we saw that the representationalist assumption of a distinction between “thick” and “thin” versions of semantic properties prevented Matthew Simpson from seeing the possibility of a global semantic deflationism.

We have also seen, as directly above, that an alternative, expressivist approach to these areas can solve these problems for the neo-Carnapian. I therefore feel justified in concluding that representationalism, speaking generally, problematizes metaphysical deflationism; the neo-Carnapian, if their deflationary project is to be successful, has to reject representationalism, and adopt the alternative metasemantic methodology, namely a global expressivism.

This is, at present, an inductive conclusion. I hope that in the preceding thesis I have managed to argue it convincingly, despite the limited data set provided, by demonstrating at each point how the relevant problems arise. However, such a conclusion will remain tentative unless we can give some sort of general account of *why* representationalism is incompatible with the neo-Carnapian project. This is a tricky question to address, but I will attempt to begin to do so here.

The neo-Carnapian approach to deflationism states that there is no objectively correct metaphysical theory, such as the activity of metaphysics is usually understood to be pursuing. That is, we can gain e.g.: ontological insights “easily” in the course of the ordinary use of our language, but metaphysical theories that go beyond these easily obtained facts are in error, and the debates between them pseudo-debates. Hence there can be no coherent conception of such a metaphysical theory being “correct”, as such theories cannot meaningfully be contrasted.

This fact makes the neo-Carnapian approach inimical to the key assumptions of representationalism. Representationalism assumes that description is the paradigm for language use, and that the key metric for a language’s success is therefore descriptive accuracy, in a metaphysical sense. This last part is the point of disagreement; the representationalist holds that metaphysical theories are capable of being more accurate or less accurate depending on how they represent the world as being, and how similar this representation is to how the world actually is.

For instance, on the representationalist paradigm, it is possible to have a language whose terms uniformly fail to refer, despite their having the linguistic form and function of referring terms and being warranted as successfully referring by the language's reference-establishing procedures, even in an ideal state of full information. A representationalist theory of reference will want to point to a certain sort of substantive relation, *R*, such that a term 't' refers to something in the world, *t*, if and only if instances of 't' stand in *R* to *t*. The language's success at referring will then be judged based on how closely its procedures for establishing reference track *R*, because this will determine how many of its purported referring terms actually stand in *R* to some object. We can replace *R* with a relatively broad class of relations, the better to account for different ways of referring to different sorts of entities, for instance. However, the relational analysis here means that it has to be possible for systematic error to occur, i.e.: where our linguistic tools for determining reference assure us that a certain sort of expression stands in *R* to some sort of entity, where this is not in fact the case.

Essentially, this amounts to the idea that there are criteria for the correctness of a metaphysical theory, including those embedded in languages, that go beyond the criteria available to ordinary language users. Again, such criteria can be as simple or underspecified as the bald requirement that the theory must reflect the actual metaphysical structure of the world. Nevertheless, they must be lurking somewhere in the background in order, for instance, for metaphysical debates to be construed as real disagreements, with a hope of each side being either correct or incorrect in their assertions.

Neo-Carnapian deflationism is committed to the rejection of this representationalist idea of metaphysical correctness. In Carnap's own terms, the question which theory or language-form accurately reflects the structure of the world as it really stands is perhaps the archetypal "external" question. Hence the Carnapian must reject the idea of such a criterion of correctness, however it is outlined. This applies, *mutatis mutandis*, for the neo-Carnapian variations we have investigated above; Hirsch, by way of a clear example, cannot make sense of the idea of the candidate languages to be compared for their practical usefulness being potentially more or less accurate to the world, lest his deflationary program collapse. It follows that the (neo-)Carnapian is committed to there being no substantive, relational account of semantic properties, such as truth and reference, that generalizes across all languages. This is because such an account would create the possibility that one language (or the terms thereof) could satisfy these relations more fully

than another, for example, if its procedures for establishing reference track the actual reference relation in more cases than a contrasting language.

The possibility of such systematic error as representationalism guarantees is incompatible with neo-Carnapianism. We can see this, first, in that it directly contravenes Carnap's Principle of Tolerance, which provides the guiding light for the neo-Carnapian school; representationalism makes languages beholden to an external constraint – of standing in the correct, substantive semantic relations – for their success. Hence the languages from which the neo-Carnapian wants to draw their deflationary metaphysical conclusions, according to representationalism, are not guaranteed to yield those conclusions, as they are not guaranteed to be successful.

Essentially, global expressivism amounts to a deflationism about semantic properties, this being the view that no substantive metaphysical account of semantic properties can be cogently given. This means that, in deriving a metasemantic requirement from the neo-Carnapian's metaphysical deflationism, we are actually taking their methodology and extending it to the domain of metasemantics, showing, in fact, that it has to be so extended if that methodology is to work. This is a similar move to the one made we saw made by Price in Ch.2, in arguing that there is no viable *local* expressivism, and that expressivism, if it is to work at all, has to be extended to all regions of language, including the language of semantics. Indeed, since global expressivism is effectively a metaphysical deflationism extended to semantic properties and relations, the two moves may be identical, and merely presented differently.

There are a few ways, then, of putting the point that representationalism is incompatible with neo-Carnapian deflationism. One is that representationalist criteria for languages violate Tolerance. Another is that, as we have seen previously in chapters 3 and 4, from representationalist assumptions about a particular area of discourse, we can derive substantive metaphysical conclusions that subvert the deflationist's aims. And yet another is to point out that, as with expressivism, a deflationism that is merely local is unstable – the only way to create a deflationism that is stable is to globalise it to all areas, including semantics. Carnapian deflationism in metaphysics thus requires questioning and rejecting our non-deflationary metasemantic assumptions, and pursuing the sort of global expressivism I have laid out here.



## Bibliography:

Ayer, A.J. 1952, *Language, Truth, and Logic*, New York: Dover Publications

Barker, S. 2007, *Global Expressivism*, University of Nottingham (internet publication)

- 2015, "Expressivism About Reference and Quantification Over the Non-Existent Without Meinongian Metaphysics", *Erkenntnis*, Vol 80, pp215-234
- 2020, "Global Expressivism", Ch. 21 in *The Routledge Handbook of Metametaphysics*

Blackburn, S. 2010, *Practical Tortoise Raising: and other philosophical essays*, Oxford University Press

- 2013, "Pragmatism, All or Some?", in Price, H. *Expressivism, Pragmatism and Representationalism*, Cambridge: Cambridge University Press
- 2015, "Blessed are the Peacemakers", *Philosophical Studies*, vol. 172, pp843-853

Boghossian, P. 1989, "The Rule-Following Considerations", *Mind*, vol. 98, no. 392, pp507-549

- 1990, "The Status of Content", *The Philosophical Review*, vol. 99, no. 2, pp157-184

Brandom, R. 1984, "Reference Explained Away", *The Journal of Philosophy*, vol. 81, no. 9, pp469-492

- 1994, *Making It Explicit*, Cambridge MA: Harvard University Press

Brenner, A. 2018, "Easy Ontology, Application Conditions and Infinite Regress", *Analysis*, vol. 78, no. 4, pp605-614

Button, T. 2013, *The Limits of Realism*, Oxford: Oxford University Press

- 2020, "Deflationary metaphysics and ordinary language", *Synthese*, issue 197, pp33-57

Carnap, R. 1928, *The Logical Structure of the World: Pseudoproblems in Philosophy*, London: Routledge and Kegan Paul

- 1937, *The Logical Syntax of Language*, London: Routledge



- 1947, *Meaning and Necessity*, London: University of Chicago Press
- 1950, "Empiricism, Semantics and Ontology", *Revue Internationale de Philosophie*, vol. 4, no.11, pp20-40

Creath, R. 1991, "Every Dogma Has Its Day", *Erkenntnis*, vol. 35, no. 1, pp347-380

Dorr, C. 2005, "What we Disagree about when we Disagree about Ontology", in *Fictionalism in Metaphysics* (ed. Eli, M.), Oxford: Clarendon

Dreier, J. 2018, "The real and the quasi-real: Problems of distinction", *Canadian Journal of Philosophy*, vol. 48, nos. 3-4, pp532-547

Eklund, M. 2006, "Metaontology", *Philosophy Compass*, vol. 1, no. 3, pp317-334

- 2013, "Carnap's Metaontology", *Noûs*, vol. 37, no.2, pp229-249

Fodor, J. 1998, *Concepts: Where Cognitive Science Went Wrong*, Oxford: Oxford University Press

Gettier, E. 1963, "Is Justified True Belief Knowledge?", *Analysis*, vol. 23, no. 6, pp121-123

Golub, C. 2017, "Expressivism and Realist Explanations", *Philosophical Studies*, vol. 174, no. 6, pp1385-1409

Grice, H. P. & Strawson, P. F. 1956, "In Defence of a Dogma", *The Philosophical Review*, vol. 64, no. 2, pp141-158

Hawthorne, J. 2009, "Superficialism in Ontology", in *Metametaphysics* (eds. Chalmers, D. & Wasserman, R.), Oxford: Clarendon

Hirsch, E. 1982, *The Concept of Identity*, New York: Oxford University Press

- *Dividing Reality*, New York: Oxford University Press
- 2011, *Quantifier Variance and Realism: Essays in Metaontology*, Oxford: Oxford University Press

- 2016, "Three Degrees of Carnapian Tolerance", in *Ontology after Carnap* (eds. Blatti, S. and Lapointe, S.), New York: Oxford University Press

Hirsch, E. and Warren, J. 2019, "Quantifier Variance and the Demand for a Semantics", *Philosophy and Phenomenological Research*, vol. 98, no. 3, pp592-605

Horwich, P. 1998a, *Meaning*, Oxford: Oxford University Press

- 1998b, *Truth* 2nd ed., New York: Oxford University Press

Husserl, E. 1970, *The Crisis of European Sciences and Transcendental Phenomenology: An Introduction to Phenomenological Philosophy*, Evanston: Northwestern University Press

Köhler, S. 2018, "Expressivism, meaning, and all that", *Canadian Journal of Philosophy*, vol. 48, nos. 3-4, pp337-356

Korman, D.Z. 2019, "Easy Ontology without Deflationary Metaontology", *Philosophy and Phenomenological Research*, vol. 99, no. 1, pp236-243

Kripke, S. 1980, *Naming and Necessity* (revised and enlarged edition.), Oxford: Blackwell

Legg, C. and Giladi, P. 2018, "Metaphysics – Low in Price, High in Value: A Critique of Global Expressivism", *Transactions of the Charles S. Peirce Society*, vol. 54, no. 1, pp64-83

Leng, M. 2016, "Naturalism and placement, or, what should a good Quinean say about mathematical and moral truth?", *Proceedings of the Aristotelian Society*, vol. 116, no. 3, pp237-260

Macarthur, D. 2010, "Wittgenstein and Expressivism", pp81-95 in Whiting, D. *The Later Wittgenstein on Language*, New York: Palgrave

Miller, A. 2003, "The Significance of Semantic Realism", *Synthese*, vol. 136, no. 2, pp191-217

Millikan, R. 1984, "Language, Thought, and Other Biological Categories: New Foundations for Realism", Cambridge MA: MIT Press

Mitchell, K. 2014, "Pragmatism, Joint-Carving, and Ontology", *The Monist*, vol. 97, no.4, pp571-591

Price, H. 1988, *Facts and the Function of Truth*, Oxford: Blackwell

- 1992, "Metaphysical Pluralism", *The Journal of Philosophy*, vol. 89, no. 8, pp387-409, reprinted in Price 2011 as Ch.2
- 1997, "Naturalism and the Fate of the M-Worlds", *Proceedings of the Aristotelian Society, Supplementary Volumes*, vol. 71, pp247-267, reprinted in Price 2011 as Ch.6
- 2009a, "Metaphysics after Carnap: The ghost who walks?", in *Metametaphysics* (eds. Chalmers, D. & Wasserman, R.), Oxford: Clarendon Press
- 2009b, "The Semantic Foundations of Metaphysics", in *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson* (ed. I. Ravenscroft), reprinted in Price 2011 as Ch.12
- 2011, *Naturalism Without Mirrors*, New York: Oxford University Press
- 2013, *Expressivism, Pragmatism and Representationalism*, Cambridge: Cambridge University Press
- 2015, "From Quasi-Realism to Global Expressivism – And Back Again?", in *Passions and Projections: Themes from the Philosophy of Simon Blackburn* (eds. Johnson, R. and Smith, M.), Oxford: Oxford University Press
- 2019, "Global Expressivism by the Method of Differences", *Royal Institute of Philosophy Supplement*, vol. 86, pp133-154

Putnam, H. 1987, "Truth and Convention: On Davidson's Refutation of Conceptual Relativism", *Dialectica*, vol. 41, no. 1/2, pp69-77

Quine, W. 1951a, "Two Dogmas of Empiricism", *The Philosophical Review*, vol. 60, no. 1, pp20-43

- 1951b, "On Carnap's Views on Ontology", *Philosophical Studies*, vol. 2, no. 5, pp65-72

- 1957, "The Scope and Language of Science", *The British Journal for the Philosophy of Science*, vol. 8, no. 29, pp1-17
- 1960, *Word and Object*, Cambridge MA: MIT Press
- 1969, *Ontological Relativity and Other Essays*, London: Columbia University Press

Richardson, A. 2007, "Carnapian Pragmatism", in *The Cambridge Companion to Carnap* (eds. Friedman, M. & Creath, R.), Cambridge: Cambridge University Press

Ryle, G. 1945, "Knowing That and Knowing How", *Proceedings of the Aristotelian Society*, New Series, vol. 46, pp1-16

Schaffer, J. 2009, "The Deflationary Metaontology of Thomasson's *Ordinary Objects*", *Philosophical Books*, vol. 50, no. 3, pp142-157

Schiffer, S. 2003, *The Things We Mean*, Oxford: Oxford University Press

Sider, T. 2009, "Ontological Realism", in *Metametaphysics* (eds. Chalmers, D. & Wasserman, R.), Oxford: Clarendon

- 2011, *Writing the Book of the World*, New York: Oxford University Press

Simpson, M. 2018, "Solving the problem of creeping minimalism", *Canadian Journal of Philosophy*, vol. 48, nos. 3-4, pp510-531

- 2020a, "Creeping Minimalism and Subject Matter", *Canadian Journal of Philosophy*, vol. 50, no. 6, pp750-766
- 2020b, "What Is Global Expressivism?", *The Philosophical Quarterly*, vol. 8, no. 278, pp140-161

Sosa, E. 2001, "Epistemology and Primitive Truth", in *The Nature of Truth: Classic and Contemporary Perspectives*, ed. M. Lynch, Cambridge MA: MIT Press

Sova, H. 2019, "A Non-substantial Meta-semantics for Global Expressivism", *Acta Analytica* vol. 34, pp505-514

Spelke, E. S., Kestenbaum, R., Simons, D. J., and Wein, D. 1995, "Spatiotemporal continuity, smoothness of motion and object identity in infancy", *British Journal of Developmental Psychology*, vol. 13, pp113-142

Steinberger, F. 2014, "How Tolerant Can You Be? Carnap on Rationality", *Philosophy and Phenomenological Research*, vol. 91, issue 3, pp645-668

Strawson, P.F. 1948, "Truth", *Analysis*, vol. 9, no. 6, pp83-97

Thomasson, A. 2007, *Ordinary Objects*, New York: Oxford University Press

- 2009, "Answerable and Unanswerable Questions", in *Metametaphysics* (eds. Chalmers, D. & Wasserman, R.), Oxford: Clarendon Press
- Thomasson, A. 2014, "Deflationism in Semantics and Metaphysics", in *Metasemantics: New Essays on the Foundations of Meaning* (eds. A. Burgess and B. Sherman), New York: OUP
- 2015, *Ontology Made Easy*, New York: Oxford University Press
- 2017, "Husserl on Essences: A Reconstruction and Rehabilitation", *Grazer Philosophische Studien* vol. 94, pp436-59
- 2019, "What Can Phenomenology Bring to Ontology?", *Res Philosophica*, vol. 96, no. 3, pp1-18
- 2020, *Norms and Necessity*, Oxford: Oxford University Press

Tiefensee, C. 2016, "Inferentialist Metaethics, Bifurcations, and Ontological Commitment", *Philosophical Studies*, vol. 173, pp2437-2459

Topey, B. 2019, "Linguistic convention and worldly fact", *Philosophical Studies*, vol. 176, issue 7, pp1725-1752

Van Inwagen, P. 1990, *Material Beings*, London: Cornell University Press

Warren, J. 2016, "Sider on the epistemology of structure", *Philosophical Studies*, vol. 173, pp2417-2435

Williams, M. 2013, "How Pragmatists can be local expressivists", in Price 2013, *Expressivism, Pragmatism and Representationalism*, Cambridge: Cambridge University Press, pp128-144

- 2015, "Knowledge in Practice", in *Epistemic Evaluation: Purposeful Epistemology* (eds. Henderson, D. K. and Greco, J.), Oxford: Oxford University Press, pp161-185

Wittgenstein, L. 1922, *Tractatus Logico-Philosophicus*, London: Routledge

- 1950, *Philosophical Investigations*, Oxford: Blackwell

Zalabardo, J. 2019, "The Primacy of Practice", *Royal Institute of Philosophy Supplements*, vol. 86, pp181-199