



Through thick and thin: seamless metaconceptualism

Christine Tiefensee¹

Received: 7 March 2022 / Accepted: 26 September 2022 / Published online: 7 February 2023
© The Author(s) 2022

Abstract

One major insight derived from the moral twin earth debate is that evaluative and descriptive terms possess different levels of semantic stability, in that the meanings of the former but not the latter tend to remain constant over significant counterfactual variance in patterns of application. At the same time, it is common in metanormative debate to divide evaluative terms into those that are thin and those that are thick. In this paper, I combine debates about semantic stability and the distinction between the thin and the thick by presenting a new seamless inferentialist account of thin and thick evaluative terms which, despite subsuming them under the same metasemantic analysis, can nevertheless account for their varying levels of semantic stability. According to this position of ‘seamless metaconceptualism’, thin and thick evaluative terms do not belong to different categories, but are both understood as metaconceptual devices which do not differ in kind, but in scope. By providing the same analysis for both thin and thick terms, seamless metaconceptualism not only entails that the latter cannot shoulder the philosophical work that some have attributed to them, but also removes much of their surrounding intrigue.

Keywords Inferentialism · Thick terms and concepts · Semantic stability · Disagreement · Meaning · Moral Twin Earth

1 Introduction

One major insight derived from the moral twin earth debate is that evaluative and non-evaluative descriptive terms appear to possess different degrees of semantic stability. At the same time, it is common in metanormative debate to divide evaluative

✉ Christine Tiefensee
c.tiefensee@fs.de

¹ Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt, Germany

terms into those that are thin and those that are thick. In this paper, I combine debates about semantic stability and the distinction between the thin and the thick by presenting a new seamless inferentialist account of thin and thick evaluative terms which, despite subsuming both under the same analysis, can nevertheless account for their varying levels of semantic stability.

Starting with the moral twin earth debate's focus on semantic stability,¹ let us say that a term is semantically stable to the extent that its meaning tends to remain constant over significant actual or counterfactual variance in patterns of application. Then in relation to non-evaluative descriptive terms, this range of variance appears to be quite strictly limited. To see why, assume that one community of competent speakers consistently applies the term 'fish' only to aquatic, cold-blooded and gill-bearing animals, whereas another additionally applies 'fish' to water-resident mammals. Assume further that a speaker of the second community asserts that whales are fish, whereas a speaker of the first asserts that whales are not fish. Then it seems that the two speakers do not substantively disagree about the nature of whales. Rather, they appear to be talking past one another, in that both communities seem to use the term 'fish' with different meanings, albeit diverging only slightly in their prevailing patterns of application.

In relation to evaluative terms, the range of variance in application patterns that is compatible with sameness of meaning appears to be considerably wider. For instance, assume that one community of competent speakers consistently applies 'morally right' only to actions which are based on universalisable maxims, whereas another applies this term only to actions which are optimific. Assume further that a speaker of the second community asserts that sacrificing the few for the sake of the many is morally right, whereas a speaker of the first asserts that doing so is not morally right. Then it does not seem as if the speakers' dispute were merely verbal; rather, they appear to be locked in substantive moral disagreement. But if so, both communities must use 'morally right' with the same meaning despite significantly diverging in their prevailing applications of this term.²

Both examples seek to elicit intuitions about sameness of meaning on grounds of intuitions about the existence of substantive disagreement. Assume that this 'disagreement test' and the semantic intuitions it invokes are defensible. Then they show that while even relatively slight variance in a descriptive term's application patterns tends to signal variance in meaning, it does not do so in the case of evaluative terms. Consequently, evaluative terms are semantically stable to an extent that descriptive terms are not.

¹ The literature on the moral twin earth argument is extensive. For its original formulation, see Horgan/Timmons (1991, 1992). For further examples of the semantic intuitions that it seeks to elicit and their critical discussion, see Merli (2002), McPherson (2013), Dowell (2016), Dunaway/McPherson (2016), Väyrynen (2018), Köhler (2021). Here, I modify Dunaway/McPherson's (2016: 641) understanding of semantic stability.

² There are forms of (moral) disagreement which do not necessitate sameness of meaning, such as metalinguistic negotiations (see Plunkett/Sundell (2013), Väyrynen (2013: 68–69)). Here, I understand disagreement in terms of "canonical disputes", which concern "disagreement over the truth or correctness of literally expressed content" (Plunkett/Sundell 2013: 7). In contrast to metalinguistic negotiations, canonical disputes do presuppose sameness of meaning.

So far, these considerations rehearse central observations of the moral twin earth debate. Next, let us add another layer to them by adopting a more nuanced approach to the evaluative. For, as a quick glance at metanormative debate shows, many agree that evaluative notions should not be understood as forming one uniform set, but that these notions divide into those that are thin and those that are thick. Amongst the thin, we find terms such as ‘good’, ‘bad’, ‘ought’, ‘desirable’ and ‘morally objectionable’. Amongst the thick are terms such as ‘brave’, ‘rude’, ‘frugal’ and ‘cowardly’. How exactly to conceive of the difference between the thin and the thick, in turn, is a matter of dispute. Some take thin and thick terms to differ in kind. Others suggest that they differ only in degree, such that evaluative terms should be understood as being arranged along a continuum, with thin terms such as ‘morally good’ and ‘morally desirable’ being located on one end, thick terms such as ‘cruel’ and ‘selfish’ on the other, and thicker terms such as ‘just’ and ‘autonomous’ somewhere in-between.³

No matter which of these positions is adopted, we can see that once the more fine-grained distinction between the thin and the thick is in view, it also invites us to develop a subtler approach to semantic stability. For, just as evaluative terms differ in thickness, they also appear to vary in degrees of semantic stability. We have already said that in relation to thin evaluative terms such as ‘morally right’, even extensive variance in application patterns does not necessarily signal variance in meaning. However, in relation to thick terms, this picture starts to change. To elaborate, assume that one community of competent speakers consistently applies the term ‘courageous’ only to people who stand up against moral wrongs even if doing so exposes them to harm, whereas another applies ‘courageous’ only to people who show unwavering support for their friends and family even if this risks being to their own detriment. Assume further that a speaker of the first community asserts that someone defending a stranger against a group of bullies was courageous, whereas a speaker of the second asserts that this person was not courageous. Then it seems that the two speakers are not morally disagreeing. Rather, they appear to be talking past one another, with the first meaning something along the lines of ‘brave’ when using ‘courageous’ and the second meaning something akin to ‘loyal’.

Again, assume that the disagreement test and its elicited intuitions are on the right track. Then we can see that in relation to thick evaluative terms, the range of variability in application patterns that is compatible with sameness of meaning seems to be much smaller than in the case of thin evaluative terms. But if so, thick evaluative terms appear to be semantically less stable than thin evaluative terms.⁴ The question now is: how come?

The standard view of thin and thick terms appears to have an answer ready at hand. According to this account, thin and thick evaluative terms differ in that the former have only evaluative meanings, whereas the latter combine evaluative and non-

³ For the former view, compare Williams (1985: 151–154) and Hare (1963: 24); for the latter, compare Scheffler (1987: 417–418) and Väyrynen (2013: 6–7). For discussions of this matter, compare Väyrynen (2021: § 5.2) and Kyle (2016: § 4). Väyrynen (2008: 391) takes the ‘continuum view’ to be the standard position.

⁴ Or, against the background of the continuum view: The thicker an evaluative term is, the less stable it appears to be.

evaluative, descriptive meanings.⁵ At the same time, we have already said that terms with descriptive meanings are semantically less stable than thin evaluative terms. But if so, proponents of the standard view can explain that since thick terms comprise evaluative and descriptive meanings, and since descriptive meanings are semantically less stable, it is their descriptive component which explains why thick evaluative terms come with a lower degree of semantic stability than thin evaluative terms.⁶

The standard view thus promises to explain thin and thick evaluative terms' different levels of semantic stability by offering different analyses of these terms. In this paper, I will reject this approach by presenting an alternative account which, based on an inferentialist metasemantics, does not treat thin and thick evaluative terms differently, but which provides the exact same inferentialist analysis for both. I call this account 'seamless metaconceptualism'. According to this position, thick evaluative terms do not combine separate evaluative and descriptive components. Nor do thin and thick evaluative terms belong to different categories. Rather, both thin and thick evaluative terms are to be understood as metaconceptual devices which do not differ in kind or degree, but in scope. Presenting the outlines of seamless metaconceptualism will be my first objective in this paper.

Seamless metaconceptualism has several important implications. Firstly, by providing the same analysis for both thin and thick terms, it entails that the latter cannot pose any distinctive puzzles *qua* thick terms. To elaborate, stimulated by the standard view sketched above, much of the attention and intrigue surrounding thick terms has been driven by the question of how exactly these terms combine evaluation and description. For instance, can their evaluative and descriptive components be separated? If they are separable, are thick terms semantically reducible to a combination of thin evaluation and nonevaluative description? Does the grasp of thick terms presuppose a prior grasp of thin terms? Is evaluation semantically built into thick terms, or do they convey evaluation only pragmatically? Which role, if any, does evaluation play in fixing the descriptive component of thick terms?⁷ By treating thick terms in exactly the same way as thin terms, seamless metaconceptualism removes this intrigue. It does so, firstly, by showing that accounting for the differences between thin and thick terms does not require us to subsume these terms under different analyses; secondly, by closing most of the questions triggered by the standard view through treating thin and thick terms alike; and, finally, by entailing that if questions still remain open, these questions will have nothing specifically to do with thickness.

Secondly, seamless metaconceptualism implies that thick evaluative terms cannot shoulder the philosophical work that some attribute to them. More specifically, since thick terms receive the same analysis as thin terms, their existence cannot be used against metaethical expressivism. Nor can it be invoked so as to erode the distinc-

⁵ As a small sample of characterisations of this standard view, see Eklund (2011: 25), Väyrynen (2013: 2), Roberts (2013: 677), Kyle (2016) and Kirchin (2017: 1–2).

⁶ This is too quick, as this answer does not address the more fundamental question of why descriptive meanings are semantically less stable than evaluative meanings. For a comparative study on the semantic stability of descriptive and evaluative terms, see Köhler (2021).

⁷ For excellent overviews of the debate about the thin and thick, see Väyrynen (2013: ch. 1), Kyle (2016), Kirchin (2017: part I) and Väyrynen (2021).

tion between fact and value or undermine the objectivity of morality, say.⁸ Put more bluntly: according to seamless metaconceptualism, differentiating between the thin and the thick is philosophically not particularly significant.⁹

Thirdly, since seamless metaconceptualism denies that thick evaluative terms include a descriptive component, it can no longer appeal to such a component in order to explain these terms' lower levels of semantic stability. Instead, it generates a new puzzle, namely: If thin and thick evaluative terms are to be given the same inferentialist analysis, how can it be that some of these terms are semantically more stable than others? The second objective of this paper will be to provide a solution to this puzzle. In a nutshell, my aim will thus be to present a seamless metaconceptualist account of both thin and thick evaluative terms which, despite subsuming these terms under the same metasemantic analysis, can nevertheless account for their varying levels of semantic stability.

The paper is structured as follows. Section 2 provides a synopsis of those inferentialist tenets that are key for my purposes. Section 3 explains how both thin and thick evaluative terms are to be embedded within a seamless inferentialist framework by providing a metaconceptualist analysis for both. Once this metaconceptual approach is in place, we will see why these terms' varying degrees of semantic stability are no longer shrouded in mystery but instead follow straightforwardly. A brief discussion of seamless metaconceptualism's merits will round off the paper in Section 4.

Before I start, though, a few 'housekeeping' remarks are in order. Firstly, I will adopt a broad construal of the term 'evaluative', according to which the evaluative is to contrast with the descriptive and not with the normative or deontic. Since my interest lies with thin and thick evaluative terms, I will ignore descriptive terms from hereon. Secondly, my focus will be firmly fixed on thin and thick evaluative terms that fall within the practical domain, rather than domains such as aesthetics and epistemology. Throughout this paper, 'normativity' will thus be used as shorthand for 'practical normativity'. Thirdly, it may have been noticed that my discussion has so far been marked by some ambiguity, in that talking about stability in *meaning* can be read in one of two ways. According to the first, accounting for different degrees of semantic stability requires us to explain why speakers who significantly diverge in their dispositions to apply a thin term such as 'morally right' still count as *competently* encoding the same concept MORALLY RIGHT, whereas competence with a thick term such as 'courageous' tolerates considerably less variability in speakers' application dispositions. According to the second, accounting for varying levels of semantic stability demands that we explain why the *extension* of thin concepts such as MORALLY RIGHT tends to remain constant over variance in patterns of competent concept application, while such variance tends to indicate shifts in the extension of

⁸ To give but one example, see McDowell (1981: 201-2) for the disentangling argument against expressivism. For thick terms' alleged impact on the fact/value dichotomy, see Putnam (1981: 127, 128); for their alleged bearing on objectivity, see Williams (1985).

⁹ As a seamless metaconceptualist, I thus join Väyrynen (2013: 10) in this assessment. In contrast to Väyrynen, though, I reach this conclusion not by locating evaluation outside the meaning of thick terms, but by conceiving of thick terms in exactly the same way as of thin terms.

thick concepts such as COURAGEOUS.¹⁰ To use standard terminology, the first reading thus asks for provision of a theory of semantic competence, whereas the second requires presentation of a theory of reference determination.¹¹ Although both tasks are closely related, they should not be conflated. One reason why is that using a term *competently* does not entail applying it *correctly*. More specifically, assume that one speaker asserts that only optimistic actions are morally right, whereas another asserts that only good-willed actions are morally right. Then although both speakers may very well be equally competent with the term ‘morally right’, in that neither of them makes a semantic mistake in their respective uses of this term, this does not entail that both or, indeed, either of them is also guaranteed to be saying something true. Put differently, speakers can be semantically competent but morally mistaken. While a comprehensive account of meaning must, therefore, comprise both a theory of semantic competence and a theory of reference determination, here I will limit myself to spelling out the former while leaving development of the latter for another paper. Finally, as this previous comment indicates, this contribution should be understood as being programmatic in nature. More precisely, my aim in this paper will be to put seamless metaconceptualism ‘on the table’ by starting to sketch, in relatively broad brushstrokes, its basic contours while leaving the refinement of its features and further components for another occasion. This said, it is high time for the first stroke of the brush.

2 Background: Inferentialist metasemantics

According to the inferentialist metasemantics I favour, all normative terms—deontic and evaluative, thick and thin—are metaconceptual devices.¹² In relation to normative terms that fall within the practical domain, I further take these metaconceptual devices to operate on language exit transitions. To gain a better grasp of this inferentialist approach, let me say a little more about the distinction between conceptual and metaconceptual vocabularies on the one hand and different kinds of language moves on the other.

According to inferentialism, the meaning of a sentence *S* is determined by the rules of inference that govern it.¹³ On the one hand, these rules concern commitments that can be correctly inferred from *S*. Call these outputs of *S*’s correct assertion its *downstream* consequences. If these downstream consequences concern further com-

¹⁰ Compare Manley (Ms.) and Dunaway/McPherson (2016: 640, 641), with the latter explaining that “a term is semantically stable to the extent that its referent tends to remain unchanged over counterfactual variance in semantically significant properties of the term.”

¹¹ For instance, see Schroeter/Schroeter (2003: 192) for this terminology.

¹² Brandom (2008: 89) might agree but does not generally distinguish between evaluative and deontic notions. In contrast, Chrisman (2016: 214) defends inferentialist metaconceptualism about deontic notions but hints at advocating a form of descriptivism about evaluative terms. However, also compare his (2018). See also footnote 15.

¹³ Compare Murzi/Steinberger (2017) for an excellent short introduction into inferentialism more generally and Brandom’s view more particularly. For the following summary, see Brandom (2000: 63, 194; 2008: 46), Chrisman (2016: 203; 2017: 269) and Sellars (2007).

mitments that we undertake within the language game, we will be making an *intra-linguistic transition*, say by inferring the claim that Charles is William's father from the claim that William is Charles' son. If, in turn, these downstream consequences concern practical commitments, we will be making a *language exit transition*, say by inferring the practical commitment 'I shall babysit the boys' from the claim that Charles needs help with the children.¹⁴

At the same time, asserting *S* not only commits us to *S*'s downstream consequences, but also exposes us to the potential demand of showing our entitlement to *S*. On the other hand, then, the inferential rules governing *S* concern the sentences and extra-linguistic circumstances from which *S* can properly be inferred. Call the inputs or antecedents of *S*'s correct assertion *S*'s *upstream* antecedents. If these upstream antecedents concern further commitments that we have undertaken within the language game, we will again be making an intra-linguistic transition, say by explaining that William is Elizabeth's grandson because Elizabeth is Charles' mother. If, in turn, they concern extra-linguistic circumstances, we will be making a *language entry transition*, through which we respond to our non-linguistic environment by undertaking commitments within the language game, say by asserting 'There are Charles and William!' upon spotting them in the playground.

Some of these inferential transitions will be '*committive*' or 'commitment-preserving', whereas others will be '*permissive*' or 'entitlement-preserving'. Committive inferences are those where commitment to their respective antecedents also entails commitment to their consequents. For instance, if I am committed to William being older than Harry, I must also commit to Harry being younger than William. Entitlement-preserving inferences are those where commitment to their respective antecedents permits, but does not oblige us to commit ourselves to their consequents. If I am committed to William being at least as old as Harry, I am permitted to hold that William and Harry are the same age, but am also permitted to hold that William is older than Harry.

In order to make assertions, we now need two main kinds of vocabularies. On the one hand, we require terms that enable us to make claims about the world and how to act in it. These are *first-order conceptual* vocabularies. 'I saw Charles and William', 'William is older than Harry' and 'I shall help Charles look after the boys' are all typical examples of first-order conceptual statements. On the other hand, we need vocabularies that enable us to talk about and make explicit the implicit inferential commitment/entitlement structures in which assertions and linguistic expressions stand. These are *second-order or metaconceptual* vocabularies. 'If William is Harry's older brother, then William and Harry have the same parents' is a paradigm example of a statement using metaconceptual vocabulary. By allowing us to articulate conditionals such as this, logical vocabulary is not used to assert additional claims about the world and how to act in it, but to explicate the inferential relations within which various statements stand.

¹⁴ Here, 'inference' is given a broad reading, which includes both transitions between contentful claims and transitions between claims and non-linguistic items, such as objects and actions. I follow Brandom (2000: 85) in using 'shall' to "express the significance of the conclusion as the acknowledging of a practical commitment. ('Will' would be used correspondingly to express a doxastic commitment to a prediction.)".

I have already stated that following seamless metaconceptualism, both thin and thick concepts are metaconceptual devices that neither differ in kind nor degree, but in scope. Let me explain next what this account involves.

3 Seamless metaconceptualism: through thick and thin

Most inferentialists agree that moral vocabulary functions as a metaconceptual device that operates on language exit moves.¹⁵ At the same, what exactly this metaconceptual role involves is a matter of some dispute.

According to Brandom (2000: 89), normative “vocabulary plays the same expressive role on the *practical* side that *conditionals* do on the *theoretical* side”, in that a sentence such as ‘I ought to help Charles’ makes explicit endorsement of the material propriety of transitioning from the commitment that Charles needs help to the practical commitment that I shall help Charles. Following Chrisman (2017: 269), ‘morally ought’ is a metaconceptual device that alters the evaluation base of embedded sentences, in that adding ‘morally ought’ to a sentence such as ‘I help Charles’ shifts the focus of this sentence’s evaluation from the actual to morally best worlds. Moving from deontic to evaluative terms, I (2021: 126) have suggested that the metaconceptual function of moral evaluative terms concerns the systematisation of legitimate language exit transitions, where based on the method of reflective equilibrium, these systematisations ground explanations of why transitions such as ‘Charles needs help with the children, so I shall help Charles’ are legitimate.

In an attempt to keep matters as simple as possible, I will gloss over many of these nuanced disagreements and also ignore important differences between deontic and evaluative vocabularies. Instead, I will limit my focus to spelling out seamless metaconceptualism’s core idea, which consists in combining Brandom’s thoughts on the explication of inferential proprieties with a reinforced focus on patterns, systematisations or ‘bundles’ of inferential relations.¹⁶ Given my interests in this paper, I will do so by measuring my account against the following three criteria:

<i>[Individuation]:</i>	Seamless metaconceptualism must specify how to <i>individuate meaning</i> by telling us when two terms t_1 and t_2 encode the same concept C and when they fail to do so.
<i>[Specificity]:</i>	Seamless metaconceptualism’s account of concept individuation must be sufficiently <i>specific</i> so as to distinguish between the contents of different moral and non-moral normative concepts, such as MORALLY RIGHT, JUST, COURAGEOUS and INSTRUMENTALLY RATIONAL, say.

¹⁵ Many inferentialists do not explicitly distinguish between deontic and evaluative vocabularies. Chrisman (2016), who does make this distinction, arguably defends metaconceptualism about the former but not the latter. At the same time, his account of ‘ought’ is not limited to language exit transitions.

¹⁶ We can find focus on patterns very clearly in Brandom’s thinking. Compare: “The idea is that ... broadly normative or evaluative vocabulary ... is used to make explicit in assertible, propositional form the endorsement of a pattern of material practical inferences” (Brandom 2000: 89–90; see also Chap. 2, sec. V). My aim here is to bring patterns even more strongly to the fore.

[*Stability*]: Seamless metaconceptualism's account of concept individuation must account for varying levels of *semantic stability* by explaining why thin terms such as 'morally right' are semantically more stable than thick terms such as 'courageous'.

Let us look at these in turn and start with [*Individuation*].

3.1 [*Individuation*]: meanings and functions

In relation to [*Individuation*], inferentialists already have the outlines of an answer at hand, namely: Two terms t_1 and t_2 encode the same concept C iff t_1 and t_2 stand in the same material, meaning-constituting inferential relations.¹⁷ The insertion of 'meaning-constituting' is important. For arguably, not all inferences in which a term features are relevant for its meaning.¹⁸ How to identify the privileged subset of meaning-constituting inferential relations, in turn, is open to debate. The proposal I favour is one which explicitly links meaning-constitution with a term's *function*, namely by explaining that only those inferential relations are constitutive of its meaning which are required for the term to realise its specific function (Köhler, 2021: § 4, 630; Sciaraffa, Ms.). To give a toy example, take the descriptive colour term 'blue' and assume that the function of descriptive vocabularies is to enable us to track and communicate about our environment. Assume further that in response to looking at my blue jumper, I assert the sentence 'This jumper is blue'. Assume finally that since blue is my favourite colour, I also move from my commitment to the jumper's being blue to the practical commitment 'I shall wear this blue jumper'. In this toy case, the term 'blue' features both in a language entry transition and a language exit transition. However, only the language entry transition is required to fulfil the colour term's function to keep track of our environment. Accordingly, while the language entry transition is relevant for determining the meaning of 'blue', the language exit transition is semantically irrelevant.

This response doubtless requires further refinement. For instance, we need to clarify how the notion of 'function' is to be understood so as to ensure that functions can fulfil the role that this account assigns to them.¹⁹ Moreover, stricter constraints must be imposed on meaning-constitution so as to distinguish between the meanings of terms whose meaning-constituting inferences partly overlap (sticking to our toy example, compare the meaning-constituting inferences of 'blue' with those of 'coloured').²⁰ The basic idea, though, should hopefully be sufficiently clear: Two terms t_1 and t_2 will encode the same concept C iff t_1 and t_2 stand in the same meaning-constituting material inferences, where only those material inferences are meaning-constituting which are required for t_1 and t_2 to realise their specific functions.

¹⁷ Analogously, two concepts C_1 and C_2 have the same content iff C_1 and C_2 stand in the same material, meaning-constituting inferential relations.

¹⁸ Brandom (1994: 485) appears to disagree by adopting a holistic approach to meaning-constitution, which holds that meaningful communication requires us to "navigate and traverse differences in points of view, to specify contents from different points of view." See Köhler (2021: 621–622) for discussion.

¹⁹ For more on this question, see Köhler (2022) and Sciaraffa (Ms).

²⁰ For instance, compare Wedgwood's (2007: 87) completeness constraint in this context.

As these considerations show, the inferentialist account of concept individuation puts functions centre-stage. Given as much, it neatly sets up our response to [*Specificity*]. For, just as we have appealed to a vocabulary's function in order to deal with [*Individuation*], we can exploit observations about moral vocabulary's *specific* function in order to tackle [*Specificity*]. Let me explain.

3.2 [*Specificity*]: seamless metaconceptualism

So far, we have examined how to individuate meaning more generally by considering when two terms t_1 and t_2 encode the same concept C . In order to meet [*Specificity*], we now need to explain how to individuate normative and moral concepts more specifically.

In order to discharge this task, seamless metaconceptualism pursues the following nested approach. Normative concepts differ from descriptive concepts in that the former are *metaconceptual* whereas the latter are conceptual. Normative concepts differ from other metaconceptual concepts (such as logical concepts) in that they make explicit endorsement of proprieties of language exit transitions. Moral concepts differ from other forms of normative concepts (such as prudential concepts) in regard to the *kind of entitlement* carried by the language exit transitions that they make explicit. Finally, specific moral concepts differ from others in regard to the *antecedents* of the inferences that these respective concepts explicate.

These differences will not be the only contrasts that can be found between these classes of concepts. Yet, they highlight the importance of three kinds of patterns, or 'inferential bundles', for seamless metaconceptualism: firstly, the bundle of language exit transitions; secondly, the bundle of language exit transitions which carry a specific kind of entitlement; and, finally, the bundle of language exit transitions that carry a specific kind of entitlement and share specific antecedents.

The first, most general bundle of language exit transitions helps us distinguish normative concepts from other non-normative metaconceptual concepts which do not (primarily) operate on language exits. In this context, the contrast between normative and logical concepts is specifically intriguing.²¹ However, since examining this contrast would take us too far afield, I will set it aside for now. Instead, let us zoom in on the bundle of language exits and consider how to differentiate between different classes of normative concepts even though all of them share the metaconceptual function of making language exits' inferential proprieties explicit.

For a moral normative concept, let us choose MORALLY RIGHT as our example; for a non-moral normative concept, let us focus on the concept INSTRUMENTALLY RATIONAL.²² We already know that since both MORALLY RIGHT and INSTRUMENTALLY

²¹ Arguably, the most important question in this context is why we need normative concepts to explicate proprieties of language exits if logical concepts could be used for the exact same purpose. For the outlines of an answer, see my thoughts on the usefulness of moral concepts that feature later in this section.

²² It may have been noted that almost all of the terms and concepts I consider in this paper carry some qualification, in that I speak of 'morally right' rather than 'right', 'instrumentally rational' rather than 'rational' and (later) 'distributively just' rather than 'just'. The reason for including these qualifications is that 'instrumentally rational' arguably means something different than 'substantively rational', just as 'distributively just' arguably means something different than 'retributively just'. This, in turn, entails that

RATIONAL are normative, they fulfil the same coarse-grained function of explicating endorsement of language exits' proprieties. Accordingly, if we want to identify differences in their contents despite this shared feature, we need to attain a more fine-grained understanding of their respective functions that reveals how they diverge. Such functional differences could take a number of forms. For instance, two concepts could have the function of making *different* language exit transitions explicit, or different *kinds* of language exits, or different *elements* and *characteristics* thereof, or all of the above. As we will see shortly, seamless metaconceptualism invokes all of these functional variations for different individuation purposes. For the purpose of distinguishing between concepts such as MORALLY RIGHT and INSTRUMENTALLY RATIONAL, it is differences in the entitlement-structures of the explicated language exits that do the individuating work.

To elaborate, let us shift focus from the general bundle of language exit transitions to different sub-bundles of language exits that manifest different kinds of entitlement structures.²³ For instance, take the exit transition 'Learning through the night is the only way to do well in the exam, so I shall stay up and learn all night'. Then on the one hand, we can see that the antecedent of this inference concerns the means to an end. On the other, we find that only those who share the end of wanting to do well in the exam will be entitled to make this inference.²⁴ Whether or not this inferential move preserves entitlement thus depends on whether or not the person making it has a specific end or preference.

Compare this with a language exit transition such as 'Charles is in a real emergency situation, so I shall help him'. The material propriety of this transition does not depend on the person making it having a specific end or preference. Rather, this transition is entitlement-preserving for anyone, no matter who engages in it.

Following seamless metaconceptualism, non-moral normative concepts such as INSTRUMENTALLY RATIONAL bundle inferences of the former kind, whereas moral concepts such as MORALLY RIGHT bundle those of the latter. More specifically, seamless metaconceptualism thus offers the following rough characterisations of MORALLY RIGHT'S and INSTRUMENTALLY RATIONAL'S fine-grained functions:

MORALLY RIGHT	MORALLY RIGHT makes explicit endorsement of language exit transitions which are entitlement-preserving for anyone, irrespective of preferences and social status, and whose antecedents are unrestricted.
---------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

the fine-grained functions of these pairs must also differ in certain respects. I leave investigation of how they differ (such that 'distributively just' and 'retributively just' have different meanings) and how they are alike (such that 'distributively just' and 'retributively just' still refer to different forms of justice) for another occasion.

²³ The following considerations closely follow Brandom (2000: 91, Chap. 2, sec. V). Here, I bracket further potential differences relating to the commitments and entitlements of *speakers* making assertions, such as 'It would be rational for Harry to learn through the night' or 'It was morally right not to tell Charles the truth'. For more on this issue, see Brandom (2000: Chap. 2).

²⁴ We could expand these thoughts further. For instance, if the antecedent of a language exit transition concerns a non-necessary means to some end, we can argue that the transition is entitlement-preserving for those who have the end. If, in turn, its antecedent concerns a necessary means, the transition will not only be entitlement-preserving but also commitment-preserving for those who have the specific end.

INSTRUMENTALLY
RATIONAL

INSTRUMENTALLY RATIONAL makes explicit endorsement of language exit transitions which are entitlement-preserving only for those who have specific preferences, while their antecedents are restricted to means/end considerations.

Hence, although INSTRUMENTALLY RATIONAL and MORALLY RIGHT share the same course-grained function of explicating proprieties of language exits, they have different fine-grained functions in explicating language exits that carry different kinds of entitlement. It is exactly this difference in fine-grained functions that grounds their membership in different classes of normative concepts. More precisely, just as Brandom declares that “[d]ifferent patterns of inference should be understood as corresponding to different sorts of norms or pro-attitudes”, we can declare that different classes of normative concepts should be understood as deriving from different fine-grained functions in making patterns of different kinds of inferential proprieties explicit.

This does not mean that concepts such as INSTRUMENTALLY RATIONAL and MORALLY RIGHT cannot operate on the same language exit transitions. For instance, take the inference ‘Ethical banking is the best way to combat climate change, so I shall move my account to an ethical bank’. Insofar as I take this inference to be entitlement-preserving for anyone, I can explicate its material propriety by asserting that moving my bank account to an ethical bank is morally right. If in addition, I have the desire to combat climate change, I can also explicate that this inference is entitlement-preserving for me by asserting that it is instrumentally rational for me to move my bank account. Hence, what distinguishes moral and non-moral classes of normative concepts is not that these concepts necessarily operate on different inferences, but that they explicate inferences’ different kinds of entitlement-structures, namely preference-dependent entitlement-structures in the case of concepts such as INSTRUMENTALLY RATIONAL and preference-independent entitlement-structures in the case of concepts such as MORALLY RIGHT.

So far, we have explained that normative and non-normative metaconceptual concepts differ because they explicate different *kinds* of inferences, whereas moral and non-moral normative concepts differ because they explicate different kinds of inferential *proprieties*. As a last step, we still need to explain how moral concepts differ from one another. I will do so by suggesting that moral concepts differ because they explicate different *inferences*.

To elaborate, we already know that all moral concepts share the same fine-grained function of making explicit endorsement of preference-independent proprieties of language exit transitions. Consequently, we also know that appeal to different kinds of inferences and types of entitlement-structures can no longer help us distinguish between these concepts’ different contents. Rather, doing so requires us to develop an even finer-grained understanding of their respective functions. In order to arrive at such an understanding, I suggest that we must invoke the *antecedents* of the specific inferences whose proprieties different concepts make explicit.²⁵

²⁵ This is not to deny that for certain moral concepts, such as BLAMEWORTHY or PRAISEWORTHY, inferential consequents rather than antecedents might be more relevant for content individuation.

To see how, let us choose MORALLY RIGHT, DISTRIBUTIVELY JUST and COURAGEOUS as our examples and consider the following proposed specifications of their fine-grained functions:

MORALLY RIGHT	MORALLY RIGHT makes explicit endorsement of language exit transitions which are entitlement-preserving for anyone, irrespective of preferences and social status, and whose antecedents are unrestricted.
DISTRIBUTIVELY JUST	DISTRIBUTIVELY JUST makes explicit endorsement of language exit transitions which are entitlement-preserving for anyone, irrespective of preferences and social status, and whose antecedents concern the properties of distributions of some entity (resources, status, rights, opportunities etc.).
COURAGEOUS	COURAGEOUS makes explicit endorsement of language exit transitions which are entitlement-preserving for anyone, irrespective of preferences and social status, and whose antecedents concern an agent's risking his own harm for the sake of goods greater than the evil of that harm. ²⁶

Next, consider inferences such as 'God ordained this distribution of social statuses, so we shall maintain this distribution', 'Redistributing money maximises utility, so we shall redistribute money', or 'This distribution of resources would benefit the worst-off, so we shall implement this distribution.' Then by having the function to bundle and explicate these inferences, DISTRIBUTIVELY JUST enables us to assert that maintaining or bringing about a specific distribution is just.

Similarly, consider moves such as 'William jumped first despite risking his life, so I shall praise William', 'Elizabeth stood her ground against Andrew despite incurring his wrath, so I shall congratulate her', or 'Harry started a new life despite not knowing what this life would hold, so I shall applaud Harry'. Then by bundling and explicating inferences such as these, COURAGEOUS allows us to assert that William, Elizabeth and Harry are courageous.

Finally, return to MORALLY RIGHT. According to seamless metaconceptualism, the antecedents of inferences whose proprieties MORALLY RIGHT explicates are not restricted.²⁷ Rather, they can concern reactions to adversity and the redistribution of money, just as much as they can pertain to making Charles happy, the universalisability of maxims or moving one's account to an ethical bank, say. Given as much, MORALLY RIGHT can be used to explicate any of these inferences by allowing us to assert that it is morally right to move my bank account, bring about a certain resource distribution or visit Charles, say.

Why are there all these different concepts? Put differently, why do we engage in these various 'bundling exercises'? The short answer is: because it is useful to do so. For instance, how resources, rights and opportunities are distributed is hugely significant for how well or badly we fare in life; DISTRIBUTIVELY JUST is a concept that helps us discuss exactly these distributional questions. The willingness to suffer harm for the sake of the greater good is a character trait that is important for the way we live and prosper together; COURAGEOUS is a concept that helps us talk about this trait. How

²⁶ I adopt this formulation from Elstein & Hurka (2009: 527).

²⁷ This is too radical, as the antecedents of the inferences explicated by MORALLY RIGHT are (at the very least) limited to actions. Given the programmatic nature of this paper, the functional specifications suggested here should thus not be read as the last word.

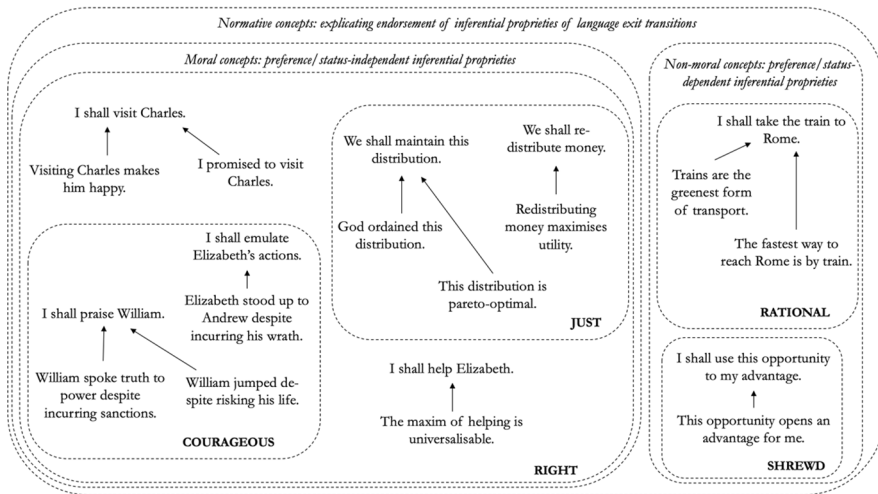


Fig. 1 Seamless metaconceptualism and nested bundles of inferences

an agent acts and makes decisions on the basis of her preferences and commitments is important in interpreting, predicting and responding to her actions; INSTRUMENTALLY RATIONAL is a concept that helps us articulate these matters. In addition, having specific normative concepts to discuss and articulate our different concerns enables us to explicate and talk about the relations within which values stand. For instance, Elizabeth's actions might be just but not generous, William might be courageous but sometimes irrational, taxation might promote justice but reduce freedom, etc.²⁸ Having various normative and moral concepts, rather than just one or two, thus increases the expressive capacities of our language in order to meet the practical needs that arise in the context of living and deliberating together.

Figure 1 above provides a simplified schematic representation of this nested approach to concept individuation, where concepts are printed in small bold capitals, arrows indicate inferential moves from antecedents to practical commitments, and the dotted lines delineate the bundles of inferences on which the respective concepts operate.

Consequently, whereas MORALLY RIGHT operates on all materially good, universally entitlement-preserving language exit transitions, DISTRIBUTIVELY JUST and COURAGEOUS concern only specific subsets thereof. On the one hand, this entails that in contrast to different classes of normative concepts, which can operate on the same inferences while making different entitlement-structures explicit, moral concepts will vary in content only if they operate on different subsets of inferences.²⁹ On the

²⁸ In footnote 21, I mentioned that the usefulness of having these different moral terms also explains why we need moral language in addition to logical vocabulary. To see this point, just try to express the statements made in this sentence by use of logical vocabulary alone.

²⁹ There will be cases where the bundles of inferences on which moral concepts operate are closely related, such as in relation to KIND and FRIENDLY, say. In addition, the boundaries of these inferential bundles are

other hand, it shows why metaconceptualism is seamless. According to the proposal developed here, thin and thick evaluative concepts do not differ in kind. Nor are thin evaluative concepts purely evaluative, whereas thick evaluative concepts combine evaluation and description. Rather, thin and thick moral concepts perform the *same* metaconceptual function on the *same* kind of language transitions manifesting the *same* patterns of entitlements. Their only difference concerns the *range* of inferential antecedents that can feature within the inferences that these concepts explicate. In relation to thin concepts such as MORALLY RIGHT and MORALLY GOOD, this range of antecedent conditions will hardly be restricted at all. In relation to thicker concepts such as DISTRIBUTIVELY JUST and AUTONOMOUS, greater constraints will apply. And in relation to thick concepts such as COURAGEOUS and SELFISH, the range of antecedent conditions will be heavily restricted. This, then, is why thin and thick concepts do not differ in kind, or degree, or types of contents, but only in scope.

Having dealt with [*Individuation*] and [*Specificity*], let us finally tackle [*Stability*] and thus the question of why thin and thick moral terms show different levels of semantic stability, despite receiving the same metaconceptual analysis. Luckily, with [*Individuation*] and [*Specificity*] in place, [*Stability*] follows straightforwardly.

3.3 [*Stability*]: seamless analysis, varying stability

Recall from earlier that [*Stability*] demands that we explain why thin terms such as ‘morally right’ are semantically more stable than thick terms such as ‘courageous’. Recall also that we have further specified this demand as requiring us to explain why speakers who significantly diverge in their dispositions to apply a thin term such as ‘morally right’ still qualify as competently encoding the same concept MORALLY RIGHT, whereas competence with a thick term such as ‘courageous’ tolerates considerably less variability in speakers’ application dispositions. Finally, remember that since seamless metaconceptualism denies that thick concepts possess descriptive content, it cannot appeal to such content in order to explain their relative semantic instability. What could play this role instead? The answer is relatively obvious, namely: varying levels of restrictions that are imposed on the range of eligible inferential antecedent conditions.

To elaborate, start with the thin term ‘morally right’ and return to a case where one speaker applies this term to actions that are optimific, whereas another applies it to actions that are based on universalisable maxims. Then as long as both speakers employ ‘morally right’ so as to explicate their endorsement of universally entitlement-preserving language exits that fall within the specific bundle of inferences on which MORALLY RIGHT operates, they will qualify as competent users of ‘morally right’. The fact that one of these speakers employs ‘morally right’ to explicate her endorsement of inferences such as ‘Lying to Charles is optimific, so I shall lie to Charles’, whereas the other uses this term to explicate his endorsement of inferences such as ‘Only the maxim to tell Charles the truth is universalisable, so I shall tell Charles the truth’ does not affect this result. Nor would this result change if they

most likely to be vague, in that there will be borderline cases where it is indeterminate whether or not a specific inference falls within a specific bundle.

employed ‘morally right’ to explicate endorsement of inferences such as ‘Today is Tuesday, so I shall lie to Charles’ or ‘I clasped my hands three times, so I shall tell Charles the truth’. Rather, since the input conditions of transitions on which MORALLY RIGHT operates are not restricted, speakers can employ ‘morally right’ to explicate their endorsement of inferences featuring a vast range of extremely varied antecedents and yet qualify as competently encoding the same concept MORALLY RIGHT. Put in a nutshell, ‘morally right’ is thus semantically highly stable because the range of inferential antecedents is hardly restricted.

This story starts to change once we turn to thicker concepts, such as DISTRIBUTIVELY JUST.³⁰ To see why, assume that one speaker applies ‘distributively just’ to resource distributions that maximise the position of the worst-off, whereas another applies ‘distributively just’ to distributions that have been ordained by God. Then on the one hand, we can echo the explanations just given by explaining that as long as both speakers use ‘distributively just’ in line with its metaconceptual function of making explicit endorsement of universally entitlement-preserving language exit transitions whose antecedents concern the properties of distributions, they will competently encode the same concept DISTRIBUTIVELY JUST. And again, we can point out that this result will remain unchanged over significant variations in inferential antecedents. For instance, besides concerning properties such as *being ordained by God* and *benefitting the worst-off*, these antecedents could concern properties such as *being distributed equally*, *being the result of a coin toss*, *tracking hard work*, *being the result of free transactions* or *having come to me in a dream*, to name but a few.

On the other hand, though, there are limits. If a speaker uses ‘distributively just’ to explicate her endorsement of the inference ‘Charles is in an emergency situation, so I shall help him’, she will be guilty of misapplying this term. Similarly, if she employs ‘distributively just’ to explicate her endorsement of the transition ‘William spoke truth to power, so I shall praise him’ or ‘Harry did not hurt me intentionally, so I shall forgive him’, she would display incompetence with its meaning. More generally, then, we can see that since the range of input conditions of the inferences explicated by the term ‘distributively just’ is restricted to antecedents that concern distributional properties, competence with its meaning (and thus with the concept DISTRIBUTIVELY JUST) will tolerate only those variations in application dispositions that fall within this restricted range. Variations that fall outside of this range, in turn, will signal semantic incompetence. Consequently, although a term such as ‘distributively just’ is semantically relatively stable, its semantic stability is nevertheless limited. More precisely, it is as limited as the range of antecedents that can feature in the inferences which ‘distributively just’ has to the function to explicate.

By now, it should be clear how to account for the relative semantic instability of thick terms. We have already explained that the range of inferential inputs explicated by thick terms such as ‘courageous’ is tightly restricted. Consequently, the boundaries within which variations in application dispositions of ‘courageous’ are compatible with competent language use will be just as restricted, such that even relatively minor deviations will amount to misuses of this term. For instance, this will include varia-

³⁰ Here, I follow advocates of the continuum view in agreeing that evaluative terms can vary in thickness. Compare footnotes 3 and 4.

tions excluding inferential antecedents which concern people who expose themselves to harm for the sake of strangers. It will also include variations involving inferential antecedents which concern people who are willing to stand up against moral wrongs, but only if the risks involved in doing so are practically non-existent. In short, thick terms such as ‘courageous’ are thus semantically highly unstable because the range of eligible inferential antecedents is tightly restricted.

Hence, we can see that semantic stability and the restrictions imposed on the range of eligible inferential antecedents are inversely correlated: the more restricted the range of inferential antecedents is, the less stable the corresponding terms and concepts will be. Since thin and thick terms and concepts are distinguished on the basis of the restrictions imposed on eligible inferential antecedents, this also entails that the thicker a term or concept is, the more unstable it will be. As a result, thin and thick terms and concepts will vary in semantic stability despite fulfilling the same metaconceptual function on the same kind of universally entitlement-preserving language exit transitions. Consequently, this paper’s mission is accomplished: we have explained why thin and thick moral terms and concepts possess different levels of semantic stability despite receiving the same metaconceptual analysis.

4 Through thick and thin, for better or for worse?

As these considerations show, seamless metaconceptualism’s appeal to nested inferential patterns that bundle antecedent conditions of universally entitlement-preserving language exit transitions serves various purposes. We have met three of them above:

Firstly, appeal to nested inferential patterns enables us to individuate moral terms and concepts by distinguishing between their different fine-grained contents without taking thick terms and concepts to possess separate evaluative and descriptive components. Rather, since thin and thick moral terms and concepts differ only in scope but not in kind, both are exclusively evaluative. Secondly, focus on the bundling of inferential antecedent conditions allows us to account for the intuition that thickness and semantic stability do not describe binary phenomena, but are arranged along a genuine continuum. Thirdly, taking evaluative concepts to operate on inferential bundles that vary in scope enables us to explain why thin and thick terms and concepts vary in semantic stability, despite receiving the same metaconceptual analysis. Let me add a fourth purpose here, namely: by taking thin and thick terms and concepts to explicate bundles of inferential relations that are nested within one another, seamless metaconceptualism also allows us to account for connections between these moral terms and concepts, in that it straightforwardly explains why being kind is also good, why being just is also right and why being selfish is also bad, say.³¹

³¹ Arguably, this is not where seamless metaconceptualism’s advantages end. To give the slightest of hints here, it arguably also promises to cater for more specific linguistic data about thick terms and concepts. For instance, as Pekka Väyrynen (2013: 62) has argued, negations such as ‘Demonstrating for free speech within well-functioning liberal democracies is not courageous’ are not naturally heard as denying that being courageous is morally good, but as denying that engaging in such demonstrations is sufficient to qualify as courageous. Väyrynen argues that accounts which do not locate evaluation within the semantic

Needless to say, these advantages will need to be spelt out more thoroughly. Still, we can see that its explanatory potential makes seamless metaconceptualism a very attractive position. Hence, having put this metaconceptualist account on the table, we can now explore its potential in full.

Acknowledgements I am grateful to Pekka Väyrynen, Sebastian Köhler and Irina Schumski for extremely helpful discussions and written comments on an earlier draft. I also thank audiences in Leeds, Amsterdam, Innsbruck, Berlin and Bonn for their helpful feedback.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Brandom, B. (1994). *Making it explicit*. Harvard University Press
- Brandom, B. (2000). *Articulating reasons*. Harvard University Press
- Brandom, B. (2008). *Between saying and doing*. Oxford University Press
- Chrisman, M. (2016). *The meaning of 'ought'*. Oxford University Press
- Chrisman, M. (2017). Conceptual role accounts of meaning in metaethics. In T. McPherson, & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 260–274). Routledge
- Chrisman, M. (2018). Two nondescriptivist views of normative and evaluative statements. *Canadian Journal of Philosophy*, 48, 405–424
- Dowell, J. (2016). The metaethical insignificance of moral twin earth. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics, Vol. 11* (pp. 1–27). Oxford University Press
- Dunaway, B., & McPherson, T. (2016). Reference magnetism as a solution to the moral twin earth problem. *Ergo*, 3, 639–679
- Eklund, M. (2011). What are thick concepts? *Canadian Journal of Philosophy*, 41, 25–49
- Elstein, D., & Hurka, T. (2009). From thick to thin: Two moral reduction plans. *Canadian Journal of Philosophy*, 39, 515–36
- Hare. (1963). *Freedom and reason*. Clarendon Press
- Horgan, T., & Timmons, M. (1991). New wave moral realism meets moral twin earth. *Journal of Philosophical Research*, 16, 447–465
- Horgan, T., & Timmons, M. (1992). Troubles for new wave moral semantics: The 'open question argument' revived. *Philosophical Papers*, 21, 153–175
- Kirchin, S. (2017). *Thick evaluation*. Oxford University Press
- Köhler, S. (2021). Normative disagreement: A functional account for inferentialists. *Philosophical Studies*, 178, 617–637

content of thick terms are better suited to explain this linguistic datum than semantic construals of thick terms. However, by providing a fine-grained differentiation between those 'courageous'-inferences that the speaker of this negation endorses and those that she rejects, seamless metaconceptualism arguably has the resources at hand to provide a convincing explanation of this linguistic datum despite defending a semantic account of thick terms. Since showing as much would go beyond the scope of this paper, turning this mere hint at an explanation into a fully worked-out account must be left for another occasion.

- Köhler, S. (2022). What is (neo-)pragmatists' function? *Australasian Journal of Philosophy*. <https://doi.org/10.1080/00048402.2022.2034904>
- Kyle, B. G. (2016). Thick concepts. *Internet encyclopedia of philosophy*, <https://philpapers.org/archive/KYLTC.pdf>
- Manley, D. (Ms.). Moral realism and semantic plasticity. (Ed.). unpublished manuscript, http://www-personal.umich.edu/~dmanley/manley/David_Manley_files/moral+realism+and+semantic+plasticity.pdf
- McDowell, J. (1981). Non-Cognitivism and rule-following. In his (1998) *Mind, value and reality*. Harvard University Press: 198–219
- McPherson, T. (2013). Semantic challenges to normative realism. *Philosophy Compass*, 8, 126–136
- Merli, D. (2002). Return to moral twin earth. *Canadian Journal of Philosophy*, 32, 207–40
- Plunkett, D., & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13, 1–37
- Murzi, J., & Steinberger, F. (2017). Inferentialism. In B. Hale et al. (eds.) *A companion to the philosophy of language, 2nd Edition, Vol. 1*. Wiley Blackwell: 197–224
- Putnam, H. (1981). *Reason, truth, and history*. Cambridge University Press
- Roberts, D. (2013). Thick concepts. *Philosophy Compass*, 8, 677–688
- Scheffler, S. (1987). Morality through thick and thin: A critical notice of ethics and the limits of philosophy. *Philosophical Review*, 96, 411–434
- Schroeter, L., & Schroeter, F. (2003). A slim semantics for thin moral terms? *Australasian Journal of Philosophy*, 81, 191–207
- Sciaraffa, S. (Ms.). On Pragmatic Inferentialism. (Ed.). unpublished manuscript
- Sellars, W. (2007). *In the space of reasons. Selected essays of Wilfrid Sellars*. Edited by K. Scharp and R. Bandom. Harvard University Press
- Tiefensee, C. (2021). Metasemantics for the relaxed. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (16 vol., pp. 108–133). Oxford University Press
- Väyrynen, P. (2008). Slim epistemology with a thick skin. *Philosophical Papers*, 37, 389–412
- Väyrynen, P. (2013). *The lewd, the rude and the nasty*. Oxford University Press
- Väyrynen, P. (2018). A simple escape from moral twin earth. *Thought*, 7, 109–118
- Väyrynen, P. (2021). Thick ethical concepts. In Edward N. Zalta (ed.). *The Stanford encyclopedia of philosophy* (Spring 2021 Edition), <https://plato.stanford.edu/archives/spr2021/entries/thick-ethical-concepts/>
- Wedgwood, R. (2007). *The nature of normativity*. Oxford University Press
- Williams, B. (1985). *Ethics and the limits of philosophy*. Harvard University Press

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.