CANADIAN JOURNAL OF PHILOSOPHY, 2018
VOL. 48, NOS. 3–4, 425–446
https://doi.org/10.1080/00455091.2017.1413843

Routledge
Taylor & Francis Group

The unity of moral attitudes: recipe semantics and credal exaptation*

Derek Shiller

Independent Scholar, New York City, NY, USA

ABSTRACT

This paper offers a noncognitivist characterization of moral attitudes, according to which moral attitudes count as such because of their inclusion of moral concepts. Moral concepts are distinguished by their contribution to the functional roles of some of the attitudes in which they can occur. They have no particular functional role in other attitudes, and should instead be viewed as evolutionary spandrels. In order to make the counter-intuitive implications of the view more palatable, the paper ends with an account of the evolution of normative judgments as exaptations of the cognitive structures that underlie beliefs.

ARTICLE HISTORY Received 25 March 2017; Accepted 30 November 2017

KEYWORDS Noncognitivism; Frege-Geach problem; language of thought; moral evolution

Normative noncognitivism is a negative theory: it holds that normative judgments lack representational contents and are therefore not beliefs. Formulating a satisfying positive theory requires saying more about what moral attitudes actually are. To that end, noncognitivists have traditionally looked to the motivational capacity of moral judgments and have focused on straightforwardly predicative forms such as *insurance fraud is wrong* and *charitable donations are morally exemplary*. Moral judgments compel us to engage or refrain from engaging in actions without help from ancillary desires.[1] Though noncognitivism can be developed in a number of ways, I will treat this motivational strain as representative.

Noncognitivists are known to face challenges in making sense of moral attitudes other than straightforwardly predicative moral judgments. The Frege-Geach problem, which addresses logically complex moral judgments, is the most well-known challenge. However, noncognitivists require characterizations of a range of other attitudes.

CONTACT Derek Shiller dcshiller@gmail.com

*This paper benefitted from comments by Adam Thompson, Derek Baker, David Faraci, Tristram McPherson, William Melanson, and two anonymous referees, as well as audiences at the Philosopher's Cocoon and Minds Online conferences.

In this paper, I will sketch a systematic strategy for characterizing moral attitudes. My approach differs from most extant noncognitivist views in that I don't aim to rationally vindicate the behavior of our moral attitudes or earn the right to realist practice and discourse. Instead, I aim only to provide an account of what distinguishes these attitudes. I hope that once this account is fleshed out, their behavior will be shown to be understandable.

In short, I propose that moral attitudes can be divided into primary and derivative classes. The primary moral attitudes admit of traditional noncognitivist characterizations. The derivative moral attitudes are identified by non-semantic similarities with members of the primary class. Namely, they involve the same constituent components. This dichotomy allows for a 'recipe semantics' of moral attitudes, which characterizes them by their inclusion of components that have distinctive functional roles in a restricted set of contexts.

The resulting noncognitivist view shares some of the spirit of error theory: I will suggest that the derivative moral attitudes are accidental byproducts of the cognitive structures underlying the primary moral attitudes. The parallels between moral and propositional attitudes result from their syntactic commonalities. Moral concepts enjoy the same syntactic freedoms as other concepts, even if they don't have coherent functions in all of the contexts in which they occur.

I will begin in Section 1 by cataloguing the attitudinal complexities facing noncognitivists. In the second, I will present an account of our propositional attitudes to frame my proposal about moral attitudes. In the third, I will explain how a recipe semantics could be used to characterize moral attitudes. In the fourth, I will conclude with a speculative account of the cognitive evolution of moral judgments that makes a recipe semantics plausible.

1. A menagerie of moral attitudes

Cognitivists hold that moral judgments, such as the sort expressed by (A), can be characterized as beliefs with a certain moral content.

(A) It is wrong to collect trophies of endangered species.

The judgment expressed by (A) takes a proposition about the distribution of a certain property (*wrongness*) as its object. Cognitivists disagree about the nature of this property, but they agree that the world must meet definite conditions in order for (A) to be true.

In contrast, noncognitivists hold that such judgments have no special representational contents. There is no property that we ascribe to acts in judging them wrong. Rather, in judging an act wrong, we adopt a conative attitude towards it. According to this approach, moral judgments are *moral judgments* not because of their representational content, but because of their motivational force.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

An account of straightforwardly predicative forms of moral judgments won't tell us much about the nature of logically complex forms. Consider (B), which employs 'wrong' under the scope of a negation.

(B) It is not wrong to collect trophies of extinct species.

The attitude expressed by (B) is not motivational in the same way as the attitude expressed by (A), so an account of straightforwardly predicative moral judgments will not carry over to negational judgments. We are not moved to act solely because we believe that something is *not* wrong.

Negation is the only most basic form of logical complication.[2] The same problem also arises for conjunction and disjunction, as in the attitude expressed by (C), and quantification, as in the attitudes expressed by (D) and (E).

(C) Human beings either have a special moral status or it is wrong to test cosmetics on animals.
(D) Everything you did today was more wrong than the thing you did before it.
(E) Most killings are deeply morally wrong.

Quantification hasn't received the same attention as conjunction, disjunction, and negation, but it is equally problematic.

In addition to ordinary quantification, moral judgments can also take plural and generic forms, such as the attitudes expressed by (F) and (G).

(F) The wrongs that the colonists did to the natives were more varied and numerous than the wrongs that the natives did to the colonists.
(G) Wrongs are a stain upon the moral character of decent people.

Geach was especially interested in conditionals, since they play an important role in moral arguments:

(H) If lying is wrong, then getting your little brother to lie is wrong.

Although there are truth-functional interpretations available for these conditionals, they are widely thought to be inadequate. The indicative conditionals found in moral judgments like that expressed by (H) are subject to the same considerations that motivate non-truth-conditional interpretations of indicative conditionals in other contexts. This means that the problems created by conditionals are not quite like the problems created by logical operators.

Moral judgments can take the form of subjunctive conditionals, as expressed by (I), and can involve tense, as expressed by (J).

(I) If it had been wrong to bring children into this world, we would not have done it.
(J) Abortion is as wrong today as it was fifty years ago.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

Noncognitivists must also account for the existence of mixed normative judgments. Mixed normative judgments, such as the attitudes expressed by (K) and (L), combine different flavors of normativity within a single attitude.

(K) You ought (epistemically) to know when you did wrong.
(L) You should (rationally) keep track of what you shouldn't (morally) do.

Finally, noncognitivists owe an explanation of moral attitudes other than judgment. Not only can we judge that actions are wrong, but we can wish that they were permissible, regret that we acted immorally and hope that we chose the right thing, wonder whether utilitarianism or deontology is correct and be more confident in one than the other, find some moral principles intuitive, imagine fictional scenarios where different moral principles hold, and suppose novel moral principles for the sake of arguments.

We must make sense of these other attitudes just as we must make sense of moral judgments. An account of what it is to judge that an action is wrong will not tell us what it is to hope that we did the right thing or wonder whether what we did was obligatory. It makes no sense to suppose ordinary conative attitudes for the purposes of arguments (though we can suppose that *we have* these conative attitudes).

Although the problems involved in logical and other semantic contexts are different from the problems created by other moral attitudes, the threat they pose is the same. Noncognitivists would be best served by a general account of the nature of these attitudes.

2. Context and meaning determination

Since moral attitudes are so similar to propositional attitudes, any inquiry into the former ought to start with the latter. In this section, I will present some assumptions about the cognitive structure of ordinary propositional attitudes in preparation for the recipe semantics that I provide in the next.

On the view to be assumed, our propositional attitudes involve relations to cognitive representations whose representational contents are determined by the representational properties of their parts. Those parts can be individuated by means of non-semantic properties and can receive their representational content from their function in special contexts. I will describe each of these ideas in turn.

2.1. The representational theory of mind

According to the Representational Theory of Mind (RTM)[3], ordinary propositional attitudes involve relations to mental representations whose contents are propositions. The type of attitude that results from tokening a mental representation depends upon the way in which it is tokened. For instance, to believe a

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

proposition might be to token a representation of that proposition in a certain way and to use representations tokened in that way in deciding how to act to satisfy one's goals.

Those who accept RTM need a psychosemantic account of the associations between mental structures and propositions. I will assume a teleological theory (e.g. Dretske 1995; Millikan 1984; Papineau 1984): roughly, mental representations have their representational contents by virtue of being designed to represent that content. This theory won't differ from its competitors on points relevant to my argument, and so readers may substitute their own preferred psychosemantics wherever appropriate.

2.2. Compositional representations

Representations can be either structured or unstructured. Structured representations are composed of constituents with their own representational contents. Natural language sentences are structured, since both sentences and words represent. Maps are also structured, since they have sub-regions with representational content (Blumson 2012). I will assume that mental representations are structured and I will refer to the ultimate constituents of individual mental representations as 'concept tokens'. The bare term 'concept' will be reserved for types of concept tokens.

Structured representations are typically compositional in the sense that the representational contents of the parts help to determine the representational contents of the whole. The representational contents of concept tokens do not determine how they (as neurophysiological entities) can be combined, even if their contents will have a say in whether the combinations they enter into are meaningful (Stich 1983). Semantic properties simply aren't the right sort of property to play that role, any more than the representational content of a LEGO head explains why it fits snugly onto a LEGO torso and not a LEGO horse.[4]

Non-semantic – 'syntactic' – properties enable the combination of some concepts and not others. As a result, there may be syntactically viable combinations of concept tokens that are semantically incoherent.

2.3. Orthographic identities

By supposition, concept tokens acquire meanings from their designed functions. The designed functions of concepts depend upon the their behavior in some range of contexts, and so it is necessary that concept tokens have identities that are preserved across contexts in order to ground their representational contents. Quine's (1960) radical interpreters could rely on the phonological properties of 'gavagai' to re-identify the word in different contexts. Their task would have been impossible had the sound of the word varied arbitrarily with each use. The same goes for cognitive semantics: representational contents could not

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

be assigned to the components of an attitude if those components could not be re-identified. The difficulty is not merely epistemic; patterns of use must underlie semantic values.

Since they are used to assign representational contents, the properties that identify concept tokens from context to context must be non-representational. Given the analogy with written language, I will refer to such properties as 'orthographic'. Orthographic properties might be closely associated with syntactic properties, but they need not be the same.[5]

2.4. Spandrel contexts

Our concepts may or may not have a designed function in every context in which their syntactic properties permit them to occur. Phenotypical traits often outstrip the requirements of their designed function. Natural selection exerts imprecise forces, and adaptations that are helpful in one context may also have effects elsewhere. This means that many phenotypic traits are byproducts – 'spandrels' – of selection for other traits.

Suppose that that we identify the concept *horse* as a concept whose designed function is to track the properties of horses by figuring into representations of horses. It isn't part of the designed function of the concept *horse* to distract us when we are bored, but its tokens can be employed to this end. This use need not count against the concept's designed function, as the concept may not have been designed *not* to be used as a distraction. Instead, its function may be given by its contributions in the context of certain kinds of attitudes, such as beliefs and desires, and it may have no functions in others, such as daydreams. If this is the case, then these other contexts are spandrel contexts for the concept.

3. Recipe semantics for moral attitudes

In Section 5, I developed a view on which propositional attitudes involve mental representations. I suggested that individual mental representations can be identified across contexts by virtue of their orthographic properties. I now propose extending this idea to moral attitudes. I will start by broadening the notion of mental representation to allow for non-representational variants. Then I will provide a recipe semantics for daydreams to serve as an analogue for my view about moral attitudes. Finally, I will turn to my proposed recipe semantics of moral attitudes.

3.1. Presentations

The orthographic properties of non-mental representations can be divided into resemblance classes. Written words, for instance, are individuated by their shapes. These shapes are distinctive, but they share more in common with each

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

other than they do with spoken words. Written English sentences also share features (across different typefaces and handwriting styles) that distinguish them from sentences in Arabic or Greek. An 'orthographic class' is a set composed of entities whose orthographic properties resemble each other and that includes a system of representations as a subclass. There are orthographic classes for English sentences, road maps, stock tables, and bar graphs.[6]

Membership in an orthographic class is determined by the orthographic properties of the potential member and the corresponding properties of other members. Representationality is not a prerequisite for membership. Many members of orthographic classes have no representational significance. Take this stanza from Lewis Carroll's Jabberwocky:

> 'Twas brillig, and the slithy toves
>
> Did gyre and gimble in the wabe:
>
> All mimsy were the borogoves,
>
> And the mome raths outgrabe.

These sentences belong to orthographic classes that also contain English sentences, but since many of the words have no established meaning, they do not represent.

Any member of an orthographic class is a 'presentation'. Representations are often presentations, but presentations need not be representations. The sentences in Jabberwocky are non-representational presentations.

Non-representational mental presentations are conceptually coherent: they are presentations that fail to meet contingently-satisfied conditions for representing. If minds use representations in their propositional attitudes, they may also use non-representational presentations in other attitudes. Not every attitude that acts something like a belief has earned a semantic value. This opens up the possibility that moral attitudes involve non-representational presentations in just the way that propositional attitudes involve representations.

3.2. The presentational theory of mind

According to RTM, propositional attitudes are relations to mental representations. If beliefs involve relations to mental representations, moral judgments plausibly involve relations to mental presentations. Beliefs and moral judgments are quite similar (Horgan and Timmons 2006), and so it would be surprising if they were implemented in fundamentally distinct ways.

Now my proposal: *straightforwardly predicative moral judgments are attitudes that involve relations to complex structured presentations that contain non-representational concepts whose proper function is not to represent but instead to influence action (in that context)*. These concepts are moral concepts and they are syntactically compatible with representational concepts in the same way

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

as representational concepts are with each other. Since the designed function of neither the whole presentation nor its constituent moral concepts is to represent, neither is representational. But since moral concepts convey an action-directing function in the context of straightforwardly predicative moral judgments, they are also not functionless.[7]

3.3. A recipe theory of daydreams

Philosophers of mind have traditionally focused on attitudes, such as belief and desire, that play an important role in influencing behavior. Not all of our attitudes are like this. Daydreams are not. This means that many strategies for characterizing propositional attitudes cannot be fully generalized. Contra functionalism, we cannot make sense of what it is to daydream about a particular situation in terms of what that daydream does in our mental life. In order to understand these states, we may need to approach them as byproducts of other cognitive faculties.

I propose that the representational properties of daydreams ought to be understood through their parts and how those parts are used elsewhere. We can't identify a concept token as a concept of a horse, for instance, solely by its deployment in daydreams. Instead, we must look to how the same sort of concept token is used in other cognitive contexts, such as perceiving horses, having beliefs about horses, and organizing intentions regarding behavior toward horses.

A daydream gets its content by employing concepts that play important roles in other contexts. These concepts are recognized by the orthographic properties of their tokens. A concept token of a horse in a daydream represents a horse only insofar as it was designed to have a certain representational content in other contexts. The fact that a dream is of a horse and not a teapot results from the fact that the same concept tokens employed in the dream have the function of keeping track of the properties of horses on other occasions.

3.4. A recipe semantics for moral attitudes

In Sections 1 and 2, I suggested that straightforwardly predicative moral judgments involve a special kind of mental presentation. In the present, I will explain how to extend this characterization to handle more complex moral attitudes.

Moral concepts are characterized by non-representational proper functions that operate in the context of straightforwardly predicative moral judgments. In order for a concept token to be a token of a moral concept, it must be designed for a certain use in certain contexts. Given that concept tokens can be recognized in a variety of contexts based on their orthographic properties, and given that they need not have functions in all of the contexts in which they occur, it is easy to systematically extend this account to handle other moral attitudes.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

The general strategy is to characterize attitudes with recipes for constructing them – i.e. in terms of the ingredients that must be combined to produce the attitude. The ingredients may be characterized by their designed function in other contexts. There need not be anything distinctive about these attitudes beyond the parts from which they are made. They may lack a function and have no important role in cognition.

Just as with daydreams, the concepts that populate our moral attitudes may get their contents from other contexts. In particular, moral concepts may be characterized by their role in straightforwardly predicative judgments and have no particular semantic characters in general.[8]

3.5. Hope

Perhaps the most promising applications of this approach are to moral attitudes other than judgments. Noncognitivists must make sense of what it is to have moral hopes, to be uncertain about moral issues, and to make moral suppositions for the purposes of arguments. With a recipe semantics, we can easily explain what makes these attitudes moral attitudes.

According to RTM, propositional attitudes involve relations to mental representations. By assumption, differences in kinds of propositional attitudes correspond to differences in the relationships taken to mental representations. To believe something and to desire it involve taking different relations to the same content. Just as with concept tokens, these relations need an orthographic identity in order to support a complex pattern of use. Each belief involves an orthographically identifiable relation to an orthographically identifiable presentation, each desire involves a distinct orthographically identifiable relation to an orthographically identifiable presentation, and so on.

If hopes are a basic propositional attitude, then we can expect that there will be a separate orthographically distinctive property that characterizes the relations we have to the representations we hope to be true. Even if hoping is characterized by its functional role within our cognitive lives, the realizer playing that role may be distinctive in a non-semantic way.

If our minds employ moral presentations and we relate to them in the same orthographically distinctive way that we relate to the representations that we hope, then we can count as hoping them as well, even though they lack representational content and do not have quite the same functional role as ordinary hopes. A moral hope is created by combining a presentation including a moral concept with the orthographic relation special to hope. There need be no deeper essence for moral hopes.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

3.6. Negation

Negational moral judgments are the sorts of judgments that we typically express with negations of moral predications. They appear to be logically inconsistent with straightforwardly predicative moral judgments. For instance, the judgment *immigration restrictions are not wrong* is a negational moral judgment that appears to be inconsistent with the judgment *immigration restrictions are wrong*. This inconsistency has been thought to create substantial problems for noncognitivists (Schroeder 2008; Unwin 1999), and has received a lot of attention in recent years (e.g. Baker and Woods 2015; Schwartz and Hom 2015; Shiller 2016; Sinclair 2011).

Noncognitivists have tried to characterize negational attitudes by their relation to their unnegated counterparts: they are states that are in some way incompatible (Blackburn 1988; Gibbard 2003; Horgan and Timmons 2009). This incompatibility might be explained by the characteristic functional role of the attitudes. For instance, the functional roles of straightforwardly predicative and negational moral judgments might be to direct us to realize mutually unsatisfiable states of affairs.

Critics have cast doubts on whether noncognitivists can locate attitudes with the relevant sort of incompatibility. The chief worry is that the incompatibility must be essential to the attitudes and must accommodate the distinction between logical and non-logical forms of inconsistency (van Roojen 1996).

Negational moral judgments are not the strongest candidates for a recipe semantics, and adopting a recipe semantics would not preclude us from characterizing straightforward predicative moral judgments and their negations in ways that account for attitudinal inconsistency. However, noncognitivists can also make use of a recipe semantics to characterize negational moral judgments and doing so helps explain the appearance of inconsistency.

On a recipe semantics, negational moral judgments might be characterized as attitudes that involve presentations with a special sort of negation concept. A negational moral judgment involves a presentation that includes a negation concept token as a constituent. Though the negation concept is designed to play a certain function in the case of representations, it need not play that specific function in other presentations and so it may have no special function or representational content whatsoever in the context of moral presentations. What identifies the moral judgment as a negational moral judgment is the presence of a constituent with a specific function in other contexts.

It is plausible that the characteristic function of negation concepts is to invert representations: a representation governed by a negation concept has the opposite satisfaction conditions as its unnegated content. This function is restricted to representations. The concept may not have this function in all of its syntactically viable contexts. In the context of moral judgments, in which there are no satisfaction conditions to invert, the negation concept may have

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

no more function than 'not' in 'it was not brillig'. Instead, it simply feeds another component into the presentation.[9]

Characterizing an attitude is one thing. Accounting for its behavior is another. This paper is specifically focused on offering a schema for characterizing moral attitudes, but the viability of these characterizations will depend in part on how they can account for the behavior of the attitudes. When it comes to negational moral judgments, the key feature to be explained is the appearance of inconsistency.[10]

The appearance of inconsistency probably does not depend on an implicit grasp of the semantics of moral attitudes. On the contrary, the appearance may result from the presence in moral attitudes of the same non-semantic mechanisms that generate genuine inconsistency in other situations.

Even in its spandrel contexts, the presence of the negation concept may color our introspective view of the attitude. Since it generates inconsistencies in so many other contexts, the negation concept could produce the appearance of inconsistency in moral attitudes and thereby lead us to treat the attitudes in which negation concept tokens occur as spandrels in the same way we treat the attitudes in which they play their proper role. In other words, we may project the inconsistency-generative appearance of negational concepts from their primary contexts to their spandrel contexts.

We can extend this explanation to handle the logicality of attitudinal inconsistency. Noncognitivists and their critics have realized that it is helpful to attribute some structure to moral attitudes (Baker and Woods 2015; Schroeder 2008) in order to explain logical inconsistency. The recipe semantics achieves this through attributing syntactic structure and the presence of distinctive concepts in the cognitive presentation. Moral attitudes often appear inconsistent because they have the form of inconsistent attitudes, and the considerations that deprive them of representational meanings are subtle.

This structure even permits a shallow sense of logical validity. While moral claims may only be truth-evaluable in a deflationary sense themselves, any substitution of representational concepts for moral concepts in a 'valid' moral argument would produce a truth-preserving argument. Moral arguments can have a valid form, since the structure of their logical concepts can guarantee representational truth preservation. Of course, moral arguments are importantly different from their propositional analogues, but their logical form might explain why they are compelling in the same way. It is unsurprising that we should be inclined to view moral arguments as akin to non-moral arguments, especially given our tenuous introspective grasp on their semantics.[11]

3.7. Comparisons with existing views

By adopting a recipe semantics, we avoid the need to attribute representational contents or essential functional characters to attitudes other than

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

straightforwardly predicative judgments. Straightforwardly predicative moral judgments establish the orthographic properties that other attitudes must share in order to count as moral attitudes. There is no semantic core to moral attitudes in all their guises. They are moral attitudes because a subset of the attitudes with the relevant orthographic properties (i.e. containing a orthographically identified moral concept) play the role characteristic of straightforwardly predicative moral judgments.

It is consistent with this approach that many kinds of moral attitudes have representational contents or functional characters – all that the recipe semantics requires is that these properties are not what qualifies them as moral attitudes. However, once we secure the morality of moral attitudes without recourse to a shared content or functional role, there is little reason to demand contents or functional roles for any particular attitudes. This makes my approach very different from traditional approaches to the problem, and for attitudes other than straightforwardly predicative moral judgments, it may more closely resemble error theory (Joyce 2001; Mackie 1977). The attitudes themselves are not mistaken in the way error theorists suppose, but our higher-order judgments about them are deeply mistaken.

3.7.1. Traditional expressivism

Traditional noncognitivists such as Simon Blackburn (1988) and Allan Gibbard (2003) have developed elaborate views with the aim of vindicating aspects of moral practice – of earning the right to realist forms of practice and discourse. The prominence of the Frege-Geach problem has led to a focus on moral argumentation and the logical relations between our attitudes. The standard line is to interpret apparently logically inconsistent attitudes as involving some sort of opposing and mutually unsatisfiable commitments.

The view about negational moral judgments suggested in Section 2 takes a very different tact. Instead of aiming to vindicate the relations between moral attitudes, it provides an explanation of how we might be led to treat them as we do. It is plausible, for reasons presented below, that some of our attitudes play no important roles in our psychology. If these aspects of moral psychology do not need intelligible rationales, than we should be satisfied with non-rationalizing explanations of moral practice.

The key point is that in denying moral attitudes special semantic contents, we need not surrender any of the cognitive machinery responsible for our use of (and intuitions about) propositional attitudes. Propositions are abstract objects we use to help us think about neurological processes. Strictly speaking, they do not contribute to the functioning of the underlying neural machinery. We impose propositions on top of this machinery in order to fit it within our folk psychological and normative schemes, but if our schemes break down, it is more likely because of the limitations of the abstracta than the machinery.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

This view will fair poorly if our goal is to support the greatest number of pre-theoretical intuitions, for it contradicts a common perspective on morality. However, error theories play a special role in philosophical methodology: intuitions that can be effectively explained away are not intuitions whose satisfaction counts for much. If we can effectively explain our moral intuitions by means of syntactic features of moral attitudes, then we should let the semantic cards fall where they may.

3.7.2. Hybrid expressivism

Forms of hybrid expressivism (Ridge 2014; Schroeder 2013; Toppinen 2013) that marry cognitive and noncognitive states have recently become popular. On these views, moral judgments involve both beliefs and noncognitive states. Judging that an act is wrong, for instance, might be analyzed into believing that the act has a certain non-moral property and disapproving of all acts with that property. Hybrid views have the advantage of capturing the behavioral similarity of moral attitudes with both cognitive and noncognitive states within a traditional psychological framework.

Hybrid expressivism may not differ so much from the view on offer in terms of the kinds of cognitive structures that it attributes to moral judgments. Both views will allow that belief-like attitudes take part in moral judgments. They will disagree, however, primarily[12] on the meta-semantic question of whether those belief-like structures have full representational contents. The motivation for the recipe semantics is the thought that moral concept tokens lack representational contents because they fail to meet the conditions necessary to represent. Hybrid expressivists must deny this.

The primary drawback to hybrid views is that they require representational contents for our moral beliefs. In order to accommodate the diversity of moral views, the most promising versions of hybrid expressivism have been forced to adopt a kind of relativism in which the representational contents of moral judgments (the parts involved in beliefs) are relativized to the judge's particular moral standards (Ridge 2014).

In many moral judgments, it is not hard to find reasonable representational contents to assign. The problem is that the approach is committed to finding a non-normative content of this sort for every moral judgment: in order to judge an action wrong, an individual must judge it to have some particular negatively regarded property. Similarly, it must be impossible to hope that an action was right without hoping that it had some specific positively regarded property.[13]

The attribution of relativized contents is unnecessary. The primary value of assigning contents to these attitudes is that it allows them to be fit within a familiar rationalizing psychological picture. It is not worth going to great lengths to read robust representational contents into moral attitudes, given the potential of non-semantic explanations to achieve the same ends. If we don't need to attribute representational contents to understand our attitudes, we shouldn't go

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

out of our way to do so. So, if the project of explaining the similarities between moral judgments and beliefs in terms of their syntactic properties pans out, then we don't need to assign representational contents.

4. Belief, exaptation, and moral judgments

The value of this recipe semantics depends on how plausible it is that some of our concepts appear in spandrel contexts. While it is coherent for concepts that have a function in one context to appear in spandrel contexts, it may still count against a theory to be forced to this conclusion.

We are able to combine concept tokens in the ways that we do because of their syntactic properties. The semantic properties of our concepts, which depend partly on their historical or counterfactual behavior, cannot themselves fully explain how it is that we are able to use them to produce attitudes with the characteristic behaviors of propositional attitudes.

Cognitivist and noncognitivists alike require explanations of the forms our attitudes take. Cognitivists face the challenge of explaining how we can form diverse beliefs. Jerry Fodor has argued for a solution invoking the structure rather than the content of our attitudes (Fodor 1998, 2008; Fodor and Pylyshyn 1988). Fodor explains our systematic and productive attitudinal capacities in terms of a cognitive system of recombinable units. The purported functional differences provide little reason to think that the challenge of accounting for all of our attitudes will look much different for noncognitivists. The mechanisms that determine unit combinability are orthogonal to the properties that convey these units their semantic values.

Nevertheless, it might be doubted that our moral concepts have syntactic properties that allow them to be combined into attitudes without functions. How is it that we're able to hope that we did the right thing, if the concept *right* doesn't make any specific semantic contribution to the attitude?

The evolutionary history of moral attitudes can make these syntactic properties unsurprising. What follows is one story of how this history may have gone. It isn't the only viable story, but it can still ward off serious worries about postulating spandrel moral contexts.

4.1. Coordinating social expectations

Our tendency to moral behavior is surely the product of evolutionary forces.[14] Full-fledged moral attitudes play an important role in regulating our behavior, and the behaviors they produce contribute to our fitness. It is fitter for us to reciprocate acts of kindness and cruelty and to act altruistically toward our kin.[15] However, full-fledged moral attitudes are not necessary to produce this sort of behavior. Moral sentiments suffice. We could have evolved to feel grateful,

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

vindictive, protective, and magnanimous without having anything that looked quite like moral judgments.

Moral judgments are complex and sophisticated attitudes. On the surface, they appear to rely on the same mechanisms that underlie our capacity to formulate complex representations, even if they are not representational. If moral sentiments are capable of driving moral behavior by themselves, why do we also have the capacity to judge actions to be right or wrong?

One explanation is that moral judgments evolved to enable us to coordinate social expectations.[16] They not only play a role in guiding moral behavior, they also increase social flexibility. They give us the cognitive resources necessary to formulate, assess, and communicate norms that can regulate our sentiments. We can voice the moral rules we agree to obey, and learn the rules others favor following. Furthermore, if we are flexible and open to influence by others, we can ensure that the rules that govern our community are a compromise of the interests of its members. This makes it easier to accommodate changing social structures and power dynamics.

In the rough and rapidly changing world of the upper palaeolithic, moral flexibility would have been advantageous. Our species spread out quickly to environments as different as Ice Age Europe and Polynesia (Henn, Cavalli-Sforza, and Feldman 2012). The new ways of life required for survival in these environments created novel social dynamics in an evolutionary instant. Furthermore, the cognitive advances that allowed for this expansion also forced our ancestors to deal with new interpersonal issues relating to property, contracts, debts, and punishments, and fueled the change from small and relatively unspecialized egalitarian family communities to large hierarchical societies of highly specialized strangers (Boehm 2012). Fixed moral sentiments would have been a handicap.

4.2. Exaptation

A system of normative attitudes for coordinating attitudes would not have sprang up overnight. Chimpanzees and bonobos, our closest relatives, may have some rough analogues to moral sentiments, but they have nothing clearly recognizable as judgments about morality. They recognize some social expectations and are able communicate with each other about their observation in very basic ways. As far as we know, they are not capable of formulating, evaluating, and communicating anything as complex as social rules (von Rohr, Claudia, and van Schaik 2011). Our nearest cousins might have some sense of reciprocity, permission, and obligation, but they do not reason, discuss, or negotiate them in the way that we do.

Our ancestors almost certainly would have started making cognitively sophisticated moral judgments only after their split with our nearest cousins some six million years ago, and the greatest cognitive changes most likely came in the last

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

fifty to two hundred thousand years. In that period of time, cultures sprang up, art and religion developed, tool use drastically advanced, long distance trade commenced, and sophisticated forms of language probably evolved.[17]

It makes sense that moral attitudes (as opposed to our moral sentiments) were largely a product of this time period, and if so, they must have evolved extremely quickly. Our pre-moral ancestors would have had much to gain by coordinating social expectations about behavior. This would have produced an evolutionary pressure to allow for such flexibility. This pressure would have set to work on a cognitive system with a stock of propositional attitudes including beliefs and desires. Moral judgments might have emerged *sui generis*, but it seems far more likely that they would have been spun off from existing attitudes. It is easier to co-opt – 'exapt'[18] – existing structures than to build them from scratch: there is evidence (Anderson 2010) that many aspects of cognition make use of the resources of older cognitive functions. Beliefs and desires seem like the two best candidates for the original source material of moral attitudes, and there are some reasons to favor the former over the latter.

Moral judgments might have arisen as exaptations of beliefs because of the advantages they afford: perhaps the structural flexibility and logical relations available to beliefs were helpful in moral reasoning and discourse. If so, then it is possible that logical connectives had an important function and that logically complex attitudes need a richer characterization than is provided by a recipe semantics.[19] However, while this might explain why it is better for us to have moral judgments that are structurally similar to beliefs, it cannot explain how a trend initiated in that direction. Why were incipient moral attitudes formed from beliefs, rather than desires?

The primary reason to favor the hypothesis that moral attitudes developed from beliefs is that beliefs about social regulations are the best candidates for precursors to moral judgments. Our ancestors were probably thinking and conversing about collective expectations for at least as long as they have been reasoning about norms. The states that play the most similar functional roles in promoting social coordination are beliefs about social roles and expectations.[20] Evolutionary forces pushing toward non-representational moral judgments would have likely moulded these states.[21]

Our pre-moral ancestors might have started with ordinary beliefs about social rules and regulations: they judged actions according to their social appropriateness within their community and expressed these judgments with simple language. Generation by generation, they became increasingly disposed to form attitudes that functioned not to simply represent social standards, but to motivate them to adopt these standards. Gradually, the set of standards they were motivated to adopt might have come to differ from those that were accepted by their society. The divergence between what was accepted by individuals and what was accepted by other members of their society could have grown to a point where the attitudes fell outside of the range of representationality,

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

especially if they came to have a use in influencing those standards, and not just reflecting them.

Our concept tokens employed in tracking rules and social regulations combine normally with negation, conjunction, quantification, generics, conditionals, tense, and other attitudes. *If our moral concepts exapted from these representational concepts, then they could easily continue to share some of their syntactic properties.*

This is supported by the rapidity of the development of moral judgments. It is quite possible that we only started internalizing social regulations as moral judgments in the last fifty thousand years (This is especially plausible if moral attitudes are the result of cultural evolution rather than biological evolution). Retaining the capacity for superfluous attitudes does us no harm, and selection has not yet weeded these attitudes from our cognitive repertoire. Even if they have no function, they may still prove useful.

This is far from the only story to be told about the origins of our moral attitudes, and it may well be false in many of its details. However, its plausibility undermines concerns about whether syntactic properties could allow moral concepts to appear in spandrel contexts. The fact that moral judgments display the same syntactic properties as beliefs without having a clear function is not a major problem for the view.

5. Conclusion

In this paper, I have argued that noncognitivists can provide a recipe semantics for a variety of moral attitudes. Moral judgments involve moral concepts. Moral concepts are characterized by the roles they play in the context of moral judgments. Moral concepts can be identified in spandrel contexts by their orthographic properties. The other moral attitudes, including both complex moral judgments and non-judgment attitudes, can be characterized in terms of their orthographic similarities to straightforwardly predicative moral attitudes.

Many of our moral attitudes may be byproducts. The logical relations that we attribute to them need no underlying rationale, and the appearance that such attributions are deserved may instead result from the historically accidental syntactic viability of logical concepts in moral attitudes.

This thesis does not entail that our attitudes are mistaken. Nothing should hold us back from putting cognitive spandrels to good use. There is also nothing wrong with treating attitudes as if they were logically inconsistent, and so we need not change our moral practices, even if we discover that some of the relations we intuit between moral attitudes are not supported by anything essential to the attitudes themselves.

Despite not requiring any change to our moral practices, the proposal does give up something of our folk conception of moral psychology. Insofar as we seek to substantiate the folk conception of morality in our metaethics, this must

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

be seen as a drawback. While the lack of particular explanatory meaning-determined rationales for the behaviors of moral attitudes may be disconcerting, they can be accounted for by the rapid evolutionary development of ordinary moral judgments. If we are willing to leave some parts of our folk psychology behind, the resulting picture may be the best path forward for noncognitivists to simultaneously understand the diversity and unity of moral attitudes.

Notes

1. A.J. Ayer (1936) and Simon Blackburn (1984) treated moral judgments as a special sort of approval and disapproval. Charles Stevenson (1937) described them as a species of interest or partiality. Allan Gibbard (2003) compared them to plans. Mark Timmons and Terry Horgan (2006) suggested that they are a special non-representational species of belief. Mark Schroeder (2008) proposed (without endorsing) that they are attitudes of favoring or disfavoring attributions of blame. These views all emphasize the pressures moral judgments exert on us to act in certain ways. Some of the remaining difficulties involved in providing an adequate characterization of straightforwardly moral judgments are explored in David Merli (2008).
2. The question of logical complexity, first raised in the work of Peter Geach (1965) and John Searle (1962), has taken a prominent place in discussions of noncognitivism. Though it is typically presented under the inclusive label of *The* Frege-Geach Problem, logical complexity poses a variety of different issues. One part of the problem – the part that I focus on in this paper – is to make sense of just what sort of attitudes logical complex moral judgments are. The other parts of the problem involve explaining why logically complex attitudes relate to each other in the ways that they do, how such attitudes have logical relations such as inconsistency and entailment, and how they can figure into rationally compelling arguments. The challenges involved in all of these parts of the problem are explored in depth by Mark Schroeder (2008). I have discussed them elsewhere (Shiller 2016), and the solution I present there for this second part of the Frege-Geach problem is compatible with the solution I present here for the first part of the Frege-Geach problem.
3. This term comes from Jerry Fodor, who uses it to describe a constellation of theories (1998). I will restrict its usage to the view defined above (a view that is found, for instance, in Field [1978]): a theory that is distinct from the Language of Thought Hypothesis, which makes additional assumptions about the structure of mental representations, and the Computational Theory of Mind, which makes additional assumptions about the way that representations are handled in deliberation and reasoning.
4. Of course, representational properties are correlated with viable combinations of both concepts and LEGOs, but these properties are not explanatory except in a unificationist sense. If anything, representational contents are determined in part by what combinations are possible, rather than the other way around.
5. Friedemann Pulvermüller (2002) presents a theory that illustrates the intended difference between orthographic and syntactic properties. On his proposal, individual concepts are implemented in the brain by networks of functionally entwined neurons that selectively respond to relevant stimuli. These functional webs are themselves connected to collections of sequence detectors that act as

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

grammatical categories. The fact that a given web is connected to a collection of detectors typical of nouns explains why it can be put into subject position in thoughts. In this theory, concepts are distinguished by orthographic properties (dispositional firing patterns of a functional web), and how they associate with each other depends upon their distinct syntactic properties (connection strengths of the web with specific sequence detectors).

6. Orthographic classes are plentiful and many will be vague. There need not be a clear cut answer as to whether or not sentences in English and Chinese fall into the same orthographic class. They fall into some of the same orthographic classes and not others.
7. We can characterize the function of moral concepts in a variety of different ways, so this proposal is consistent with many analyses of straightforwardly predicative moral attitudes.

   It might be objected that my proposal will restrict moral judgments only to those creatures who employ presentations. I am sympathetic to this worry, and I will allow that having moral judgments may not require having presentations. Nevertheless, it is not misguided to limit our focus to creatures like us. Our way of having moral attitudes involves taking attitudes toward presentations with moral concepts.
8. Are concepts meaningless in their spandrel contexts? I want to resist the urge to say that moral concepts mean anything different in spandrel contexts than in the contexts in which they have a function. Representational concepts are not meaningless in their spandrel contexts precisely because they borrow their meaning from their functional contexts. For both moral and representational concepts, there is only one fundamental bearer of semantic values between the two kinds of contexts, and whatever semantic values those concepts have, they have because of their behavior in the functional contexts.

   That said, I doubt that it is possible to assign an object to count as a moral concept's meaning in both spandrel and functional contexts in the way that is possible for representational concepts. If it takes a meaning object to be meaningful, then moral concepts are not meaningful in any of the contexts in which they occur.
9. The fact that an attitude lacks a function or a meaning does not entail that it is useless. The attitude may have begun its life as a spandrel and subsequently come to be put to good use. The present value of a concept may be tangential to the properties that provide it with a coherent meaning.
10. Negated moral judgments and their unnegated counterparts may genuinely be inconsistent in the shallow sense in which two attitudes are inconsistent when we are inclined to treat them as disagreeing and perceive them as incompatible social commitments. We try not to adopt such 'inconsistent' attitudes, and expect others to expect this of us. This shallow sense of inconsistency often arises from, and cannot explain, the appearance of something deeper. Nevertheless, recognizing that such shallow forms of inconsistency may persist mitigates the unintuitiveness of denying authenticity to deeper senses. Noncognitivists have been happy to embrace minimalist interpretations of truth and content, I see little further cost to also adopting shallow interpretations of inconsistency.
11. If the contents of our moral concepts are determined by population-level regularities of concept use (Schroeter 2014), then we can have no special introspective insight into their meanings, and it is highly plausible that we might be misled by the forms of our attitudes.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

12. They may also disagree about whether the motivational force of the moral attitude is provided by the belief-like structure or by a separate structure. The most straightforward way of interpreting hybrid expressivism within RTM analyzes moral judgments into two structures, but it is also possible to treat them as one structure implementing two attitudes.
13. I expand on this line of criticism in (Shiller 2017).
14. I do not mean to imply that most aspects of morality are genetic. In fact, the proposal developed here is quite a natural fit for a cultural revolution approach (Powell, Shennan, and Thomas 2009; Sterelny 2011) to explaining human behavioral changes of the last hundred thousand years.
15. Joyce (2007) provides an in depth overview of the ways that stereotypically moral behavior is fitness enhancing.
16. Gibbard (1990), Sinclair (2012), and Björnsson and McPherson (2014) advocate for a similar perspective on the role of morality. Michael Tomasello (Tomasello 2016; Tomasello et al. 2012) has also developed a two-step process in the evolution of morality in which cultural aspects of moral psychology arose with the dawn of culture, long after the appearance of cooperation-inducing components of moral psychology.
17. The possible range of dates for the evolution of cognitively modern humans spans the emergence of anatomically modern humans ~160 ka (d'Errico and Stringer 2011) and the great dispersal out of Africa ~50 ka (Klein 2008).
18. Gould and Vrba (1982) present the concept and argue that it plays an important role in evolution, including in cognition.
19. Thanks to an anonymous reviewer for raising this idea.
20. Tomasello (2016) hypothesizes that early normative judgments focused on proper role-playing in complex cooperative activities. If so, they surely trailed beliefs about what roles individuals could productively play.
21. Furthermore, is unlikely that young children are able to discern representational from noncognitive concepts. They must form attitudes in response to moral instruction very early in life, and their mature attitudes grow out of these early ones. Whatever attitudes they ultimately come to have, children probably begin classifying actions under moral categories in much the same way that they classify actions under other categories, and this probably involves the structures of belief. Later on, they adopt the subtleties of moral concepts that rob them of their representational contents, but they don't need to fundamentally refigure their existing attitudes. In this case, perhaps, ontology recapitulates phylogeny: states that start out primarily as representations of social categorizations become something more closely tied to motivation.

Notes on contributor

Derek Shiller is a software engineer in New York City. He writes on issues relating to consciousness, ethics, and normative attitudes.

References

Anderson, Michael. 2010. "Neural Reuse: A Fundamental Organizational Principle of the Brain." *Behavioral and Brain Sciences* 33 (4): 245–266.
Ayer, Alfred J. 1936. *Language, Truth and Logic*. London: Victor Gollancz.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

Baker, Derek, and Jack Woods. 2015. "How Expressivists Can and Should Explain Inconsistency." *Ethics* 125 (2): 391–424.
Björnsson, Gunnar, and Tristram McPherson. 2014. "Moral Attitudes for Non-Cognitivists: Solving the Specification Problem." *Mind* 123 (489): 1–38.
Blackburn, Simon. 1984. *Spreading the Word*. Oxford: Oxford University Press.
Blackburn, Simon. 1988. "Attitudes and Contents." *Ethics* 98 (3): 501–517.
Blumson, Ben. 2012. "Mental Maps." *Philosophy and Phenomenological Research* 85 (2): 413–434.
Boehm, Christopher. 2012. *Moral Origins: The Evolution of Virtue, Altruism, and Shame*. New York: Basic Books.
Dretske, Fred. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
d'Errico, Francesco, and Chris Stringer. 2011. "Evolution, Revolution or Saltation Scenario for the Emergence of Modern Cultures?" *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1567): 1060–1069.
Field, Hartry. 1978. "Mental Representation." *Erkenntnis* 13 (1): 9–61.
Fodor, Jerry. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
Fodor, Jerry. 2008. *Lot 2: The Language of Thought Revisited*. Oxford: Oxford University Press.
Fodor, Jerry, and Zenon Pylyshyn. 1988. "Connectionism and Cognitive Architecture." *Cognition* 28 (1–2): 3–71.
Geach, Peter. 1965. "Assertion." *The Philosophical Review* 74 (4): 449–465.
Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
Gould, Stephen J., and Elisabeth Vrba. 1982. "Exaptation – A Missing Term in the Science of Form." *Paleobiology* 8 (1): 4–15.
Henn, Brenna, Luigi Cavalli-Sforza, and Marcus Feldman. 2012. "The Great Human Expansion." *Proceedings of the National Academy of Sciences* 109 (44): 17758–17764.
Horgan, Terry, and Mark Timmons. 2006. "Cognitivist Expressivism." In *Metaethics after Moore*, edited by Terry Horgan and Mark Timmons, 255–298. Oxford: Oxford University Press.
Horgan, Terry, and Mark Timmons. 2009. "Expressivism and Contrary-Forming Negation." *Philosophical Issues* 19 (1): 92–112.
Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
Joyce, Richard. 2007. *The Evolution of Morality*. Cambridge, MA: MIT Press.
Klein, Richard. 2008. "Out of Africa and the Evolution of Human Behavior." *Evolutionary Anthropology: Issues, News, and Reviews* 17 (6): 267–281.
Mackie, John L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin Books.
Merli, David. 2008. "Expressivism and the Limits of Moral Disagreement." *Journal of Ethics* 12 (1): 25–55.
Millikan, Ruth. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
Papineau, David. 1984. "Representation and Explanation." *Philosophy of Science* 51 (4): 550–572.
Powell, Adam, Stephan Shennan, and Mark Thomas. 2009. "Late Pleistocene Demography and the Appearance of Modern Human Behavior." *Science* 324 (5932): 1298–1301.
Pulvermüller, Friedemann. 2002. *The Neuroscience of Language: On Brain Circuits of Words and Serial Order*. Cambridge: Cambridge University Press.
Quine, Willard V. O. 1960. *Word and Object*. Cambridge, MA: MIT Press.
Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press

von Rohr, Rudolf, Judith Burkart Claudia, and Carel van Schaik. 2011. "Evolutionary Precursors of Social Norms in Chimpanzees: A New Approach." *Biology and Philosophy* 26 (1): 1–30.
van Roojen, Mark. 1996. "Expressivism and Irrationality." *The Philosophical Review* 105 (3): 311–335.
Schroeder, Mark. 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
Schroeder, Mark. 2013. "Tempered Expressivism." In *Oxford Studies in Metaethics 8*, edited by Russ Shafer-Landau, 283–314. Oxford: Oxford University Press.
Schroeter, Laura. 2014. "Normative Concepts: A Connectedness Model." *Philosophers' Imprint* 14 (25), 1–26.
Schwartz, Jeremy, and Christopher Hom. 2015. "Why the Negation Problem Is Not a Problem for Expressivism." *Noûs* 49 (4): 824–845.
Searle, John. 1962. "Meaning and Speech Acts." *The Philosophical Review* 71 (4): 423–432.
Shiller, Derek. 2016. "A Primitive Solution to the Negation Problem." *Ethical Theory and Moral Practice* 19 (3): 725–740.
Shiller, Derek. 2017. "The Problem of Other Attitudes." *American Philosophical Quarterly* 54 (2): 141–152.
Sinclair, Neil. 2011. "Moral Expressivism and Sentential Negation." *Philosophical Studies* 152 (3): 385–411.
Sinclair, Neil. 2012. "Metaethics, Teleosemantics and the Function of Moral Judgements." *Biology and Philosophy* 27 (5): 639–662.
Sterelny, Kim. 2011. "From Hominins to Humans: How Sapiens Became Behaviourally Modern." *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1566): 809–822.
Stevenson, Charles L. 1937. "The Emotive Meaning of Ethical Terms." *Mind* 46 (181): 14–31.
Stich, Stephen P. 1983. *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.
Tomasello, Michael, Alicia Melis, Claudio Tennie, Emily Wyman, and Esther Herrmann. 2012. "Two Key Steps in the Evolution of Human Cooperation: The Interdependence Hypothesis." *Current Anthropology* 53 (6): 673–692.
Tomasello, Michael. 2016. *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
Toppinen, Teemu. 2013. "Believing in Expressivism." In *Oxford Studies in Metaethics 8*, edited by Russ Shafer-Landau, 252–282. Oxford: Oxford University Press.
Unwin, Nicholas. 1999. "Quasi-Realism, Negation and the Frege-Geach Problem." *The Philosophical Quarterly* 49 (196): 337–352.

https://doi.org/10.1080/00455091.2017.1413843 Published online by Cambridge University Press