

A Dual Aspect Account of Moral Language

CAJ STRANDBERG

University of Gothenburg and Lund University

Abstract. It is often observed in metaethics that moral language displays a certain duality in as much as it seems to concern both objective facts in the world and subjective attitudes that move to action. In this paper, I defend The Dual Aspect Account which is intended to capture this duality: A person's utterance of a sentence according to which ϕ ing has a moral characteristic, such as " ϕ ing is wrong," conveys two things: The sentence expresses, in virtue of its conventional meaning, the belief that ϕ ing has a moral property, and the utterance of the sentence carries a generalized conversational implicature to the effect that the person in question has an action-guiding attitude in relation to ϕ ing. This account has significant advantages over competing views: (i) As it is purely cognitivist, it does not have the difficulties of expressivism and various ecumenical positions. (ii) Yet, in spite of this, it can explain the close, "meaning-like," connection between moral language and attitudes. (iii) In contrast to other pragmatic accounts, it is compatible with any relevant cognitivist view. (iv) It does not rest on a contentious pragmatic notion, such as conventional implicature. (v) It does not imply that utterances of complex moral sentences, such as conditionals, convey attitudes. In addition, the generalized implicature in question is fully calculable and cancellable.

1. Introduction

It is a common observation in metaethics that moral language has a certain duality: On the one hand, moral language appears theoretical since it concerns objective facts that might make our moral views correct or incorrect. On the other hand, it appears practical since it has an intimate connection to subjective attitudes that move to action. Much of the discussion in metaethics is concerned with devising a theory that accommodates both these features. If we are to believe the traditional tale, there are two main positions with corresponding difficulties. While cognitivism can explain the former feature but has difficulties accounting for the latter, the converse holds true for expressivism. Recently, views have been presented which are explicitly designed to encompass both features. Ecumenical views have been suggested according to which moral sentences express both non-cognitive states and beliefs, and cognitivist views have been suggested according to which the latter feature can be explained in pragmatic terms. However, all these views have met with serious objections.

In this paper, I propose a purely cognitivist and pragmatic account that I think has significant advantages over competing views. The model I propose explains the practicality of moral language in terms of Paul Grice's notion of *generalized conversational implicatures*. This notion is quite commonly appealed to in different areas of philosophy and linguistics, but generally unexplored in this part of metaethics.

In the next section, I spell out the relevant notion of the practicality of moral language. In Section 3, I provide a brief overview of relevant metaethical positions and spell out how this paper contributes to the debate. It is observed that expressivism has a number of difficulties which motivate that we consider a purely cognitivist view combined with a pragmatic story of the practicality of moral language. In Section 4, I provide a short reminder of Grice's theory of implicatures. In Section 5, I argue that the two most prominent pragmatic accounts, proposed by David Copp and Stephen Finlay, suffer from important problems. In the main part of the paper, Section 6 through 13, I develop The Dual Aspect Account. Moral conversations have dual purposes: to communicate moral beliefs and to influence action. Correspondingly, a person's utterance of a sentence to the effect that an action has a moral characteristic conveys two things: First, such a sentence, in virtue of having a certain conventional meaning, expresses the belief that the action has a moral property. Second, an utterance of this type of sentence carries a generalized conversational implicature to the effect that the person in question has a certain action-guiding attitude in relation to the action. This account, it will be argued, captures the duality of moral language, and particular, offers a model of the various mechanisms that make moral language practical. The Dual Aspect Account has important advantages over competing pragmatic views of the practicality of moral language. It does not presuppose any contentious pragmatic notion, such as conventional implicature, and it does not imply that utterances of complex moral sentences, such as conditionals, necessarily convey attitudes (in contrast to Copp's view). It is able to explain the intimate, "meaning-like," connection between moral language and attitudes, and it does not presuppose any particular cognitivist theory (in contrast to Finlay's view). Indeed, it is compatible with any purely cognitivist view of the relevant kind. Moreover, the conversational implicature in question is calculable and cancellable in Grice's sense.

2. The Practicality of Moral Language

Moral language has been claimed to be practical in a number of respects, but the one that is pertinent here distinguishes moral language from other parts of language.¹ One feature that seems to be characteristic of moral language is that we take for granted that

¹ See e.g. Smith (1994), 6; Blackburn (1998), 1–4; Finlay (2004), 205–207, and Railton (2006), 203.

certain uses of it are accompanied by attitudes. Thus, then notion in question can be characterized as follows:

The Practicality of Moral Language: It is generally the case that if a person S employs a sentence to the effect that ϕ ing has a certain moral characteristic, such as “ ϕ is wrong,” we presume that S has a certain action-guiding attitude in relation to ϕ ing.²

That is, we generally assume that a person who employs a moral sentence of this type has a certain attitude towards the action in question, *even* if we do not possess any special information to the effect that she *has* that attitude. In particular, we do so even though we do not have any special information about her or the situation in which she finds herself. Thus, when a person uses a moral sentence of this type, we automatically infer that she has the attitude without reflecting on the particularities of the circumstances at issue. The existence of such a connection between moral language and attitudes is the default understanding; in case there is any exception to it, this is something that calls for explanation. In this characterization, relevant uses of different parts of moral language (e.g. “good,” “ought,” “right”) may be connected to different kinds of attitudes (e.g. approval, states of exhortation or command, and various forms of motivation), and possibly to a number of such attitudes.

Non-moral language is usually not practical in this way. To take an example, we might come to believe that a person who says “It’s raining” wants an umbrella. However, this is because we have access to particular information—about the situation, say, or what kind of person she is—that gives us grounds for this assumption.

One important reason of why moral language is practical is presumably that it thereby serves a major function: it enables us to regulate one another’s behaviour. That is, moral language is generally utilized to get us to act, or to get us to refrain from acting, in certain ways so as to adjust our various actions in relation to one another.³ A very general way of affecting other people’s behaviour is to let them know what our attitudes towards it are. Assume that it generally is the case that when a person says that an action is wrong, we infer that she has a negative attitude towards it being performed.

² Cf. Finlay (2004), 212. I use “moral characteristic” in a way that is neutral between expressivism and cognitivism.

³ See e.g. Gibbard (1990), part IV, and Blackburn (1998), Ch. 1.

This might make us less willing to do the action or to contribute to it being carried out. The underlying explanation for our responses is presumably very basic, for example that we need to avoid conflicts with other people and need to cooperate with them.⁴

3. Metaethics and the Practicality of Moral Language

Consider a moral sentence of the type “ ϕ ing is wrong.” It has a conventional semantic meaning that is determined by the meaning of its various parts. Assume that a person who understands its meaning asserts or accepts it. It is then plausible to suppose that she is in a state of mind that has a content which answers to the meaning of the sentence. This might be put by saying that, in virtue of its meaning, the sentence expresses such a mental state. A moral sentence can thus be understood to express the minimal mental state that a person has to be in, in order for it to be consistent with the conventional meaning of the sentence to assert or accept it. I will refer to such a mental state as “moral judgment.”⁵ We can now distinguish four types of metaethical positions.

Pure expressivism: A sentence of the type “ ϕ ing is wrong” expresses a moral judgment which consists of (1) a non-cognitive state, but which (2) does not involve any belief.

Following Michael Ridge, there are two kinds of ecumenical views.

Ecumenical expressivism: A sentence of the type “ ϕ ing is wrong” expresses a moral judgment which consists of (1) a non-cognitive state, and (2) a belief. However, (3) the moral sentence is not true by virtue of the belief in (2) being true.

Ecumenical cognitivism agrees with (1) and (2), but states that (3) the moral sentence is true if the belief referred to in (2) is true.⁶

These three positions entail that moral judgments involve a non-cognitive state of some kind. Non-cognitive states, such as desires, are mental states that have the

⁴ One explanation of why moral language has this function might in turn be that it enables us to coordinate our actions so as to benefit mutually; see e.g. Gibbard (1990), 61–68.

⁵ See e.g. Ridge (2003), 563–574; Kalderon (2005), Ch. 2, and Schroeder (2008), 86–116.

⁶ Blackburn (1998) and Gibbard (2003) defend pure expressivism. For defences of ecumenical views, see Ridge (2006), 302–336; Barker (2002), 268–279; Boisvert (2008), 169–203, and Eriksson (2009), 8–35.

function to move to action, and they thus exemplify what I call “action-guiding attitudes.”

Lastly, there is the traditional type of cognitivism:

Pure cognitivism: A sentence of the type “ ϕ ing is wrong” expresses a moral judgment which consists of (1) a belief, but which (2) does not involve any non-cognitive state.

There is also another kind of view that is important in this context: internalism. The kind of internalism that is relevant here maintains that a sentence of the type “ ϕ ing is wrong” expresses a moral judgment such that if a person S holds that judgment, then S is motivated not to ϕ . There are purely expressivist, ecumenical, and purely cognitivist versions of internalism.⁷ However, according to any of these views, being motivated amounts to having an action-guiding attitude: a state that moves to action.

We can now see that (i) pure expressivism, (ii) ecumenical expressivism, (iii) ecumenical cognitivism, and (iv) internalism each entails:

The Attitude Implication Thesis (The AI Thesis): A sentence of the type “ ϕ ing is wrong” expresses a moral judgment such that if a person S holds that judgment, then S has an action-guiding attitude of some kind in relation to ϕ ing.

This is an abstract formulation of *The AI Thesis*, and it exists in a number of versions. Conditional versions of it maintain that the necessary connection holds only if the person in question satisfies a particular condition, such as rationality or normality, where this qualification does not render the thesis trivial.⁸

Metaethical views that entail *The AI Thesis* seem well suited to account for the practicality of moral language since they maintain that a moral sentence, in virtue of its very meaning, expresses a moral judgment that implies an action-guiding attitude.

⁷ According to the latter view, a moral judgment consists in a belief which motivates to action, either by consisting in a cognitive action-guiding attitude or by giving rise to a separate action-guiding attitude of some kind. See e.g. Nagel (1970), Ch. 1; Shafer-Landau (2003), Ch. 5, and Wedgwood (2007), Ch. 1. As I characterize “action-guiding attitude,” it can consist in either a non-cognitive or a cognitive attitude.

⁸ See e.g. Korsgaard (1996 (1986)), 311–334; Dreier (1990), 6–26, and Smith (1994), esp. Ch. 4. For defences of externalism, see e.g. Brink (1989), Ch. 3; Sayre-McCord (1997), 267–292, and Svavarsdóttir (1999), 161–219.

Indeed, this is one of main reasons to adopt any of these metaethical views. It can be maintained, in this view, that the reason of why we presume that a person who employs a moral sentence has a certain attitude is simply that we know, because we know the meaning of the sentence, that it expresses such a judgment. In contrast, metaethical views that do not entail *The AI Thesis* might seem far less suited to explain this essential feature of moral language.

However, it is well-known that the various metaethical views that entail *The AI Thesis* have a number of serious problems.⁹ In this paper, I will however limit myself to arguing that *any* version of *pure cognitivism* which *denies The AI Thesis* can explain the practicality of moral language by employing a certain pragmatic account. The main challenge for this approach is to provide a pragmatic story of the tight, “meaning-like,” connection between moral language and attitudes as it manifests itself in the practicality of moral language.

4. Grice’s Notion of Implicatures

Paul Grice argues that we are justified to assume that people who are involved in conversations with one another conform to the cooperative principle: “Make your conversational contribution such as is required at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.”¹⁰ He identifies four maxims, recognition of which enables us to comply with this principle.¹¹ The basic idea is that on the presumption that we conform to the cooperative principle and the maxims, our utterances can convey information that is not part of the conventional meaning of the sentence we utter: “conversational implicatures.”¹²

⁹ Pure expressivism has a number of potential weaknesses, such as the Frege-Geach problem and difficulties to account for moral normativity. In Strandberg (forthcoming) I argue that expressivism is committed to an unreasonably strong version of internalism. Mark Schroeder argues in a recent article that ecumenical views also face serious challenges (Schroeder (2009), 257–309). For instance, he argues that ecumenical expressivism cannot solve the Frege-Geach problem, and that other ecumenical positions have serious drawbacks of their own, such as difficulties to account for moral motivation, the expression relation, and attitude-ascriptions. Internalists face well-known problems, such as to account for amorality and to formulate a non-trivial version of their view. Cognitivist internalism has the well-known problem of being incompatible with the contention that beliefs cannot motivate.

¹⁰ Grice (1989a (1975)), 26.

¹¹ The maxims of *quantity* (make your contribution neither less nor more informative than is required to fulfil the mutually accepted purpose of the conversation); *quality* (say what is true and not what is false); *relation* (make your contribution relevant in view of the purpose of the conversation), and *manner* (make your contribution perspicuous, avoiding ambiguity and the like).

¹² Utterances involve not merely spoken sentences but intentional production of sentences in accordance with their meaning (see e.g. Boisvert (2008), 175).

As far as I understand, Grice does not provide a general definition of conversational implicature. However, this notion can perhaps be characterized as follows:¹³ A person S's utterance of a sentence *s* to the effect that *p* conversationally implicates *q*, provided that the assumption that S thinks *q* is required in order to make S's utterance of *s* consistent with the supposition that (i) S adheres to the cooperative principle and the maxims so as to fulfil the mutually accepted purpose of the conversation; (ii) S knows the meaning of *s*; (iii) S is aware of the context of the utterance, and (iv) S believes that those she converses with are able to recognize that the assumption that she thinks *q* is required to understand her as conforming to the cooperative principle and the maxims. Grice argues that there are different criteria that need to be fulfilled if something is to qualify as a conversational implicature, two of which are particularly important: it should be possible to *calculate* it and to *cancel* it.¹⁴

There are two types of conversational implicatures in Grice's view. Assume that an utterance of a sentence *s* to the effect that *p* conversationally implicates *q*. It is a *particularized conversational implicature (PCI)* in so far as an utterance of *s* does not standardly implicate *q*. The utterance carries the implicature only provided that the context in which it is uttered has certain specific features, in other words features pertaining to the individual context in question, which support this understanding of the utterance. Thus, a *PCI* needs specific contextual features to arise. It is a *generalized conversational implicature (GCI)* in so far as an utterance of *s* *standardly* implicates *q*. The utterance carries the implicature even if the contexts in which it is uttered do not have any specific features. Thus, a *GCI* does not need such contextual features to arise—indeed, it arises unless there are any specific circumstances which defeat it—and

¹³ Grice provides an account of a person implicating something by uttering a sentence (Grice (1989a (1975)), 30–31). With the present formulation, it is an utterance of a sentence which implicates something. This is also how *GCIs* are understood in the recent literature (see e.g. Levinson (2000), 22–23). Following Grice, I understand “think” in a way that leaves open for other mental states than beliefs. Grice's characterization of implicature has been criticized by e.g. Harnish (1998 (1976)), 256–267; Neale (1992), 528–529, and Davis (1998), Ch. 1.

¹⁴ It is calculable if it is possible to construe an argument that takes into considerations (i) through (iv) and concludes that the utterance carries the implicature. Grice also suggests other criteria (see Grice (1989a (1975)), 31, 39), but they are more contentious; see e.g. Sadock (1998 (1978)), 315–331, and Nunberg (1981), 199–222.

thereby constitutes the default understanding of the utterance.¹⁵ However, in spite of this, *GCI*s come into existence as a result of the same mechanisms as *PCIs*.¹⁶

In Grice's view, there is also another sense in which a sentence *s* to the effect that *p* can implicate *q*. It is a *conventional implicature (CI)* in so far as an utterance of *s* implicates *q* in virtue of the conventional meaning of a term that is part of *s*, but *q* does not affect what is said in Grice's sense.¹⁷ A *CI* does not need the cooperative principle and the maxims to arise and this kind of implicature is not calculable or cancellable. Whereas the notions of *PCI* and *GCI* are well-established, the notion of *CI* is more contentious.¹⁸

5. Alternative Pragmatic Accounts: Copp and Finlay

In Grice's view, *GCI* is the pragmatic notion that is most useful in philosophy. He argues that there are cases where it might be thought that a certain piece of information is part of a sentence's meaning, although it is a matter of a *GCI*, the reason being that utterances of such a sentence regularly carry the implicature across various contexts.¹⁹ Numerous philosophers and linguists have accordingly made use of this notion in various areas where what belongs to the meaning of a given expression is a matter of dispute.²⁰ In this light, it is slightly surprising that it has not been employed in the debate on the practicality of moral language to any significant extent.²¹ There are, however, two prominent cognitivist and pragmatic accounts that make use of other Gricean concepts.²² Although I cannot make full justice to the depth and originality of these ideas here, I will make some points which motivate that we consider an alternative pragmatic account.

¹⁵ Grice (1989a (1975)), 37, and Grice (1981), 185. As this connection is standardized, it is plausible to assume that it is not, strictly speaking, individual utterances that carry *GCI*s, but rather *types* of utterances. Cf. Levinson (2000), 22–26.

¹⁶ Cf. Harnish (1998 (1976)), 267; Neale (1989), 535–535; Carston (2002), 111; Recanati (2003), 302, and Horn (2004), 6.

¹⁷ See e.g. Grice (1989 (1975)), 25–26.

¹⁸ See esp. Bach (1999), 327–366.

¹⁹ See e.g. Grice (1989a (1975)), 37–38, and Grice (1981), 185.

²⁰ For a brief overview, see Levinson (2000), Ch. 2.

²¹ However, see footnote 29 below on Copp's first formulation of his account.

²² Neither Copp's nor Finlay's views entails *The AI Thesis*. In this paper, I will not comment on Bar-On and Chrisman's account since it is not pragmatic in nature (Bar-On and Chrisman (2009), 133–167). However, one important difference between their view and mine is that I provide a detailed explanation of why moral utterances convey attitudes.

David Copp argues that a person's use of a sentence of the type "ϕing is wrong" carries a *CI* to the effect that she has an action-guiding attitude in relation to ϕing. Accordingly, he thinks that this fact is explained by the conventional meaning of the moral term that is part of the sentence.²³ In this respect, Copp maintains, moral terms are analogous to pejoratives. For example, in his view an utterance of a sentence of the type "X is a wop" carries a *CI* to the effect that the speaker feels contempt for Italians.

It is regularly observed that *CI*s are generally inherited by utterances of complex sentences, such as conditionals.²⁴ That is, if an utterance of a sentence carries a certain *CI*, then an utterance of a complex sentence, where the first sentence occurs embedded, carries that *CI* as well. Consider "but", one of Grice's main examples of a term whose conventional meaning gives rise to a *CI*.²⁵ In his view, an utterance of the sentence "It's cold but sunny" carries a *CI* to the effect that there is a contrast between it being cold and it being sunny. An utterance of a complex sentence where this sentence is embedded carries the *CI* in question as well. Thus, the utterance "If it is cold but sunny, I'll take a walk" conveys the aforementioned contrast. There is a straightforward explanation of why *CI*s are inherited in this manner: A term makes the same contribution to the meaning of the sentences in which it is involved, irrespective of whether these sentences occur freestanding or embedded. In consequence, since a *CI* is a function of the conventional meaning of a particular term, an utterance of a complex sentence in which a sentence involving such a term is embedded carries the *CI* in question as well.²⁶ The same appears to hold for pejoratives; for example, an utterance of the sentence "If Mario is a real wop, he'll have pasta for dinner" conveys the speaker's disrespect for Italians.

Importantly, in contrast to what Copp's view predicts, moral utterances do not follow this pattern. If a person utters the sentence "It's wrong to hit one's children," it is plausible to think that her utterance conveys that she has a certain attitude towards hitting one's children. However, a person's utterance of the sentence "If it's wrong to hit one's children, it's wrong to hit other people's children" does not necessarily convey any attitude. In discussing belief ascriptions, Copp recognizes a related difficulty for his

²³ Copp (2009), 175, 184. Copp adopts the notion of a "conventional simpliciture" (instead of "implicature"), a complication which, although interesting, is not of immediate relevance here.

²⁴ See e.g. Karttunen and Peters (1979), 33–48, and Barker (2003), 1–33.

²⁵ Grice (1989b (1975)), 46.

²⁶ Cf. Potts (2005), 28.

view that moral terms are analogous to pejoratives.²⁷ He responds by distinguishing between two kinds of semantic rules: As regards pejoratives, there is a semantic rule saying that one is not to use *a single term*, such as “wop,” in *any* sentence unless one has the relevant attitude. As regards moral terms, there is a semantic rule saying that one is not to use *an entire sentence of a certain type where a moral term is involved*, for example a sentence of the type “ ϕ ing is wrong,” unless one has the relevant attitude in relation to ϕ ing.²⁸

The fact that uses of a sentence carry a certain *CI* is, as noted, explained by the conventional meaning of a particular term, such as “but,” which is part of the sentence. However, the explanation of why uses of a sentence of the type “ ϕ ing is wrong” convey an attitude does not seem to be that the sentence involves a moral term with a certain conventional meaning. The reason is that if a person utters a complex sentence where this sentence occurs embedded, her utterance does not standardly convey any attitude. This strongly suggests that in the moral case it is not a matter of a *CI* at all. Furthermore, the fact that only utterances of certain types of sentences in which a moral term is involved convey the attitude in question indicates that it is a person’s *use* of a *particular type of moral sentence* that explains that her utterance conveys the attitude, and *not* the *conventional meaning* of the *moral term* that is part of the sentence she uses. This indicates in turn that the fact that moral utterances convey attitudes should not be explained by semantic rules, but rather by purely pragmatic factors. Copp thinks we have to adopt an explanation that refers to the conventional meaning of moral terms so as to account for our intuition that if a person utters a sentence of the type “ ϕ ing is wrong” without having any attitude in relation to ϕ ing, she uses the moral term in a linguistically inappropriate manner. However, I will argue below that my pragmatic account can explain this intuition.²⁹

²⁷ For a criticism of Copp’s account, see Finlay (2005), 10–20, and Bar-On and Chrisman (2009), 154–157. Schroeder (2009), 304–306, argues that moral terms and pejoratives are not analogous.

²⁸ Copp (2009), 187–188. Copp might reply that there are pejoratives that work as moral terms. He argues (in correspondence) that the utterance “If Smith is a jerk, he won’t return the book” does not convey any attitude towards anyone. However, I have difficulties to see that this is a relevant comparison. While the utterance in question does not convey contempt for Smith, it does convey contempt for a certain group of people: “the jerks.” In this way, “jerk” functions in the same way as other pejoratives. An utterance such as “If Mario is a real wop, he’ll have pasta for dinner” does not convey contempt for Mario, but for a certain group of people: Italians. Moreover, the important point is not that moral terms do not function as pejoratives, but that they do not function as *CIs*. Perhaps pejoratives do not work as *CIs* either; indeed, there are some reasons for this view (see Hom (2008), 416–440).

²⁹ In an earlier paper, Copp made use of the notion of conversational implicature as part of his view (Copp (2001), 31–33): A person’s use of a sentence of the type “ ϕ ing is wrong” carries a conventional

It might be objected that it seems that there are cases where *CI*s are not inherited by utterances of complex sentences in the way described above. Moreover, there seem to be cases where conversational implicatures (*PCIs* or *GCI*s) are inherited by such utterances.³⁰ However, there are important differences between *CI*s and conversational implicatures here. First, it is generally agreed that *CI*s standardly are inherited by utterances of complex sentences. The reason is that *CI*s are explained by the conventional meaning of certain terms. Accordingly, the cases in need of explanation are those where *CI*s are *not* inherited. Moreover, according to a plausible explanation of why there are cases where *CI*s are not inherited, they still retain their identity as implicatures. They contribute to utterances of complex sentences in a systematic way by virtue of the conventional meaning of the relevant terms, and the reason why a particular utterance of a complex sentence does not carry a *CI* is explained by certain “intervening” factors.³¹ Second, although it is agreed that there seem to be cases where conversational implicatures are inherited, it is generally agreed that there is nothing in this notion that implies that this must be so.³² The reason is that conversational implicatures are explained by pragmatic factors, and not in terms of meaning. If anything needs to be explained, it is the cases where conversational implicatures *are* inherited. Moreover, such an explanation seems to offer itself: Whether a conversational implicature is inherited or not depends on what it is about the utterance of the embedded sentence that explains that it carries the implicature in the first place.³³ In section 10, I will explain why the relevant *GCI* is not inherited in moral cases.

It might further be objected that there is a way of skirting my objections to Copp’s account. Daniel Boisvert has put forward a novel version of ecumenical cognitivism according to which moral terms are analogous to pejoratives. In his view, a person’s

implicature (*CI*) and a conversational implicature to the effect that she subscribes to a certain moral standard. Later he maintained that he had a *GCI* in mind (Copp (2009), 192). This view differs from that proposed in the present paper. First, as far as I see, Copp’s suggestion stems from his controversial view of the meaning of moral sentences, which he shares with Finlay. Second, he provides only a rough statement of this part of his view. Third, as we will see below, his view is quite different from the account I propose. Finally, Copp has recently abandoned this view on the grounds that the conversational implicature in question would not be calculable (Copp (2009), 192–193).

³⁰ For discussions of these two types of cases, see e.g. Recanati (1989), 295–329; Green (1998), 61–90; Barker (2003), 1–33; Recanati (2003), 299–332, and Carston (2004), 65–100.

³¹ See e.g. Barker (2003), 19–20, 23–24.

³² For discussion, see e.g. Recanati (1989), 321–325; Green (1998), 73–80; Recanati (2003), 299–332, and Carston (2004), 73–78. As Green points out, not all *GCI*s are inherited by utterances of complex sentences.

³³ See e.g. Levinson (2001), 157–164.

utterance of a sentence such as “If it’s wrong to hit one’s children, it’s wrong to hit other people’s children,” intuitively conveys an attitude, not towards any particular action, but towards actions that have the property of being wrong *in general*.³⁴ However, I question the plausibility of this idea. First, it is reasonable to doubt whether Boisvert’s intuition is representative; I, for one, do not feel that someone who uses such a sentence needs to have any attitude towards actions that are wrong *in general*. This doubt is reinforced by our responses to uses of pejoratives and moral terms, respectively. Assume that a person says “If Sam is a n****r, he’s of African origin” and then adds “But I don’t have anything against black people.” A reasonable response is “But if you don’t dislike black people, *why* do you use that hideous word?” Without an appropriate explanation, we might even think that she is not fully aware of its meaning. However, in case a person says “If it’s wrong to give money to charity, it’s wrong to give money to Oxfam” and then adds “But I don’t care whether such actions are wrong or not,” I do not think we would question her use of “wrong;” rather, she merely seems to make some kind of inference. Second, it is difficult to see *why* we should have acquired the device of using moral terms in this way. It is easy to see why we use a sentence of the type “ ϕ ing is wrong” to convey a negative attitude towards particular actions, since we then might influence people not to perform them. However, it is difficult to understand why *each time* we use “wrong” we should want to commit ourselves to having a negative attitude towards *all* actions that are wrong, including the innumerable actions we have not even considered.

Stephen Finlay defends a view that rests on a theory of the meaning of moral sentences according to which a sentence of the type “ ϕ ing is wrong” expresses a belief to the effect that ϕ ing does not conform to a social moral standard, where this standard might vary in different contexts.³⁵ Finlay argues that a person’s use of such a sentence carries a conversational implicature that conveys that she subscribes to a moral standard where “subscribe” refers to an action-guiding attitude. Thus, a moral sentence such as “It’s wrong to hit one’s children” expresses, when fully spelled out, that such actions do not conform to a certain moral standard. Assume that a person uses a moral sentence in which the specification of the relevant standard is omitted, but that such a standard is presumed in the context in question. In Grice’s view, the person’s audience is

³⁴ Boisvert (2008), 185–186.

³⁵ Finlay (2005), 1–20; cf. Finlay (2004), 205–225. For criticism of Finlay’s view, see Copp (2009), 190–193, and Bar-On and Chrisman (2009), 153.

conversationally licensed to take her utterance to convey any information that is needed to understand it as an intelligible contribution to the conversation. Finlay thinks that such an utterance “will therefore conversationally implicate that the speaker possesses some attitude, just in case her possessing that attitude is (part of) the context that her audience must presuppose in order to be able to identify the omitted standard-relativization.”³⁶

Finlay’s account rests on a contentious theory of the meaning of moral sentences that seems open to standard objections to relativist views as regards belief ascriptions.³⁷ However, even if we grant this theory, it is difficult to see that he succeeds in showing that the conversational implicature in question is calculable. It can be accepted that, in this view, a person’s utterance of a sentence such as “It’s wrong to hit one’s children,” in which reference to a moral standard is omitted, implicates that such actions *do not meet the standard that is presupposed in the context in question*. However, it is unclear why it should implicate that the person has any particular *attitude* in relation to these actions. Finlay maintains, as we have seen, that the audience needs to assume that she has a certain attitude in order to be able to identify the moral standard to which she implicitly refers. However, it is difficult to understand why this needs to be the case. One reason is that, in this view, it seems that the audience should be able to identify the standard in question simply by being aware of what standard is contextually presumed.³⁸ Finlay argues that adherence to a moral standard is often socially demanded, which gives rise to an expectation that people who make moral utterances have attitudes which conform to the standard in question. However, it seems that a person’s moral utterance might convey that she has a certain attitude in cases where a particular moral standard is not socially demanded, as there is no agreement in the social setting in question over the appropriate standard.³⁹

³⁶ Finlay (2005), 12.

³⁷ See Schroeder (2009), 284–287, and Bar-On and Chrisman (2009), 155–156.

³⁸ On the other hand, in cases where no moral standard is contextually presumed, it is difficult to see why the assumption that a person has a certain attitude would make her audience conclude that she has a certain standard in mind. The idea is perhaps that a person would not utter a sentence that makes implicit reference to a certain standard unless she has an attitude that conforms to the standard, because in that case her utterance would lack in relevance. However, there is nothing in Finlay’s account to explain *why* this should be the case. Besides, this does not explain why the audience has to assume that she has a certain attitude in order to identify what standard she embraces.

³⁹ In an earlier paper, Finlay defended a slightly different view on the meaning of moral sentences, according to which a sentence of the type “ ϕ ing is wrong” would express a belief to the effect that ϕ ing fails to satisfy a (somehow) specified interest (Finlay (2004), 205–223). He argues that a person’s utterance of a moral sentence, where the specification of the interest is omitted, conversationally

Nevertheless, even if we grant that Finlay succeeds in calculating the conversational implicature he has in mind, his account is not apt to explain the practicality of moral language. In his view, a person's moral utterance implicates that she has a certain attitude by virtue of the fact that a given moral standard is presumed in the context in question. This indicates that the conversational implicature at issue is a *PCI*, since the implicature arises by virtue of particular information being presupposed in the context: that a particular standard is taken for granted. Imagine, however, that a person says "It's wrong to hit one's children" in a context where no particular moral standard is presumed, perhaps because it is controversial in that context what the correct standard is. Yet, it seems that her utterance conveys that she has a negative attitude towards such actions. Indeed, we do not need to make any particular assumptions about her, or the situation in question, to arrive at this conclusion. This suggests that an account formulated in terms of *PCI* is ill suited to account for the "meaning-like" connection between moral language and attitudes that is evinced by the practicality of moral language. This notion means that we *presuppose* a person who employs a sentence to the effect that an action has a moral characteristic to have a certain attitude in relation to the action. More particularly, we think this *even* if we do not have access to any specific information about her or the situation, for instance that a certain moral standard is presumed.

The discussion in this section can now be summarized. *On the one hand*, Copp's account in terms of *CI* might be able to establish a close, "meaning-like," connection between moral language and attitudes, since it employs the conventional meaning of moral terms, but it runs into difficulties because what explains this connection does not seem to be the conventional meaning of moral *terms*, but rather *uses* of certain *types of moral sentences*. *On the other hand*, Finlay's account in terms of *PCI* does not have this latter problem, but fails to explain the very close connection between moral language and attitudes as it manifests itself in the practicality of moral language. In view of this, there are reasons to investigate an account in terms of *GCI*. There are grounds to think that such an account would be able to avoid these difficulties since it explains the

implicates that she has a certain interest or attitude in relation to the action at issue. Again, a person's utterance of a sentence such as "It's wrong to hit one's children" may implicate that such actions fail to satisfy a given interest, but it is difficult to see why it should implicate that she has any particular interest or attitude in relation to it. Finlay suggests that it does so in contexts where the interest is not only presumed, but where it also is presupposed that the audience shares it. Yet, consider a case where a person utters this sentence in a context where people do not share the interest in question. Her utterance would still convey an interest or attitude. Finlay suggests that this can be explained as a rhetorical effect, but it is difficult to see that this can explain *all* such cases.

connection in terms of uses of certain types of moral sentences, at the same time as it implies that there is a standardized connection between moral language and attitudes. Moreover, such an account should ideally be neutral towards different cognitivist views, and it should be possible to demonstrate how the implicature in question can be calculated.

6. The Dual Aspect Account

In what follows, I will develop the following view:

The Dual Aspect Account (DAA): A person S's utterance of a sentence of a type according to which ϕ ing has a certain moral characteristic, such as " ϕ is wrong," conveys two things: (i) The sentence expresses, in virtue of its conventional meaning, the belief that ϕ ing has a moral property. (ii) An utterance of this type of sentence carries a generalized conversational implicature, *GCI*, to the effect that S has a certain action-guiding attitude in relation to ϕ ing.

In other words, an utterance of an assertive moral sentence conveys both a belief and an attitude, but while the former is explained by the conventional meaning of the sentence, the latter is explained by a *GCI* that is carried by an utterance of the sentence.

7. Dimensions of Conversational Purposes

A basic thought in Grice's theory is that an utterance carries a conversational implicature relative to a mutually accepted purpose of the conversation in which it occurs. We may start by observing that purposes of conversations may vary along at least three dimensions. (i) They might vary in number; in particular, a conversation might have more than one purpose. Grice explicitly claims that a conversation might have a "set of purposes" and thus acknowledges that a conversation might have more than one aim.⁴⁰ (ii) They might vary as regards their content. Many conversations have as an obvious purpose to exchange beliefs on a certain subject matter. This is what Grice focuses on in most of his examples where "giving and receiving information" is at issue.⁴¹ However, he also recognizes that the aim of a conversation might have other

⁴⁰ Grice (1989a (1975)), 26.

⁴¹ Grice (1989a (1975)), 30; see also 28.

contents, and he especially mentions “such general purposes as influencing or directing the actions of others” or “influencing or being influenced by others.”⁴² (iii) They might vary in generality. The purpose of a conversation can vary from being quite particular, in the sense of being the aim of few conversations of a certain kind, to being quite general, in the sense of being the aim of a great number of conversations of a certain kind.

These distinctions can now be applied to moral conversations: (i) A moral conversation might have more than one purpose; (ii) these purposes might have different contents, and (iii) the kind of conversation that moral conversations constitute might have these purposes more or less generally.

8. Dual Purposes of Moral Conversations

I contend that it is reasonable to assume that moral conversations have two purposes, both of which are general. These purposes require thorough discussion, but in the present paper I am merely able to provide a basic outline of them.

First, moral conversations generally have as a purpose to communicate what the participants in such a conversation believe about matters that bear on the moral issue they are conversing about. Most importantly, in moral conversations we let one another know that we believe—and hence hold true—that some actions have certain moral properties.⁴³

According to pure cognitivism, this purpose is intimately connected with the conventional meaning of moral sentences, since this view states that moral sentences, in virtue of their conventional meaning, express beliefs, but not any action-guiding attitudes.

Importantly, *DAA* does not entail any particular cognitivist view. However, as we will see, it presumes one essential but uncontroversial claim: A sentence of a type according to which ϕ ing has a moral property entails the existence of a moral reason in relation to ϕ ing. Accordingly, a sentence of the type “ ϕ ing is wrong” entails that there is a moral reason not to ϕ . This kind of entailments has its source in the fact that if an action has a certain moral property, it follows that there exists a moral norm, or

⁴² Grice (1989a (1975)), 28 and 30. Cf. Harnish (1998 (1976)), 257.

⁴³ The beliefs we communicate may of course relate to other moral matters. It might be objected that the purpose of this type of conversations is not that of communicating beliefs but, in Grice’s words, that of “giving and receiving information.” However, as the basic way of transmitting information is to let other people know about our beliefs this purpose implies the one considered above.

standard, which justifies the action from a moral perspective.⁴⁴ Corresponding claims hold for other kinds of reasons and norms. According to an influential view, a reason to perform an action consists in a fact which “counts in favour” of doing it and, similarly, a reason not to perform an action consists in a fact which “counts against” doing it.⁴⁵ Thus, “ ϕ ing is wrong” entails that, *from the perspective of morality*, there is a fact which counts against ϕ ing. This claim is uncontroversial, especially as it does not entail that moral reasons constitute a subclass of normative reasons, where normative reasons are understood in terms of rationality.⁴⁶

Second, moral conversations generally have as a purpose to influence behaviour; in particular, they generally have as a purpose to contribute to people performing or not performing certain actions.⁴⁷

Most evidently, there are cases where someone aims to influence other people’s behaviour (moral advice, blame, commendation, persuasion, supervision, upbringing, etc.) and where people aim to influence one another’s behaviour (moral arguments, etc.). There are also cases where a person is open to having her own behaviour influenced by others without trying to influence their behaviour in return, and cases where people converse about moral issues without having any predetermined moral opinions when they enter the conversation, but where such opinions might evolve as it proceeds.

Three things are especially important to note here. First, as is clear from what Grice writes, we need not be consciously aware of the purposes of the conversations in which we are involved. When it comes to moral conversations, this contention is particularly plausible given that it is a general purpose of such conversations to influence behaviour, since it means that this purpose permeates contexts in which we converse about moral issues. In other words, it is reasonable to think that this purpose is habitual, and not in the foreground of our moral deliberation. Thus, the assumption that moral language has this purpose works as a tacit notion in moral conversations—a purpose on which we usually do not reflect, but which governs our understanding of moral utterances. Second, even if we *are* consciously aware that it is a purpose of a certain moral

⁴⁴ See e.g. Brink (1992), 1, 8–9; Smith (1994), 95–96; Copp (1995), 190, and Shafer-Landau (2003), 166.

⁴⁵ See e.g. Scanlon (1998), 17.

⁴⁶ In contrast, see e.g. Smith (1994), esp. Ch. 5, and Shafer-Landau (2003), Ch. 8.

⁴⁷ Cf. Copp (2001), 32.

conversation to influence people in some respect, this need not equate to aiming to influence their behaviour. For example, we might take a conversation to have the purpose to influence people's beliefs and attitudes, and the influence the conversation has on them in this respect might then influence their behaviour indirectly. Third, the subject under discussion in a moral conversation is not necessarily the behaviour that is the object of influence. Suppose, for example, that someone argues that certain actions are wrong in a situation where none of the participants are in a position to perform such actions. Nevertheless, this might influence their behaviour, for example by making them reluctant to perform similar actions. Similarly, although the behaviour we seek to influence is ongoing or future behaviour, we might evoke past or hypothetical behaviour, since doing so might affect the behaviour we are interested in.

The view that moral conversations generally have as a purpose to influence behaviour is, as far as I understand, widely agreed on. Moreover, it finds support in what I said earlier. We observed that an essential feature of moral language is that it is practical in a certain sense. We also observed that it is widely accepted that moral language thereby fulfils the essential function of regulating people's behaviour. Granted that moral language serves to regulate behaviour in this way, it is reasonable to assume that moral conversations generally have as a purpose to influence behaviour. Put abstractly: if this purpose is successfully carried out, and people who are involved in various moral conversations actually succeed in influencing one another's behaviour, the overall result will be that people regulate one another's behaviour and, hence, that moral language fulfils this function. As it is generally acknowledged that moral language has this function, it is plausible to assume that it has the mentioned purpose. Conversely, if it were not a general purpose of moral conversations to influence behaviour, it is difficult to see how moral language would be able to fulfil this function, since moral language in that case would work in a way that is not reflected in our purpose in using it.⁴⁸

⁴⁸ It is important to note that metaethical positions that do entail *The AI Thesis* also have to assume the second purpose of moral conversations in order to account for the function of moral language. To see this, suppose that *The AI Thesis* is correct, but that it is not a general purpose of moral conversations to influence behaviour. In that case, utterances of moral sentences that express judgments which imply action-guiding attitudes would not serve the function to regulate behaviour. The basic reason is that we would not identify these utterances as saying things that are entirely relevant in moral conversations. In particular, the utterances would not completely accord with the cooperative principle and the maxim of relation, and they would thus not contribute to moral conversations in an appropriate manner. As a result, they could not serve the mentioned function. In this light, it is not surprising that expressivists stress that the purpose of moral language is to influence behaviour (see e.g. Blackburn (1998), 1–4).

We have seen that according to pure cognitivism, the first general purpose of moral conversations is intimately connected to the conventional meaning of moral sentences. In contrast, assuming that pure cognitivism is combined with a denial of *The AI Thesis*, the second general purpose of moral conversations does not seem to be connected to the meaning of moral sentences. However, as Grice's discussion of conversational implicatures bears witness, there are numerous cases where conversations have purposes that are not reflected in the meaning of the sentences uttered in them.

9. First Basis of the Dual Aspect Account

There are two bases of the second part of *DAA*—the part intended to account for the practicality of moral language—which I will discuss in turn in this and the following section.

First, we have seen that moral conversations have dual purposes: to communicate beliefs about moral matters, and to influence behaviour. Second, in describing the practicality of moral language, we saw that one general way to affect people's behaviour is to let them know about our attitudes towards it. Finally, we have seen that a sentence of the type “ ϕ ing is wrong” entails that there is a moral reason not to ϕ , and hence, that from the perspective of morality, there is a fact that counts against ϕ ing.

Now, assume that a person who is involved in a moral conversation which has this purpose utters a sentence of the type “ ϕ ing is wrong,” for example “It's wrong to hit one's children.” How would we conceive of her utterance? It seems that in order for us to make adequate sense of it, we need to assume that she wants that such actions are not performed, or at least that she has a negative attitude towards them being performed. There does not seem to be any *point* in uttering a sentence that entails that there is a moral reason not to hit one's children in a moral conversation which has as a mutually accepted purpose to influence behaviour unless she wants that such actions are not carried out. More precisely, since one purpose of the moral conversation in question is to influence behaviour, and a person utters a sentence which entails that there exists a moral reason not to hit one's children, we understand her as wanting that such actions are not performed, because if she does not have this attitude, it seems that her utterance would lack in *relevance*.⁴⁹

⁴⁹ The fact that a person wants an action not to be performed is compatible with the fact that she also does want it to be performed, and that she wants another action to be performed. Whether this is the case might be indicated by her further utterances or by the context of the conversation in question.

This can be explicated in terms of Grice's theory, and in particular the maxim of relation: make your contribution so as to be relevant in order to fulfil the purposes of the conversation. In Grice's view, we can expect people to contribute to a conversation in a way that fulfils its purposes, especially in a way that abides by the cooperative principle and the maxims. Now, in order to understand the person in the example as doing so, we have to understand her as wanting that the kind of actions under considerations are not performed. More specifically, in order to understand her as contributing to the moral conversation in a way that is *relevant* given its purpose to influence behaviour, we should understand her in that way. Thus, her utterance conversationally implicates that she wants that the actions in question are not performed. However, this holds only with the proviso that she does not make an additional utterance, or the context contains information, that cancels the implicature.

We might formulate the first basis for *DAA* as follows: A person's utterance of a sentence of the type " ϕ ing is wrong" conversationally implicates that she wants that ϕ is not performed, since this assumption is required in order to make her utterance consistent with the supposition that she (i) adheres to the cooperative principle and the maxims, especially the maxim of relation; (ii) knows the meaning of the sentence she utters, above all that it entails that there is a moral reason not to ϕ ; (iii) is aware of the context of the utterance, and (iv) (see above).⁵⁰

Grice stresses that it should be possible to calculate a conversational implicature. We may do so in the following manner. The person has uttered a sentence to the effect that certain actions are wrong. There are no grounds for believing that she is not observing the cooperative principle; that is, there are no reasons to believe that she does not try to contribute to the moral conversation in a way that fulfils its mutually accepted purposes, which, among other things, is to influence behaviour. In particular, there are no reasons to believe that she does not conform to the maxim of relation, and hence, that she utters something that is not relevant given this purpose. Moreover, there are no grounds to think that she is not aware of the meaning of the sentence she utters. In particular, there are no grounds to think that she does not realize that the sentence she utters entails that there is a moral reason not to do the actions under consideration, with the result that there is fact that, from the perspective of morality, counts against performing them. Given these assumptions, she would not have uttered the sentence

⁵⁰ Since (iv) is not directly relevant for the present discussion, I will presume it in what follows.

unless she wants that such actions are not carried out. More precisely, in view of the mentioned assumptions, it is plausible to think that she realizes that if she utters this sentence, her audience understands her as wanting that such actions are not performed. And so, if she does not wish to be understood in this way, she would not have uttered the sentence, at least not without modifying it. She has not uttered anything to suggest that she approves of such actions being carried out, or that she is neutral towards them being carried out, and there is nothing in the context to indicate this. Therefore, she wants that such actions are not performed.

Thus far, I have maintained that a person's utterance of a sentence of the type "ϕing is wrong" carries a conversational implicature to the effect that she wants that ϕ is not performed. It is now time to note that this is a *GCI* rather than a *PCI*. In Grice's view, an utterance carries a *GCI* in so far as it standardly carries the implicature; the implicature does not need any special contextual factors to arise, but carries it in the absence of special circumstances that defeat it. I have argued that moral conversations generally have as a purpose to influence behaviour. It follows that this holds true in the various contexts in which people are involved in such conversations. Consequently, an utterance of a sentence that instantiates "ϕing is wrong" carries the implicature I identified earlier across such contexts. Thus, it can be justifiably claimed that such an utterance standardly carries the implicature. The following is especially worth observing. Given that moral conversations generally have the purpose to influence behaviour, they have this purpose even in the absence of specific contextual features. In other words, they have this purpose even if the context in question lacks special features which provide particular evidence that this *is* the purpose of the conversation.⁵¹ Thus, an utterance of a sentence that instantiates "ϕing is wrong" carries this implicature even if the context in which it is uttered lacks such special features, and the implicature does not need them to arise. On the contrary, such an utterance carries the mentioned implicature unless there are any special circumstances that defeat it.

We might end this section by considering an objection. In arguing for the first basis of *DAA*, I employed the maxim of relation. However, it has been claimed that this

⁵¹ This does not entail that *GCI*s have to be entirely independent of contextual features. They are independent of *specific* contextual features—features which merely occur in certain specific contexts. This is compatible with there being some standing features that are needed if a *GCI* is to arise.

maxim cannot be involved in the production of *GCI*s, only of *PCI*s.⁵² Whether an utterance is relevant in the pertinent sense depends, as we have seen, on the purposes of the conversation. Now, it might be thought that what purposes a conversation have always varies with the context of the conversation. If this assumption is correct, it follows that the maxim of relation cannot be involved in the production of *GCI*s. However, this objection overlooks the possibility that some kinds of conversations can have certain purposes *generally*. In my view, this is precisely the case with moral conversations: they generally have as a purpose to influence behaviour. Consequently, moral conversations generally have this purpose in the various contexts in which people are involved in such conversations. This purpose of moral conversations thus does not vary with context in a manner that is problematic for *DAA*.⁵³

10. Second Basis of the Dual Aspect Account

Moving to the second basis of *DAA*, we may start by observing that there are various background presumptions concerning moral beliefs and utterances, according to which certain stereotypical features attach to them. As a result, there is a tacit notion of these phenomena which is present in moral conversations and which governs our understanding of moral utterances.

According to one background presumption, people who utter sentences of the type “ ϕ ing is wrong” generally want that ϕ is not performed. This should not come as a surprise given what was said in the last section; since such utterances carry a *GCI* to this effect, it is generally expected that people’s moral utterances standardly are connected to their attitudes in this manner. There are also other background presumptions that are relevant. According to one of them, people who think that a certain type of actions are wrong generally want that such actions are not carried out. There are presumably a number of psychological and social explanations for this phenomenon. One is moral upbringing. From early childhood we are taught not to do what is wrong in such ways that we are socialized into society and avoid performing these actions. Another

⁵² See Levinson (1983), 127; Levinson (2000), 74, and Davis (1998), 149 (but see Levinson (1983), 107). Several authors claim, or imply, that the maxim of relation can be involved in the production of *GCI*s; see e.g. Bach and Harnish (1998 (1979)), 698; Neale (1992), 534–535; Carston (1995), 221, and Horn (2004), 13.

⁵³ It should be clear that the first basis of *DAA* is not circular. One part of the explanation of why moral utterances conversationally implicate attitudes is that moral conversations have as a purpose to influence action. However, it is not part of the explanation of why moral conversations have this purpose that moral utterances actually implicate attitudes. What explains that moral conversations have this purpose is rather that if it is fulfilled, we can use moral language to regulate one another’s behaviour.

explanation may be that people are widely assumed to share certain psychological traits, such as empathy, which give rise to the mentioned expectation.⁵⁴

Consider again a person who is involved in a moral conversation and says “It’s wrong to hit one’s children.” We have seen that, according to pure cognitivism, moral conversations generally have as a purpose to exchange beliefs about matters that bear on moral issues, and that this purpose is intimately connected to the conventional meaning of moral sentences. As I try to develop a purely cognitivist account of the practicality of moral language, I will assume that the conversation under consideration has this purpose. However, for reasons that will be evident below, I will not assume that it has as a purpose to influence behaviour. Yet, in view of the mentioned background presumptions, we would presumably understand the person’s utterance as conveying that she wants that such actions are not performed.

This, I contend, can be explicated in terms of Grice’s notion of the second maxim of quantity: “Do not make your contribution more informative than is required.”⁵⁵ Stephen Levinson argues that this maxim is involved in the production of *GCI*s in a way which is relevant to the present inquiry.⁵⁶ A corollary of this maxim, Levinson suggests, is that one need not provide information about what can be taken for granted in the context of a conversation, such as information that is part of the background presumptions regarding stereotypical features. He succinctly summarizes this notion in the heuristic “What is expressed simply is stereotypically exemplified.”⁵⁷ Levinson argues that the fact that people conform to the second maxim, and that they share various background presumptions, gives rise to a number of *GCI*s.⁵⁸ Applying this

⁵⁴ See e.g. Brink (1989), 49; Boyd (1988), 215–216; Svavarsdóttir (1999), 183–187, and Shafer-Landau (2003), 159–160. Another explanation concerns normative reasons. *DAA* does not presuppose that a sentence of the type “ ϕ ing is wrong” entails that there is a normative reason not to ϕ . However, it is plausible to assume that we regularly believe that we have strong normative reasons not to do what is wrong. It is also plausible to think that, if we are rational, we are motivated to refrain from performing actions we think there are strong normative reasons not to do. In so far as we generally are rational, it can thus be expected that we are regularly motivated not to do what we believe is wrong.

⁵⁵ Grice (1989a (1975)), 26. The first maxim of quantity says “Make your contribution as informative as is required,” which Grice evidently understands as “Don’t make your contribution *less* informative than is required.”

⁵⁶ Levinson (2000), 37–38, and 112–134. See also Harnish (1998 (1976)), 267; Horn (2004), 12–17, and Jaszczolt (2005), 55–58. Cf. Searle (1969), 141–146.

⁵⁷ Levinson (2000), 37.

⁵⁸ For instance, Levinson thinks that “John said ‘Hello’ to the secretary and then he smiled” implicates “John said ‘Hello’ to the female secretary and then John smiled” (Levinson (2000), 117). The underlying assumption at work here is that secretaries are generally women.

reasoning to the present inquiry, the result is that a moral utterance carries the *GCI* I identified earlier.

To see this more clearly, return to the person who is involved in a moral conversation and says “It’s wrong to hit one’s children.” As Grice suggests, we can expect people to adhere to the cooperative principle and the maxims, such as the second maxim of quantity. Now, it may be maintained that in order to deem her as doing so, we should understand her as wanting that such actions are not performed. In order to take her as following the principle of not providing more information than is required, we should understand her utterance in such a way that there is no reason for her to submit additional information which modifies her utterance. For instance, we ought to assume that there is no reason for her to provide information to the effect that what usually is presumed to hold true does not apply in this particular case. As a result, we understand what she says in such a way that there is no exception to the background presumptions regarding moral beliefs and utterances. Consequently, her utterance conversationally implicates that she wants that such actions are not performed. Again, this holds on condition that she does not make a further utterance, or the context contains information, which cancels the implicature.

We might formulate the second basis for *DAA* as follows: A person’s utterance of a sentence of the type “ ϕ ing is wrong” conversationally implicates that she wants that ϕ is not performed, since this assumption is required in order to make her utterance consistent with the supposition that she (i) adheres to the cooperative principle and the maxims, in particular the second maxim of quantity; (ii) knows the meaning of the sentence she utters; (iii) is aware of the context of the utterance, in particular the background presumptions regarding moral beliefs and utterances, and (iv) (see above).

The conversational implicature is calculable in the following manner. The person has uttered a sentence to the effect that certain actions are wrong. There are no grounds to believe that she is not observing the cooperative principle. In particular, there are no reasons to believe that she does not adhere to the second maxim of quantity, and hence, that she does not comply with the principle of not saying more than is required. Moreover, there are no reasons to believe that she is unaware of the context of the utterance, such as the background presumptions regarding moral beliefs and utterances. Given these assumptions, she would not have made the utterance unless she wants that these actions are not carried out. Put more accurately, she would not have made the utterance without providing appropriate additional information if she does not have that

attitude. More precisely, in view of the mentioned presumptions regarding moral beliefs and utterances, it is reasonable to presume that she realizes that if she utters this sentence without adding information to the effect that there is some exception to these presumptions, her audience will understand her utterance in such a way that these presumptions hold true, and thus that she wants that these actions are not carried out. So, if she did not wish to be understood in this way, she would not have uttered the sentence without due modification. She has not said anything which indicates that she does not accept this understanding of her utterance, and there is nothing in the context which suggests this. Therefore, she wants that the actions in question are not performed.

Moreover, once again it is plausible to think that the conversational implicature in question is a *GCI*. According to the second basis of *DAA*, the fact that a person's moral utterance carries the conversational implicature in question depends on there being certain background presumptions regarding moral beliefs and utterances. There are, as indicated, reasons to think that these presumptions are quite widely shared. Consequently, they are normally presumed in the various contexts in which people are involved in moral conversations. It follows that an utterance of a sentence of the type "ϕing is wrong" carries a *GCI* to the effect that she wants that ϕ is not performed. Thus, such an utterance standardly carries this implicature. The following is especially worth observing. It was implied above that the fact that people who are involved in moral conversations hold certain background presumptions is part of the context of such conversations. However, the presumptions regarding moral beliefs and utterances are commonly accepted; although they are part of the context, they are not special contextual features in the relevant sense—they are not features that merely pertain to individual contexts. Rather, they are standing contextual features of moral conversations. For this reason the moral utterance carries the mentioned implicature even in the absence of special contextual features, and such features are not needed for the implicature to arise. On the contrary, such an utterance carries the implicature unless there are special circumstances that defeat it.

We saw earlier that advocates of pure cognitivism should maintain that moral conversations have the first purpose (to exchange moral beliefs). However, they are not committed to the presence of the second purpose (to influence behaviour). In arguing that there is a further basis for thinking that a moral utterance carries a certain *GCI*, I therefore assumed that the moral conversation in question had the first purpose without presupposing that it had the second purpose as well. We have seen that a person's

utterance of a sentence of the type “ ϕ ing is wrong” might carry the *GCI* I identified earlier *even if* it is not a purpose of the moral conversation in question to influence behaviour. The basic reason is that the background presumptions about moral beliefs and utterances still are in place. This is an important result since it is plausible to think that there might be cases where a moral conversation does not have as a purpose to influence behaviour but where a moral utterance still carries the *GCI* in question.⁵⁹

We are now in a position to explain why the *GCI* under consideration is not inherited by utterances of complex sentences. According to *DAA*, it is utterances of a type of sentence that attributes a moral property to something, such as “ ϕ ing is wrong,” which carry a *GCI* relating to attitudes, and *not* utterances of *any* sentence where a moral term, such as “wrong,” is involved. Above all, in this view an utterance of a complex sentence, such as a conditional, where a sentence which involves a moral term occurs embedded need not carry such a *GCI*.⁶⁰ The basic reason is this. On the first basis of *DAA*, part of the explanation of why an utterance of a sentence such as “It’s wrong to hit one’s children” carries the *GCI* in question is that a sentence of this kind entails that there is a *moral reason* in relation to these actions. However, a complex sentence where such a moral sentence is embedded—such as “If it’s wrong to hit one’s children, it’s wrong to hit other people’s children”—does not entail any moral reason. As a result, an utterance of such a sentence does not carry the *GCI* in question. A further reason is that, on the second basis of *DAA*, moral utterances carry such a *GCI* because there are certain background presumptions regarding moral beliefs and utterances where these beliefs and utterances correspond to the mentioned type of moral sentence. Consequently, *DAA* does not have the difficulty that afflicts Copp’s view.

11. Explaining the Practicality of Moral Language

Recall the practicality of moral language: It is generally the case that if a person employs a sentence to the effect that ϕ ing has a certain moral characteristic, we presume

⁵⁹ One example might be cases where the moral conversation concerns behaviour that the participants cannot influence. Another might be cases where the moral issue under discussion is very abstract.

⁶⁰ In commenting on *GCI*s, Grice writes that “one can say that the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature” (Grice (1989a (1975)), 37). However, there is nothing in Grice’s account that commits him to the view that *GCI*s are connected to uses of single terms rather than uses of entire utterances. On the contrary, as Robyn Carston argues, what Grice says here is misleading given his official view that it is utterances that carry implicatures (Carston (1995), 229). Grice characterizes *GCI*s by saying things as “Anyone who uses a sentence of the form *X is meeting a woman this evening* would normally implicate [...]” (Grice (1989a (1975)), 37), thereby indicating that *GCI*s are carried by utterances of certain types of sentences.

that she has a certain attitude in relation to ϕ ing; more particularly, we do so even if we do not have any special information which indicates that she has that attitude.

According to *DAA*, a person's utterance of a sentence of the type " ϕ ing is wrong" carries a *GCI* to the effect that she wants that ϕ is not performed, which means that such an utterance standardly carries this implicature. To borrow a phrase from linguistics, we can perhaps say that there exists a *convention of usage* such that utterances of this type of sentence convey that speakers have the mentioned attitude.⁶¹ This indicates that *DAA* can explain the practicality of moral language: Given that there is this standardized connection between moral utterances and attitudes, it should come as no surprise that we strongly expect a person who makes such an utterance to have the attitude in question.

However, *DAA* is also able to provide a more thorough explanation of the practicality of moral language. In Grice's view, if an utterance of a sentence carries a *GCI*, utterances of such a sentence standardly carries this conversational implicature. This means that the utterance carries such an implicature even if the context in question does not have any special features—in other words, features merely pertaining to individual contexts—that support this understanding of it. The two bases of *DAA* that I identified in the preceding sections suggest basically two reasons of why this is the case as regards utterances of moral sentences: (i) moral conversations generally have as a mutually accepted purpose to influence behaviour; and (ii) there are certain background presumptions regarding moral beliefs and utterances. These two considerations, I suggested, constitute *tacit notions* that are present when we are involved in moral conversations. Thus, I argued that these notions govern our understanding of a moral utterance independently of the special contextual features of the case at hand. As a consequence of these notions, we assume that a moral utterance carries the implicature under consideration even where the case in question does not have specific contextual features which support this understanding of it. This readily explains the practicality of moral language. Hence, we generally presume that a person who employs a sentence of the type " ϕ ing is wrong" wants that ϕ is not performed, although we do not have any special contextual information about the particular case at hand which indicates that she *has* this attitude. The basic reason is that, as a result of the tacit notions present in moral

⁶¹ See e.g. Searle (1979 (1975)), 49; Morgan (1998 (1978)), 644–651; Bach and Harnish (1998 (1979)), 701–703; Bach (1998 (1995)), 713, and Levinson (2000), 21–27.

conversations, the utterance conversationally implicates that she has this attitude, even if the context in which it is uttered does not have any special features which support this understanding of it.⁶² Moreover, since moral language is practical in this way, *DAA* is able to explain how moral language can serve the function of regulating behaviour.

DAA is also able to explain why metaethical views that entail *The AI Thesis* are so tempting.⁶³ According to *The AI Thesis*, a sentence of the type “ ϕ ing is wrong” expresses a moral judgment such that if a person holds it, she has a certain attitude in relation to ϕ ing. This view offers an attractive explanation of the practicality of moral language: we know—merely because we know its meaning—that the sentence expresses such a judgment, and we therefore presume that an utterance of it is accompanied by the attitude in question. However, Grice maintains that we might be misled to think that a sentence has a certain meaning when it actually is a matter of a *GCI*, since an utterance of the sentence standardly carries a given implicature in that case. In a similar manner, *The AI Thesis* might appear plausible because utterances of a sentence that ascribes a moral characteristic to something, such as “ ϕ is wrong,” standardly carry a *GCI* to the effect that the speaker has a certain attitude in relation to ϕ ing. However, *DAA* can provide a more thorough explanation of the attraction of *The AI Thesis*. According to *DAA*, a person’s utterance of a sentence of the type “ ϕ ing is wrong” carries the implicature I identified earlier, even if the context of the utterance does not have any special features that support this understanding of it. This makes it possible for us to infer *immediately* that she has the attitude in question, because we do not need to attend to any special contextual features to reach this understanding of her utterance. This fact might make us believe that it is *sufficient* to know the meaning of a moral sentence to conclude that an utterance of it is accompanied by a certain attitude. Thus, it might lead one to embrace a metaethical view that entails *The AI Thesis*.⁶⁴

⁶² Assume that a person utters a sentence of the type “ ϕ ing is wrong” without having any negative attitude in relation to ϕ ing and that the *GCI* I have identified is not cancelled. If the moral conversation at issue has both the purposes considered above, her utterance violates both the maxim of relation and the second maxim of quantity. If the moral conversation lacks the second purpose, it only violates the latter maxim.

⁶³ Similar considerations explain why an account in terms of *CI* may seem appealing.

⁶⁴ Advocates of metaethical positions that entail *The AI Thesis* might want to object that *DAA* does not provide an explanation of the intimate connection between moral *thoughts* (judgements) and attitudes. (Cf. Bar-On and Chrisman (2009), 154–155.) However, I think *DAA* readily can explain this notion. First, we have already seen that *DAA* is able to explain why we might be led to think that a moral sentence, in virtue of its meaning, expresses a judgment which implies an attitude. Clearly, this also provides an explanation of why we think that there is an intimate connection between moral judgments and attitudes. Moreover, according to the second basis of *DAA*, a reason of why a moral sentence carries a *GCI* is that

12. Calculability and Cancellability

In Grice's view, conversational implicatures have to fulfil two essential criteria, and these criteria consequently also apply to *GCI*s, in spite of their being standardized.

First, it should be possible to *calculate* a *GCI*. We have already seen that *DAA* satisfies this requirement. It should be stressed, though, that even if it is possible to calculate a *GCI*, we need not do so in order to recognize it. It is generally acknowledged that *GCI*s are *short-circuited*: since such an implicature is standardized, it is possible to recognize it immediately without calculating it.⁶⁵ Thus, in order to see that a person who utters a sentence of the type “ ϕ ing is wrong” wants that ϕ is not performed, we need not calculate the implicature in the manner described above in Sections 9 and 10, since the connection between moral utterances and attitudes is standardized. This is so because the two tacit notions that are present in moral conversations make it possible for us to infer immediately that such an utterance carries the relevant implicature without attending to specific contextual features.

Second, it should be possible to *cancel* a *GCI*.⁶⁶ Assume that utterances of a certain type of sentence carry a *GCI*. The *GCI* is *explicitly cancellable* if and only if a person who makes such an utterance can make an additional utterance that ensures that the original utterance does not carry the *GCI* in that particular case, and this is compatible with the conventional meaning of the sentence. Assume that someone says “It's wrong to hit one's children,” but then adds “But I don't want people to stop hitting their children.” I think there might be cases where we accept that after the second utterance she has informed us that although she actually does think that hitting one's children is wrong, she does not have the mentioned attitude towards such actions. However, it

there are certain background assumptions regarding the connection between moral judgments and attitudes. In this way, our notion of the intimate connection between moral thoughts and attitudes already is part of *DAA*. Accordingly, these background assumptions can be appealed to in the required explanation. Finally, as Timothy Williamson has pointed out, it is plausible to think that conversational implicatures often take place in thought. Roughly put, in our thinking we follow the conversational implicatures of the utterances we would have used to communicate our thoughts. This suggestion is especially plausible as regards *GCI*s because they are standardized. Accordingly, a person who thinks “She ate some of the cookies” is apt to think that she did not eat all the cookies, since an utterance of such a sentence standardly carries an implicature to this effect. Likewise, a person who thinks “Capital punishment is wrong” is apt to have a negative attitude towards this type of penalty in view of the *GCI* I have described above. (Cf. Copp (2009), 200.)

⁶⁵ See e.g. Morgan (1998 (1978)), 651–652, and Bach and Harnish (1998 (1979)), 701–702. This is sometimes called “default reasoning” (see e.g. Bach (1984), 37–58, and Jaszczolt (2005), Ch. 2).

⁶⁶ Grice (1989a (1975)), 39. Finlay argues that a moral implicatures are cancellable: Finlay (2005), 13–15.

might be that we accept this only after a prolonged discussion in which she explains why she does not have the attitude in spite of making her moral utterance. Undoubtedly, we would presumably find her utterances puzzling, and require an explanation of *why* she made the utterance when in her the next breath she says that she does not want that people stop hit their children. It should be easy to see, however, that our response is explicable in view of the standardized connection between moral utterances and attitudes.

Assume again that utterances of a certain type of sentence carry a *GCI*. The *GCI* is *contextually cancellable* if and only if the context of the utterance can be such that it means that the utterance does not carry the *GCI* in that particular case, and this is compatible with the conventional meaning of the sentence. I think it is an advantage of *DAA* that it makes clear under what conditions this is possible. There are two tacit notions present in moral conversations that govern our understanding of moral utterances: moral conversations have as a general purpose to influence behaviour, and there are certain background presumptions regarding moral beliefs and utterances. We have seen that even if the first notion is absent, a moral utterance might still implicate an attitude by virtue of the second notion being present. This means that in order for the implicature under consideration to be contextually cancelled, the context has to be such that *both* these tacit notions are prevented from governing our understanding of the utterance. I think there are such cases, although they may be quite rare, given the firm grip these notions have on our understanding of moral utterances.⁶⁷

Copp maintains that if someone says, for instance, “Capital punishment is wrong,” but does not have a negative attitude towards such actions, we would feel that her use of the term is “linguistically inappropriate.”⁶⁸ He therefore thinks that it is incompatible with the conventional meaning of the moral term to make such an utterance without having the attitude in question. As a consequence, he concludes that moral implicatures are not cancellable in Grice’s sense.

Let us accept that we would feel that a person’s use of the mentioned type of moral sentence would be linguistically inappropriate if she lacks the relevant attitude. It is unclear, however, why this intuition needs to be explained by the conventional meaning

⁶⁷ For example, there might be moral conversations that do not have the purpose to influence behaviour (perhaps because the moral issue under discussion is very theoretical) and moral conversations where the participants leave the background presumptions out of account (perhaps because they are discussing hypothetical situations where no such presumptions are in place).

⁶⁸ Copp (2009), 195–197. Advocates of *The AI Thesis* might argue in a similar manner.

of a moral term. One reason to be suspicious of such inferences is that there are well-known examples in the philosophical and linguistic literature that might evoke such responses, but where the received view is that this is not to be explained by the meaning of any particular term. Indeed, according to an influential view, many of these examples should be understood in terms of *GCI*s.⁶⁹

Moreover, *DAA* offers a very plausible account of why we might be led to think that what explains why utterances of a certain type of moral sentences convey attitudes is the conventional meaning of the moral term that is part of these sentences. *DAA* maintains that utterances of a sentence of the type “ ϕ ing is wrong” carry a *GCI* to the effect that the speaker wants that ϕ is not performed, which in turns mean that such utterances standardly carry this implicature. We can thus infer automatically, without considering particular contextual features, that a person who uses such a sentence has the attitude. Similarly, a *GCI* is short-circuited in the sense that although it is possible to calculate, we need not do so in order to identify it. As there is such a remarkably close correlation between, on the one hand, utterances of a certain type of moral sentences which involve a moral term and, on the other hand, a certain attitude, it is easy to come to think that it is the conventional meaning of the term itself that explains that these utterances convey the attitude. This might in turn make one believe that it is a matter of a *CI*. In Grice’s words: “[I]t is all too easy to treat a generalized conversational implicature as if it were a conventional implicature,” since “the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature.”⁷⁰ For the same reason it is easy to be misled into thinking that if a person utters a sentence such as “Capital punishment is wrong,” without having a negative attitude towards this type of penalty, her use of “wrong” is linguistically inappropriate.

13. Conclusion

In this paper, I have defended The Dual Aspect Account (*DAA*): A person’s utterance of a type of sentence according to which ϕ ing has a moral characteristic, such as “ ϕ ing is

⁶⁹ For an overview of such examples, see e.g. Levinson (2000), esp. Ch. 2. For instance, if someone says “I ate some of the cookies,” but it turns out that she ate all the cookies, we might want to say that her use of “some” is linguistically inappropriate. Similarly, if someone says “I’ll have pizza or pasta for lunch,” but it turns out that she has decided to have both, we might respond in a similar manner as regards her use of “or.” On the standard view, this is explained by the fact that utterances of sentences where these terms (“some” and “or”) are involved carry *GCI*s. See e.g. Grice (1989 (1975b)), 45–47.

⁷⁰ Grice (1989 (1975)), 37. Cf. Grice (1981), 185.

wrong,” conveys both the belief that ϕ ing has a certain moral property and that she has a certain attitude towards ϕ ing. However, whereas the belief is conveyed by means of the conventional meaning of the sentence she uses, the attitude is conveyed by means of a generalized conversation implicature (*GCI*). This account, I argued, might be adopted by any pure cognitivist who denies that there is a necessary, non-trivial, connection between moral judgments and attitudes, so as to account for the practicality of moral language. I have also argued that it has significant advantages over competing accounts of this essential feature of moral language.

Finally, *DAA* provides a general formula that generates systematic explanations of the practicality of moral language. According to *DAA*, there are five explanatory features that interact so as to generate a *GCI*: Grice’s notion of *GCI*; the conventional meaning of the type of moral sentence at issue; the two distinct purposes of moral conversations, and various background presumptions about moral beliefs and utterances. Moral sentences of the relevant type differ in their meanings. The different meanings of these sentences may interact with the other four features in the manner I have outlined above. As a result, utterances of these sentences carry different *GCI*s that communicate different attitudes. Thus, *DAA* provides systematic explanations of why utterances of different moral sentences carry *GCI*s which communicate different attitudes.⁷¹

References

- Bach, K. and R. Harnish (1998 (1979)), “Indirect Acts and Illocutionary Standardization,” in A. Kasher, ed. *Pragmatics*, vol. IV, London: Routledge, pp. 682–711.
- Bach, K. (1984), “Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice,” *Pacific Philosophical Quarterly*, vol. 65, pp. 37–58.
- (1998 (1995)), “Standardization Revisited,” in A. Kasher, ed. *Pragmatics*, vol. IV, London: Routledge, pp. 712–722.
- (1999), “The Myth of Conventional Implicature,” *Linguistics and Philosophy*, vol. 22, pp. 327–366.
- Barker, S. (2000), “Is Value Content a Component of Conventional Implicature?,” *Analysis*, vol. 60, pp. 268–279.
- (2003), “Truth and Conventional Implicature,” *Mind*, vol. 112, pp. 1–33.
- Bar-On, D. and M. Chrisman (2009), “Ethical Neo-Expressivism,” in R. Shafer-Landau, ed. *Oxford Studies in Metaethics*, vol. 4, Oxford: Oxford University Press, pp. 133–167.

⁷¹ I am grateful to colleagues in Gothenburg and Lund for helpful discussions of earlier versions of this paper. Special thanks are due to David Copp, Stephen Finlay, Gunnar Björnsson, John Eriksson, Cathrine Felix, Ragnar Francén, Ingmar Persson, and Wlodek Rabinowicz for valuable comments. My work on the paper was supported by a research grant from The Bank of Sweden’s Tercentenary Foundation.

- Blackburn, S. (1998), *Ruling Passions*, Oxford: Clarendon Press.
- Boisvert, D. (2008), "Expressive-Assertivism," *Pacific Philosophical Quarterly*, vol. 89, pp. 169–203.
- Boyd, R. (1988), "How to Be a Moral Realist," in G. Sayre-McCord, ed. *Essays on Moral Realism*, Ithaca: Cornell University Press, pp. 181–228.
- Brink, D. (1989), *Moral Realism and the Foundations of Ethics*, Cambridge: Cambridge University Press.
- (1992), "A Puzzle About the Rational Authority of Morality," *Philosophical Perspective*, vol. 6, pp. 1–26.
- Carston, R. (1995), "Quantity Maxims and GCI," *Lingua*, vol. 96, pp. 213–244.
- (2004), "Truth-Conditional Content and Conversational Implicature," in C. Bianchi, ed. *The Semantics/Pragmatics Distinction*, Stanford: CSLI, pp. 65–100.
- Cholbi, M. (2009), "Moore's Paradox and Moral Motivation," *Ethical Theory and Moral Practice* vol. 12, pp. 495–510.
- Copp, D. (1995), "Moral Obligation and Moral Motivation," in J. Couture and K. Nielsen, ed. *On the Relevance of Metaethics*, University of Calgary Press, pp. 187–220.
- (2001), "Realist-Expressivism: A Neglected Option for Moral Realism," *Social Philosophy and Policy*, vol. 18, pp. 1–43.
- (2009), "Realist-expressivism and Conventional Implicature," in R. Shafer-Landau, ed. *Oxford Studies in Metaethics*, vol. 4, Oxford: Oxford University Press, pp. 167–202.
- Davis, W. (1998), *Implicature*, Cambridge: Cambridge University Press.
- Dreier, J. (1990), "Internalism and Speaker Relativism," *Ethics*, vol. 101, pp. 6–26.
- Eriksson, J. (2009), "Homage to Hare: Ecumenism and the Frege-Geach Problem," *Ethics*, vol. 120, pp. 8–35.
- Finlay, S. (2004), "The Conversational Practicality of Value Judgment," *The Journal of Ethics*, vol. 8, pp. 205–223.
- (2005), "Value and Implicature," *Philosophers' Imprint*, vol. 5, pp. 1–20.
- Gibbard, A. (1990), *Wise Choices, Apt Feelings*, Oxford: Clarendon Press.
- (2003), *Thinking How to Live*, Cambridge, Mass.: Harvard University Press.
- Green, M.S. (1998), "Direct Reference and Implicature," *Philosophical Studies*, vol. 91, pp. 61–90.
- Grice, P. (1981), "Presupposition and Conversational Implicatures," in P. Cole, ed. *Radical Pragmatics*, New York: Academic Press.
- (1989a (1975)), "Logic and Conversation," *Studies in the Ways of Words*, Cambridge, Mass.: Harvard University Press, pp. 22–40.
- (1989b (1975)), "Further Notes on Logic and Conversation," *Studies in the Ways of Words*, Cambridge, Mass.: Harvard University Press, pp. 41–57.
- Hare, R.M. (1952), *The Language of Morals*, Oxford: Clarendon Press.
- Harnish, R. (1998 (1976)), "Logical Form and Implicature," in A. Kasher, ed. *Pragmatics*, vol. IV, London: Routledge, pp. 230–314.
- (2008), "The Semantics of Racial Epithets," *Journal of Philosophy*, vol. XX, pp. 416–440.
- (2004), "Implicature," in L.R. Horn and G. Ward, eds. *The Handbook of Pragmatics*, Oxford: Blackwell, pp. 3–28.
- Jaszczolt, K.M. (2005), *Default Semantics*, Oxford: Oxford University Press.
- Karttunen, L. and S. Peters (1979), "Conventional Implicature," in C. Oh and D.A. Dinneen, eds. *Syntax and Semantics*, vol. 11, New York: Academic Press, pp. 1–56.
- Korsgaard, C. (1996 (1986)), "Skepticism About Practical Reason," *Creating the*

- Kingdom of Ends*, Cambridge: Cambridge University Press, pp. 311–334.
- Levinson, S. (1983), *Pragmatics*, Cambridge: Cambridge University Press.
- (2000), *Presumptive Meanings*, Cambridge, Mass.: The MIT Press.
- Morgan, J.L. (1998 (1978)), “Two Types of Convention in Indirect Speech Acts,” in A. Kasher, ed. *Pragmatics*, vol. IV, London: Routledge, pp. 639–657.
- Nagel, T. (1970), *The Possibility of Altruism*, New York: Oxford University Press.
- Neale, S. (1992), “Paul Grice and the Philosophy of Language,” *Linguistics and Philosophy*, vol. 15, pp. 509–559.
- Nunberg, G. (1981), “Validating Pragmatic Explanations,” in P. Cole, ed. *Radical Pragmatics*, New York: Academic Press, pp. 199–222.
- Potts, C. (2005), *The Logic of Conventional Implicatures*, Oxford University Press: Oxford.
- Railton, P. (2006), “Moral Factualism,” in J. Dreier, ed. *Contemporary Debates in Moral Theory*, Oxford: Blackwell, pp. 201–219.
- Recanati, F. (1989), “The Pragmatics of what is Said,” *Mind and Language*, vol. 4, pp. 295–329.
- (2003), “Embedded Implicatures,” *Philosophical Perspectives*, vol. 17, pp. 299–332.
- Ridge, M. (2003), “Non-Cognitivist Pragmatics and Stevenson’s ‘Do so as well!’,” *Canadian Journal of Philosophy*, vol. 4, pp. 563–574.
- (2006), “Ecumenical Expressivism: Finessing Frege,” *Ethics*, vol. 116, pp. 302–336.
- Sadock, J.M. (1998 (1978)), “On Testing for Conversational Implicature,” in A. Kasher, ed. *Pragmatics*, vol. IV, London: Routledge, pp. 315–331.
- Sayre-McCord, G. (1997), “‘Good’ on Twin Earth,” *Philosophical Issues*, vol. 8, pp. 267–292.
- Scanlon, T.M. (1998), *What We Owe to Each Other*, Cambridge, Mass.: Harvard University Press.
- Schroeder, M. (2008), “Expression for Expressivists,” *Philosophy and Phenomenological Research*, vol. LXXVI, pp. 86–116.
- (2009), “Hybrid Expressivism: Virtues and Vices,” *Ethics*, vol. 119, pp. 257–309.
- Searle, J. (1969), *Speech Acts. An Essay in the Philosophy of Language*, Cambridge: Cambridge University Press.
- (1979 (1975)), “Indirect Speech Acts,” *Expression and Meaning. Studies in the Theory of Speech Acts*, Cambridge: Cambridge University Press, pp. 30–57.
- Shafer-Landau, R. (2003), *Moral Realism. A Defence*, Oxford: Clarendon Press.
- Smith, M. (1994), *The Moral Problem*, Oxford: Blackwell.
- Strandberg, C. (forthcoming), “Expressivism and Dispositional Desires,” *American Philosophical Quarterly*.
- Svavarsdóttir, Sigrún (1999), “Moral Cognitivism and Motivation,” *The Philosophical Review*, vol. 108, pp. 161–219.
- Wedgwood, R. (2007), *The Nature of Normativity*, Oxford: Oxford University Press.
- Williamson, T. (2009), “Reference, Inference and the Semantics of Pejoratives,” in J. Almog and P. Leonardi, eds. *The Philosophy of David Kaplan*, Oxford: Oxford University Press.