

Reasons, inescapability and persuasion

Neil Sinclair¹

Published online: 29 February 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract This paper outlines a new metasemantic theory of moral reason statements, focused on explaining how the reasons thus stated can be inescapable. The motivation for the theory is in part that it can explain this and other phenomena concerning moral reasons. The account also suggests a general recipe for explanations of conceptual features of moral reason statements.

Keywords Expressivism · External reasons · Inescapability · Moral reasons · Williams

1 Introduction

I am sure that I do not understand the idea of a reason for acting, and I wonder whether anyone else does either (Foot 1972: 156).

The fact that people are starving is reason to feed them. That Prince is on tour provides a reason for fans to buy tickets. These statements of reasons are as commonplace as their philosophical treatment is puzzling. In this paper I outline a new metasemantic theory of such statements, focused initially on the moral case, but with expansive ambitions.¹

¹ For the distinction between semantics and metasemantics see Lewis (1970: 18–19). This paper therefore shares the methodological approach of Pérez Carballo and Santorio (forthcoming), Ridge (2014), Silk (2013) and others, who take expressivism to be a meta-semantic, and not semantic, thesis. As Ridge explains: ‘...semantics assigns literal meanings...to meaningful units of language. Meta-semantics, by contrast, explains that in virtue of which a given word, morpheme or sentence has the meaning it does’

✉ Neil Sinclair
neil.sinclair@nottingham.ac.uk

¹ Department of Philosophy, University of Nottingham, Nottingham NG7 2RD, UK

1.1 Practical reason statements

First, I need to delineate my subject matter. I begin with an artificially simple schema, hoping to extract from there the raw materials required for a more comprehensive account. Consider sincere utterances of sentences with the form:

(1) F is a reason for A to ϕ in C

I shall call such utterances ‘reason statements’ and say that they state reasons.² Each part of (1) requires elaboration.

1.1.1 A, C, ϕ and ‘is’

‘A’ refers to an agent or set or agents and ‘C’ gives a description of the situation in which the reason statement applies. For most purposes these can be treated as a single variable, since there is a certain arbitrariness in whether a condition is offered as part of a description of an agent (‘Starving people have reason to eat’) or as part of their circumstances (‘People have reason to eat when starving’). Hence in what follows I will only add ‘in C’ when exposition demands. ‘ ϕ ’ refers to an action or type of action. Such statements are sometimes called statements of practical reasons, to contrast with epistemic reasons, where ϕ is a belief, and evaluative reasons, where ϕ is a feeling (Skorupski 2009: 114–117). The ‘is’ should be taken to be that of predication rather than identity, as is shown by its common synonyms here: ‘provides’ and ‘constitutes’ (Olson 2009).

1.1.2 ‘...reason...’

The term ‘...reason...’ is here used in its ‘objectively’ or ‘standardly’ normative sense (Scanlon 1998: 19), the sense in which F justifies A’s ϕ ing or to some degree makes A’s ϕ ing right or obligatory. The normative is that which is fraught with ought, as Sellars (allegedly) put it. This is in contrast to explanatory and subjectively normative senses of the same term (Schroeder 2007: 10–15). Only in the normative sense does assent to a reason statement commit one to holding that an action has some consideration that stands in its favour.³ Such considerations needn’t be taken to be decisive: the reasons at stake here are *contributory* in the sense that

Footnote 1 continued

(2014: 8). However, the paper is in substantive disagreement with the hybrid descriptive/expressive account of reason statements given by Ridge, which takes such statements to express both beliefs concerning standards and endorsements of such standards (Ridge 2014: 37, 123–123). The account offered here eschews such beliefs.

² This paper is therefore concerned with ‘reasons’ only in the narrow sense that they are the subject matter of reason statements. Many other phenomena are also labelled ‘reasons’—such as the considerations that we weigh in deliberation—but such phenomenon are not the immediate concern. For an account reasons as inputs into deliberation compatible with the view defended here, see Schroeder 2007 ch. 8. See also Sect. 1.1.3.

³ Note also that my concern is with statements of reasons not rationality. For this distinction see Hooker and Streumer (2004: 72–73) and Kolodny (2005).

there can be more than one consideration in favour of A's ϕ ing, some considerations can favour A's ϕ ing whilst others favour A's not- ϕ ing, and the strength of a consideration in favour of A's ϕ ing can be amplified, attenuated, or even undermined by further features of the case (which may or may not be reason-giving themselves; cf. Dancy 2004: 15, 38–43).

1.1.3 *F*

I call *F* the 'reason-giver' in a reason statement, although these are standardly labelled 'considerations' (for example in Dancy 2004). Reason-givers are sometimes called 'reasons', but this can be misleading since it fails to distinguish between those things that provide reasons and what is stated by complete reason statements. Reason-givers are often taken to be (sets of) facts or true propositions (Schroeder 2007: 20; Skorupski 2009: 115) but here I take them to be intensional facts. That is, I take *F* to be a placeholder for terms that both refer to a particular part of the fabric of the world and offer a particular mode of presentation of that world-bit. On this view the fact that the Queen is offended is distinct from the fact that Elizabeth is offended, even though Elizabeth is Queen. That reason-givers are intensional facts is supported by the thought that sometimes the normative salience of a given world-bit can be affected by the manner in which it is presented. For example, there are situations in which the fact that the Queen is offended is reason to seek a Royal Pardon, but the fact that Elizabeth is offended is not, even though Elizabeth is Queen.⁴ Henceforth when I use the term 'fact' and its cognates, I have in mind intensional facts.

Finally, in the type of case I am concerned with, *F* is a placeholder for a non-normative term or phrase, such as 'That Elizabeth is offended'. Of course, many intelligible reason statements do involve normative locutions in the *F*-position ('The fact that literacy is good is a reason to promote it') but I prefer to give an account of these cases by first understanding cases involving non-normative reason-givers.

1.1.4 *Other reason statements*

I shall say little about variations on the schema given in (1), such as reason statements that exclude some of the variables ('A has a reason to ϕ ') or include extra variables such as weight ('*F* is a strong reason for A to ϕ '). Schroeder (2007: 15–21) has convincingly argued that the former are less basic than statements that fit schema (1). For reason statements that involve weight something more needs to be said, but it is better left until after an account of the simpler form is given. I give a brief account in Sect. 4.1.

⁴ For further reasons for taking reason-givers to be intensional facts see Suikkanen (2012).

2 Moral and non-moral reason statements

Within the class of reason-statements formalised by (1), common understanding recognises various types. Some reasons are decidedly moral, others less so. The theory I offer is first and foremost an account of moral reason statements, so I need to say something about this distinction.

Consider the following scenario:

The Last Glass. This stuff is water. Anna wants water to drink and believes that this stuff is water. Giving water to Brian would save Brian from death by dehydration.

Here, at least two reason statements are plausible:

- (1a) The fact that this stuff is water is a reason for Anna to drink it.
- (1b) The fact that this stuff is water is a reason for Anna to give it to Brian.

Let us say that (1a) is a non-moral reason statement, stating a non-moral reason, and (1b) is a moral reason statement, stating a moral reason. This distinction is sometimes, but not always, made explicit by talk of ‘moral’ reasons. But what distinguishes such reasons?

I cannot provide a complete answer to this question here, but I will assume that one necessary feature of moral reasons is their inescapability (Foot 1972; Joyce 2000, 2006: 61–64). To say moral reasons are inescapable is to say (roughly) that when they apply they do so regardless of any particular interest, project, desire, motivation, value or commitment of the agent whom they apply to, or equivalently that when they apply they do not apply because of any interest project, desire, motivation, value or commitment of the agent whose reasons they are. Williams (1981) calls these states elements of agents’ ‘subjective motivational sets’ and I shall sometimes refer to them generally as ‘desires’.⁵ More precisely, then, my assumption is that F provides a moral reason for A to ϕ only if F’s status as a moral reason-giver is not in any way dependent on any element, D, of A’s motivational being served or promoted by A’s ϕ ing.⁶

This assumption can draw inspiration from the work of Kant, who took inescapability to be a mark of the categorical imperative: ‘...that which represent[s] an action as objectively necessary in itself, without reference to another end’ (Kant 2007: 530). In Williams’ terms (1981: 101), there are some reasons for A to ϕ —‘external reasons’—which do not imply that ‘A has some motive that will be served or furthered by ϕ ing’. The assumption is also not without some warrant. In *The Last Glass*, when we judge that the fact that this stuff is water

⁵ See Schroeder (2007) for precedent. One way of distinguishing such states by their distinctive world-to-mind direction of fit (Humberstone 1992).

⁶ See Schroeder (2007: 110–113) for analyses of the ‘promotion’ relation. I take inescapability to be necessary but not sufficient for a reason to count as moral. Reasons of etiquette (Foot 1972) and personal excellence (Copp 1997: 92) also appear to be inescapable. Inescapability is also compatible with other suggested necessary features of moral reasons, such as universality (Schroeder 2007: 103–108) and overridingness (Copp 1997; Dancy 2004: 43).

is a moral reason for Anna to give it to Brian, we do not withdraw the statement on discovering that Anna currently has no desire served by doing so. Thus inescapability appears to be necessary for a reason to count as moral.⁷

Inescapability can be construed as a particular type of mind-independence. Call a motivational state of A that would be promoted by A's ϕ ing a 'D*-state'. Inescapability is the claim that moral reasons are independent of the 'D*-state' bit of the relevant agent's mind. More precisely, this independence is captured by the following two negated (strict) conditionals:

(I1) It is not the case that: If A lacks a D*-state, then there is no F such that F is moral reason for A to ϕ .

(I2) It is not the case that: If A has a D*-state then there is an F such that F is moral reason for A to ϕ .⁸

(I1) captures the thought that moral reasons to ϕ do not necessarily disappear in the absence of desires promoted by ϕ ing. Schroeder (2007: 106) calls this 'strong modal status'. (I2) captures the thought that moral reasons to ϕ are not consequent on agents having desires promoted by ϕ ing. I call this 'super strong modal status'. Together, these amount to the claim that for moral reasons, it is not because A's motivational set contains an element that would be promoted by her ϕ ing that F is a moral reason for A to ϕ (when it is). Moral reasons are D*-state independent. At least, so I will assume.

I will assume, further, that this type inescapability is a conceptual truth about moral reasons (Joyce 2000: 471). By conceptual truth I mean one such that it is constitutive of competence with the relevant concept that one recognises that truth in one's deployment of that concept. This assumption is not unreasonable. If someone claims that you have a moral reason to ϕ , but insists that this reason is purely dependent on your having some desire promoted by ϕ ing, we might question whether they have understood the concept of a moral reason. At the very least, we might ask what they take the distinction between moral and non-moral reasons to be.

In what follows, therefore, I will heed these (necessarily brief) remarks and assume inescapability to be a conceptually necessary feature of moral reasons. Part of my remaining task is to consider whether an expressivist theory can accommodate both this inescapability and its conceptual status.⁹

⁷ A further argument derives the inescapability of moral reasons from the inescapability of moral oughts. See Joyce (2000: 464–467).

⁸ These conditionals capture what Jenkins (2005) calls 'modal independence'.

⁹ Note that, in the final analysis, the claim that moral reasons are necessarily inescapable is not a necessary part of my argument. The account in Sect. 4 can, alternatively, be taken as an account of a type of reason statement that conceptually satisfies inescapability, with it being a further issue whether moral reasons are necessarily inescapable.

3 Motivating (interest in) the expressivist theory

3.1 A trilemma

The claim that moral reasons are inescapable is one pillar in the following trilemma that has done much to organise debates about the nature of reasons:

- (1) Necessarily, moral reasons are inescapable.
- (2) There are moral reasons.
- (3) Necessarily, if F is a reason for A to ϕ , then A has some desire that would be promoted by her ϕ ing.

Each claim is well-supported by argument, yet any two provide an argument against the third.

(1) is the claim defeasibly argued for above: moral reasons are D^* -state independent. Note that it is not necessary to understand this claim as describing the behaviour of a reified category of moral reasons, picked out by the predicate ‘...is a moral reason for...to...’. (1) can also be understood, in a metatheoretically neutral sense, as a claim concerning the correct application of the moral reason predicate. On this view, (1) is the claim that the correct application of this predicate does not depend on A having some desire that is promoted by her ϕ ing. Further, this understanding of (1) is preferable in the current context, since, unlike the alternative, it begs no questions about the semantic function of the moral reason predicate. Hence, in what follows I understand (1) and similar claims in this ontologically non-committing sense.

(2) can also be supported by argument. Charles Manson had a moral reason not to murder his victims; so there are moral reasons. More generally ordinary language users assume that some agents possess moral reasons. The same people use moral reason statements in ways well explained by the truth of this assumption. There is thus a strong presumptive case in favour of there being moral reasons. Note that to say that there are moral reasons in this sense is just to say that there are some situations where the predicate ‘...is a moral reason for...to...’ is correctly applied. Further, a true statement is simply one that involves the correct application of a predicate. So another way of reading (2) is as the claim that (positive) moral reason statements are sometimes true.

(3) is the Humean Theory of Reasons. As Hume might have said, but didn’t: Our reasons are the slaves of our passions. The Humean view is defended in some form by Schroeder (2007), Joyce (2000), Goldman (2009), Markovits (2011) and Williams (1981), among others. This crowd is not without ammunition. One type of argument in their favour is the ‘wedge’ argument: at least some normative reasons (such as Anna’s reason to drink this stuff) depend on the desires of the agents whose reasons they are in the way hypothesised by Humeanism. Hence, if Humeanism turns out to be broadly extensionally adequate for our considered talk of reasons it is preferable on grounds of simplicity (Schroeder 2007). Another type of argument lays down an explanatory challenge to those who would deny Humeanism (Finlay 2006: 4, following Williams 1989: 39). If our reasons are not bound by our desires,

what does explain the intuitive constraints on what reasons there are (for example, the constraint against there being reasons to perform impossible actions)? A third type of argument mirrors the type of worries that Mackie had about ‘objectively prescriptive’ moral properties. If normative reasons are not based on desires then (if they are explicable at all) they would seem to be based on, or identical with, normative entities that are metaphysically mysterious, epistemically troubling and motivationally ‘magnetic’ in ways distinct from all other type of entity (Goldman 2009: 9–28). While none of these arguments are conclusive they create a strong defeasible case for Humeanism. Those who reject Humeanism need to take these challenges seriously.

Given the trilemma, there are two types of Humeans. First, those who use (2) and (3) in an argument against (1). For example, Schroeder (2007: 101–122) rejects (1) by rejecting super strong modal status. Second, those such as Joyce (2000) who use claims (1) and (3) in an argument against there being moral reasons. To mitigate the implausibility of this error-theory, Joyce supplements it with an account of why although, strictly speaking, there are no moral reasons, it might make sense to go on behaving and speaking as if there are. Joyce (2000: 472–474) argues that though no moral reason statements are true, the making of some such statements can be part of worthwhile attempts to change the motivational sets of others so that they come to acquire new desires, and hence new reasons, that they did not previously have. Likewise, Williams talks about ‘optimistic’ internal reason statements (1981) and the usefulness of ‘proleptic invocation’ (1989; compare Finlay 2006: 12–13). On such views (positive) moral reason statements are strictly-speaking false (or at least not true) but pragmatically useful.

The last argument here is an example of a well-known ‘preservationist’ strategy which takes a set of target statements to be strictly speaking false (or not true) yet sometimes useful. This strategy faces a well-known problem, sometimes referred to as the objection from an alternative norm (Wright 1992: 10; Blackburn 1993: 149–150). The strategy accepts that some of the target statements serve some legitimate purpose, such as, in this case, proleptic invocation. It also accepts that these statements serve this purpose by asserting falsehoods (or not by stating truths). We might say that the purpose served is part of the pragmatics of the making of the statements (the use they are put to when made) but not part their semantics (their distinctive, truth-conditional, meaning). This raises a question: If some legitimate purpose is served by making these statements, why is that purpose not reflected in their semantic truth-evaluable content? If there is a norm governing what makes some of the statements appropriate and others not, why is that norm not reflected in the truth-conditions of the statements themselves, so that those that satisfy the norm can be rightly considered true? In short, why is the norm of usage recognised in the pragmatic post-script of the Joyce/Williams story not a norm of truth for the statements concerned?

This objection is not conclusive, in either the general or particular cases (Daly and Liggins 2010: 216–220). It may be that the purpose served by asserting false (or not true) statements cannot be well-served in any other way, or that it could be served in other ways, but through historical accident it is not. But here my interest is not to refute preservationism but to illuminate an under-explored possibility. What

this objection suggests is a theory of moral reason statements where the practical purposes of making such statements is part of the semantics of those statements themselves, and not merely part of the pragmatics of their use. The expressivist account is one such theory.^{10,11}

In sum, Humeans come in two varieties. Those who accept (2) and (3) and reject inescapability. And those who accept (1) and (3) and deny that there are any moral reasons. However, given that each of (1), (2) and (3) can be supported by argument, it is worth considering turning the triangle on the other side, by accepting (1) and (2) and rejecting Humeanism. Such a view will be only be plausible to the extent that it can both explain the inescapability of moral reasons and successfully neuter the arguments in favour of the Humean Theory. In what follows, I first take inspiration from Joyce and Williams to help elaborate an expressivist theory of moral reasons. I go on to show how this theory can explain inescapability and provide the materials to resist the common arguments for Humeanism. Although by no means complete, I take this work to establish the resulting view as worthy of further consideration.

3.2 First lacuna

Expressivist theories of reason statements, such as the one I offer here, have been conspicuously absent from the existing philosophical literature on the nature of reasons. They do not appear in general surveys (e.g. Wiland 2012) and are neglected by opponents (e.g. Schroeder 2007; Skorupski 2010). In this section I put forward a hypothesis about the source of this neglect, namely an under-exposed truth-maker assumption. Exposing this assumption helps illuminate the logical space in which my preferred theory operates.

The existing literature on reasons (moral and non-moral) recognises several non-Humean options which do not see all reasons as ultimately grounded in agents' desires. Such views may accept a Humean account of *some* reasons, but will insist that it is not universally applicable. One theory is the Moorean view that the predicate '...is a reason for...to...' refers to a three-place reason-relation, the obtaining of which is a sui generis normative fact (see Olson 2009 for discussion). Another is the Arsitotelian view that facts about reasons to act are reducible to facts about the value of states of affairs actions bring about, where these facts are themselves irreducibly normative (Gaut 1997; Raz 1999). A third, Kantian, view

¹⁰ Joyce (2000: 473) quickly rejects this option, citing the Frege–Geach problem. But that problem is not particular to the expressivist account of moral *reason* statements. It is worthwhile to consider the particular details of that account should the general problem be solvable (indeed, such details may help discern a general solution).

¹¹ I do not claim that the expressivist theory is *entailed* by the claim about the practical purposes of making reason statements. Such a move would commit what Searle (1969: 139) would later labelled the 'speech-act fallacy' but which was first hinted at by Dewey (1945: 702–3). My point is simply that it is worth considering a view which, as Dewey puts it, takes 'emotional factors' to be 'an inherent part of the judgment'. My hypothesis is that such a view can explain, at least as well as competitors, some of the interesting phenomenon concerning moral reasons (such as inescapability) and that it is this explanatory adequacy that gives it such plausibility as it has. I take this to be a particular instance of Blackburn's (1984: 170) general reply to Searle's point.

takes facts about reasons to be reducible to facts about the preferences or motivations of rational agents, where these are not constrained by the particular motivational set one starts with but only by the demands of ‘pure practical reason’ (Korsgaard 1996).

It is interesting to note a common substantial assumption in these popular theories: that reason statements (including moral reason statements), when true are made true by the instantiation of a distinct worldly relation—the *reason relation*. On the Humean view, this relation is partly constituted by the desires of actual or idealised agents; on the Moorean view it is a sui generis normative relation; on the Aristotelian view it is partly constituted by sui generis value-facts; on the Kantian view it is partly constituted by the preferences of rational agents. Hence these theories assume that reason statements are what Skorupski (2006) calls *pictorial*: they depict a certain reason-state-of-affairs as obtaining in the world and are true when the state of affairs they depict really does obtain (This is not to say that these theories are always put forward as analyses of the concept ‘reason’: see Schroeder 2007: 61–83). This assumption is so ingrained in existing debates that Schroeder, in his defence of a Humean Theory, claims that any ‘perfectly general explanatory theory of reasons’ must have the following form:

For all propositions, r , agents x , and actions a , if r is a reason for x to do a , that is because r , x , and a stand in relation R (2007: 50).¹²

It is interesting to consider the possibilities that open up when this substantial truth-maker assumption is rejected. Skorupski has questioned the assumption’s motivations. One is the thought that since reason statements express propositions that can be true or false, they require substantive truth-makers when they are true. But, as Skorupski notes, this assumes a less-than compulsory ‘correspondence’ model of propositions (2006: 32–34). Another reason why people may be drawn to the assumption is the terminological preference for framing debates in this area as debates about the ‘nature of reasons’, apparently taking for granted that reasons are *things* in the world. I noted above that talk of (moral) reasons can be understood in an ontologically non-committal way, as talk about the correct application of reason predicates. By failing to recognise this possibility discussions of reasons often sleep-walk into the substantial truth-maker assumption. But once this assumption is exposed it is clearly question-begging when deployed as an axiom that frames the range of possible metasemantic theories of reason predicates (since not all predicates denote world-bits). Hence, the only non-question begging starting point here is our making of (moral) reason statements, with it being a further question how the content of those statements is understood (depictive or otherwise).

Skorupski’s own view seems to leave the important question of *content* unaddressed. It claims that reason statements are assertions of normative propositions that can be true or false, but denies that such propositions need to be made true by worldly states of affairs. Here it is unclear how reason statements

¹² See also Dancy (2004: 29): ‘We will...be trying to understand the nature of a certain relation, not the nature of a certain monadic property such as goodness’.

acquire their particular semantic content, if it is not derived from their depicting reason-relations.¹³ An interesting alternative, therefore, is to agree with Skorupski in rejecting the substantive truth-maker assumption while giving a further, non-depictive, account of the semantic content of (some) reason statements. The expressivist theory is one such account.

To emphasise, all this is not to say that we have reason to *reject* the truth-maker assumption, just that we do not have sufficient reason to accept it as framing axiom of the debate. It may be that the best metasemantic theory of reason statements is one that incorporates the truth-maker claim; or it may be that the best theory is one that rejects it, but includes (as, it seems to me, Skorupski's view does not) equally substantive claims about the content-determining functional role of moral reason statements (such as the expressivist commitments I elucidate below). The current point is simply that we cannot motivate either of these views by framing the debate with assumptions that render the competitors invisible.

3.3 Second lacuna

As well as being neglected by opponents, expressivist accounts of (moral) reasons have been left underdeveloped in the hands of expressivists whose primary focus is on other normative concepts (for brief remarks on the topic by expressivists, see Blackburn 1998: 257; Gibbard 1990: 163). This neglect is potentially problematic, since it is commonly accepted, for example, that 'A (morally) ought to ϕ ' entails 'A has (moral) reason to ϕ ' (e.g. Shafer-Landau 2003: 190–214) and that 'A's ϕ ing is morally wrong' entails 'A has moral reason not to ϕ '. Expressivists need to address these potential problems; the theory that follows is part of an attempt to do so.

4 Expressivism about moral reason statements

4.1 The expressivist theory

In Sect. 3.1 I suggested there was space for a theory of moral reason statements that incorporates their persuasive role into their semantic content. In Sect. 3.2 I suggested there was space for a theory of reasons that rejects the assumption that all reason statements are depictive of reason-relations. In Sect. 3.3 I suggested a need for expressivists to extend their account beyond the more extensively studied normative concepts. These considerations come together in the following theory.

The expressivist theory is initially characterised by the following two claims, with an important third to follow:

- (a) When a speaker sincerely utters a sentence of the form 'F is a moral reason for A to ϕ ' she expresses the belief that F.

¹³ This mirrors Williams' (1989) criticism that defenders of external reasons do not 'offer any content for external reason statements'.

And

- (b) When a speaker sincerely utters a sentence of the form ‘F is a moral reason for A to φ ’ she expresses moral approval of A’s being motivated (to some extent) by awareness of F to φ .

Here (a) captures the uncontroversial depictive content of moral reason statements. One cannot assert ‘F is a moral reason for A to φ ’ without asserting that F, and on one plausible view of assertion, if an agent makes a sincere assertion ‘it follows that she has a belief whose content can be captured by means of the sentences used’ (Wright 1992: 14).¹⁴ Thus even the expressive theory of moral reason statements admits that they have some depictive content; what it denies is that they depict moral reason-relations.

According to (b) moral reason statements express a particular kind of approval in the same way that, according to traditional expressivist views, statements like ‘ φ is right’ express a particular kind of approval of φ . The notion of ‘expression’ is the same as that employed in expressivist accounts of other normative statements and though this notion may prove problematic, there is nothing additionally puzzling about its use here.¹⁵ One point, however, deserves emphasis. Although (a) and (b) provide an account of the meaning of individual speech acts, what determines that meaning (according to the account) is not the idiosyncratic attitudes of the agent performing those acts. That is, (a) and (b) should not be tied to a generally discredited Lockean theory of meaning. Such a theory may perhaps apply to wincing and other primitive ‘venting’ of emotion, but it is implausible in the case of linguistic meaning (as we have here) primarily because constancy in meaning can be preserved across idiosyncrasies in corresponding ideas or attitudes. Instead (a) and (b) should be understood as involving a type of ‘expression’ determined by social conventions, which are in some sense grasped by both speaker and hearer and embedded in the wider linguistic community of which both are a part. When (b) claims that that reason statements express a particular type of moral attitude, therefore, it should be understood as saying that the public conventions between speakers and hearers that define the meaning of predicates determine that such statements expresses such attitudes.

Claim (b) also assumes that some attitudes are distinctively moral. Explaining this distinctiveness is an important part of any expressivist theory. But again this is a general issue for expressivists, not specific to the particular case of moral reason statements.¹⁶ What is distinctive about (b) is the claim it makes about the *first-order object* of the distinctively moral approval expressed by moral reason statements, and

¹⁴ The claim that reason statements are ‘factive’ in this way is questioned by Suikkanen (2012: 599), but this argument can be countered by providing a more fine-grained account of the relevant reason-givers.

¹⁵ See Schroeder (2008) for some problems and Blackburn (2001) for an account.

¹⁶ For what it is worth, my preferred view is that moral attitudes are stable emotionally ascended states of approval and/or disapproval partly constituted by dispositions to encourage and/or sanction the attitudes and behaviours of others (Sinclair 2014). I take this account to be broadly in line with those suggested by Blackburn (1998: 8–21) and Gibbard (1990). See also Björnsson and McPherson (2014). An interesting alternative is given by Chrisman (2008).

it is this object that (according to the expressivist theory) distinguishes the attitudes expressed by such statements from the attitudes expressed by other moral statements. According to standard versions of expressivism about wrongness, for instance, the first-order object of the attitude expressed by the statement that ‘Murder is wrong’ is a particular type of action, viz. murder. According to expressivism about moral reasons, in contrast, the first-order object of the attitude expressed by the statement that ‘F is a moral reason for A to ϕ ’ is a particular type of motivational sensibility, namely one that takes awareness of F as an input and generates some motivation to ϕ as output. I shall call moral approvals of motivational sensibilities ‘structured approvals’. (b) claims that moral reason statements express structured approvals.

Since it is a distinctive feature of the present account, it is worth dwelling on the objects of structured approvals, viz. motivational sensibilities. A motivational sensibility can be defined by a function from input in cognition or perception to an output in motivation. In the sense at issue here, to say that A is motivated (to some extent) by awareness of F to ϕ is not to say that A ϕ 's. Rather, it is just to say that A's awareness of F gives rise in A to some motive, of some strength, to ϕ . This motive might not lead to action. To have some motive to ϕ is, roughly, to be disposed to ϕ (or intend to ϕ) absent competing motives and independently specifiable psychological maladies, such as depression and weakness of will.¹⁷ The strength of a motive is therefore determined by its interactions with other motives, actual and hypothetical, in the issuance of action. For example, a strong motive to ϕ is one that results in the agent possessing it ϕ ing in a large range of close-by scenarios where it is coupled with other, contrary, motives. A weak motive to ϕ , by contrast, is one that results in the agent ϕ ing only in a small number of close-by cases, where the agent has few or no other competing motives. The notion that motives can have differing strengths is common in (folk) psychology and not an idiosyncratic commitment of the expressivist theory.

One consequence of this notion of strength of motive is the availability of a related account of weight for moral reasons. A natural corollary of (b) is the view that the claimed weight of a reason is directly proportional (though not semantically reducible) to strength of the motive approved of by the state of structured approval expressed by the relevant statement. So, for example, to think that F provides a *heavy-weight* reason for A to ϕ is to approve of A's awareness of F giving rise in A to a *strong* motive to ϕ . The expressivist account can therefore respect the general thought that there is some connection between the weight of reasons and the strength of motives they demand.¹⁸

More generally, it is easy to see how the present account can account for the *contributory* nature of normative reasons. To judge that F is a reason for A to ϕ is to express moral approval of A's being motivated (to some extent and in a certain range of situations) by awareness of F to ϕ . Suppose one also approves of A's being

¹⁷ I intend to remain neutral on the question of whether a motive is a state that properly ends in ϕ ing or intending to ϕ , but for convenience will speak only of the former.

¹⁸ See Setiya (2014: 229). Of course, an account of weight needs to do more much more than this (Schroeder 2007: 123–129) but these details must wait for another occasion.

motivated (to some extent and in a certain range of situations) by awareness of G to φ . Sometimes, the cases where one approves of A being to some extent F -motivated to φ will overlap with the cases where one approves of A being to some extent G -motivated to φ and one will approve of A being motivated *to a greater extent* to φ by F and G ; the reasons have ‘added-up’. In another type of case, one can approve of A being F -motivated to φ and approve of A being H -motivated *not* to φ ; the reasons are in conflict. For example, in a range of situations I may approve of my children being sufficiently motivated to tell lies (when it is necessary to save a life, for example) but still, when they do, I would rather they did so only reluctantly. Furthermore, the type of structured approval which is (according to the present account) expressed by reason judgements can be more or less fine-grained. For example, I may approve of A ’s being motivated (to some extent) by awareness of F to φ only in the range of cases where not- I obtains. This is to treat I as a *silencer* of what would otherwise be a reason to φ . In a different range of cases, one may morally approve of A being *somewhat* moved by F to φ when J is absent but approve of A being moved by a greater (or lesser) degree when J is present. This is to treat J as an *amplifier* (or *attenuator*) of the initial reason. The expressivist account of reason statements presented here, therefore, seems well-placed to capture the way in which reason-givers and other conditions can interact.¹⁹

Before moving onto the final part of the expressivist account, it is worth dwelling briefly on potential objection to (b). One’s approval of A being F -motivated to φ may arise in an odd way. Suppose, for example, that an evil demon offers Barbara ten dollars if she is moved to some small degree by awareness of green objects to pick them up. Suppose that I , being generally well-disposed towards Barbara and happy for her to make some easy money, consequently approve of Barbara being so moved. In such a case, I would not want to say that the fact that an object is Green is reason for Barbara to pick it up and yet I have the attitude which, according to the expressivist account, is expressed by such a statement. Something must have gone wrong with the account.

This problem is based in the more general one of marking the distinction between the ‘right-kind’ and ‘wrong-kind’ of reasons for possessing certain mental states, such as beliefs, intentions, motivational sensibilities or approvals thereof. According to a recent suggestion, ‘however this distinction works, it must work by somehow associating each attitude (or activity) with a distinctive kind of benefit which generates the right kind of reason with respect that that activity’ (Schroeder 2012: 485). On this (reductive) view of the ‘right kind of reasons’ each attitude or activity of forming such attitudes is associated with a distinctive goal or purpose, and when the benefits of holding (or not holding) such an attitude are appropriately connected

¹⁹ Contra Dancy (2004: 57–60) who argues that Gibbard’s expressivist account of reasons cannot capture the contributory nature of reasons. The above point about the ‘fine-grain’ of structured approvals dispenses with Dancy’s objection that expressivism cannot account of the ‘organicity’ of reasons. On the expressivist view, to judge that F is a reason for A to φ in C is to express approval of a certain (type of) agent being moved in a certain direction by a certain feature *in a certain range of cases* (which may or may not be well-defined in the mind of the judge). It does not follow, therefore, that when an agent holds that F provides a reason for A to φ , she must hold that F provides such a reason whenever it appears, no matter what other features it is bundled with.

to that purpose, there are reasons of the right kind to have that attitude. For example, according to Schroeder the distinctive purpose of *intending* is to ‘close off deliberation, in order to allow us to co-ordinate and control our own actions across time and make decisions at times at which we have more available cognitive resources’ (2012: 483). Hence reasons of the right kind to intend are those deriving from the benefits of forming a given intention having to do with this distinctive role. In Kavka’s (1983) toxin puzzle, for instance, the benefits of forming the intention to drink the poison arise whether or not the intention plays this ordinary role, so these are not ‘right-kind reasons’ to form the intention.

Consider, then, the distinctive purpose of the attitudes I am calling motivational sensibilities. Plausibly, their distinctive purpose is to (be one of possibly many states which) mediate the relation between features of agents’ situations, their motivations and (ultimately) actions, so that agents can respond in action to the various situations they find themselves in. Hence reasons of the right kind to form a motivational sensibility are those deriving from the benefits of forming such a sensibility having to do with this distinctive, ultimately action-guiding role. If so, Barbara has the wrong kind of reason to form the motivational sensibility which takes awareness of green colour as an input and generates a motive to pick up as an output (the ‘Green-Pickup Sensibility’), since the benefits of her having this state arise whether or not it plays its ordinary role (of mediating the connection between worldly features and action). The expressivist can then claim that moral reason statements express approval of motivational sensibilities, where the basis of that approval relates to the right kind of reason to have such sensibilities. In the case of Barbara, although I approve of her possessing the Green-Pickup Sensibility, the basis of that approval does not relate to the ‘right kind of reason’ to have such a sensibility, rather it relates to a reason of the wrong kind (the reason given by the evil demon’s financial incentive). In general, then, and to avoid the objection, the expressivist can modify (b) to the claim that when a speaker sincerely utters a sentence of the form ‘F is a moral reason for A to φ ’ she expresses moral approval of A’s being motivated (to some extent) by awareness of F to φ , where the basis of this approval relates to the ‘right kind of reasons’ to have such a motivational sensibility, and where the right kind of reasons for possessing such a sensibility are understood (reductively) as those which derive from the benefits/costs of possessing such a certain state vis-à-vis its distinctive purpose of mediating the relation between worldly features and action. Although I think that, ultimately, this more nuanced understanding of (b) is required in order to give a fully satisfactory account of moral reason statements, for ease of exposition in what follows I will ignore this (important) complication.

In addition to (a) and (b), the expressivist theory makes the following claim, about the practical purpose of the type of speech act in question:

- (a) The function of the expression in (a) and (b) is the mutual co-ordination of attitudes and actions through the distinctive mechanisms of moral discussion.

According to expressivism about moral reason statements, then, such statements express particular combinations of beliefs and moral attitudes for the practical

purpose of mutual co-ordination of attitudes and actions through moral discussion. It is important to note that the notion of function or practical purpose at stake here (like the notion of expression which it modifies) is communal or anti-individualist. To say that moral reason statements express attitudes and that that expression has a particular co-ordinating function (or equivalently, to say that moral reason statements express attitudes for a particular co-ordinating purpose) is to say that such statements are part of a convention-governed expressive practice and that the overall function of that practice is to help agents co-ordinate in a particular way. It is consistent with this view that some token moral reason statements do not fulfil this function or are made for distinct individual purposes, so long as they are part of a conventional practice which, overall, does have this function or purpose. The function of the overall practice (as I shall explain below) helps determine the particular norms governing the individual speech-acts involved and it is these norms, rather than those deriving from idiosyncratic speaker intentions, that, according to expressivism, determine the speech-acts' meaning. As Blackburn (1995: 49) puts it: 'Subjective states are not the authority: we are entered into a conventional system, and we can no more escape its norms than we can make the coins we tender worth what we think they are worth'. In what follows, therefore, talk of the function of moral reason statements should be understood as the function of the overall practice of making such statements.

Note also that this claim about practical function is not peculiar to expressivism about moral reason statements, although the application of it to this context is novel. All plausible expressivist views supplement the claim that moral statements express attitudes with an account of the practical purpose of that expression. On the most common version of this claim, the practice of which such statements are a part is a distinctive linguistically infused mutual co-ordination device through which competing parties can negotiate towards (and thence maintain) mutually beneficial and stable patterns of attitudes and actions (cf. Blackburn 1998, Gibbard 1990). A particular instance of this functionality was described by one early expressivist:

When you tell a man that he ought not to steal, your object is not merely to let him know that people disapprove of stealing. You are attempting, rather to get *him* to disapprove of it. Your ethical judgment has a quasi-imperative force which, operating through suggestion and intensified by your tone of voice, readily permits you to begin to influence, to modify, his interests... Thus ethical terms are instruments used in the complicated interplay and readjustment of human interests. (Stevenson 1963: 16–17).

While modern expressivists might not agree with Stevenson's implication that moral statements fulfil their practical function in part by describing the contours of societal approval, or with the implicit assumption that the function involved is derived from individual speaker intentions, they would agree with his claim that having some co-ordinative function is distinctive of moral statements. As Blackburn puts it: 'The ethical proposition gets its identity as a focus for practical thought, as people communicate certainties, insinuations and doubts about what to value' (1998: 50).

The particular application of this general idea to the case of moral *reason* statements generates a view similar to Williams' and Joyce's claims about the pragmatic function of moral reason statements. According to the general claim, moral statements express attitudes in part for the persuasive purpose of altering the attitudes and actions of others.²⁰ In the case of reason statements, the primary purpose of this type of expressive speech-act is not persuade agents to ϕ , but to persuade them adopt a *sensibility* that takes recognition of F as an input and generates some motive to ϕ as an output (and hence, in the particular circumstance where F obtains, to persuade them to be moved by recognition of F in the direction of ϕ).²¹ For Williams and Joyce, with their implicitly depictive model of moral reason statements, this type of persuasive influence proceeds via the stating of falsehoods (or non-truths). For expressivists, this persuasive influence is part of the meaning of the statements themselves.

One pleasing feature of this type of expressivism is that it can explain why we might want the particular mode of persuasive co-ordination made possible by moral reason statements. Roughly, it will be helpful for groups of co-ordinating agents to agree, not only to shared patterns of *action* regarding new and recurring types of situations, but to shared patterns of *reaction* to particular features of such situations. For example, a process of agential co-ordination might have the function of getting those agents to coalesce around encouraging and performing a particular type of action, say, the looking after of elderly relatives. Judgements of rightness and wrongness, whose first-order objects are actions, seem well-suited to this role. But it can be useful for a group of agents to coalesce not just around performing and encouraging certain *actions*, but on possessing and encouraging certain types of *motivational sensibility*. For example, a discursive practice might aim to influence others so that, like us, they feel some motive to alleviate others' pain. This is not necessarily to encourage them to always act so as to alleviate pain; it is consistent with this motivational profile to, for example, leave sado-masochists to their own devices. But it is to encourage agents to feel *some* motive to alleviate observed pain. Since how we get along with others depends not just on how they act in given situations, but also on how they *react* to the particular features of those situations, it makes sense to have a particular mode of co-ordination aimed towards harmony of motivational sensibilities alongside a mode aimed towards harmony of action. The downside of the former is that it is less determinate than the latter: we may achieve harmony in having *some* motive to alleviate others' pain, for instance, without achieving harmony in how to deal with sado-masochists. The upside is that it allows a deeper and more fine-grained mode of co-ordination.²² That there are distinct

²⁰ Another important part of co-ordination is the reinforcement of existing attitudes and patterns of action. For ease of exposition, I henceforth ignore this aspect.

²¹ A secondary purpose is to persuade others to adopt the same sensibility.

²² See Regan (1980) for the view that co-operation requires sensitivity to others' motivational sensibilities rather than prediction of their behaviour. If Regan's arguments are correct, there is some reason to think that the type of co-operation aided by moral reason statements is more fundamental than the type aided by other types of moral statements. On the expressivist account, this in turn can help explain why the concept of a (normative) reason is more fundamental than other normative concepts.

advantages and disadvantages of both modes of co-ordination plausibly explains why actual moral practice involves both mechanisms.

According to expressivism about moral reason statements, then, when a speaker sincerely utters ‘F is a moral reason for A to ϕ ’ she asserts that F and (conventionally) expresses moral approval of awareness of F giving rise, in A, to a motivation to ϕ . The primary (again, conventionally enshrined) purpose of this type of expression is persuasive: to persuade A to recognise F and, subsequently, to be moved by that recognition in the direction of ϕ . Moral reason statements function to guide action by alerting agents to features of their situation which (if all goes well) will move them in the direction suggested.

4.2 Expressivism and inescapability

Can the preceding account explain the inescapability of moral reasons? The key to doing so is to recognise that only statements that obey the constraint of inescapability can play the persuasive co-ordinating role which the expressivist takes to be (partly) definitive of moral reason statements.

First, as previously argued, the inescapability claim is best understood as the claim that the moral reason predicate is not correctly applied on the basis of D*-states. Second, according to expressivism, this predicate expresses structured approvals for a particular persuasive purpose, viz. the co-ordination of motivational sensibilities described above. Third, a group of agents could not employ such a structured-approval-expressing predicate to play this particular purpose unless its correct application was, quite generally, D*-state independent. To see this last claim, consider a group of agents for whom the meaning-determining conventions surrounding the use of a structured-approval-expressing predicate determine that that predicate *is* correctly applied purely on the basis of D*-states. Members of such a group would not be able to (semantically) correctly say that F is a moral reason for A to ϕ unless A had some desire promoted by ϕ ing and, further, they would be (semantically) correct to make such a reason statement whenever A had some desire promoted by their ϕ ing. It is clear that the (semantically) correct deployment of such a predicate could not be helpful in *altering* or *re-orientating* A’s motivational sensibility towards ϕ ing, since these agents could only correctly deploy the predicate when A’s motivational sensibility is already disposed (in virtue of her D*-state) in that direction.²³ But the purpose of *altering* others’ motivational sensibilities is an essential part of the persuasive co-ordinating role outlined above. Thus, a structured-approval-expressing predicate can only play this persuasive role if the meaning-determining conventions surrounding the use of such a predicate determine that its correct application is not D*-state dependent (i.e. that moral reasons are not Humean reasons). Finally, because (expressivists claim) playing this particular persuasive purpose is (partly) definitive of the concept of a moral reason, what is necessary for a predicate to play this persuasive role is also necessary for a

²³ I admit that such a predicate could serve a distinct but related practical purpose, that of *directing* an existing sensibility in ways it is already somewhat disposed to go. See Sect. 4.3.

collective use of that predicate to count as deployment of the moral concept. Thus, if convention-governed use of a structured-approval-expressing predicate is to count as the deployment of the concept of a moral reason, it must be that the predicate is not (semantically) correctly applied purely on the basis of D*-states. Thus inescapability is a conceptual truth of moral reasons.

In short, a set of moral reason statements must be semantically governed by the correctness conditions enshrined in inescapability because those statements' essential co-ordinating role is in part to *change* targets' motivational sensibilities, not reflect them.

It is interesting to note the potential for this type of expressivist explanation to extend to other features of moral reasons. Note the general form. First, the phenomenon to be explained is given a metatheoretically neutral interpretation, as a claim concerning the correct application of the moral reason predicate. Second, the expressivist gives an account of the (communal, conventionally enshrined) practical function of that predicate: the mutual co-ordination of motivational sensibilities. Third, the expressivist explains how the norms of application of the predicate can therefore be understood as arising from the norms of application that must apply to predicates if they are to function in this particular co-ordinating way. Finally, because the concept of a moral reason is defined by this co-ordinating role (morality is essentially practical) the norms laid down by the need to co-ordinate are thereby shown to be the norms which must apply if agents' use of that predicate is to count as deployment of the concept of a moral reason. That is, the norms engender conceptual truths about moral reasons.²⁴

4.3 The arguments for Humeanism

The above provide grounds for optimism that expressivists can deploy their account of the practical role played by moral reason statements to explain some of the features of such statements. But, as noted at the outset, the expressivist theory will only be plausible to the extent that it can also defang some of the motivations for the rival Humean Theory. So how might the expressivist reply to the arguments of Sect. 3.1?

It is helpful to consider the second type of argument first. That argument challenged those who denied Humeanism to explain where the constraints on what moral reasons there are might come from, if not from the limits set by the existing motivations of agents. Given my general strategy, this can be recast as the challenge to explain where the norms of correctness for applying the moral reason predicate might come from, if not the function of accurately depicting some (Humean) reason-relation. The expressivist answer is clear: the application of the moral reason predicate is part of a particular co-ordinating practice, and the needs of that practice partially limit when this predicate is (in)appropriately applied. In particular, given the needs of this practice, agents engaging in it must be bound by the standard that

²⁴ For a further example, concerning the conceptual truth of moral mind-independence, see Sinclair (2008).

the predicate is not appropriately applied purely on the basis of D*-states. Note that these constraints do not provide a completely determinate account of correct application for the moral reason predicate—that, after all, is in part a matter of substantive normative debate, which, on the expressivist view, is a matter of how to co-ordinate. But these constraints do help explain the general contours of a complete theory of moral reasons: whatever that theory is, it should allow that there can be moral reasons in the absence of D*-states.²⁵

The first argument in favour of Humeanism relied on the premise that some reasons are desire-based in the way hypothesised by the Humean Theory. On my account, these are a particular type of non-moral reason, which Copp (1997: 92) helpfully labels ‘self-grounded’. The argument then seeks to drive a wedge from this type of reason to all reasons. A complete response to this argument requires the expressivist to offer a companion theory of non-moral reasons (including self-grounded reasons) to go with their theory of moral reasons. I cannot provide all the details of such an account here, but I can provide some grounds for optimism for thinking that such an account is possible. First, note that this more comprehensive theory of reasons needs to do two things: (1) give a general account of the role of reason statements, such that it can explain the similarity between moral and non-moral reason statements (for instance, both are canonically stated as ‘F is a reason for A to ϕ ’) and (2) explain the difference(s) between moral and non-moral reason statements. There is reason to think that the expressivist account can fulfil both desiderata. First, it is open for the expressivist to claim that all (normative) reason statements play some role in the practice of mutual co-ordination of action and reaction described above. Such statements, expressivists can claim, possess a unity of persuasive function. But second and more important, expressivists can give an account of the differences between types of reason statements in terms of different ways of contributing to that function. For example, the expressivist might claim that self-grounded reason statements are made for the purpose not of *altering* motivational sensibilities, but of *directing* existing ones. Such statements may even play this role by depicting the relations Humeans take to truth-make reason statements. After all, to be informed that, because of F, ϕ ing will promote one of my desires, is likely to move me in the direction of ϕ ing (such depictive statements may, further, conversationally implicate the speaker’s approval for being so moved, given Grice’s maxim of conversational relevance). On this view, *directing* and *altering* motivational sensibilities are two distinct yet complementary parts of the

²⁵ The constraints on the correct application of the moral reason predicate that I elucidate are perhaps more accurately labelled constraints of *competence*: constraints one must generally abide by in order to be considered a competent user of that predicate at all. These can be distinguished from strict constraints of correctness, that is, constraints you must abide by in order to earn my assent to your use of the predicate. For this helpful distinction, see Cantwell (2013: 1382). It follows that the expressivist position here is not committed to any controversial thesis about the normativity of meaning: whether or not *competence* with a given piece of language is a fundamentally normative matter is independent of whether the *correct* application of it is a normative matter. In general, the expressivist will hold that to think a predicate *correctly* applied is to think it appropriate given the (conventionally enshrined) standards which constitutively govern that predicate. These standards can yield meaning which is descriptive (a matter of how the world is) or normative (a matter of how to respond to the world, of how to live).

mutual co-ordination of action and reaction, and thus there is, at a general level of description, a unified practical function that both self-grounded and moral reason statements serve (a function which can explain their syntactic similarity). Nevertheless, the different ways in which those statements fulfil that function gives rise to their differences (and reflects an underlying semantic difference). If this is right, the expressivist can resist the wedge from self-grounded reasons to all reasons by subsuming all reason statements under the category of 'practical persuasive devices' and giving distinct accounts of how different types of statements fulfil this function.²⁶ Although the details of such an account need filling out, there is no reason as yet to think that they cannot be provided.

The final, Mackian argument in favour of the Humean Theory is easier to deal with. By understanding moral reason statements as expressive of attitudes rather than depictive of reason-relations, the expressivist faces no difficulties about the ontological, epistemological or psychologically 'dynamic' status of such relations.

5 Conclusion

There are two types of unfinished business. First like any expressivist theory, the above account needs to say something about general objections to expressivist accounts of any area of normative language, such as the Frege–Geach problem. Though these are potentially serious problems there is no reason to think the particular application of expressivism to reason statements either exacerbates or ameliorates them. General problems require general solutions, so solving these problems is not a particular burden of the current version of expressivism.

The second type of unfinished business is more particular. As a theory of moral reason statements, this account needs to explain all of the uncontroversial phenomena concerning those statements. Here I have briefly tried to explain weight, contributoriness, and inescapability. But a complete theory would need to explain much else besides, such as: the fact that reason statements are partly non-extensional contexts; the connection between reason statements and other types of normative statements (such as judgements of wrongness and ought); the connection between reason statements, the considerations with which we deliberate and motivation; the apparent irreducibility of the concept of a reason. Nevertheless, the foregoing arguments provide some grounds for optimism that such explanations can be given. Those grounds are, first, the existence of a recipe for explaining features of moral reasons which can first be understood as providing constraints on the correct application of the moral reason predicate (see Sect. 4.2). It is an

²⁶ The 'bifurcated' account just given is for illustrative purposes only. An alternative is a thorough-going expressivist option which takes all reason statements to express structured approvals, but holds that self-grounded reason statements express structured approvals that are themselves sensitive to elements of the target agent's motivational set. On this latter view, the claim that there are some self-grounded reasons expresses approval of some agents being motivated by facts appropriately related to some elements of their subjective motivational set. One might worry that on the bifurcated account there is no common category of 'reasons' that serve as inputs into deliberation. But this would be mistaken – for both types of judgement the reason-givers are (intensional) facts. See footnote 2.

interesting and substantively open question just how generally palatable this recipe is. Second, and more generally, by locating the role of moral reason statements within a broader framework of a co-ordinating practice, the expressivist provides additional resources to understand the complexities of such statements (such as their dimension of weight). In so far as these resources are available the hitherto neglected expressivist treatment of reason statements seems to represent a research programme well worth engaging with.

Acknowledgments Work for this paper was supported by the Arts and Humanities Research Council (grant number AH/J006394/1). I would also like to thank seminar audiences at the Universities of Manchester and Lingnan.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bjornsson, G., & McPherson, T. (2014). Moral attitudes for non-cognitivist: Solving the specification problem. *Mind*, 123(489), 1–38.
- Blackburn, S. (1984). *Spreading the word*. Oxford: Clarendon Press.
- Blackburn, S. (1993). *Essays in quasi-realism*. Oxford: Oxford University Press.
- Blackburn, S. (1995). The flight to reality. In R. Hursthouse, G. Lawrence, & W. Quinn (Eds.), *Virtues and reasons: Philippa Foot and moral theory*. Oxford: Clarendon Press.
- Blackburn, S. (1998). *Ruling passions*. Oxford: Oxford University Press.
- Blackburn, S. (2001). Reply. *Philosophical books*, 42, 27–32.
- Cantwell, J. (2013). First order expressivist logic. *Erkenntnis*, 78, 1381–1403.
- Chrisman, M. (2008). Expressivism, inferentialism, and saving the debate. *Philosophy and Phenomenological Research*, 77(2), 334–358.
- Copp, D. (1997). The ring of Gyges: Overridingness and the unity of reason. *Social Philosophy and Policy*, 14, 86–106.
- Daly, C., & Liggins, D. (2010). In defence of error theory. *Philosophical Studies*, 149(2), 209–230.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Oxford University Press.
- Dewey, J. (1945). Ethical subject-matter and language. *Journal of Philosophy*, 42(26), 701–712.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(1), 1–20.
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *Philosophical Review*, 81, 305–316.
- Gaut, B. (1997). The structure of practical reason. In B. Gaut & G. Cullity (Eds.), *Ethics and practical reason*. Oxford: Oxford University Press.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge MA: Harvard University Press.
- Goldman, A. (2009). *Reasons from within*. Oxford: Oxford University Press.
- Hooker, B., & Streumer, B. (2004). Procedural and substantive rationality. In A. Mele & P. Rawling (Eds.), *The Oxford handbook of rationality*. Oxford: Oxford University Press.
- Humberstone, L. (1992). Direction of fit. *Mind*, 101, 59–83.
- Jenkins, C. (2005). Realism and independence. *American Philosophical Quarterly*, 42(3), 199–209.
- Joyce, R. (2000). The fugitive thought. *Journal of Value Inquiry*, 34, 463–478.
- Joyce, R. (2006). *The evolution of morality*. Cambridge, MA: MIT Press.
- Kant, I. (2007). Groundwork of the metaphysics of morals. In: R. Shafer-Landau (Ed.) *Ethical theory: An anthology* (Mary J Gregor and extracted, Trans.). Oxford: Blackwell.
- Kavka, G. (1983). The toxin puzzle. *Analysis*, 43, 33–36.
- Kolodny, N. (2005). Why be rational? *Mind*, 114, 509–563.

- Korsgaard, C. (1996). *The sources of normativity*. Cambridge: Cambridge University Press.
- Lewis, D. (1970). General semantics. *Synthese*, 22(1–2), 18–67.
- Markovits, J. (2011). Why be an internalist about reasons? In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 6). Oxford: Oxford University Press.
- Olson, J. (2009). Reasons and the new non-naturalism. In S. Robertson (Ed.), *Spheres of reason*. Oxford: Oxford University Press.
- Pérez Carballo, A. & Santorio, P. (Forthcoming). Communication for expressivists. *Ethics*.
- Raz, J. (1999). *Engaging reason: On the theory of value and action*. Oxford: Oxford University Press.
- Regan, D. (1980). *Utilitarianism and co-operation*. Oxford: Clarendon Press.
- Ridge, M. (2014). *Impassioned belief*. Oxford: Oxford University Press.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge, MA: Harvard University Press.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. (2008). Expression for expressivists. *Philosophy and Phenomenological Research*, 76, 86–116.
- Schroeder, M. (2012). The ubiquity of state-given reasons. *Ethics*, 122(3), 457–488.
- Searle, J. (1969). *Speech acts*. Cambridge: Cambridge University Press.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies*, 167(2), 221–235.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Silk, A. (2013). Truth conditions and the meanings of ethical terms. *Oxford Studies in Metaethics*, 8, 195–222.
- Sinclair, N. (2008). Free thinking for expressivists. *Philosophical Papers*, 37(2), 263–287.
- Sinclair, N. (2014). On standing one's ground. *Analysis*, 74(3), 422–431.
- Skorupski, J. (2006). Propositions about reason. *European Journal of Philosophy*, 14, 26–48.
- Skorupski, J. (2009). The unity and diversity of reasons. In S. Robertson (Ed.), *Spheres of reason*. Oxford: Oxford University Press.
- Skorupski, J. (2010). *The domain of reasons*. Oxford: Oxford University Press.
- Stevenson, C. L. (1963). *Facts and values*. New Haven: Yale University Press.
- Suikkanen, J. (2012). Reasons-statements as non-extensional contexts. *The Philosophical Quarterly*, 62(248), 592–613.
- Wiland, E. (2012). *Reasons*. London: Continuum.
- Williams, B. (1981). Internal and external reasons. In R. Harrison (Ed.), *Rational action*. Cambridge: Cambridge University Press.
- Williams, B. (1989). Internal reasons and the obscurity of blame. In William J. Prior (Ed.), *Reason and moral judgment, logos* (vol. 10). Santa Clara, CA: Santa Clara University.
- Wright, C. (1992). *Truth and objectivity*. Harvard: Harvard University Press.