



Norm-expressivism and regress

Attila Tanyi

To cite this article: Attila Tanyi (2017) Norm-expressivism and regress, South African Journal of Philosophy, 36:3, 362-376, DOI: [10.1080/02580136.2017.1294642](https://doi.org/10.1080/02580136.2017.1294642)

To link to this article: <http://dx.doi.org/10.1080/02580136.2017.1294642>



Published online: 24 Aug 2017.



Submit your article to this journal [↗](#)



Article views: 3



View related articles [↗](#)



View Crossmark data [↗](#)

Norm-expressivism and regress

Attila Tanyi

Department of Philosophy, University of Liverpool, UK; and
Collegium Helveticum, ETH/University of Zurich, Switzerland
attila.tanyi@liverpool.ac.uk

This paper aims to investigate Allan Gibbard's norm-expressivist account of normativity. In particular, the aim is to see whether Gibbard's theory is able to account for the normativity of reason-claims. For this purpose, I first describe how I come to targeting Gibbard's theory by setting out the main tenets of quasi-realism cum expressivism. After this, I provide a detailed interpretation of the relevant parts of Gibbard's theory. I argue that the best reading of his account is the one that takes normativity to be carried by a controlled, coherent, comprehensive set of norms. Finally, I present a potential obstacle to Gibbard's approach: the regress problem. The idea is to examine the structure of the non-cognitive state expressed and find it inadequate due to the possibility of an infinite regress in the justification of the norms whose acceptance it contains. I then end the paper with some concluding remarks.

Quasi-realism, norm-expressivism and reasons

Reasons are no doubt important. They are important in our everyday life: we customarily refer to reasons when we explain our actions, as well as when we try to justify what we plan to do, or what we did. Reasons are perhaps even more important in contemporary analytic philosophy. In a recent book, Thomas Scanlon (2014, 1–2) points out that reasons have become the focus of philosophical research in at least two ways. First, there is currently more interest in normativity and reasoning spelled out in terms of reasons than in morality. Secondly, there is more interest in reasons than in motivation – that is, in what reason one has to do what, e.g. morality or prudence demands rather than what moves one to so act.

The present article adds to this growing literature by focusing on an influential line of thought about normativity: the non-cognitivist. Any inquiry of this sort must answer three basic questions. First, it must say something about why it approaches normativity from the standpoint of reasons. Is normativity exhausted by reasons? Not everyone would agree. Second, it must explain why non-cognitivists would want to account for normativity at all. Early non-cognitivists such as Ayer or Stevenson were willing to abandon the right to ethical (or, at least, moral) knowledge and truth; so why would they and their followers care about normativity? Third, even if they do, it is still not obvious what sort of approach they are going to take. There are many theories in the offing, and it is not at all clear which version they are endorsing (or should endorse). In this introductory section, my aim will be to answer these three questions.

We can get the first question quickly out of the way. My discussion in this paper does not require the controversial thesis that normativity is exhausted by reasons. All it requires is that reasons are part of normativity and I think no one would doubt this claim.¹ In any case, my focus in this paper, Allan Gibbard, clearly endorses the connection between normativity and reasons. The response to the second inquiry needs a bit more spelling out, but it rests on an equally simple response: the idea of quasi-realism. Quasi-realism – advocated by philosophers such as Simon Blackburn and Gibbard – is the enterprise of quasi-objectifying non-cognitive attitudes so that they mimic strict

¹ For the debate see Broome (2000), Dancy (2004), Scanlon (2014). The main competitors are the concepts of ought and rationality and so the debate circles around the question whether we can derive one from the other.

factual beliefs. That is, the aim is to work out a non-cognitivist semantic proposal that fits well with the features of ordinary practical discourse by applying metaphysics and epistemology that are consistent with a non-cognitivist position. Concerning the latter part, quasi-realists endorse the two defining tenets of non-cognitivism: non-descriptivism (ethical² claims are not, fundamentally, attempts to describe the world; Gibbard 1990, 7–8; 2003, 7; Blackburn 1984, 167; 1998, 49) and non-factualism (there are no ethical facts that ethical claims can track; Gibbard 1990, 23; 2003, 182; Blackburn 1984, 170–171; 1998, 50). Turning now to their “realist” side, quasi-realists, unlike early non-cognitivists, refuse to draw the conclusion that led many to label non-cognitivism a nihilist, sceptical position. That is, they claim to have room for ethical truth, knowledge, objectivity and normativity in their theory (Gibbard 1990, 154–155, 218; Blackburn 1984, 171; 1993, 149–152).

However, quasi-realism operates on a level different from what we need: it is the enterprise of explaining why our ethical discourse has the shape it does. It thus sets the target, but it does not say how we can reach it; it provides the frame that needs to be filled in. Therefore, to complete the picture, we have to find the missing content – and thereby answer the third inquiry above. In this paper my candidate for this role will be *expressivism* – the view that ethical claims and thoughts *express* non-cognitive states of mind – as it is developed in Allan Gibbard’s 1990 book *Wise Choices, Apt Feelings*. My reason for this exclusive focus is simple: Gibbard’s norm-expressivism is clearly the most developed non-cognitivist account of normativity (understood in terms of reasons) at hand.³

There are of course many critical discussions of Gibbard’s theory. In line with some of them (Jackson & Pettit 1998; Smith 2001; 2002a; 2002b), I will argue that the expressivist account of the mental state expressed in normative claims is inadequate because it cannot preserve the normativity of those claims. But my rendering of this problem will be, to my mind, novel. I will examine the *structure* of the complex mental state expressed to see whether it is suitable to account for the normativity of reason-claims. But to do this, we first need to have a plausible expressivist account at hand. Let us therefore turn to Gibbard’s norm-expressivism and its account of normativity (that is, of reasons).

The norm-expressivist account of normativity

Gibbard’s norm-expressivism fits perfectly with the construction of quasi-realism cum expressivism. When we analyse the rationality of an act, Gibbard says, what we are dealing with is not what it is to *be* rational, but what it is to *judge* it to be rational (Gibbard 1990, 8).⁴ And our judgments about the rationality of actions, i.e. about what the agent ought to do, are not fundamentally attempts to describe the world: they primarily express desires, not beliefs (*ibid.*, 112–117). More precisely, they are expressions of the agent’s acceptance of norms that permit the act. Moreover, though for Gibbard the basic normative notion is “rationality”, he admits that his analysis can just as well be put in terms of reasons (*ibid.*, 160).⁵ In what follows I will do just that and will speak of “reason” instead of “rationality”.

With these remarks in mind, let us now turn to specifics. The norm-expressivist proposal is the following. The statement:

(R) *P* is a reason for *A* to ϕ

2 I use “ethical” as a placeholder term here: I intend it to be broader than “moral” and encompassing both the deontic and the evaluative.

3 Non-cognitivist authors are typically concerned with morality or with value concepts; no detailed analysis of practical reason can be found in their writings (Blackburn [1998] says more about normativity, but he clearly endorses expressivism: parallels to his theory will emerge during my discussion). Gibbard’s (2003) new version of his account could be a candidate, but I prefer to focus on his older theory for two reasons. One, there are questions about the non-factualism of the new version – in particular, whether the concept of the thing to do picks out an ethical property. See Hawthorne (2002, 169–170) and cf. Copp’s (2001) realist expressivism. Two, I agree with Silverstein (2016, 5n) that focusing on the concept of “the thing to do” makes Gibbard’s project incomplete because as an aim of practical reasoning, “the thing to do” is only a formal object in Velleman’s (2000) sense.

4 Gibbard also speaks of the rationality of feelings and emotions, but I restrict my discussion to the rationality of action.

5 Gibbard presents both concepts as equal to the concept of being the *most* reason for the agent to do something. See Gibbard (1990, 160) and (2003, 8n).

means

(NE) *A* expresses her acceptance of norms that say to treat *P* as weighing in favour of φ -ing

where NE primarily expresses *A*'s non-cognitive state, not her belief. Note further that there are *two* notions of “weighing in favour of” in play here. One is normative, and we get to it if we paraphrase R to say “*P* weighs in favour of *A*'s φ -ing”: this is what needs to be explained. The other is psychological, giving the content of the mental state expressed in NE and understood as forming norms: this is what explains (Gibbard 2003, 189–190).⁶

Of course, this raises a further question: what sort of psychology are we referring to? For an answer, we need to tackle the following questions:

1. What is a norm?
2. What is it to accept a norm?
3. Why did people come to accept norms?

Start with the first question. For Gibbard, every norm is an imperative or prescription in the general form: “Act like that!” Such an imperative “gives the rule a sophisticated observer could formulate. The imperative is a formulation of a pattern, which, in effect, controls the organism's behaviour” (Gibbard 1990, 70). Take the example of Joe. Let us suppose that there is a reason for him to climb Mount Everest because that would fulfil his desire. In Gibbard's formulation this would look something like this. The claim, “Joe has reason to climb Mount Everest” expresses the imperative “Take any means that would conduce to the fulfilment of your desires!”. By picturing norms as imperatives, Gibbard is able to avoid potential pitfalls. On the one hand, imperatives are not statements that can be believed: there is no way to give a cognitivist interpretation of imperatives. On the other hand, imperatives need not use any normative concept: hence norm-expressivism avoids circularity.

Turn now to the second and third questions. Gibbard's answer connects with his answer to the third question.⁷ It starts with a contrast. “If a norm is simply an imperative”, he says, “the real psychological question is what it is to internalize it.” And then he gives his solution: “A norm prescribes a pattern of behaviour, and to internalize a norm is to have a motivational tendency of a particular kind to act on that pattern” (1990, 70). But internalising a norm is not the same as accepting it. Animals can internalise norms, but they cannot accept them. The specialty of *accepting a norm* can be discerned in language. Language allows for shared evaluation, that is, for normative discussion in which we converse about what to do. This, Gibbard claims, is mutually fitness enhancing because it helps coordination and rehearsal. Moreover, because it has these biological functions, the discussion must influence our acting. This effect – Gibbard calls it normative governance – is that in discussing such issues we come to accept norms for the situation (ibid., 73).

In doing so, he continues, we come to avow that norm in a variety of ways: we may, for instance, simply evince an emotion or we can express an emotively charged hypothetical decision in words – in any case we express our motivation to be governed by the norm. “*Accepting a norm*”, says Gibbard, “is whatever psychic state, if any, gives rise to this syndrome of avowal of the norm and governance by it” (ibid., 75, original emphasis). Moreover, our avowal is sincere for two reasons. On the one hand, it is spontaneous: it takes place in contexts where the mechanisms of self-censorship are relaxed. On the other hand, it is responsive to demands for consistency: in avowing a norm we are always prepared to shift our position when it is shown to be inconsistent. In sum, to accept a norm “is in part to be disposed to avow it in unconstrained normative discussion, as a result of the workings of demands for consistency in the positions one takes in normative discussion” (ibid., 74).

⁶ It is thus not obviously the case that norm-expressivism helps itself to a concept of reason as Scanlon (1998, 58) and Dancy (2004, ch. 2) claim; instead, it attempts to give a genuine analysis of the concept via psychology. However, as I think they would point out, it is unclear how one can provide such analysis of a *relation* that “favouring” clearly is. I set aside this difficulty here, however, in order to be able to move on.

⁷ Or, at least, it gives part of his answer. What I do not mention in the text is the first part of the story according to which biological adaptation requires us to coordinate. This then leads to the emergence of our linguistic, emotional and normative capacities in the way told in the text. See Gibbard (1990, 61–68).

Strictly speaking, this *is* Gibbard’s account of normativity. According to him, normativity is motivational force, a form of *endorsement* our normative judgments carry (ibid., 10). And now we see how he understands endorsement: it is expressing our acceptance of norms where norms are understood as imperatives, and acceptance is seen as a motivational syndrome of avowal characteristic of normative discourse. I do not want to question this analysis, so if we also accept Gibbard’s claim that this is an account of normativity, then the rest of this paper will be about something else. It will be about what Gibbard calls “normative authority”, the objective validity we attach to our normative claims: that it is right what we are saying and therefore those who are addressed (including ourselves) should accept it (ibid., 155). It would still be important to account for this aspect of our normative life because normative authority, as Gibbard readily admits, is an essential part of the meaning of normative concepts (ibid., 153). But it will not any more be about normativity.

I, however, would like to go further than this and stick to my original claim that I am concerned with normativity. This is because I think that what Gibbard calls normative authority is in fact a *part* of what we mean by normativity. We can see this if we recall that Gibbard’s norm-expressivism is meant to give an analysis of the concept of normative reason. And when we say that some reason is normative for someone (including ourselves), we do not merely endorse a claim and guide ourselves by it in normative discussion. We also imply that one can be wrong about the reasons one has – hence that there are grounds for our claim, which one is or is not right about. (In fact, Gibbard himself calls these reasons “grounding” reasons; ibid., 162). But if we stop at the suggested reading of normativity, we end up with claims that involve no such grounds, no such criterion of correctness, but mere normative guidance understood motivationally. To get more than this, we need to be able to claim that our norms are for some reason right or correct. Hence what Gibbard calls normative authority is included in our concept of normativity and I will go on speaking about and analysing it as such.

Justification and higher-order norms

The best way to see how correctness appears in norm-expressivism is to go back to the distinction between accepting and internalising norms. When one merely internalises a norm, says Gibbard, one is “in the grip of that norm”. One is acting in accordance with the norm, though one thinks it “makes no sense” to do so; in short, one is weak-willed (ibid., 60). Take Gibbard’s example of the subjects of Milgram’s famous experiment who were told to administer painful electric shocks to another subject. These subjects, Gibbard tells us, experienced a conflict of norms: the norm of non-harm on the one hand, and the norms of following directions, cooperating, on the other. Nevertheless, the subjects were clear on which norms they accept (former) and which norms they merely internalise (latter) (ibid., 58–59). To put it another way, the subjects thought they had no reason to harm others even if that would require them not to follow orders. They were thus rejecting these “peculiarly social” norms as, at least in the given situation, *unjustified*. Since, however, they internalised them, they were moved, in fact, overwhelmingly moved, to act on them – but they saw their motivation as alien and thus lacking any normative force. We thus get the following result: when people accept a norm, they endorse it as justified; when they merely internalise a norm, they reject it as unjustified. In the first case, their claim expressing their acceptance of a norm has normative force, in the second case, it does not.

This analysis can be extended to the perspective of the observer who criticises another agent. Take Gibbard’s example of the girl who thinks she has reason to starve herself to death for the sake of a slim figure. Let us also suppose that there are norms in operation so that she is not merely overwhelmed by a powerful craving (ibid., 165). The observer, however, thinks the opposite: he says that the girl has no reason to starve herself to death for a slim figure even if she has norms that tell her to do so. What does he do? According to Gibbard, he can make three kinds of criticism: contextual, Socratic and fundamental (ibid., 175). It is very probable that in this case he does neither of the first two: both would require that he accept the girl’s norms, which he evidently does

not.⁸ What he does, instead, is to claim “her premises are crazy”, that is, he “claims to be ‘seeing’ something that she does not: that the norms she accepts just do not make sense”. He is claiming, Gibbard says, “a kind of fundamental authority, an authority that does not stem from any common acceptance of more fundamental norms. [He is] claiming that her norms are crazy, so that [his] own are, in a fundamental way, more to be trusted than hers” (ibid.). In other words, he claims that her reason has no normative force because the norm it is based on is *unjustified*.

In both cases then, normativity accrues to reason-claims only if the norms they refer to are justified. The challenge for norm-expressivism, therefore, is to find place for the justification of norms. Gibbard’s solution is to appeal to other norms that can justify these norms. We thus arrive at the territory of *higher-order norms*. According to Gibbard, higher-order norms have two characteristics. First, they tell everyone to act according to some lower-order norm (ibid., 171). Second, they leave no slack: they leave no alternatives open so that if one does not follow the norm, one cannot remain rational (ibid., 169).⁹ The picture that emerges in the end is similar to what Blackburn calls the “the staircase of practical and emotional ascent” (Blackburn 1998, 9; cf. Gibbard 1990, 169n). It is a staircase where the first steps are made of our ordinary preferences, likes, and dislikes; then follow our first-order norms that govern these desires; and then come our higher-order norms that require everyone to share our norms and the simple desires they govern. Again, it seems that we have reached the point where we can stop and declare, now with respect to my broader understanding of normativity, that we reconstructed the norm-expressivist account of normativity.

But we cannot stop. For, no matter how compelling this picture is, it does not solve our original problem. This can be seen from both perspectives. Take Milgram’s subjects. What we now further suppose is that they have higher-order norms that tell them to accept those first-order norms they find alien. Perhaps they were so educated to accept everything society wants from them, including administering electric shocks to others. But they may well claim that such kind of social education is no more than mere indoctrination and they want no part of it. In other words, they may say that the concerned higher-order norms are just as well unjustified and alien to them as the first order-norms whose acceptance they demand. Similarly, the observer need not accept the higher-order norms, which, we may suppose, the anorexic girl accepts to back up her position. He can find them outrageous and demand their elimination. But then, evidently, we are back where we started: we need to ground higher-order norms. How shall we do it?

Justifying higher-order norms

One might think that this question leaves Gibbard without an answer. But it does not. He approaches the problem from afar, but most of what he discusses (such as “conversational demands”, or a lengthy explanation why the hearer should listen to the speaker’s criticism) say nothing about the justification of higher-order norms. For his real response, we have to turn to the third-person case. “I as a speaker”, says Gibbard, “do not simply demand; I claim to have a basis of my demands... Speakers distinguish browbeating demands from reasonable demands, and in normal conversation a speaker confines himself to demands he thinks reasonable” (ibid., 173). How does the speaker do this? The idea must be similar to what we have already seen: it must be that the speaker has a story, told in higher-order norms, which tells why his norms are “better” than those of the hearer. But his story cannot involve just *any* higher-order norms because that would lead us to the problem we started out with. Hence emerges Gibbard’s answer: the norms on which the speaker bases his claim

8 Contextual authority concerns cases when the audience accepts claims made by the speaker just because the speaker accepts it. The reason for this is that both the speaker and the audience share the same basic norms. Socratic authority is similar in the sense that the ground for the claim is based on premises the agent accepts. The difference is that in this case hearers accept what the speaker says not because the speaker accepts it, but on the basis of claims they were prone to accept if they properly reasoned from their premises. See Gibbard (1990, 174–175). Both these constructions can be interpreted as demands for justification, but fundamental criticism takes us to the very heart of the issue; hence my choice to use Gibbard’s example in the text.

9 Here the comparison is with what Gibbard calls existential commitments. The difference between them and norms of rationality, Gibbard says, lies in the higher-order norms that govern their acceptance. In particular, higher-order norms governing the acceptance of existential commitments leave some slack: they do not require, only permit these commitments, leaving at least one incompatible alternative open to the agent. See Gibbard (1990, 169).

in the case of fundamental authority are norms that *govern the acceptance* of other norms, both lower- and higher-order kind. He suggests that

[w]e are trying to decide what norms to accept as governing our thought. We are trying to decide what norms to accept for letting one's judgment depend on what others think. We are thus trying to settle on norms for the acceptance of norms (ibid., 177).

Gibbard groups these higher-order norms in two sets. In one set, we find epistemic norms or norms of warrant, in the other, we find norms of rationale. *Epistemic* norms, as the attribute suggests, are norms for the recognition of lower-order norms. They determine generic qualifications for ideal normative judgment by warranting the acceptance of a lower-order norm (ibid., 181). And the way they do this is to circumscribe a set of potentially competent normative judges by listing content-fixed (depending on who makes the judgment) and content-neutral (not depending on who makes the judgment) properties of them (ibid., 195–197, 208). Examples for epistemic norms include the norm of coherence, but also norms that go beyond sheer coherence. A typical and widely accepted criterion is given by the higher-order norm of reflective equilibrium, which tells agents to accept norms that would be arrived at after vividly considering all relevant facts and philosophical arguments. Another higher-order norm can require people to follow norms that would be accepted by them in a cool hour of thinking when they are not influenced by turbulent affections and the like (ibid., 168–170).

In contrast, norms of *rationale* are non-epistemic in character because their aim is not to tell how to recognise a good judgment and thus determine who counts as a good judge. Their primary objective is to tell what norms to accept: they aim to give a rationale for accepting a certain set of first-order norms (ibid., 213). If a higher-order norm, for instance, tells agents to accept only those lower-order norms whose acceptance enhances the sense of meaning in life, it qualifies as a norm of rationale. It says nothing about how to recognise those norms; instead, it gives directions as to what gives the norms their point (ibid., 218). Epistemic norms, on the contrary, are silent regarding the rationale for ordinary norms or, rather, they are neutral on the issue. Anything, for instance, that comes out of reflective equilibrium counts as warranted regardless of its contribution to the sense of meaning in one's life.

These two types of norms provide the agent with the possibility of justifying his other norms, both lower- and higher-order. Of course, he may also find, based on these norms, that certain of his norms lack justification. This is what happens to Milgram's subjects. In the given situation, they think that the higher-order norms that tell them to always follow certain "peculiarly social" norms are unjustified because they were acquired through the wrong epistemic means, namely through, what they take to be, social indoctrination. And, similarly, when the observer criticises the anorexic girl, he refers to the epistemic norms and norms of rationale that warrant and give the point of the lower-order norms he, the speaker, accepts and thus justifies his normative judgment. He gives what Gibbard calls his "epistemic story" outlining openly the grounds – that is, the higher-order norms – for his claim. Such a story specifies the conditions that the speaker's judgment has to fulfil, that is, it tells the story of why the speaker is a competent judge and why the girl is not.¹⁰

The CUC set of norms

One step is still missing. Gibbard notes that people can choose among communities of judgment. They can decide whom to include and whom to exclude and thus arrive at communities of judgment, which overlap with other communities within a college of loose public groups (ibid., 204, 234). *Within* the community, members have a set of epistemic norms as well as norms of rationale that determine what norms one can refer to as authoritative in normative discussion. But then the question naturally presents itself: what happens in the normative discussion that takes place *between* groups, between communities of judgment? The problem that first strikes the eye is that when it comes to fundamental criticism, the epistemic stories of different groups may on certain restricted topics be

10 This is why Michael Smith's (2001, 107–114; 2002b) argument against non-cognitivism is in the end unsuccessful. As Blackburn (2002, 175–176) in his response to Smith points out, the alleged belief about coherence will not be a "belief" on non-cognitivism exactly because coherence is a substantive standard (a higher-order norm, to invoke Gibbard's terminology).

so deeply contradictory that the chance of reaching a decision reduces to zero (ibid., 198). Though convergence on norms is possible, the attitude of groups toward each other will not be so favourable. Gibbard draws up two possible scenarios, both relying on some version of standpoint-dependence.

First, communities may say nothing about why they do not accept the judgments of other communities, while maintaining epistemic stories for their members. Gibbard calls this *parochialism* (ibid., 208). What parochial communities claim is that though there could be non-members who meet their requirements as to what an ideal normative judgment consists in, they still do not qualify as potentially normative judges just because they do not belong to the community: simply because of who they are (ibid., 216). Second, communities can also have a story about why they do not regard non-members as competent judges even if they meet their epistemic requirements – other than the mere fact that they do not belong to the community. They can say that norms depend on the features of the specific community and thus differ from community to community (ibid., 210). They can base this claim on some higher-order norm of rationale that tells them how to treat the ethos of one's community, or how to treat one's commitment to the community and so on. This is Gibbard's version of *relativism* or, more precisely, partial relativism, because though the higher-order norms invoked have relativistic effects, their status is non-relative (ibid., 210, 214).

Undoubtedly, parochialism and relativism are deep problems for norm-expressivism (and non-cognitivism, in general). Both make it impossible to construct an account of truth and of normativity. While *within* a group there are things that one has reason to do being backed by a hierarchy of lower- and higher-order norms the group accepts, *between* groups one cannot claim the same given the lack of a shared system of norms. However, here I will set aside Gibbard's long and complex treatment of the problem because I am more interested in its upshot than how he gets there. Let me just register that his argument follows Blackburn's idea of transcending the "tree structure" of sensibilities and Rawls's political conception of justice (Blackburn 1984, 198–210; Rawls 1993, Lecture I, IV). The idea is that relativism and parochialism can be overcome by transcending disagreement between communities either by considering the costs of isolation and repression or by fostering mutual respect through toleration. Gibbard proposes higher-order norms to do the job: norms of accommodation to cater for the former, and norms of toleration to accomplish the latter (Gibbard 1990, ch. 12–13).

The result that emerges completes the norm-expressivist "staircase of emotional and practical assent". First, we now know all the steps. They are: ordinary non-cognitive states of mind (e.g. ordinary desires and preferences); lower-order norms that govern these preferences; higher-order norms that tell everyone to accept the lower-order norms; norms of accommodation and toleration that help transcend relativism and parochialism; and norms of rationale and warrant that govern the acceptance of all other norms. Second, with some ingenuity we can also say something about how the steps are related to each other. Namely, we can define a set of norms that would result from taking all possible opportunities for improvement through norms of coherence, warrant, rationale, accommodation and toleration. This set will be *controlled* due to the operation of epistemic norms and norms of rationale that include and exclude norms in and from the set. Further, as a result of transcending relativism and parochialism, hence differences among groups, the set will also be *unique*. Finally, the set will be *comprehensive* because, in construing it, all available norms are considered (including those the agent could be aware of with some effort or will). The result is a controlled, unique, and comprehensive set of norms: a *CUC set of norms*. Saying that a claim is normative is to say that the norm whose acceptance it expresses is a member of this set. This is the norm-expressivist theory of normativity.¹¹

The regress problem

Let us retrace our steps. We began with two examples that both required some account of the justification of the norms whose acceptance our reason-claims express. Take first Milgram's subjects. The explanation of what is going on there is this. Standing on the ground of a CUC set of norms, they declare that neither the lower-order, "peculiarly social" norms, nor the higher-order

11 The original idea comes from Blackburn (1984, 197–198, 233–248). But he calls it the CCC set of norms where the second C stands for "coherent". I, however, list coherence among the ways to control the set, and emphasise instead the uniqueness of the set.

norms that demand their acceptance are justified: they are not members of the set. As a result, the claims that express those norms are non-normative, and if they act on them, that is because they are in the grip of those norms and not because they accept them. Similarly, the observer who criticises the anorexic girl refers to a CUC set of norms claiming that the girl's norms, be they lower or higher-order, are not members of the set. Hence her norms lack justification, and, consequently, her claim lacks normative force.

In what follows, I will question the adequacy of this construction. I will pursue a simple idea: *the regress problem*. We saw that norm-expressivism must have an account of the justification of norms in order to preserve the normativity of reason-claims. But then an Achilles and the Tortoise-type of regress appears to threaten. Recall how we have proceeded so far. We started with simple, lower-order norms; then invoked higher-order norms that demand their acceptance; after this, we turned to norms of rationale, warrant, accommodation, and tolerance that govern the acceptance of these norms and thus justify them. But what justifies *these* norms? Gibbard, for instance, throughout his discussion, gives a crucial role to norms of coherence. But one can ask why one should be coherent. As Blackburn (1998, 117; cf. Gibbard 1990, 156–157) at one point reminds us: sometimes it is virtuous to be drawn in two directions, and sometimes we do not take the advice of the ideally coherent. But if one questions coherence, where can he turn to for justification? This is the question – and without an answer, the observer cannot criticise the anorexic girl, nor can Milgram's subjects explain why they are in the grip of “peculiarly social” norms.¹²

There are three general ways to respond to a regress problem. To claim that the regress does not start; to admit that it starts, but show that it can be stopped; or to accept both that it starts and that it is infinite, but argue that it is not vicious. I will proceed in the following order. I begin with the first and third reactions and then end with a discussion of the second strategy (which requires the investigation of techniques adopted from epistemology: coherentism and foundationalism, in particular).

The first and third response

To be successful, norm-expressivist advocates of the first response must show that justification is “contagious” in a particular way. We can distinguish two categories of contagion for justification (Radzik 1999, 41). Certain kinds of justification can be conferred only by things that are themselves in possession of that sort of justification. For instance, many epistemologists think that only those beliefs can justify other beliefs that are themselves justified. But other kinds of justification are not like that: in their case, the justifier need not itself be justified. Justification here is understood like a disease of which the justifier is a “mere carrier”: it can pass justification onto other “organisms” without itself having it (or at least needing it). The challenge the norm-expressivist faces is then twofold: he must show that this kind of justification is unproblematic and that he can appeal to it. I think there is sufficient ground to doubt both claims.

As to the first, the analogy with the “mere carrier of disease” version seems to break down when it comes to issues of normativity. For it is difficult to see how something that is not itself justified can confer justification on other things when what is at stake is the normativity of that thing. There appears to be a conceptual tension here, the resolution of which at least requires some explanation, something the norm-expressivist does not give. But even if we set aside this problem, we still stumble into the second challenge. As we saw, in giving an account of normativity, norm-expressivism appeals to norms and this suggests that it is norms that play the role of justifier in the theory. But norms stand in need of justification, this, if my account in the previous section was right, is encoded into our concept of accepting a norm; hence the regress gets underway. The only alternative would be to claim that norms are justified by the facts, natural or non-natural, which they track in the world. But even if we set aside the controversial issue whether these facts themselves are in need of justification, we still cannot come over the non-factualism of norm-expressivism. As

12 To my knowledge, this problem has not received detailed discussion in the literature. Campbell (1993, 314) mentions this difficulty in his review of Gibbard (1990), but does not dwell on it further. Railton (1992, 963n) appears also to note a similar objection in a footnote of his review of Gibbard but says nothing further.

a result, the norm-expressivist must appeal to factors other than facts and these can only be norms – and then the regress immediately gets off the ground.

To some, however, this brief rebuttal may sound overly hasty. For we have another approach that appears to go along the lines of the first response: Peter Railton's idea of default trust. Railton's (2004, 185–186) proposal is that the regress loses ground because most of our rational, intentional activity is unmediated by recognition or judgment that would go together with demands for justification. What goes on there instead, he suggests, is that agents express their default trust, their blind, unquestioning obedience in the authority of their memory, senses, thoughts and the like (ibid., 190, 191). There are two ways to understand this claim. One focuses on Railton's additional claim that memory, senses etc. have only defeasible authority, and thus interprets his proposal as a version of foundationalism. I will deal with this idea later. The other interpretation, I think, stands closer to Railton's intentions and fits the present context well: here the idea is simply that memory, senses etc. do not stand in need of justification. Applied to the context of norm-expressivism, Railton's claim would then be that at least some of the agent's norms are endowed with default authority: just as we do not question our senses, memory or thoughts, we do not demand justification for these norms either.¹³

There can be serious questions asked about the plausibility of Railton's proposal (Tanyi 2006), but here it is enough to point out its inconsistency with norm-expressivism. Turn first to the rest of Railton's argument. Driven by the concern that something lacking justification cannot itself confer justification, Railton (ibid., 191–196) aims to show that belief and desire are not only attitudes of default trust, but also are normative attitudes whose function is to track truth and value respectively. But while norm-expressivism can accommodate the first part of the claim concerning belief, it is clearly inconsistent with the second part concerning desire: its non-factualism forces it to deny any such claim. Second, Gibbard's account of self-trust just does not fall in line with Railton's reading. On Gibbard's (1990, 179) view, when one tries to decide whether to give authority to one's past, present, or hypothetical judgments, one is deciding about accepting a higher-order norm that accords authority to one's own normative judgments. Consequently, the authority of these judgments is not default, but acquired through their relation to higher-order norms. And though Gibbard exempts one's present fundamental judgments from this account, he does so because he thinks that one does not accept these judgments on anyone's authority including one's own (ibid. 177–178, 227).¹⁴ Until these inconsistencies are overcome, Railton's argument cannot even be taken seriously as an attempt to save the first response.

Norm-expressivism must look for other ways to deal with the regress. One such attempt is given by the third response: to accept that there is a regress as well as that it is infinite, but then show that the regress is not vicious. There are two ways to back up this claim. The idea that does the work in the first case is that what counts is not the *process* of justification, i.e. the agent's performing the steps of justification, but the *status* of being justified, i.e. the claim's having a place in an infinite chain of justification. Like in geometry: perhaps there is one theorem of geometry only if there is an infinity of theorems, but this is not to say that there is one theorem only if an infinity of theorems has been *proven* to be such (Copp 1995, 43). This response clearly builds on a particular understanding of practical justification, what we might call *third-person* justification. It claims that the agent need not have access to the full justification of his norms; he need not be able to comprehend it.

However, third-person justification is not the only alternative; in fact, there are reasons to favour another, *first-person* justification. Briefly, these are the following. First, questions of practical justification are primarily questions the agent asks about himself: *he* wants to know whether this or

¹³ For other passages that confirm the second reading, see Railton (2004, 177, 185–187).

¹⁴ Gibbard says that our present fundamental judgments express our acceptance of norms that we find self-evident in the sense that we cannot doubt them. But, he continues, our accepting them in this way is not the same as our trusting them. For the fact that we find something self-evident is not our ground for accepting it; instead, our finding it so *consists in* accepting it in this way. See Gibbard (1990, 177–178). Railton, on the other hand, takes self-evidence to be an indication of self-trust. He says: "Even if I say 'Some things are simply *self-evident* in the light of reason', I now see this is whistling in the dark: 'self-evidence in the light of reason' is just default trust in what seems evident to me, garnished with parsley. If I cannot even imagine an alternative to *p*, then to say that *p* is 'impossible' is just to trust my powers of imagination and command of *p*'s content" (2004, 187). It may be that these positions can in the end be reconciled, but that certainly needs further philosophical work. And even when this happens, the other discrepancies described in the text will still remain.

that norm is justified. Hence the answer to his question had better be accessible to him, had better be something he can grasp and appreciate the import of (Korsgaard 1996, 16–17). Second, attributions of justifications try to pick out agents as acting conscientiously, i.e. in a responsible, blameless manner. Therefore, agents should at least be given the opportunity to detect the justification of their norms. Perhaps they will not always seize this opportunity, but if they did they could find the missing justification (Radzik 1999, 39). Finally, third, normativity is a regulative notion. By “regulative” it is meant that the justification and thus the normativity of reason-claims is something people can usefully think about in order to decide what to do. If thinking about normativity is supposed to be of any use to agents when they are trying to make choices, they had better be able to tell what is normative and what is not (ibid.).

The second way of showing that the regress is not vicious is by making use of an idea pursued by indirect consequentialism. It is the claim that utilitarian agents need not be aware of their true moral motives, but can act on their personal motives. Similarly, by extension to the sphere of justification, a norm-expressivist might claim that people should not be aware of the norm-expressivist account of reasons (hence the regress), but can pursue their own, probably cognitivist theory (without regress). However, by analogy, the same objections that beset indirect consequentialism are also effective against the norm-expressivist response. The problem is well articulated both by John Rawls’ publicity requirement and Bernard Williams’ transparency condition (Rawls 1971, 130; Williams 1973, 128; 1985, 101–102). They emphasise the counter-intuitive nature of the claim. To paraphrase, most of us would think that a norm-expressivist theory that requires widespread ignorance of its account of reasons is self-defeating and thus unacceptable. For its distinctive contribution to ethical theory is exactly this account (remember quasi-realism and its aspirations!). Therefore, when norm-expressivism does not require people to be aware of the doctrine, it abandons the basis on which its own foundation, as a distinctive ethical theory, is laid.¹⁵

The second response: foundationalism

I thus conclude that the regress starts, and if nothing comes in its way, it will be viciously infinite. But is there really nothing that may come in its way? To see this, we need to examine the second strategy: perhaps there is a way to meet the regress head on. Epistemology provides us with two ideas that we can apply to the present case: foundationalism and coherentism. In what follows I take them up in turn.

Foundationalism in the sphere of norm-expressivism claims that a CUC set of norms contains foundational elements. At one place, Gibbard appears to suggest that agents have “very highest-order norms, norms governed by no further norms”. At another place, he speaks of the observer’s “seeing” norms as “self-evidently valid”, as “having independent credibility” (Gibbard 1990, 170, 176–177). And in his subsequent exposition he says that to accept a norm as self-evidently valid is “to accept it without being prompted, even on challenge, to doubt it, or to appeal to something else in its support” (ibid., 177). This suggests the following picture of a foundational norm: it is a norm that (i) is capable of justifying other norms, (ii) is not itself justified on the basis of another norm, but is freestanding, and (iii) is itself justified. And the challenge to the norm-expressivist is to give an account of (ii)–(iii), to construct what epistemologists call a theory of privilege for norms.

Gibbard says nothing more on the issue; hence we need to act on our own. There are six major attempts in the offing. The most straightforward option for norm-expressivists is to claim that certain norms are insusceptible to rational or reasonable doubt. They are like Descartes’ *cogito*: just as one cannot ask the question “Why should I believe that I exist?” because the evidence for it is present in the very thinking of the thought, one cannot question the justification of certain norms. The problem with this approach, however, exists on its face. To say that it is unreasonable

15 There is no doubt more to be said here. I find discussions surrounding moral error theory instructive in this regard. There is a crucial question in that debate, too: what happens if indeed moral error theory is true? What do we do then? Should we believe the theory and discard our moral talk (abolitionism) or should we consider morality as a mere fiction (fictionalism), or perhaps try to conserve traditional moral talk despite endorsing error theory (conservationism)? It seems to me that in the present scenario norm-expressivism has the same three options and they are equally problematic. For a good discussion (that, however, attempts to defend conservationism), see Olson (2014).

or irrational to doubt the validity of these norms is to say that one *ought not* or has no *reason* to do so. But, on the one hand, such a move presupposes what it aims to explain, namely the notion of reason, rationality and normativity. And, on the other hand, since on norm-expressivism reason-claims express the acceptance of norms, it makes reference to further norms to back the allegedly foundational norms. In either case, the regress continues (cf. Radzik 2000, 638).

Perhaps norm-expressivists do better if they shift attention to the practical side. For instance, they can claim that there are certain norms whose *motivational efficacy* is universal: any agent who is aware of the norm is also moved to apply it in action. Hence, from a strictly practical point of view, the norm is unquestionable and thus foundational. This approach can be interpreted in two ways, but neither delivers the goods. Turn to the first version. The claim here is that certain norms never run out of motivational force: perhaps the agent is not moved decisively, but *some* motivation is always produced. On norm-expressivism, however, there is a difference between merely internalising a norm and accepting it as justified. Both phenomena are motivational, but only acceptance carries normativity; hence, in order to know whether a motivational effect of a norm is normative, we first have to have an account of justification at hand and not the reverse. And to this it does not do to reply that the agent need not be aware of the justification of the norm that invariably moves him; for then the present approach becomes indirect with all its problems with transparency. The second version restricts the circle of norms depicted on the first interpretation: norms must have *inevitable* motivational force such that agents cannot but act on them. But these “practical necessities” are not normative. They allow no room for failure, and a dictate that one cannot fail to do cannot be normative (Williams 1995, 53; Korsgaard 1997, 229). Hence the “practicality” attempt fails to make its point.

The next two approaches appeal also to ideas familiar from discussions on other subjects. According to the first, a norm may be foundational in the sense that it is *unintelligible* to question it: it is foundational by virtue of meaning. However, this raises the worry that such a norm would be too thin or insubstantial to serve as the basis of normativity. The most likely candidate is a norm that is *formal* in David Velleman’s sense. On Velleman’s (2000, 176) account, “winning”, for instance, is the formal object of a competitive game; and the “quarry” is the formal object of hunting. By analogy, the object of reasoning would be set by the foundational norm in the same way. But, as Velleman (*ibid.*, 175) points out, every enterprise that has a formal object must have a substantive object as well. It must do so because otherwise the agent would have no criterion to recognise what is a reason and what is not. It would be like telling him to hunt for something described only as the “quarry”, or to play a game with the end only described as “winning”. In short, the enterprise would have no object at all and thus would cease to be normative for the agent.

The fourth attempt to provide us with a plausible foundational norm directs our attention to the *cost* of scepticism. Take first the theory of Christine Korsgaard. Within the confines of her theory, she argues that the justificatory regress can only be brought to a satisfactory halt if one values oneself as a human. In other words, the ultimate justification comes from a foundational norm, namely from the moral law (Korsgaard 1996, 121–123). And in response to the question “Why value my own humanity?”, she points to the *price* of such rejection. It is “complete normative scepticism”, she says: one who disvalues one’s humanity cannot act rationally, cannot have reasons for action. He may just as well commit suicide, but we cannot say even this much for there is, strictly speaking, nothing the sceptic should or must do (*ibid.*, 123–124, 129, 163). Similarly, Gibbard, at the only place where he seems to be concerned with the regress, argues that someone who rejects the norm of consistency, opts out of normative discussion altogether and becomes some sort of normative solipsist (Gibbard 1990, 289–290).

The general problem with this approach is that there is a difference between scepticism *itself* and the *non-sceptical* position, which tries to elicit a better answer by employing sceptical arguments. In other words, since there can be an alternative to both Gibbard’s and Korsgaard’s argument, therefore one who denies the conclusion they draw need not also endorse the consequence they point to. That is, one need not endorse complete normative scepticism, or opt out of normative discussion altogether. He might just be like me, looking for a better theory with entirely non-sceptical intent

while writing this paper. The price is not complete normative scepticism or normative solipsism; it is merely the rejection of Korsgaard's theory or Gibbard's norm-expressivism.

The fifth attempt emphasises the *self-justificatory* character of the foundational norm. It is such that it meets its own criterion, hence, when the demand for justification arises, it will justify itself. For instance, the instrumental norm "Take the necessary means to your end, or else give up the end!" is self-supporting because accepting it as a guide to one's actions will be part of taking the necessary means to one's ends. Similarly, the higher-order norm "Accept norms that direct you to promote or preserve what is, in fact, valuable!" is a norm that would itself direct one to what is valuable, hence it is self-justifying. The problem with this approach is that it leads to unwelcome results; it would lead us to accept norms as foundational, which are unimportant, false, or absurd (Copp 1995, 41; Radzik 2000, 643). Take Copp's example. The higher-order norm "Accept standards that can be expressed in English in fewer than twenty words!" is clearly self-supporting: in expressing, I used less than twenty words. Nevertheless, no one would claim that this norm could serve as the foundation for a CUC system of norms. As a result, the most we can say about this attempt is that it gives a necessary condition for the justification of foundational norms, but does not provide the whole of it. For that we need to look elsewhere.

We have one attempt to deal with. It points out that the responses we have considered so far all supposed that, in addition to requirements (i)–(iii),¹⁶ the foundational norm must also be incorrigible. That is, it must be unquestionable in some sense, and thus be undefeated by any other norms that may enter the normative picture. But there is an alternative: *defeasible justification*. The idea comes from epistemology where certain object-level perceptual beliefs, like "I see a brown table in front of me", are justified without appeal to other beliefs of the agent, but are justified only if the agent has no reason for thinking that his eyes are deceiving him (as noted, this is one reading of Railton's proposal as well). Accordingly, defeasible justification in the sphere of norm-expressivism claims that the foundational norm the norm-expressivist is looking for is freestanding and justified in itself, but adds that the justification it possesses is only *prima facie*: it is "innocent until proven guilty". This would put an end to the regress because we would always grant initial justification to the norm satisfying criteria (i)–(iii), even though we know that questions regarding its validity may arise. But until this happens, we take it that justification simply runs out at this point.

This response falters on the following obstacle (cf. Radzik 2000, 646–647). We saw that deflationism runs counter to people's need to demand justification of their norms. They need an answer why to stop looking for justification where they are supposed to stop, and to this the approach under consideration replies that this is because certain norms have at least initial justification. But then one can further inquire what makes these norms possess even *prima facie* support. Say all norms are justified that are actually accepted by the agent; the question is why would such an act even initially justify those norms. Agents accept many strange and absurd norms, so their actual acceptance can hardly serve that role. Of course, the norm-expressivist then can come up with another, more plausible sounding candidate, but, it seems, all candidates can be questioned. To defend his position, the norm-expressivist will in the end be forced to appeal to a general *norm*, something like "Accept norms that possess a certain property in virtue of which they are defeasibly justified!". And then he will have to justify this norm and thus continue the regress.

The second response: coherentism

If a CUC set of norms cannot include foundational elements, then perhaps members of the set draw support from other ordinary members of the set. In other words, some form of *coherentism* may be true: norms are justified because they are members of a coherent and mutually supportive set. Gibbard (1990, 176n) himself suggests that coherentism can help us understand the idea of fundamental criticism, and thus the regress problem. We also saw how important a role coherence plays in the Gibbardian picture of normative discussion. Finally, the coherentist approach comes very close to the construction of a CUC set of norms. Here, recall, the first C stands for "controlled", and one way to control the set is through coherence. Coherentism thus appears to be a natural

16 These are, to repeat, that a foundational norm is one that (i) is capable of justifying other norms, (ii) is not itself justified on the basis of another norm, but is freestanding, and (iii) is itself justified.

solution to the regress problem, one that does not import controversial views about foundational norms.

However, the fact that coherence is one means by which the set is controlled points to the deepest problem with the coherentist approach. It shows that the requirement of coherence is just another epistemic norm; it is just another way to govern the acceptance of lower-order norms. Hence it is in need of justification, and since it is not taken to be foundational, justification can only come from other norms: the regress continues. This conclusion does not change even if we further sophisticate the idea. One may add, for instance, that when testing a norm's coherence with other norms, the conditions in which the test is carried out also count. One typical condition is that the agent should be in a normal frame of mind, uninfluenced by turbulent affections when testing coherence. But this move only introduces further epistemic norms, such as the epistemic norm of "cool hour", and thus cannot stop the regress. In fact, what we get in the end is a set of epistemic norms (perhaps also including norms of rationale) framed by the general norm of coherence. But what we do not get, to repeat, is what justifies these norms; hence we are back where we started.

Recently, however, Linda Radzik put forward a theory, she calls it Reflective Endorsement Coherentism (REC), which claims to overcome this obstacle. Call the above norms "norms of reflection". Then her proposed solution takes the following form (Radzik 2002, 34). The higher-order norm, RE_1 , "Accept norms that you would endorse upon reflection!" needs justification because the agent requires explanation why he should trust his norms of reflection including, most prominently, the norm of coherence. Radzik's reaction is to point out that there is a norm that does just that. This norm, RE_2 , says "Accept the norms of reflection that you do!". Of course, RE_2 stands in need of justification as well, but Radzik thinks she can provide one. The justification for RE_2 , she points out, comes from RE_1 : one ought to accept the norms of reflection that one does because when one reflects on those norms, one endorses them, and one ought to accept norms that one endorses upon reflection. She calls this construction the *loop of reflective endorsement*, which, she claims, ties up her coherentist theory of justification.

Technically speaking, Radzik's loop solves the problem and does so in good coherentist fashion: lower-order norms are justified by RE_1 , RE_1 is justified by RE_2 , and RE_2 is justified by RE_1 . But the loop has its bad side too. The message it conveys, as Radzik readily admits, is that "reflective endorsement is like a tribunal that passes judgment on its own procedures" (ibid., 35). Norms of reflection, including the norm of coherence, are endorsed on the basis of other norms and those other norms must be ones the agent already has: his CUC set of norms. But then Radzik's solution turns RE_1 and RE_2 , the two anchors of REC, into a single *self-supporting* norm. She claims that when the agent reflects upon his norms of reflection, he endorses them on the basis they themselves provide. It is thus, properly speaking, not RE_2 that justifies RE_1 , but RE_1 that justifies itself (ibid., 36).

Yet, it is difficult to see why a self-justifying norm would do better now than it did before. Consider the situation. The agent does not trust his norms of reflection but Radzik tells him to endorse them because they justify themselves. Why would he? This is mere question begging for him. Since he does not trust his norms of reflection, he does not trust the norms they justify, including the norms of reflection themselves. Such an act of justification is nothing else but self-gratification, and this is not what he is looking for. To get around the problem, Radzik (ibid., 36–37) further claims that the self-justification of RE_1 is not automatic because the agent refers to the rest of his CUC set, not only norms of reflection. But lower-order norms get into the set exactly by norms of reflection, so if there is a conflict, then this must be because two higher-order norms conflict. And such a conflict cannot be decided on the basis of lower-order norms. The agent needs to invoke further higher-order norms and thereby continue the regress. So the choice is: self-gratification, or regress. In either case, REC fails as a theory of justification.¹⁷

17 Somewhat out of line with the rest of her position on justification, Radzik appears to appeal to the cost of scepticism as a last resort claim. She says, "an agent for whom RE_1 is not self-supporting [i.e. who does not accept the loop of reflective endorsement] would hardly count as an agent. An agent, if she is to be an agent, must be able to trust her own practices of reflective endorsement in general, even if she recognizes that she is untrustworthy in certain contexts" (Radzik 2002, 37). But, again, while this is certainly so, Radzik's attempt to secure such trust is not the only one. Hence the "sceptic", who rejects it, need not be a sceptic about agency but only about Radzik's defence of it.

I thus conclude that norm-expressivism cannot stop the regress. There is no way it can find a defensible foundational norm among the members of a CUC set, nor can it appeal to coherentism for justification. And since the regress cannot in some other way be avoided, the norm-expressivist theory of normativity is doomed. Milgram's subjects cannot in the end make a difference between accepting a norm and merely internalising it; nor can the observer criticise the anorexic girl with normative force. It is difficult to say if this conclusion can be generalised to non-cognitivism per se, but it certainly effects many versions of it. Most non-cognitivist theories would include something like an "emotional and practical staircase", and this is exactly where the regress begins to unfold.

Summary

In this paper, I have argued against an influential line of thought about normativity (which I understood, at least in part, in terms of reasons): non-cognitivism. In my interpretation, the non-cognitivist account of normativity emerges from a quasi-realist cum expressivist framework with normativity understood along Allan Gibbard's norm-expressivism. Discussion has showed that norm-expressivism needs a well-circumscribed account of the justification of norms in order to have a defensible theory of normativity. But my investigation has proved that it does not have one: a potential regress lies in the justification of norms that norm-expressivism has no resources to overcome.

Acknowledgments

This paper started life as a chapter of my doctoral dissertation back in the mid-2000s. My work in this period was strongly helped by the comments of János Kis and Krister Bykvist as well as by audiences around Europe too numerous to mention (and because I fear that my memory would not serve me well after all these years). I would like to thank all of them for help with my thinking about non-cognitivism – even though in the end I decided against including the chapter in the submitted version of my dissertation. Instead, I returned to the argument presented in the chapter only recently and decided to publish it as a separate paper. Hence the present article is the result of the combination of my work years ago and my more recent thinking about these matters. My work on the paper was supported by a EURIAS Junior Research Fellowship at the Collegium Helveticum, Zurich.

References

- Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.
- Blackburn, S. 1993. "Errors and the phenomenology of value." *Essays in Quasi-Realism*, 149–166. Oxford: Oxford University Press.
- Blackburn, S. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon Press.
- Blackburn, S. 2002. "Replies." *Philosophy and Phenomenological Research* 65(1): 164–176. doi:10.1111/j.1933-1592.2002.tb00192.x.
- Broome, J. 2000. "Normative requirements." In: *Normativity*, edited by J. Dancy, 78–99. London: Blackwell.
- Campbell, R. 1993. "Critical notice of Allan Gibbard's *Wise Choices, Apt Feelings*." *Canadian Journal of Philosophy* 23(2): 299–323. doi:10.1080/00455091.1993.10717323.
- Copp, D. 1995. *Morality, Normativity, and Society*. New York: Oxford University Press.
- Copp, D. 2001. "Realist expressivism: A neglected option for moral realism." *Social Philosophy & Policy* 18(2): 1–43. doi:10.1017/S0265052500002880.
- Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press. doi:10.1093/0199270023.001.0001.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Massachusetts: Harvard University Press.
- Gibbard, A. 2003. *Thinking How to Live*. Cambridge, Massachusetts: Harvard University Press.
- Hawthorne, J. 2002. "Practical realism?" *Philosophy and Phenomenological Research* 64(1): 169–178. doi:10.1111/j.1933-1592.2002.tb00149.x.
- Jackson, F., and P. Pettit. 1998. "A problem for expressivism." *Analysis* 58(4): 239–251. doi:10.1093/analys/58.4.239.

- Korsgaard, C. 1996. *The Sources of Normativity*. New York: Cambridge University Press. doi:10.1017/CBO9780511554476.
- Korsgaard, C. 1997. "The normativity of instrumental reason." In: *Ethics and Practical Reason*, edited by G. Cullity and B. Gaut, 215–254. Oxford: Oxford University Press.
- Olson, J. 2014. *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780198701934.001.0001.
- Radzik, L. 1999. "A normative regress problem." *American Philosophical Quarterly* 36(1): 35–47.
- Radzik, L. 2000. "Incorrigible norms: Foundationalist theories of normative authority." *Southern Journal of Philosophy* 38(4): 633–649. doi:10.1111/j.2041-6962.2000.tb00919.x.
- Radzik, L. 2002. "A coherentist theory of normative authority." *Journal of Ethics* 6(1): 21–42. doi:10.1023/A:1015887107235.
- Railton, P. 1992. "Nonfactualism about normative discourse." *Philosophy and Phenomenological Research* 52(4): 961–968. doi:10.2307/2107923.
- Railton, P. 2004. "How to engage reason: The problem of regress." In *Reasons and Value: Essays on the Moral Philosophy of Joseph Raz*, edited by R. J. Wallace, M. Smith, S. Scheffler, and P. Pettit, 176–201. Oxford: Oxford University Press.
- Rawls, J. 1971. *A Theory of Justice*. New York: Oxford University Press.
- Rawls, J. 1993. *Political Liberalism*. New York: Columbia University Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, Massachusetts: Harvard University Press.
- Scanlon, T. M. 2014. *Being Realistic About Reasons*. New York: Oxford University Press.
- Smith, M. 2001. "Some not much discussed problems for non-cognitivism in ethics." *Ratio* 14(2): 93–115. doi:10.1111/1467-9329.00149.
- Smith, M. 2002a. "Evaluation, uncertainty, and motivation." *Ethical Theory and Moral Practice* 5(3): 305–320. doi:10.1023/A:1019675327207.
- Smith, M. 2002b. "Which passions rule." *Philosophy and Phenomenological Research* 65(1): 157–163. doi:10.1111/j.1933-1592.2002.tb00191.x.
- Silverstein, M. 2016. "Reducing reasons." *Journal of Ethics and Social Philosophy* 10(1): 1–23.
- Tanyi, A. 2006. An Essay on Desire-Based Reasons Model. Doctoral dissertation. Central European University. http://politicalscience.ceu.edu/sites/politicalscience.ceu.hu/files/basic_page/field_attachment/attilatanyi.pdf (Accessed: 24/10/2016)
- Velleman, D. 2000. "The possibility of practical reason." In: *The Possibility of Practical Reason*, 170–199. New York: Oxford University Press.
- Williams, B. 1973. "A critique of utilitarianism." In: *Utilitarianism, For and Against*, edited by J. J. Smart and B. Williams, 77–151. Cambridge: Cambridge University Press.
- Williams, B. 1985. *Ethics and the Limits of Philosophy*. Cambridge: Cambridge University Press.
- Williams, B. 1995. "Moral Incapacity." In: *Making Sense of Humanity*, 46–56. Oxford: Oxford University Press.