

What is normativity?

John Skorupski
University of St Andrews

Abstract

The thesis that the concept of a reason is the fundamental normative concept is in the air. In this paper, I examine what it amounts to, how to formulate it, and how ambitious it should be. I distinguish a semantic version, according to which any normative predicate is definitionally reducible to a reason predicate, and a conceptual version, according to which the sole normative ingredient in any normative concept is the concept of a reason. Although I reject the semantic version, I examine its potential in some detail. And I claim that the conceptual version is plausible.

1. Normativity and reasons

The concept of a reason is the fundamental normative concept: this thesis is in the air.¹ One of its attractions, for me, is that it supplies illuminating terms for an old distinction: the distinction between the descriptive and the normative. It has been the fashion to deprecate any such distinction, but in my view its foundations in the critical epistemological tradition, whether empiricist or Kantian, remain as strong as ever. Putting the contrast in terms of propositions about the world and propositions about reasons, as against ‘is’ and ‘ought’, or ‘fact’ and ‘value’, is a helpful way of stating two ideas I would want to defend. The first is that normative propositions are not themselves descriptive or factual — they do not present more facts, they do a different job. The second, contrary to the non-cognitive (emotivist, prescriptivist, voluntarist etc.) strains in the critical tradition, is that

¹ I am using the word ‘normative’ broadly to contrast with ‘descriptive’ — not narrowly, as some writers use it, to contrast with ‘evaluative.’ As to the thesis, although it is in the air it is difficult to attribute definite versions of it to particular people. I have been helped in thinking about it by Gibbard (1990) and Scanlon (1998). Gibbard works with the basic notion of what is ‘rational,’ or ‘what makes sense,’ while for Scanlon the basic notion is explicitly that of a reason. (Substantively this may not be a big difference, though other substantive differences between them are indeed big).

they are nonetheless straightforward, genuine, propositions. Both ideas become, I think, a little more persuasive when we focus directly on propositions about reasons, in their full generality: reasons for belief, reasons for feeling, reasons for action — although I should emphasize that neither of them is thereby forced.

This semantic and epistemological background will not however be our topic here.² The thesis that something's being a reason for someone is the fundamental normative concept — as against for example obligation, or value or 'rationality' — is independent and can be assessed in its own right, whatever one thinks of those background issues. One can favour this thesis on more specific grounds; for example that it clears up the obscure relation of normative supervenience (which simply reduces to the reason relation) or that it opens up a plausible sentimentalist account of value in terms of reasons for feelings. And then there is the underlying, stage-setting thought that 'normativity' can be nothing more than that by which autonomous — reason-sensitive — agents steer. It must come down to their reasons for belief, feeling and action: to the normative relation between facts on the one hand and those reason-responsive acts or states on the other. The elaborations of normative discourse must issue in propositions about reasons or they are idle wheels.

But I shall not discuss these persuasive considerations (as they seem to me) either. My aim in this paper is simply to examine what the thesis amounts to, how to formulate it, and how ambitious it should be.

I should note that we cannot simply assume there is only one way of interdefining the circle of normative concepts. If it turns out that reductive analysis of normative concepts can be done in more than one way, of which reduction to reasons is one, we will not have shown that reasons are *the* fundamental normative concept. I myself do not believe it can be done in any other way, but I am only going to examine here whether a reduction to reasons is possible. Call this claim the *Reasons Thesis*. If it is sound, I leave it as a challenge to others to show that some other reduction, for example to obligation, value, or rationality is possible.

² I discuss it in Skorupski, 2006.

2. Semantic reduction and concept-possession

The simplest and sharpest version of the Reasons Thesis is explicitly semantic: any normative predicate is definitionally reducible to a reason predicate (as well as non-normative predicates if it is not itself wholly normative). Call this the *semantic thesis*.

An alternative is to put the Reasons Thesis at the level of concepts: the sole normative ingredient in any normative concept is the concept of a reason. This claim is weaker, in that the semantic thesis entails it but the converse need not hold.

To illustrate the difference, consider the moral concepts, which are obviously a major case in any general discussion of normativity. Let us allow that the basic concept here is that of moral wrongness: a moral obligation, for example, is by definition something with which non-compliance is morally wrong. And let us also allow for the sake of argument that the morally wrong is that which (absent extenuating circumstances) there is reason to blame. So that promises a reduction of moral claims to claims about reasons. But it may plausibly be objected that it cannot yield a *definition* of ‘morally wrong’, since ‘blame’ in the relevant sense should be defined as that sentiment which is appropriate — which there is reason to feel — towards the *morally wrong*. A good objection, but not the end of the story. Suppose that there is a distinctive sentiment of blame whose object is the morally wrong; suppose also that we can be given an independent characterization of this sentiment, say in terms of the actions to which it disposes, or by having it explained as the sentiment we feel when we consider some specific paradigm cases. Suppose, finally, that when we have been made familiar, in the first-person way, with the sentiment, we find ourselves able to go on spontaneously, and reasonably convergently, making confident new judgements about when there is reason to feel that specific sentiment. In that case we have everything we need to grasp the concept of moral wrongness. The morally wrong is that which is blame-sentiment-worthy in the absence of an excuse. The concept can in this sense be exhaustively captured in terms of the concept of a reason and the concept of a certain sentiment, even though ‘morally wrong’ cannot be defined in terms of ‘reason’. Furthermore if we can individuate the sentiment by these methods we can also use them to introduce a term to refer to the sentiment, say ‘BS’. The morally wrong is that which is BS-worthy

(which there is sufficient reason to respond to in that way) in the absence of an excuse.³

Colour concepts provide a partial analogy. To grasp the concept of yellow you have to be familiar in the first-person way with the sensation of yellow, and take the presence of the sensation as a reason to think it indicates a particular objective feature, while also being able to recognise cases in which that reason is defeated. If you have all that, you have what's needed to possess the concept. Nonetheless, 'yellow' is not definable in terms of 'sensation of yellow'. 'Yellow' is semantically simple, 'sensation of yellow' is complex; you need to know what the former means to know what the latter means. G. E. Moore was right about the indefinability of 'yellow' (if not about the indefinability of 'good').

The partial analogy with moral wrongness is that a being which has the blame-response, and an ability to distinguish between having it and having reason to have it, has what is needed to possess the concept of moral wrongness, even though 'morally wrong' is not definable in terms of 'what there is reason to blame'.

This is an approach that captures a concept by saying what is required for possession of the concept, rather than by the semantic route of explicit definition. Nonetheless, if the story is right we have shown that the normative concept of moral wrongness can be fully captured in terms of the normative concept of a reason, and in this sense made progress in defending the Reasons Thesis. (Attributing full possession of *either* concept, that of wrongness or that of yellowness, involves attributing reason-sensitivity to its user, in that it is a condition for grasping a concept that one is sensitive to the reasons that warrant applying it. What distinguishes moral wrongness as a *normative* concept, unlike yellowness, is that you can have the concept only if you are sensitive to reasons for the blame-sentiment itself. It always makes sense to ask whether there are reasons to feel the blame sentiment; with the sensation of yellow the analogous question makes no sense.)

If these points about the moral concepts are right, their normativity cannot be captured through a direct semantic reduction of moral predicates to reason predicates within our actual language. I shall come back to this in section 9. On either approach, however, we are

³ See Skorupski, 2005, for further discussion of blame.

going to need a definition of normative predicates in terms of reason predicates. In sections 2 — 8 I want to consider how to construct such a definition, and how far it can take us in accounting for predicates that strike us intuitively as normative. In other words, I want to explore how far we can take the semantic thesis.

3. The semantic thesis: preliminaries

This project encounters two difficulties at the outset — controversy about what terms are normative, and their indefinite extent on any view. The circle of terms that most people would recognise as normative contains a vocabulary that is diverse and wide. There are normative terms that span all fields — including epistemology, ethics and aesthetics — and whose normativity is pretty uncontroversial, at least in their most salient uses: ‘good’ and ‘bad’, ‘right’ and ‘wrong’, ‘ought’ and ‘should’. But then there are terms special to a field, as to whose normativity or otherwise there can be significant dispute. In epistemology, for example, we have ‘evidence,’ ‘a priori’, ‘valid’, and perhaps a certain sense of ‘probable’. Whether these are normative is a question of fundamental significance, lying at a philosophical crossroads. Suppose we take them as normative, and apply the semantic thesis. Evidence that p consists of accessible facts that gives one reason to believe that p . A proposition is a priori if there is outright reason to believe it, that is, on the basis of no facts. Validity is a priori truth-preservation. Probability, in one of its senses, is a matter of degrees of reason to believe. Clearly these suggestions are not philosophically neutral. A broadly ‘Kantian’ epistemology that takes normativity as a precondition of factual discourse is only just below the surface.

Still, we can distinguish, at least initially, between a dispute about whether certain terms are normative and a dispute about whether normative terms satisfy the semantic thesis. What then remains is the problem of indefinite extent. Proponents of the semantic thesis cannot be expected to consider all putative normative terms individually. The most that should be expected is that they define the ‘thin’, topic-neutral ones, and then consider representative classes of ‘thick’ or topic-specific ones. So the semantic thesis will inevitably have an element of the stipulative to it, however thoroughly it is carried

through. But if it is carried through convincingly enough, the final stipulation will seem natural and useful.

To advance the thesis, then, we need to show how certain prominent types of *prima facie* normative terms are definable in terms of reason predicates. What are these reason predicates? I think there are three. First, one can say that some particular facts are a reason, weaker or stronger, for some person at some time to ψ . Second, one can say that taking everything that counts for and against into account there is more or less strong overall reason for a person at a time to ψ . And third, one can say that some facts give a person at a time sufficient reason to ψ . That gives us the following relational predicates:

Specific reasons of degree: facts p_i are at time t a reason of degree of strength d for x to ψ — $R(p_i, t, d, x, \psi)$

Overall reasons of degree: facts p_i are at time t overall reason of degree of strength d for x to ψ — $R_o(p_i, t, d, x, \psi)$

Sufficient reasons: facts p_i are at time t a sufficient reason for x to ψ — $S(p_i, t, x, \psi)$.

I do not think these are interdefinable, but I will not argue that here. If I am wrong, the semantic thesis will say that all normative terms are definable in terms of less than three reason predicates, if I am right, then it is these three reason predicates that count for the thesis.

It is also very important for our purposes to note clearly that we don't just talk about practical and epistemic reasons, but that we also talk about reasons for feeling, for affective responses, such as gratitude, resentment and so forth. We can call them evaluative reasons. That we talk in this way is indisputable; evaluative reasons are constantly in our conversations and our minds. So ' ψ ' ranges over actions, beliefs, and feelings.⁴ What can still be debated is whether we can get a reduction here, for example of evaluative to practical reasons or epistemic to practical reasons. I am a 'trichotomist' in the sense that I believe these three kinds of reason are not reducible. Thus

⁴ With epistemic reasons, by the way, we must introduce a reference to epistemic fields: sets of facts that are discoverable by the inquirer and which give the inquirer reasons to believe this or that. Epistemic reasons are reasons relative to epistemic fields. This will not affect what follows. For more on this see Skorupski, 2006, section 2.

our reduction base consists of three irreducible reason predicates which range over three distinct kinds of state, attitude or act. With these preliminary points in hand we can set about formulating the semantic thesis.

4. The semantic thesis: normative sentences

The three reason predicates yield sentences of the form

$$\begin{aligned} &R(p_i, t, d, x, \Psi) \\ &R_o(p_i, t, d, x, \Psi) \\ &S(p_i, t, x, \Psi). \end{aligned}$$

Call such sentences *atomic* normative sentences, and define an *explicitly* normative sentence as one that is either

- (i) an atomic normative sentence

or

- (ii) a sentence which is built from sentences which include at least one atomic normative sentence, by means of the connectives of propositional logic, quantifiers, and the truth operator ('it is true that').

Now define the class of normative sentences as follows:

- a. Any explicitly normative sentence is normative.
- b. Any sentence which has a normative sentence as a *definitional* consequence is normative.
- c. No other sentences are normative.

By a 'definitional consequence' I mean something narrower than analytic consequence in the usual sense, i.e. derivability by principles of pure logic plus definitions. Definitional consequence is derivability by means of the natural deduction rules for conjunction and the universal quantifier alone, plus substitution of explicit definitions.

Before moving on let me note two points. First, our definition of normative sentences commits us to counting the negation of a normative sentence as a normative sentence. To count these wide-scope negations as normative sentences is admittedly somewhat artificial; it is not obvious, for example, that the denial that there is a reason to ψ

is itself a normative claim. One might accept the truth of such propositions on nihilist grounds. Yet the nihilist about reasons is presumably not taking a normative stance, whereas our definition attributes to him acceptance of some propositions expressed by normative sentences. This is an awkwardness. But we can say that the normative propositions a nihilist accepts are all expressible as wide-scope negations. It is in this sense that he takes no affirmative normative stance: he thinks that there are no non-negative normative truths.

Second, our definition commits us to counting a conditional with a normative antecedent or a normative consequent as a normative sentence. So consider 'If Michelangelo was not Italian, there is reason for you to take up golf'. This normative sentence can be deduced from non-normative sentences alone — in this case, from the non-normative sentence that Michelangelo was Italian. Thus we cannot maintain that no normative sentence is deducible from non-normative sentences.⁵ However, we can still maintain two truths that lie behind that thesis, and I think capture what is intended by the 'is/ought' or 'fact/value' distinction. The first truth is that there is no deductively valid argument (relying only on principles of logic) from *true* non-normative premises alone to an *atomic* normative sentence as conclusion. This holds on our account. As to the second truth: that, I suggest, is that no normative sentence can be the *definitional consequence* of non-normative sentences alone. This follows from our definition of 'normative sentence', together with another claim (which I take to be true), that the reason predicates are not themselves definable 'naturalistically', or in descriptive terms.

5. 'Ought', 'should' and 'right'

Let us say that any predicate formed from a normative sentence by dropping singular terms is normative. The project is then to show that all and only the predicates we are inclined to class as normative can be shown to be normative in this sense. A first stumbling block,

⁵ If we said that the conditional was non-normative, a normative sentence could still be deduced from non-normative sentences alone. For this conditional, together with the non-normative premise that Michelangelo is not Italian, entails the normative conclusion that you have reason to take up golf. The point goes back to Arthur Prior.

however, lies in the very varied way that many of the ordinary-language predicates that come to mind as normative are in fact used.

Take 'should.' Suppose A says 'It should be getting colder soon,' B says 'You shouldn't park here — it's a double yellow line,' and C says 'You should show some consideration for other people's feelings.' Does the word 'should' in these three statements have different senses, or the same sense but different uses? Without meaning to take a firm view about these semantic questions, I am going to talk about normative and non-normative *uses* of ordinary-language 'should,' 'ought,' etc. The question I am interested in can then be put by asking whether what a person says on a particular occasion could have been said — neither more nor less — by using a normative sentence, as defined above.⁶

Take the case of A. If A could have said exactly the same thing by saying 'There's (sufficient) reason to expect it to get colder soon' then he's using 'should' normatively. But without further context it's not clear that that *is* what he's saying. What about B's statement? Suppose B could have said the same thing by saying 'It's illegal to park on yellow lines' — is what he did say normative? Defining normative predicates in terms of reason predicates helpfully focuses the question. On a positivist view of law, at any rate, his use of 'should' is not normative by that definition. For according to legal positivism the fact that a law exists in a population is a fact about certain attitudes, dispositions and actions in that population (and perhaps their causal antecedents). If so, then 'It's illegal to park on yellow lines' has no explicitly normative sentence as a definitional consequence. If what B says could be said by saying only that, then he is not using 'should' normatively. The same goes for statements of convention — 'You should only move the king one square at a time'. If this sentence is used just to state the rules of chess then again 'should' is not being used normatively. What if it simply is not determinate what a person said? It may be simply indeterminate, for example, whether B is just saying that it is illegal to park on double yellow lines, or saying that there is sufficient reason for you not to do so, because it's illegal. In that case it is indeterminate whether A's use of 'should' is normative.

⁶ I am assuming that the reason predicates I have listed do not have the fluidity of use that I am allowing to 'should' etc. That seems right, however.

As for C, in a given context he may be saying that you have a moral obligation to consider other people's feelings. I have already suggested that moral concepts constitute an insurmountable objection to the semantic thesis. However that does not settle the question whether what C is saying, on this moral reading, *can* be said by means of normative sentence. I shall come back to the question in section 9.

'Should' has various uses, then, not all normative.⁷ The same applies to 'ought'. But 'should' and 'ought' also have salient uses that are normative. In one of these uses, what one should or ought to do is simply *any one of the things there is sufficient reason to do*. In another use, what one should or ought to do is any one of the things that one has sufficient reason to believe there is sufficient reason for one to do. Let us stick with the first, so-called 'objective' use. The practical-normative case tends to be the one that people concentrate on, but we can talk in the same normative way about what a person should believe or feel. We can consider whether a person *should* feel resentment given the way he has been treated, or whether he *should* believe something in the light of the facts available to him. Here again we may be talking about what there is sufficient reason for that person to believe or to feel.

We can define a regimented sense of 'should' and 'ought', reserved for just that normative use. First, note that in the practical case, unlike the epistemic and evaluative cases, there can be sufficient reason for you to do any one of a number of strictly incompatible things. So let us say in the practical case that you have *uniquely* sufficient reason to ψ if ψ -ing is the only thing you have sufficient reason to do. What you have uniquely sufficient reason to do may thus be stateable only in a disjunction — to α or β or χ . In the epistemic and evaluative cases you have uniquely sufficient reason to ψ if and only if you have sufficient reason to ψ . We can now define the regimented sense of 'should' and 'ought': x should/ought to ψ if and only if there is uniquely sufficient reason for x to ψ . And we can define a regimented sense of 'right'. A right thing to do, believe or feel is a thing there is sufficient reason to believe, do or feel.

⁷ Another example is its use as an auxiliary verb — 'If I should die, think only this of me.'

6. 'Good'

While 'ought', 'should', 'right' have varied uses in ordinary language, the value terms 'good' and 'bad' are more stable. And here there is a well-known reduction to reasons — what Scanlon has called the buck-passing account.⁸ According to this, if we say that y is good we imply that there are some facts or other, p_i , which are sufficient reason for some pro-act or attitude towards y . (For 'bad' it would be a con-act or attitude.) But we do not say what those facts are. Nor do we say what the relevant pro-act or attitude is. However the buck-passer holds that a list of the pro-acts and attitudes can be provided, allowing the analysis to apply to the whole range of things that can be said to be good, including good arguments, good picnic tables, good people and good music. So 'good' is a normative predicate, as are 'right', 'should' and 'ought.'

I think that the buck-passing view is correct, but it has proved controversial and needs to be defended. There is the question of what the pro-attitudes are, how to formulate the view for all cases, including, for example, functional uses of the term 'good', and how to extend it to *better*, *worse* etc. I think these questions can be satisfactorily answered; but here I want to discuss very briefly two more general objections that have been made to the buck-passing account, since we can illuminate the Reasons Thesis by considering how to answer them.

Consider the following definition of 'good F':

y is a good F in certain respects C, to degree d at t if and only if there is sufficient reason at t to pro y to degree d as an F in those respects C.⁹

Two objections may be raised to it.¹⁰ It may be objected in the first place that the equivalence fails to hold. And secondly, even if the equivalence holds it may be objected that the definition gets things the wrong way round. If there is sufficient reason to pro y as an F, that is

⁸ Scanlon, 1998, p. 96. Rabinowicz and Rønnow-Rasmussen, 2004, provide some earlier history.

⁹ I discuss buck-passing more fully in Skorupski, 2007.

¹⁰ See for example D'Arms and Jacobson, 2000 (especially section IV), Crisp, 2000, 2005.

because y is a good F . I *explain* why there is reason to pro y by pointing out that the reason is that it is a good F . If saying that it was good just *was* saying that there is sufficient reason to pro it, the objection runs, that would not be an explanation.

As to the first objection, 'evil demon' examples illustrate why the equivalence may be thought to fail. Suppose the violin performance is not good, but the evil demon will punish me with eternal torture if I fail to admire it. Is that not sufficient reason for me to admire it, even though it is not good? Similarly, suppose that some purported evidence to the effect that p is not good evidence, but the evil demon will again punish me if I fail to believe that it probabilises the conclusion that p . Is that not sufficient reason for me to believe that it does probabilise that conclusion, even though it is not good evidence?

In reply to this objection, we can apply the distinction between reasons to believe or feel on the one hand and reasons to bring it about that one believes or feels on the other. Thus: there is no sufficient reason for me to admire the performance, though there certainly is sufficient reason for me to bring it about that I admire the performance, if I can. In other words, in this case there is reason for me to bring it about that I admire something which there is no reason for me to admire. Likewise: there is no sufficient reason for me to believe that this purported evidence to the effect that p probabilises that conclusion, but there certainly is sufficient reason for me to bring it about that I believe it does, if I can. Bringing about these things is undoubtedly choice-worthy, if it is possible. So what is good in these cases, according to our definition, is the choice or policy or action of (trying to) bring it about that I admire the performance, or that I believe the evidence to be probabilising. It does not follow, on the buck-passing analysis, that the performance or the evidence itself is good.

Is this response *ad hoc*?¹¹ It does not seem to me that it is. The distinction between epistemic, evaluative and practical reasons is independently based on what kind of act — belief, feeling or action — a given reason is a reason for. So the response is just an automatic consequence of identifying what exact reason relation we are discussing. In the case of the violin performance, the fact that the evil demon has his evil plans is a sufficient reason for me to *do* something —

¹¹ As suggested in Rabinowicz and Rønnow-Rasmussen, 2004, p. 412.

namely, bring it about that I admire the performance, if I can.¹² In the circumstances, that would be a very good thing to bring about. Over and above that uncontroversial point, there is then the question of whether an *evaluative* reason relation also holds. Does the self-same fact about the evil demon stand in that *distinct* reason relation to me and a certain feeling of mine, namely, *admiring the performance*? The two relations are distinct, since their *relata* are distinct. And once they have been distinguished, there seems to me to be no case for holding that the second relation holds as well as the first.

This may be denied. Take the case of belief. Someone may agree that there is a type difference between epistemic reasons to believe and practical reasons to make yourself believe, or bring it about that you believe, but still insist that facts about the usefulness of believing that *p* are not just a reason to make yourself believe that *p* but *also* a reason to believe that *p*. Against this one can object that there is nothing self-contradictory in saying that a person has reason to *make* himself believe something, for example, that he will survive the dangerous mission, even though he has no reason to believe it. The objection seems to me decisive, but suppose that it too is denied. The denier would have to have, it seem to me, an at least partly pragmatic theory of truth. A pragmatist about truth can hold that the usefulness of believing a proposition is a priori indicative of the truth of that proposition. For him, therefore, the usefulness of believing that *p* is unproblematically both a reason to believe that *p* and a reason to believe that it is true that *p*. If, in contrast, one denies that the usefulness of believing a proposition is a priori indicative of its truth, then how can this usefulness constitute, in and of itself, a reason to believe that *p*?

Let us turn to the second objection. Does our definition get things the wrong way round? I can *explain* why there is reason to favour something by pointing out that it is good. I do not *explain* why it is good by pointing out that there is reason to pro it.

This is a weak objection. In the first place, there may be reason to favour it in one way because there is reason to favour it in some other way. Thus, the reason to *choose* this CD performance may be that it is the best performance — the one there is most reason to *admire*. We might have reasons to choose this particular CD for reasons other

¹² Rabinowicz and Rønnow-Rasmussen agree.

than the goodness of the performance: because it is cheap for example. If we are working to a tight budget the cheapest CD performance may be the best choice, even though the performance is not the best. So it has explanatory force to say that there is reason to choose it because it is the best performance.¹³

Further, how do we explain why something is a good F? We do so by pointing out the facts about it that make it good as an F. Those are the very facts that give one reason to favour it as an F. In that sense we do explain why it is a good F by stating the reasons for favouring it as an F. We do not of course explain why a thing is good by simply saying *that* there are reasons to pro it, but we do explain by saying what those reasons are.

My conclusion is that if one clearly distinguishes and acknowledges the three basic kinds of reasons — epistemic, practical and evaluative — and looks in detail at the structure of normative explanations, or justifications, what emerges is a perfectly sound defence of the buck-passing view.

7. Thick evaluative terms

We have now considered the thin or topic-neutral normative terms that make up the central circle of normative vocabulary. We also have to consider the difficult and ramified outer circle. Significantly, most terms in this outer circle are in fact evaluative, having to do with reasons for feeling this or that — I will call them ‘thick evaluative terms.’

When we say that something is good or right, we say there are reasons without saying what the facts are that constitute those reasons, or for what act or attitude they provide a reason. Thick evaluative terms effectively impose tighter restrictions. How do they do that?

One way they might do so is by being more specific about what attitude there is reason for. Consider in particular predicates of the form

$$A: (\exists p_i)S(p_i, \phi(y)).$$

¹³ You might have reason to choose a CD which has a bad performance just to illustrate in your consumer survey what poor value there is on the market. In this case the bad performance is a good choice.

where ' ϕ ' names some specific affective reaction. What they say about y is that there is sufficient reason to take attitude ϕ to y .¹⁴ They say no more than that. However, that may tell us quite a lot. From our inner understanding of that particular reaction — admiration, say, or disgust — we can work out quite a bit, indirectly, about what sort of facts those facts must be, even though we are not told what specific facts obtain.

The number of potential practical normative or epistemic normative predicates is as large as the number of action and belief contents, i.e. indefinitely large. In the evaluative case the affective *attitudes*, as well as their contents, ramify indefinitely. Furthermore, and very importantly, evaluations of type A will not all be neatly classifiable as asserting that there is sufficient reason for a *pro* or *con* attitude toward the object, and thus, as implying that the object is good or bad in some respect. Does the fact that there is reason to be disturbed or shocked by a performance, for example, make it good or bad? That may well depend on context. Disgust is generally a con-attitude. But even this need not always be evident. An object may give sufficient reason to respond with a complex of attitudes of which disgust may be an element. Yet the overall complex may justify a pro-attitude. This openness between making a specific evaluation in terms of an affective attitude and determining whether that evaluation counts towards the goodness or badness of the object is one of the things that gives critical discourse in the arts its depth, complexity and elusiveness. Specific evaluations can be indefinitely subtle, quite informative (though inexplicitly) about the facts, and linked in no neat or simple way to an assessment of their objects as good or bad.

Does ordinary language contain any terms analysable as type-A predicates? One set of candidates for the role consists of terms like 'admirable,' 'lovable,' 'enviable,' 'desirable' and so on. Arguably, 'This is admirable' means 'there is sufficient reason to admire this.' If that is right, ' ψ is admirable' is a purely normative predicate.

It may be replied that predicates of this kind — ' ϕ -able' — can have the non-normative meaning 'capable of inspiring admiration/desire/love.' Maybe they are useable in either or both ways, the

¹⁴ This is the simplest case. Where the reason is relative to an agent the agent variable can come into play: there is reason for me but not for you to desire y . Where it is temporally relative the time variable will come into play: there was reason then, but there is no reason now, to hope that y would come.

normative and the non-normative. But we need not accept this. The appearance that such predicates have non-normative use may be explained away, by reference to their epistemology and their function. The fact that something does inspire admiration, desire, love or whatever (lastingly, on reflection, after discussion) is the fundamental epistemic criterion for its being worthy of admiration, desire, love etc. Now just because the fact that p is a criterion of, that is, constitutes a defeasible warrant for, the assertion that q , it does not follow that in asserting that q we are asserting that p . Nonetheless, when a proposition is very tightly tied to a dominant criterion, as in the case of these ' ϕ -able' predicates, it is easy to think (though strictly it is incorrect) that asserting the proposition is asserting that the criterion obtains. Furthermore, given the use of these predicates in criticism and persuasion, it is very natural, if one is challenged as to whether a performance is really admirable or a house is really desirable, to fall back on the fact that many people do admire it or desire it. But neither point shows that 'admirable' or 'desirable' have a non-normative ingredient in their meaning. So the view that these are purely normative predicates can be defended.

At first sight terms such as 'frightening', 'moving', 'boring' and so on seem more clearly to fall on both sides of the normative/non-normative border line. Again, however, one can argue by reference to their epistemology and their function that the appearance of non-normative content is misleading. Take 'frightening'. Can it *mean* 'fear-causing' as well as 'fear-worthy'? It is not easy in fact to come up with an example of a use that is clearly non-normative. Consider 'These steps would be very frightening to someone with poor balance and short sight.' 'Frightening' here could mean 'fear-causing', but it could also mean 'fear-worthy'. Or again: 'Spiders are frightening to many people but there's really nothing about them to be frightened of.' Is the first occurrence of 'frightening' in this sentence a non-normative use? Not necessarily. The sentence could mean 'Spiders seem to many people to be fear-worthy but they are not really.'

In the same vein, suppose I attend a lecture on quantum physics and find it boring. Note the typical locution: *finding* something boring, amusing, irritating etc. This looks as though it might be factive, but it is not: from the fact that I found something boring it does not follow that it *is* boring. Rather, to say I found it boring is equivalent to saying that it seemed boring to me. Someone may reply 'It was not *really* boring — it is just that you do not know anything about the

subject. The experts who were there found it really interesting.' Likewise, if one says that something was genuinely moving, or really or truly moving, one makes a contrast between what merely seems moving and what really is. How does the contrast between appearance and reality work here? It is not that something can seem to move me without really moving me. Rather, the contrast is between apparent and real normativity. When one is moved by something, a musical performance say, one doesn't simply experience an emotion, one experiences it as appropriate, fitting to its object. But it can happen that this normative dimension in one's feelings strikes one as misguided. I am moved by the performance, and that emotional reaction comes packaged, so to speak, with an impression of itself as reasonable — but at a level more detached from the reaction itself I think to myself that I am being sentimental. The performance was not genuinely, really, moving; I am just a sucker for that sugary kind of violin sound. So 'genuinely moving', 'truly moving' are normative.

So far, so good. Yet many terms used in ordinary-language evaluative discourse may seem normative even though they carry no connection to an affective reaction on their face. In aesthetic evaluation they abound: 'bizarre', 'beautiful', 'sublime'. Let us examine a few cases.

Since our linguistic taxonomy of affective responses is less ramified than our taxonomy of evaluations one should expect difficulty in finding a word to characterise, for each evaluative predicate, the *specific* affective response that its application says there is sufficient reason for. Take the example of 'bizarre.' It has reference to some form of surprise. To say that something is bizarre is to say, according to the *Oxford English Dictionary*, that it is 'extravagant, whimsical, strange, odd, fantastic.' But also: 'At variance with the standard of ideal beauty or regular form; grotesque, irregular.'¹⁵ In each case, the surprise is justified by the fact that the item in question goes beyond what there is reason to expect from a 'regular' item of that kind. The whimsical, fantastic, etc. is surprising in *that* way; it gives reason for an attitude of surprise directed to *that* content. What content? Something that is bizarre. So it seems that we have the same circularity here as we found with 'morally wrong'.

¹⁵ The OED also has: 'At variance with recognized ideas of taste, departing from ordinary style or usage; eccentric ...' But recognized ideas of taste, and ordinary style and usage, may themselves be bizarre, in which case the eccentric may precisely not be bizarre. None of these, of course, are strictly *definitions*.

With the beautiful and the sublime we arrive at terms which have long been recognised as difficult to define. Many of the traditional discussions in aesthetics have in fact been concerned with what *makes* a thing beautiful (in the sense in which one might discuss what makes a thing annoying.) It seems uncontroversial that the beautiful is that which is admirable in a certain way. But what way? By being delightful, charming, graceful, by giving a certain sort of pleasure, in virtue of a certain sort of rightness of combination, of proportions, colour-combinations or whatever. So: if y is beautiful then y has properties combined in such a way as to give satisfaction or pleasure of the kind taken in this perceived *rightness* of combination, and in virtue of that pleasing combination, to give sufficient reason for admiration. Thus, the concept of the beautiful involves a double layer of normativity. (The concept of the sublime seems simpler in its logical structure: it is that which is uplifting, exalting, ennobling, awe-inspiring.) The properties of a beautiful thing have a rightness of combination, they combine as they should. And it is this rightness of combination that gives sufficient reason for delight and admiration. This account of the notion also highlights one reason why the concept of the beautiful remains a problem for aesthetic theory. What is it for properties to combine as they should, or as there is reason for them to combine? What can it mean to talk about whether properties are doing what they should — how can the normative concept get an application here? There seems to be a simulated functional assessment involved in a judgement of beauty, captured in Kant's formula of 'purposiveness without a purpose.' We can say that an object with a function is working as it should, arranged as it should be, by reference to its function. But with beauty we have the idea of rightness of arrangement without a function to underpin it.

8. Mixed predicates

There might be terms that predicate both a non-normative property and a reason constituted by possession of that property. Where 'F' is the non-normative property in question, such terms will have the form

B: S(the fact that Fy , $\phi(y)$)

Call this a mixed predicate. Are there any such predicates?

We have just suggested that various predicates which may look as though they have this form don't really. Term like 'admirable' and 'boring' are purely normative. It's true that the fact that something is boredom-causing is, epistemically speaking, a reason to believe it is boring, but that is a different point.

Then again, some predicates which may be thought to be mixed are probably not normative at all. Consider 'That was a cruel thing to do'. This means (at least) 'That was an action which showed in its motivation the presence of a non-instrumental desire to inflict suffering'.¹⁶ So the assertion makes a factual claim. But does it also, as part of its meaning, make a normative claim? That would make it a mixed predicate: e.g. 'y shows in its motivation the presence of a desire to inflict suffering, and that fact is a sufficient reason to disapprove of y'.

It is not clear to me that this is the right account. Is it self-contradictory to deny that there is reason to disapprove of cruelty? I think not: it is just false. Likewise with kindness, courage, resentfulness, laziness, arrogance, coldness and so on. These are names of character traits which we take to be virtues or vices, excellences or defects, but it does not seem to me to be part of their meaning that that is what they are. Similarly with such terms as 'adultery': there seems to be nothing self-contradictory in saying that there is nothing wrong with adultery. (One might avoid that word, meaning to avoid unintended conversational implications, or one might choose it precisely to rid the word of those implications.)

But now consider 'murder'. To call something a murder is at least to say that it's an intentional killing with 'malice aforethought'. Is it a definitional truth that there is reason to condemn murder? If it is, that's because *malicious* killing means 'killing there is reason to condemn.' Not every intentional killing is malicious, however, and it does not seem possible to narrow down the intentional killings that are malicious by a wholly non-normative criterion. A murder in an intentional killing where the intention is of such a kind that there is sufficient reason to blame the person who does the act with that intention.¹⁷ This gives us as a descriptive as well as a normative com-

¹⁶ Cruelty is more than callousness, or mere lack of care. 'It wasn't just thoughtless, brutal, stupid, reckless etc. — it was cruel.'

¹⁷ There can of course be further distinctions, say between murder and culpable homicide, or murders of various degree.

ponent. Importantly, however, it does not give us a wholly *non-normative* characterization of the fact on which the blameworthiness supervenes, so it does not fit form B. It has this structure:

x is an intentional killing of such a kind that there is sufficient reason to blame the person who does x .

Predicates with this structure have a non-normative component, but one that is not strong enough on its own to characterise the reason purely descriptively. Consider courtesy and discourtesy. Courtesy is a matter of paying due care to putting people at their ease in a social context where that is relevant. Discourtesy is giving insufficient attention to that. 'Insufficient' is normative: the attention (if any) the action showed to placing others at their ease was small enough to constitute sufficient reason for disapproval of that aspect of the action. So the predicate ' x is discourteous' has a reason predicate as a definitional consequence — 'there is sufficient reason to disapprove of x '. But it is not possible to factor out a wholly non-normative component that constitutes the sufficient reason for disapproval. We are assessing the action by reference to how much attention it showed, in its context, to putting others at their ease, and saying that it did not show enough. Saying that is assessing the action normatively. Equally if we say the action was courteous, we are assessing it in the same dimension and saying that it showed a degree of attention to putting people at their ease great enough to give sufficient reason for approval.

So what we are saying, in these two examples, is of the form

C: y has property F to a degree or in a way which is such as to give sufficient reason to $\phi(y)$.

Various other terms work like that: 'careless', 'reckless', 'negligent', 'inattentive', 'overbearing'. We characterise the degree of care, attentiveness, etc. by saying that it is such as to satisfy a normatively expressed condition. So we cannot factor out a non-normative sentence which states the fact that constitutes the sufficient reason.

I therefore agree with people who hold that there are thick evaluative terms from which a reason-giving factual component is inextricable. We may not have a way of exactly characterising just those facts that merit a particular attitude except by reference to the attitude in question. And we may not have neat vocabulary to characterise

exactly the attitude in question. However these points do not undermine the two truths which I said lay behind the *is/ought* distinction.

9. Normative content and the reason predicates

If the semantic thesis were correct, we could draw the conclusion that anything normative our actual vocabulary allows us to say could be said while using only descriptive predicates and reason predicates. All normative content could be expressed that way. Other normative predicates would be expressively, if not practically, redundant.

Does this conclusion still apply if we have to supplement the semantic thesis with the concept-possession thesis, as with the predicates ‘morally wrong’ or ‘bizarre’? Take ‘morally wrong.’ Would it be possible, without expressive restriction, not to have the term ‘morally wrong’ but only a word for a certain affective response to actions, plus a reason predicate that allowed one to distinguish between actions that provide sufficient reason for that response (‘merit’ it) and actions which do not? More generally, could terms that are not normative by the semantic thesis, but nevertheless express normative concepts — as is the case with ‘morally wrong,’ and arguably, ‘bizarre’ — be replaced without loss of expressive power by explicitly normative predicates? If so, then the conclusion would still apply.

Consider again the analogy with colour words. Would it be possible, without expressive restriction, to have no colour words but only colour-sensation words, and to talk only about the powers of objects to cause those sensations, and the grounds of those powers? This question should be distinguished from a question about the order of language-learning, or of concept-acquisition. Plausibly, the concept of a colour-sensation is further on in that order than the concept of a colour, and it might even be argued that it cannot be grasped except through a prior grip on the concept of a colour. This might explain the semantic priority of colour words in our actual language. However, this impossibility would not show that a language which had colour-sensation words but no colour words would be *expressively* restricted in comparison to our actual language. To assess the issue, we would have to consider more carefully the notion of ‘saying the same thing,’ ‘expressing the same content’.

The same goes for the moral concepts. Perhaps they lie earlier in the order of learning than the concept of a blame-sentiment, and

perhaps that explains their semantic priority in relation to terms denoting the sentiment. But again that does not show that a language which included the reason predicates, and terms for the sentiments or attitudes, including the blame sentiment, would be expressively restricted in comparison to the one we have.

This has been an illustrative rather than a comprehensive survey, but we have turned up no good objection to the Reasons Thesis — the thesis that all normative concepts are reducible to the concept of a reason. The question on which I have ended is whether a language which added only the reason predicates to its descriptive vocabulary, and contained a sufficiently ramified and nuanced vocabulary of sentiments, would be able to express any normative contents that we can express. I have found no decisive objection to that idea, and some reason to consider it plausible. Furthermore, taking the concept of a reason as basic seems to me to fit in with other persuasive philosophical considerations. All in all, it seems to me that the case for thinking the normative domain is the domain of propositions about reasons is pretty strong.¹⁸

John Skorupski
 Department of Moral Philosophy
 Edgecliffe, The Scores
 The University of St Andrews
 St Andrews, KY16 9AL
 United Kingdom
 jms2@st-andrews.ac.uk

References

- Crisp, Roger. 2000. Review of *Value ... And What Follows*, by Joel Kupperman, *Philosophy*, 75.
- Crisp, Roger. 2005. Value, Reasons, and the Structure of Justification: How to Avoid Passing the Buck. *Analysis* 65.
- D'Arms, Justin and Jacobson, Daniel. 2000. Sentiment and Value. *Ethics* 110.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Oxford University Press.
- Rabinowicz, Wlodek and Rønnow-Rasmussen, Toni. 2004. The Strike of the Demon: On Fitting Pro-attitudes and Value. *Ethics* 114.

¹⁸ I am grateful to Jonas Olson, Wlodek Rabinowicz and Toni Rønnow-Rasmussen for their helpful comments.

- Scanlon, T. M. 1998. *What we Owe to Each Other*. London: Harvard University Press.
- Skorupski, John. 2005. Blame, Respect and Recognition: A reply to Theo van Willigenburg. *Utilitas* 17.
- Skorupski, John. 2006. Propositions about Reasons. *The European Journal of Philosophy* 14.
- Skorupski, John. 2007. Buckpassing About Goodness. In *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Ed. T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson, 2007, where I also give definitions of good for *x*, good of *x*, and functional goodness. (<http://www.fil.lu.se/hommageawlodek/site/papper/SkorupskiJohn.pdf>).