

This article was downloaded by: [University of Oxford]

On: 21 January 2011

Access details: Access Details: [subscription number 773573297]

Publisher Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Australasian Journal of Philosophy

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713659165>

The Standard-Relational Theory of 'Ought' and the Oughtistic Theory of Reasons

Daan Evers^a

^a University of Oxford,

First published on: 22 February 2010

To cite this Article Evers, Daan(2011) 'The Standard-Relational Theory of 'Ought' and the Oughtistic Theory of Reasons', Australasian Journal of Philosophy, 89: 1, 131 – 147, First published on: 22 February 2010 (iFirst)

To link to this Article: DOI: 10.1080/00048401003610890

URL: <http://dx.doi.org/10.1080/00048401003610890>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

THE STANDARD-RELATIONAL THEORY OF ‘OUGHT’ AND THE OUGHTISTIC THEORY OF REASONS

Daan Evers

The idea that normative statements implicitly refer to standards has been around for quite some time. It is usually defended by normative antirealists, who tend to be attracted to Humean theories of reasons. But this is an awkward combination: ‘A ought to X’ entails that there are reasons for A to X, and ‘A ought to X all things considered’ entails that the balance of reasons favours X-ing. If the standards implicitly referred to are not those of the agent, then why would these entailments hold? After all, Humeanism says that ‘A has a reason to X’ is true if and only if A has some desire which is furthered by X-ing.

In this paper, I develop a standard-relational theory of ‘ought’ and a non-Humean theory of reasons (oughtism). Together, they explain why ‘A ought to X’ entails not only that there are reasons for A to X, but also that the balance of reasons favours X-ing. The latter explanation depends on a theory of weight, in which the weight of a reason depends on the position of a rule (standard) in an order of priorities. The theories are truth-conditional, but do not require objective normative facts for the truth of ‘ought’ judgments and judgments about reasons.

1. Introduction

Antirealists or subjectivists about normativity deny the existence of objective normative facts (in some sense of ‘objective’ which I will not attempt to delineate here). Concerning the meaning of normative statements (statements involving words like ‘ought’, ‘good’, ‘wrong’, etc.), antirealists roughly have three options: (1) they can say that normative statements purport to describe objective normative facts which do not exist. This entails that all positive and non-analytic normative statements are false. Or (2) they can deny that normative statements purport to describe anything at all. Instead, they might be expressions of approval or disapproval, much like ‘Hurray!’ is an expression (but not a description) of excitement or enthusiasm. This means that the surface grammar of normative language is deceptive. Finally, (3) they can say that normative statements do purport to describe normative facts, but not objective ones (in the sense of ‘objective’ I have not bothered to explicate). Or at least, they can say that the meaning of normative statements is not such as to require the existence of objective

normative facts in order for these statements to be true. I want to develop a version of the third option. It is a form of contextualism about normative language.

Here is one reason why the third option is not easy to develop. It is not particularly plausible that ‘ought’ statements are sensitive to the desires of agents. Take the statement: ‘Hitler ought not to have ordered the extermination of the Jews’. Few think that its truth is affected by whether Hitler had any desires that were furthered by that action. So it is, at least *prima facie*, plausible that its truth conditions do not involve Hitler’s desires. However, ‘A ought to X’ arguably entails that there are reasons for A to X. But many antirealists are Humeans about reasons: they believe that the truth of ‘A has a reason to X’ requires that A has some desire which is furthered by X-ing. But if the truth conditions of ‘A ought to X’ have nothing to do with A’s desires, then why should it entail anything about A’s reasons, if the reasons that A has *do* depend on A’s desires? This challenge is not often met. In this paper, I provide a theory of ‘ought’ and reasons which explains why ‘A ought to X’ entails not only that A has reasons to X, but also that the balance of reasons favours X-ing. My attempt to meet the challenge, however, involves abandoning Humeanism about practical reasons.

2. The Standard-Relational Theory of ‘Ought’

In this section, I outline a standard-relational theory of ‘ought’. The idea behind it is that all statements about what we ought to do are made relative to some contextually-supplied standards. In other words: there are no categorical ‘ought’s, if by that we mean a truth about what we ought to do which does not implicitly refer to a standard. This type of theory has been mentioned or developed (more or less explicitly) by authors like Mackie [1977], Wong [1984], Harman [1975], and Harman and Thomson [1996]. But so far it lacks enough semantic detail to allow a fair assessment. And it hasn’t been combined with a non-Humean theory of reasons.

The standard-relational theory of ‘ought’ says this: ‘ought’ refers to a three-place relation between an agent, an act, and a system of rules. It is the relation of being required by. This is not always true, but as a first approximation it will do.¹ This proposal comes to the following: if I say ‘A ought morally to X’, the meaning of my statement can be represented as: X is required of A by [the relevant system of] moral rules. If I say ‘A ought

¹It isn’t always true, because we can judge that things ought to be the case. There needn’t be anyone who ‘owns’ this requirement in the sense that s/he ought to bring it about that the state of affairs obtains. It is not obvious what to say about this, but I suggest that in such cases ‘ought’ refers to a two-place relation between a state of affairs and a system of rules. As will become clear, I analyse the notion of an act being required by a system in terms of rules in the system entailing that an act is (not) done by the agent. Rules entail various things in virtue of their propositional contents. Normally, these contents involve agents. E.g. the content of ‘Don’t kill!’ involves a quantifier which ranges over agents. But this needn’t be the case. There could also be ‘rules’ for states of affairs, which reflect certain ideals about the world. These rules can be represented by means of constructions like: ‘Let there be no suffering!’, whose propositional content is *that there is no suffering*. I suggest that ‘ought’ in statements like ‘It ought to be the case that p’ refers to a relation between a state of affairs and a system of non-agent-involving rules.

prudentially to X', the meaning of my statement can be represented as: X is required of A by [the relevant system of] prudential rules.²

Several things need specification. First, what are rules? Second, what is it to be required by a system of rules? Third, what determines the relevant rules? Fourth, in what way do the rules figure in the truth conditions of an 'ought' statement? I will answer these questions in turn.

First: what are rules? The kind of rules relevant here can be represented by imperatives: 'Do/don't do X!'.³ We can think of them as propositional contents which are not asserted but put forward with imperative force. The propositional content of 'Don't steal (in circumstances C)!' is *that no one steals (in C)*. However, it is not essential to think of rules as propositions put forward with some kind of non-assertoric force. We can also think of rules as propositions describing (aspects of) ideal worlds.⁴ What matters is that rules do not involve normative concepts like *ought* or *permissible*. For I want to *analyse* such concepts in terms of rules.

Second: what is it for an act to be required by a system of rules? For ease of exposition, I will stipulate that 'X' is a variable that takes both acts and omissions as values. So 'X' can both stand for the act of killing and the omission of refraining from killing. I will sometimes simply talk about *acts*, which are supposed to comprise both doing something and refraining from doing something. Now I can say that X is required by a system if and only if *conformity* to the system requires that one X-es. This can mean either one of two things:

(a) In a simple case, where the rules have no inconsistent implications with respect to X in C, conforming to the system means doing what is required by at least one rule in the system. An act (omission) is required by a rule just in case the act (omission) has some feature such that the rule entails that you do acts (refrain from doing acts) with that feature. Say that pushing a button in C will kill many people. In that case, the rule 'Don't kill in C!' requires you not to push the button, because the act has a feature (being such as to kill people) such that the rule entails that you refrain from acts with that feature in C. (I am assuming that the propositional content of the rule is something like *that no one kills in C*, which entails that you don't kill in C and thus not push the button either.)

(b) In a complex case, where the system contains rules with inconsistent implications with respect to X in C, conforming to the system means doing

²The standard-relational theory of 'ought' theoretically allows for objective standards: standards that are objectively correct. In this respect, it is neutral with respect to realism or antirealism.

³This is of course exactly what David Wong says in [1984: 37]. Not all rules can be represented by means of imperatives. For example, it is a rule of chess that one wins by checkmating one's opponent. The rules that matter for my theory can all be represented by imperatives. (Thanks to an anonymous referee for this journal.)

⁴(Part of) what makes an item with this propositional content a rule rather than, say, a report about the world, is, in the case of the speaker, the attitude which s/he takes towards the proposition (i.e. some attitude other than belief). Or, in the case of the hearer, it is his/her understanding that the proposition is not to be considered as a claim about the world. This may itself be a belief about the role of the content in the context. Because rules (and ideal worlds) are individuated with reference to the attitude taken towards propositional contents, I rely on a substantive assumption about mental states like desires, namely that they can be understood in non-normative terms. This assumption is often made by antirealists about normativity, but is not uncontroversial. I do not have the space to argue for it here. If this assumption is false, I may not without circularity be able to analyse 'ought' in terms of rules. (Thanks to an anonymous referee for this journal.)

what is required by the rule with the highest priority. For example: a system may contain both a rule of morality (which requires X in C), and a rule of prudence (which requires not-X in C). In this case, X is required of you in C by the system if and only if (1) the moral rule requires that you X in C (as explained in (a)) and (2) the moral rule has higher priority than the rule of prudence.

We can now state a general definition which covers both cases of conforming to a system of rules:

Conformity: an agent A conforms to a system of rules S by X-ing in C if and only if X-ing has a feature such that a rule in S entails that A performs acts with that feature in C and X-ing does not have a feature such that another rule in S with the same or higher priority entails that A does not perform acts with that feature in C.

What does it mean to say that a rule has higher priority than another rule? It means that there is a relevant system of higher-order rules conformity to which requires that the first rule (or the system of which it is a part) prevail. These rules can be represented by means of imperatives like: 'In C, act in accordance with the moral rule!' (but, again, that is not essential).

This sheds light on 'ought' all things considered (or rather some such judgments). Suppose that I ought morally to X, but prudentially to not-X. After some reflection, I judge that I ought to X, all things considered (henceforth I will call such judgments 'all-in "ought"s'). What is the content of this judgment?

If we want a unified semantics, we should say that this judgment is also about the requirements of rules. But it is implausible that the judgment is relative to the same moral system that is in conflict with the prudential one. If 'A ought morally to X' means that X is required of A by moral system M, then it would be surprising if (in this situation) 'A ought all-in to X' means that, all-in, X is required of A by moral system M. That would be to repeat the original moral judgment, this time with the redundant add-on that we have considered everything. (It may be possible for pragmatics to explain away the oddness. But there is at least some pressure to avoid the need for such an explanation.)

In this case, it is more plausible that the rules referred to are higher-order ones. When the speaker says that A ought all-in to X, s/he is not just indicating that s/he has considered all relevant things, but also that there is a system of rules conformity to which requires acting in accordance with the moral requirement.⁵

One may wonder whether this system was already in place prior to the speaker's reflection on the conflict at hand. And one may also wonder whether s/he was conscious of being committed to the system prior to this reflection. But I don't think that matters, so long as it is true that the speaker

⁵It is not ruled out in advance that the system contains rules of higher orders than the second. But this will be rare in practice.

is, at the time of speaking, committed to a system conformity to which requires that the moral 'ought' prevail.

I don't think this is *ad hoc*. In the situation described, the all-in 'ought' reflects a certain ordering of rules or ends which the speaker either arrives at on considering all things, or which was already in place before, and which reflection reveals to require that the moral rule prevail.⁶

Third: what determines the relevant standards? Stephen Finlay claims that 'A ought to X' means that A ought to X in order that E, where 'E' is an end which is determined by context [Finlay 2009]. In some contexts, the speaker's desires play a role in determining the end. But even in these cases, Finlay does not build those desires into the truth conditions of normative statements. The identity of the end (which is part of the truth conditions) is *determined* by the speaker's desires, but normative statements are not *about* anyone's desires. Rather, they are about acts and the ends to which the acts are instrumental.

I think this idea can be retained in the current theory of 'ought'. My answer to the third question, then, is this: the relevant standards are determined by context. Exactly how this works varies from case to case. In sincere moral 'ought's, the speaker's own commitments are involved in the identification of the standards. In inverted-commas uses, the speaker refers to standards s/he does not (necessarily) endorse. The same happens if the relevant standards are the legal requirements of a country. In some cases, we identify with people in order to advise them (as it were from their perspective). Here, the standards are (at least in part) determined by the advisee's commitments, or more precisely, by our beliefs about what those standards are. But although what the standards are determined by varies, attitudes or beliefs *themselves* never occur in the truth conditions of normative statements. Only the standards determined by them do.

Fourth: in what way do the standards determined by context figure in the truth conditions of an 'ought' statement? The standards can figure in the truth conditions in different ways: for example, their content can be described, or they can simply be referred to. I don't think it's plausible that the content of the rules is part of the truth conditions of 'ought' statements. If I say: 'You ought not to lie', I indicate that lying is forbidden by a system of rules, but I don't indicate that lying is forbidden by a system of rules part of the content of which is 'Don't lie!'.⁷

Another question about the way in which the standards figure in the truth conditions of 'ought' statements is to do with quantification over systems versus reference to a system. Take the following explication of 'ought' by David Wong, who claims that 'A ought morally to X' means that

By not doing X under actual conditions C, A will be breaking a rule of an adequate moral system applying to him or her.

[Wong 1984: 40]

⁶The above considerations apply to all-in 'ought's made in the light of conflicting requirements. But all-in 'oughts' are not always relative to higher-order rules. One can also consider what is all-in morally required. The 'all-in' indicates that we have considered all relevant features of the situation. But the 'ought' itself may be relative to a system of first-order moral rules (unless these rules themselves conflict). This shows that if any special interest accrues to all-in judgments, they are judgments made in the light of conflicting requirements.

In Wong's case, it is enough for an 'ought' judgment to be true that there is *some adequate system or other* which contains a rule which would be broken by A if s/he refrained from X-ing. So according to Wong, 'ought' judgments quantify over systems of rules. There is not a unique system of rules to which the speaker is referring. I disagree with Wong in this respect. It is more plausible that 'ought' judgments refer to a particular system of rules.

My standard-relational analysis of 'ought', then, is:

Ought: 'A ought to X in C' is true iff conformity to the contextually salient system of rules requires that A X-es in C.

Fully spelt out, this means:

Ought_{full}: 'A ought to X in C' is true iff X has a feature such that the contextually salient system of rules contains at least one rule which entails that A does acts with that feature in C and X has no feature such that a rule with the same or higher priority entails that A does not do acts with that feature in C.

On this theory, relativized 'ought' statements are treated as follows: 'A ought *morally* to X' means that conformity to the contextually salient system of moral rules requires X-ing. 'A ought *prudentially* to X' means that conformity to the contextually salient system of prudential rules requires X-ing.

To sum up: the standard-relational theory of 'ought' says that 'ought' (often) refers to a three-place relation between an agent, an act, and a system of rules. It is the relation of being required by. The system is referred to in an 'ought' statement, but its content is not described (although the system may, but needn't, be explicitly identified as moral or prudential). The kind of system referred to varies with the type of 'ought' in question. By 'type' of 'ought', I don't mean that the *word* 'ought' itself has variable meaning. The semantic value of 'ought' is always the relation of being required by. But the proposition expressed by 'A ought morally to X' is not the same as the proposition expressed by 'A ought prudentially to X' (which is what one would expect). Similarly, the proposition expressed by 'A ought to X all things considered' is (in our example) different from the proposition expressed by 'A ought morally to X' or 'A ought prudentially to X'.

The standard-relational theory of 'ought' has a number of virtues: first of all, it is unified for all non-epistemic uses of 'ought' (and may well allow extension to epistemic uses). Second, it has intuitive appeal, at least for lower-level 'ought's: it is plausible that 'A ought morally to X' is about the requirements of moral rules or standards. But since it is also plausible that the meaning of 'ought' is similar in moral, prudential, and other contexts, it is a virtue that the semantics assigns the same value to all such 'ought's (i.e. the relation of being required by). Third, it does not build the desires of either speakers or agents into the truth conditions of normative statements. Moral statements are about the requirements of rules, not about anyone's desires. Fourth: it allows us to explain why speakers are often motivated to act in accordance with their 'ought' judgments: these judgments often

concern rules to which the speaker is him/herself committed. And (in ordinary cases), such commitment will in part be a matter of motivational states. Fifth: as we will see, it is compatible with a theory of reasons which explains why 'A ought to X' entails that there are reasons for A to X, and also why 'A ought all-in to X' entails that the balance of reasons favours X-ing.

3. The Oughtistic Theory of Reasons

Recently, several authors have begun to work on theories of reasons in which (a) reasons are analysed in terms of 'ought' and (b) the strength of a reason is analysed in terms of 'ought' as well.⁷ I want to sketch a view of reasons with similar properties.

The standard-relational theory of 'ought' says that 'ought' statements are about the requirements of systems of rules. A system of rules requires X-ing if and only if conformity to the system requires X-ing. The proposal about reasons is that:

Reasons are features of acts indispensable in an explanation of why a relevant rule entails that the act is (not) done in C.

Reasons comprise both reasons for and reasons against acts. We can spell this out as follows:

Reason For: A reason for an agent A to X in C is a feature of X indispensable in an explanation of why a relevant rule entails that A X-es in C.

Reason Against: A reason for an agent A not to X in C is a feature of X indispensable in an explanation of why a relevant rule entails that A not-X-es in C.

The definition is restricted to *relevant* rules, which are rules that are contextually salient in the conversation or thought process at hand. This is necessary in order to forestall overgeneration: there is an infinite number of rules in logical space, but that does not mean that there are reasons for anything whatsoever.⁸

I have said that reasons are those features of acts indispensable in an explanation of why a relevant rule entails that the act is (not) done. I have not said that reasons are those features indispensable in an explanation of why *the system* requires the act. This is to allow for the possibility that one has reasons for acts which one ought not, on the whole, to do.

⁷Examples include Gert [2005] and Horty [2007] (or at least his view can be read along these lines). Theories of this type are intended as theories of *normative* as opposed to *motivating* reasons. Normative reasons count in favour of or against actions, whereas motivating reasons play a causal-explanatory role in the genesis of an agent's action.

⁸Or at least, there aren't reasons to do anything whatsoever relative to the system salient in the conversation. There is a sense in which there are reasons to do anything whatsoever. Like all reasons, these are relative reasons in the sense that they are reasons relative to certain possible rules.

The definition involves the notion of *indispensability*, which should be understood as follows. It may be thought that if a rule entails that A X-es, then X's being such that A is *capable* of it is part of an explanation of why the rule entails this. But it is not, intuitively, a reason for an act that one is capable of it (or at least not in most contexts).⁹ However, explanations are answers to why-questions and what counts as an adequate answer depends on what one is interested in. By (my) hypothesis, when one asks 'What reasons do I (or some agent) have to X?', one is interested in learning whether there are rules that apply to X (i.e. entail something about it). But when 'X' is an act, and one is interested in the reasons in favour of it, one is already presupposing that one (or some relevant agent) is capable of it. For reasons count in favour of things that can be done. This makes the answer 'You're capable of X' irrelevant and thus dispensable. Similarly for such answers as 'X has a spatio-temporal boundary'. Hence my definition of a reason as a feature of X *indispensable* in an explanation of why some relevant rule entails that the act is (not) done in C. (But I will sometimes abbreviate 'being indispensable in an explanation of why' to 'explaining why'.)

To illustrate the definition: if a salient rule is 'Don't lie in C!', then those features of an act which make it into a case of lying explain why the rule entails that the agent does not do it in C. And thus they are reasons to refrain from doing it. If the rule is 'Preserve life in C!', then those features of an act which make it life-preserving explain why the rule entails that the agent does it. And thus they are reasons for the action.

So my theory of reasons says that reasons for (against) an act are features of the act which explain why a relevant rule entails that the act is (not) done (in the situation). But what about the strength of reasons? I will say that a feature of an act *triggers* a rule (terminology I take from Horty [2007]) just in case the rule entails that we (not) do that act at least in part in virtue of that feature and the feature is a reason. I suggest that a feature is a stronger reason than some other feature just in case the rule which is triggered by the first has higher priority than the rule triggered by the second. For example: the fact that A will die unless I X is a stronger reason to X than the fact that B will be disappointed is a reason not to X, because the rule which tells us to preserve life is higher up in the ordering of rules than the rule which tells us to avoid disappointment.

So we have arrived at the following view of the strength of reasons:

Stronger Than: A feature (or features) F_1 is stronger reason for A to X in C than a feature (features) F_2 iff F_1 triggers a rule (or rules) higher in the order of priorities in C than the rule (rules) triggered by F_2 .^{10,11}

⁹This problem was raised by an anonymous referee for this journal.

¹⁰This formulation is compatible with a situation in which two (or more) features are individually weaker reasons than a third, though collectively stronger. They will be collectively stronger if and only if the contextually salient system of rules contains a higher-order rule which requires giving priority to both rules triggered by these features, and no rules which require giving either one of them priority.

¹¹This theory of the weight of reasons has a lot in common with the theory developed by Horty [2007]. But there are some differences. First, Horty is primarily interested in formal models of reasoning (about beliefs or action). He is not in the business of doing semantics. Second, Horty does not explain what makes it the case that a rule is higher up the order of priorities than another rule. Third, Horty only explains what makes an

This by itself does not tell us what makes it the case that a set of reasons is collectively stronger than another set of reasons. I propose to deal with that as follows:

Collective Weight: If, in C, a feature F_1 triggers a rule higher in the order of priorities than any other rule triggered in C, then the features that trigger rules prescribing the same act as the one triggered by F_1 are, in C, collectively stronger reasons than the features triggering rules prescribing incompatible acts lower in the order.

4. Explaining the Entailments

This brings me to the challenges I promised to discharge at the start of this paper. The first is to explain why 'A ought to X in C' entails that there are reasons for A to X in C (at least one reason). The second is to explain why 'A ought all-in to X in C' entails that the balance of reasons favours X in C.

The first challenge is met as follows: I have said that the truth of 'A ought to X in C' requires that the contextually salient system of rules is such that conformity to it requires that A X-es in C. In order for this to be true, there has to be at least one rule in the system which entails that A X-es in C. I have also said that a reason is a feature of an act which explains why a relevant rule entails that it is (not) done in C. But if a rule entails that A X-es in C, then X must have some feature in virtue of which that is entailed. So in all possible worlds in which A ought to X in C, there is a reason for A to X in C. Therefore: the standard-relational theory of 'ought' paired with oughtism about reasons explains why 'A ought to X in C' entails that there is at least one reason for A to X in C.¹²

We also want to show that 'A ought all-in to X in C' entails that the balance of reasons favours X in C (provided that the judgment is made in the light of conflicting requirements). I have argued that such all-in 'ought's are relative to a system of higher-order rules which determine the relative priorities among rules. If conformity to this system requires acting in accordance with the rule entailing that A X-es in C, then there must be a first-order rule which entails that A X-es in C. And so there is a feature of X which explains why this first-order rule entails that A X-es in C (there is a reason for A to X in C). But since we know (from *Ought_{full}*) that this rule is highest in the order of priorities (at least higher up than any other rule that bears on the situation), it follows that there can be no feature such that it explains why a rule higher or at the same level in the order of priorities

individual reason stronger than another individual reason. He does not explain what makes several reasons collectively weightier than other reasons.

¹²One worry: if explanation is interest-relative, can it be necessary that X has some feature which satisfies the interest? Can it be ruled out that there will only be dispensable features that explain why the rule entails that A ought to X? Yes it can, because if otherwise dispensable features are the only ones that play a role in the explanation, the rules themselves must be concerned with those features. But when that is so, and if one is interested in learning whether any rules apply to the action, then it is in the relevant sense explanatory to show that the act has those features.

entails that A not-X-es in C. Given *Collective Weight*, this means that the collective weight of the reasons in favour of X is greater than the collective weight of the reasons in favour of not-X. Therefore the standard-relational theory of ‘ought’ paired with oughtism about reasons explains why ‘A ought all-in to X in C’ entails that the balance of reasons favours that A X-es in C.

5. Some Questions and Objections

Many questions and objections can be raised about the theory of ‘ought’ and reasons described above. I will attempt to answer some of them in the rest of this paper.

5.1 Gert on the Balance of Reasons

Joshua Gert [2007] believes that reasons (can) play two distinct roles in determining the rational status of an action. The first is the requiring role: reasons can make an act rationally required (i.e. such that one ought to do it). The second is the justifying role: reasons can make otherwise irrational acts permissible (i.e. not irrational), without requiring one to do them. For example, harms to oneself ordinarily have requiring strength (they require one to prevent them). Altruistic reasons (preventing harms to others) needn’t have requiring strength, although they can justify incurring harm oneself. Gert believes two types of strength correspond to both roles: requiring strength and justifying strength [2007: 537–8].

If Gert is right, then a theory of reasons should be able to make sense of these two roles. My theory can do this as follows:¹³ the requiring strength of a feature (reason) in C is measured by the number of rules that the rule triggered by it trumps in C. If, in C, a rule is at the same position in the hierarchy as another rule, I will say that these rules are *equalized* in C. The justifying strength of a feature in C can then be thought of in terms of the number of rules that the rule triggered by the feature equalizes in C.

Gert also claims that requiring strength and justifying strength can come apart. For example, a reason can justify acts without requiring one to do them. This is also possible on my theory, since the number of rules trumped by a rule triggered by a feature may be higher or lower than the number of rules equalized by that rule.

Gert is sceptical about metaphors like ‘the balance of reasons’ and ‘the collective weight of reasons’ in so far as they suggest that there is only one kind of strength. But he does not dispute that reasons with different kinds of strength still collectively determine what one ought to do. In other words: what one ought to do is a function of the reasons pro and con. If so, we should be able to make sense of their collective force or what they require on balance.

¹³The following is an indication of how my theory would deal with these distinctions. It is not meant to capture all the details. Thanks to an anonymous referee for this journal for forcing me to consider this.

If, in a context, a reason with justifying strength cancels a requirement, there is a sense in which the reason corresponding to the latter is not stronger than the former. It is this relation between reasons that *Stronger Than* intends to capture. Some such relation should exist if what one ought to do is a function of the reasons pro and con.

My theory, then, not only makes sense of the notions of requiring strength and justifying strength, but also tells us how to determine what one ought to do in the light of different reasons (which may exemplify both types of strength): one has to check whether any of the reasons trigger rules higher in the order of priorities than any other rules. If so, then one ought to do the act which is uniquely favoured by those reasons (if any). This is intended to capture the idea that if one ought (all-in) to X, then the reasons favour X on balance.¹⁴

5.2 Reasons for Permissible Acts

Someone might object to oughtism as follows: some acts are permissible, but neither forbidden nor required. Some of those are the acts which the rules are silent about. But can't we have reasons for such acts? So perhaps we should say that reasons are those features of acts which either explain why a rule entails that the act is (not) done, or why it is neither the case that it is done nor not done. This is essentially John Broome's [2004] definition of reasons.

I think Broome's definition overgenerates, however.¹⁵ Suppose that a trivial act is an act which has no important consequences whatsoever (such as rubbing one's chin). Further suppose that the rules say that trivial acts are neither required nor forbidden. In that case, the fact that an act is trivial plays a role in an explanation of why the act is neither required nor forbidden. But it is not intuitively a reason for or against the act.

So I prefer my definition of a reason. It's true that oughtism denies that one could have reasons for acts which are permissible in virtue of the rules' silence about them (although it allows for reasons for acts which are permissible in virtue of a tie between the rules). Is this a problem?

Philosophers who think it is believe that reasons do not always ground requirements. In other words, they deny the *Reason-Ought Link*:

¹⁴Here is another way of putting it: Gert does not dispute that acts which are rationally required are such that one ought to do them. Neither does he dispute that if an act is required then it is impermissible not to do it. So if X is such that one ought to do it, we can infer at least three things: (1) that there is a set of reasons with requiring strength favouring X, (2) that there is no set of reasons with equal or more requiring strength disfavouring X and (3) that there is no set of reasons with justifying strength which makes not-X permissible. If conditions (1) to (3) are satisfied, then one ought to X, and *vice versa*: if one ought to X, then conditions (1) to (3) are satisfied. When I say that the balance of reasons favours X, I mean that these three conditions are satisfied. A theory of 'ought' and reasons should explain why this is entailed by 'A ought all-in to X' and *vice versa*. The standard-relational theory of 'ought' paired with oughtism about reasons enables us to do so, although my presentation in this paper does not put it in these terms.

¹⁵I owe the example in this paragraph to Ralph Wedgwood.

ROL: if one has a reason to X and no reason to do anything else, then one ought to X.

Oughtism says that reasons for A to X in C are features of X which explain why a relevant rule entails that A X-es in C. It follows from this view (and the standard-relational theory of ‘ought’) that if one has a reason to X and no reasons to do anything else, then one ought to X (this roughly follows from the fact that in such a situation, there is some rule requiring X and no rule requiring not-X). So my theories entail *ROL*.¹⁶

One argument against *ROL* is by Jonathan Dancy [2004]. He claims that I may have a reason to buy an ice cream (and no reason to do anything else), but that it’s odd to say that I *ought* to buy the ice cream. This, I take it, is a semantic observation. However, as argued by Broome [2004], the impression that there is something odd about saying that one ought to buy the ice cream is plausibly due to the impression that ‘ought’ is a very serious word, a word with moral connotations. But ‘ought’ isn’t intrinsically moral (or even serious). It may be that you ought to buy a pair of shoes. Some things may be required to do yourself a favour, even if no one will blame you if you don’t.

Another argument against *ROL* derives from Gert [2007]. As we’ve seen, Gert distinguishes between justifying and requiring strength (of reasons). Although my theory makes sense of the distinction conceived in terms of making permissible (justifying) versus grounding ‘ought’s (requiring), it does not cater for reasons that *merely* justify (i.e. don’t require in any context, including ones where only one reason applies).

I don’t think that is clearly wrong, however. The kind of cases that lend plausibility to Gert’s distinction are cases where an act involves some cost to oneself. Such acts are normally prohibited, but the prohibition can be lifted by benefits to others. Such benefits make it permissible to incur the cost without requiring us to do so. But now imagine there to be no cost *at all* (which is actually impossible given our limited natures). Is it clear that we are not required to prevent a harm to others even then? I don’t find that obvious (in fact I find it wrong). So we cannot be confident that there are any reasons with only justifying strength.

5.3 *Wrong Order of Explanation*

It seems unlikely that one could explain the greater weight of a feature in terms of a higher-order rule which arbitrates between (potentially) conflicting requirements. Doesn’t this higher-order rule itself depend on prior insight into the relative significance of the conflicting features?

¹⁶They share this entailment with a popular view which says that if the balance of reasons favours X, then one ought to X. If the balance of reasons favours X, then X is required and not-X impermissible. This kind of view can allow for permissible acts for which one has reasons, but only if the balance of reasons does not favour one particular act. So this view also entails that if one has a reason to X and no reason to do anything else, then one ought to X. In other words, this view also entails that what reasons do is, most fundamentally, to require.

But, if so, then I haven't explained what the weight of a reason consists in.¹⁷

The answer to this objection emerges when we remind ourselves that I am engaged in a particular project. It is the project of doing normative semantics on the assumption that there are no objective normative facts. Given normative antirealism, the normative significance of a feature derives most fundamentally from the stance taken towards it by a subject (at least in cases where the identification of the relevant standards proceeds via the commitments of the subject¹⁸). For example, the fact that A cares about harm to others is what gives such harm normative significance for A. This is a metaphysical claim: a claim about the generation of normative significance. It entails that reasons are always (in a sense) derivative: features of the world matter (count as reasons) only because we care about some things. But if caring to some degree is metaphysically rock-bottom in the generation of normative significance, caring to different degrees is metaphysically rock-bottom in an explanation of relative significance. These are metaphysical claims. My suggestion is that these phenomena find semantic expression in claims about higher-order rules and their requirements.

The idea is this: the fact that F_1 is weightier than F_2 is a reflection of one's preference for F_1 over F_2 in case of conflict (which preference is rock-bottom¹⁹). This preference explains one's commitment to a rule which requires that standards furthering F_1 take priority. Statements about relative weight can then be understood as statements about the priorities determined by such rules.

5.4 Complicated Rules

The next set of questions I want to discuss concerns the nature of the rules to which I claim speakers are referring with 'ought' statements. The questions I want to focus on are relevant mainly when the relevant rules are ones to which the speaker is him/herself committed.

(1) Most plausible rules are rules that hold only *ceteris paribus*. For example, it is implausible that it is always bad to lie, and most people know that. So hardly anyone is committed to the rule 'Don't lie!'. Most people are committed to something like: 'Don't lie, other things being equal!'. But that rule does not entail that one does (not) lie if an act is an instance of lying. It only does so if other things are equal. However, spelling out what that amounts to makes the rule impossibly complicated. And it is implausible that anyone is committed to such a complicated rule. So a view where 'ought' is related to what such a rule entails is implausible.

¹⁷This objection was pressed by an anonymous referee for this journal.

¹⁸In cases where this is not so, the higher-order standards are determined by the subject's beliefs about the standards in force. These beliefs may themselves be based on information about what tends to take priority (e.g. in court rulings or etiquette).

¹⁹Or at least some such preferences will be rock-bottom for antirealists. Others may be derivative (depend on preferences which are themselves rock-bottom).

This type of objection is based on the observation that most people allow exceptions to rules. However, in order to allow for that, it is not necessary to build the exceptions into the content of the rules themselves. It suffices if we have a structure which explains why certain rules (which can themselves be simple) sometimes are to be heeded and sometimes not. And I have offered just such a structure: a rule wins out in a context just in case it is (in that context) higher in the order of priorities than any other (relevant) rule with inconsistent implications. In other contexts it may be lower than another rule, in which case we allow for an exception.

(2) My answer to the first question naturally raises the second: does my view not commit me to the claim that people tend to be committed to impossibly complicated structures of first-order and higher-order rules? Even if the first-order rules are simple, the second-order ones are not. There are indefinitely many situations in which different rules are higher up the order of priorities. But is anyone committed to higher-order rules which legislate for all such situations in advance? It seems that such a rule would have to be very complicated: ‘Let the moral system prevail over the prudential system, except if a, or b, or c, or ...’.

I agree that if the higher-order rules had to legislate for all possible situations in advance, then they would have to be very complicated. But there is no reason why this should be so. Take (again) a case where a moral rule (or system of rules) which requires A to X conflicts with a prudential rule which requires A to not-X. Now imagine that a speaker makes the judgment that A ought, all things considered, to X (i.e. respect the moral requirement). It suffices for this judgment to be true that the speaker is, at the time of speaking, committed to a rule (subscribes to an imperative) which entails that the moral rule prevail in this context. But the rule needn’t say anything more general than that.

5.5 *Normativity*

Another problem for the theory is that of normativity [Boghossian 2006: 24]: saying that X is required by some system does not seem normative. But ‘A ought to X’ does.

In order to respond to this problem adequately, we need to know what is meant by ‘normative’. What makes a statement normative? A plausible and popular answer is that a statement is normative if and only if it entails something about reasons (in this case, ‘A ought to X’ is normative if and only if it entails that there are (sufficient) reasons for A to X). But on my theory, ‘ought’ statements are normative in this sense. For I have argued that the standard-relational theory and oughtism about reasons explain why ‘A ought to X’ entails not only that there are reasons for A to X, but also that the balance of reasons favours X-ing.

If there is a problem of normativity for my theory, it must stem from something else. Some philosophers believe that normative statements cannot be reduced to non-normative statements. But what is meant by

'non-normative statement'? Presumably, a successful analysis of 'ought' statements in terms of different vocabulary makes that vocabulary itself normative, at least as used in the analysis. This must be so if the 'ought' statement is itself normative and the analysis successful.²⁰

However, anti-reductionists think that *no* analysis of claims involving words like 'ought' and/or 'reasons' can work precisely *because* the meaning of these claims has a feature, normativity, which is lacked by the meaning of claims couched in other vocabulary. However, such claims are question-begging unless the feature can be identified in some way other than by saying that the analysis does not contain words like 'ought' and 'reason'. So I believe that, in this context, the burden of proof is on the opposition.

5.6 Disagreement

The last problem I will discuss is that of disagreement. The standard-relational theory does not always allow inconsistent propositions to be expressed by speakers who respectively say 'A ought to X' and 'A ought not to X'. If the first statement refers to a system of rules accepted by one speaker, and the second to a different system accepted by the other, they are not inconsistent. The speakers merely state facts about the requirements of different systems. This problem is sometimes called the problem of disagreement: on the face of it, 'A ought to X' and 'A ought not to X' are incompatible. But the kind of view I have sketched does not always make it so.

Notice that my view does allow some genuine disagreement. Because I don't use possessive pronouns in representing the meaning of 'ought' statements, two speakers referring to the same system genuinely disagree if the one says that A ought to X and the other denies it.²¹ But even in cases where speakers do not express inconsistent propositions, it is not hard to see (1) how the appearance of inconsistency may arise or (2) why the (verbal) disagreement does not cease.

The appearance of inconsistency may arise from the supposition that one's interlocutor is referring to similar rules as oneself. It is hard to be certain that people do not share our values because normative disagreement is often explained by differences in non-normative beliefs. There is much overlap in values [Finlay 2008: 356]. But even if we are certain of a difference in rules, we can still explain why we might insist that our interlocutor is wrong. Björnsson and Finlay [unpublished ms] plausibly claim that in (some) normative matters our interest in truth is secondary to our interest in the promotion or protection of values. This explains why, in normative debates, the proposition assessed for truth or falsity is not always the one expressed by my interlocutor. It is sometimes one closely related to it. When

²⁰I am assuming that the reduction in question is analytic. Thanks to an anonymous referee for this journal.

²¹The same is true if speakers subscribe to many rules different to those accepted by the other, but (1) only a subset of the rules apply to the case and (2) these rules have the same content. In such cases, it makes sense to construe the speakers' statements as involving reference to the same subsystem of rules.

John says ‘A ought morally to X’ I don’t check whether X-ing is required by John’s standards (even if the truth conditions of the statement do involve them). Rather, I check whether it is required by mine. The explanation for this can be found in the *point* of moral discourse, which is not entirely theoretical:

Part of what it is to engage in a practice of *moral* discourse is to subscribe to a particular standard, to the exclusion of any rival, as determinative of what to do—both for one’s own conduct and others’. Here too there is a context which for us is privileged, consisting in our own standards. The truth of ought-propositions relativized to *other* standards is irrelevant to this conversational interest. To address the question of whether [John] was right to think that [X-ing] is required [of A] by [John’s] standard ... would therefore be a perverse fixation on truth for truth’s sake, in neglect of what is relevant to the purpose of moral thought and discourse.

[Björnsson and Finlay, unpublished ms]

6. Conclusion

Many more objections ought to be addressed, but space does not permit. I have outlined a standard-relational theory of ‘ought’ and an oughtistic theory of reasons. The combination of these theories meets the explanatory challenges stated at the start of this paper. The theory of ‘ought’ allows many normative judgments to be true even if there are no objective normative facts. Oughtism about reasons explains why ‘A ought to X’ judgments entail that there are reasons for A to X. Furthermore, it explains why ‘A ought all-in to X’ entails that the balance of reasons favours X-ing. I take the ability of the theories to explain these (often neglected) phenomena to be a major point in their favour. It lends further plausibility to the ancient idea that normative judgments are relative to standards.²²

University of Oxford

Received: June 2009

Revised: December 2009

REFERENCES

- Björnsson, G. and S. Finlay unpublished ms. Metaethical Contextualism Defended, URL=<<http://www-ref.usc.edu/~finlay/MCD.pdf>>.
- Boghossian, P. 2006. What is Relativism?, in *Truth and Realism*, ed. P. Greenough and M. Lynch, Oxford: Clarendon Press: 13–37.
- Broome, J. 2004. Reasons, in *Reason and Value: Essays on the Moral Philosophy of Joseph Raz*, ed. R. Jay Wallace, M. Smith, S. Scheffler and P. Pettit, Oxford: Oxford University Press: 28–55.
- Dancy, J. 2004. Enticing Reasons, in *Reason and Value: Essays on the Moral Philosophy of Joseph Raz*, ed. R. Jay Wallace, M. Smith, S. Scheffler and P. Pettit, Oxford: Oxford University Press: 91–118.

²²I would especially like to thank Ralph Wedgwood for extremely helpful comments and suggestions. I have also benefited from commentary by Natalja Deng, Barbara Vetter, Stephen Finlay, Herman Philipse and anonymous referees for this journal. Thank you all.

- Finlay, S. 2008. The Error in the Error Theory, *Australasian Journal of Philosophy* 86/3: 347–69.
- Finlay, S. 2009. Oughts and Ends, *Philosophical Studies* 143/3: 315–40.
- Gert, J. 2005. A Functional Role Analysis of Reasons, *Philosophical Studies* 124/3: 353–78.
- Gert, J. 2007. Normative Strength and the Balance of Reasons, *The Philosophical Review* 116/4: 533–62.
- Harman, G. 1975. Moral Relativism Defended, *The Philosophical Review* 84/1: 3–22.
- Harman, G. and J. Thomson 1996. *Moral Relativism and Moral Objectivity*, Oxford: Blackwell.
- Horty, J. 2007. Reasons as Defaults, *Philosophers' Imprint* 7/3: 1–28.
- Mackie, J. 1977. *Ethics: Inventing Right and Wrong*, Harmondsworth: Penguin.
- Wong, D. 1984. *Moral Relativity*, Berkeley: University of California Press.