Review Article

Thinking about Normativity: Ralph Wedgwood on ‘Ought’

Wim de Muijnck

Radboud Universiteit Nijmegen, Faculteit der Sociale Wetenschappen
Postbus 9104, 6500 HE Nijmegen, The Netherlands
w.demuijnck@planet.nl

Abstract

Ralph Wedgwood's *The Nature of Normativity* provides a theory about the semantics, metaphysics, and epistemology of normative judgments, taken to be judgments of the form ‘I ought to ϕ ’. The theory is based on the principle of Normative Judgment Internalism, and the principle that ‘the intentional is normative’. I argue, first, that by being merely about oughts, Wedgwood's account leaves out one essential constituent of normativity: value. Secondly, I argue that mainly because of this, the account faces a serious issue of relevance.

Keywords

agency, metaphysics, normativity, oughts, rationality, value

R. Wedgwood, *The Nature of Normativity* (Oxford: Clarendon Press, 2007), 296 pp. ISBN 9780199251315 (hbk). Hardback: £37.00.

Introduction

Normativity is a broad subject: it covers the wide field of issues connected with *reasons*. We might divide this field into three domains. As it is discussed in the philosophical literature, normativity includes the domain of ‘ought’ and ‘should’: of duties and obligations, prohibitions, and permissions. In brief, normativity includes *deontology*.¹ It also includes *rationality*, the consistency among our beliefs, desires, feelings, goals, actions, and identity. And as these

¹ I use this term for the sake of readability. ‘Deontology’ does not here refer to deontological theories in normative ethics, but is rather to read as ‘the domain of the deontic’.

latter things do not occur out of the blue, normativity also includes the domain of *value*, of what is good or bad, important or unimportant, to be sought or avoided—the domain of needs, goals, preferences, and interests.

These three domains—deontology, rationality, and value—are closely interconnected. They are all aspects of the classic ‘is-ought’ distinction’s right side. None can be conceived in abstraction of the others. But because nonetheless deontology, rationality, and value are not entirely the same thing, it is possible to approach normativity in different ways. This, among other things, is why normativity is a difficult subject. There are three big issues about normativity in general. First, what do normative concepts or expressions refer to? Secondly, are there normative facts of the matter? And thirdly, how could we know about such facts?

Many philosophical subdisciplines have to come to terms with normativity, some way or other: metaphysics, the philosophy of biology, the philosophy of mind and action, normative ethics, meta-ethics, aesthetics, and all the different branches of philosophy on social, political and cultural issues. Normativity is not a side issue: much of the terminology needed to even mention a subject in any of these fields is normative in character. Our reality, we might say, is deeply normative.

Because of this, we can assume that mistaken assumptions about normativity in any of its aspects will hamper progress in any of the fields just mentioned. A book called *The Nature of Normativity*, clarifying the semantics of normative discourse, the metaphysics of normative facts, and the epistemology of normative belief, could thus be of great benefit to many philosophical debates at once.

Ralph Wedgwood has written such a book. At its core is an impressive account of the semantics, metaphysics, and epistemology of ‘ought’. It is an inventive, yet carefully constructed philosophical edifice, the ripened fruit of exchanges with an impressive list of great minds (as a look at the Preface makes clear). It addresses issues across diverse philosophical subdisciplines. It presents meticulous and intricate arguments that are sensitive to subtleties and counterarguments at every turn. The writing is compact and highly abstract, but is also elegant and readable.

Although naturalistic thinking about normativity (Aristotelian, Humean, Darwinian, or otherwise) seems on the rise, Wedgwood’s approach is overtly rationalistic. Wedgwood assumes that minded beings such as us are essentially rational. He also assumes that the reasons in our minds govern our actions. Given these connections among mindedness, rationality, and agency, which are conceptual as well as constitutive, ‘ought’ is interpreted by Wedgwood as that which is rational. The approach is realistic about oughts: if there is

an abstract domain of *a priori* rational truths, and if oughts belong to that domain, then they must be real.

I have indicated above of what great use a general theory of normativity could be. Does Wedgwood give us what we need? I have a worry. It is kindled by the book's focus, which is on deontology and rationality, but not on value. My worry is that because of this, the results the book achieves face a serious issue of relevance. The results, I wish to argue, will be relevant to debates in semantics, logic, epistemology and metaphysics, but they will leave empty-handed the philosophers in the fields mentioned earlier, including ethics and social and political philosophy—the fields where normativity is most pertinent and where a theory of normativity could be the most helpful. I will try to explain my worry, below.

The Main Argument

Wedgwood's aim in *The Nature of Normativity* is a 'substantive philosophical explanation' (p. 4) of normative thought, truth, and knowledge. The book has three parts. Part I is about the semantics of normative concepts and terms, Part II is about the metaphysics of normative facts, and Part III is about the epistemology of normative beliefs.

The account is a defence of normative realism: against non-cognitivist views such as expressivism, Wedgwood will argue that normative judgments can be literally true or false, hence that there are normative facts of the matter that can be known and be referred to. It is also nonreductionistic: Wedgwood's account will itself rely on normative terminology.

There is a special reason for this. One of Wedgwood's key principles is that 'the intentional is normative' (p. 2): normativity and intentionality are so closely connected—conceptually, ontologically, and epistemically—that any philosophical explanation of the former is only possible in terms of the latter, and *vice versa*. Facts concerning 'ought'—and these are what Wedgwood's account is all about—*may* be noncircularly and nontrivially explained in terms of something else—but this something will be intentionality. And, as Wedgwood argues, intentionality in this sense cannot be a matter of straightforward 'purely psychological predicates' (p. 172). The explanation thus cannot be reductive: the required intentional terms in the metalanguage, such as 'reasoning', 'inference' and 'deliberation', are themselves essentially normative in character.

The other key principle in Wedgwood's argument, closely related to 'the intentional is normative' principle, is Normative Judgment Internalism (NJI). This is the idea that first-person normative judgments are motivating in and

of themselves. NJI says that in a normal case and insofar as the agent is rational, if the agent judges: ‘I ought to φ ’, then she is necessarily disposed to φ . Or at least, she is disposed to intend to φ . Or even more precisely, she is disposed to have ‘fully correct plans’ for what to do at a time t that logically imply p . NJI is limited to situations ‘with no relevant uncertainty’ (p. 31) and where what the agent does or does not do is ‘manifestly dependent on intention’ (p. 30).

Wedgwood initially presents this principle as a ‘compelling’ intuition (p. 11) that every account of normativity should be able to explain. This is why in Chapter 3 a ‘conceptual analysis’ approach for normative statements is rejected: even if intuitively plausible *non-normative* necessary and sufficient conditions could be specified for the truth of such statements, these would not explain NJI. Later on (chapter 4), Wedgwood uses this very principle to define the essence of ‘ought’. ‘Ought’ will then turn out to be the concept that bridges the gap between mere mental states and actions, at least those of a rational agent in normal circumstances. Thus, ‘ought’ will be just the concept that plays the regulative role in practical reasoning that is demanded by NJI.

Given the central position of NJI in Wedgwood’s argument, it is no coincidence that the entire book is about the concept of ‘ought’. In Wedgwood’s account, normativity consists mainly in what ought to be done or ought to be the case. The argument in the book is thus centred around oughts and NJI, and rationality figures as an essential precondition.

Chapters 2 and 3 are about semantical issues. Wedgwood argues here that a semantics for normative statements will have to be truth-conditional. But a truth-conditional semantics for such statements will have to be nonreductionistic, i.e., be couched in normative terms, otherwise it will not be able to explain the essential regulative role in practical reasoning that NJI entails for ought-judgments.

Is there a truth-conditional semantics of normative statements that *can* explain NJI? Yes, Wedgwood claims, and this will have to be a ‘conceptual role semantics’. Such a semantics specifies the ‘basic rules of rationality’ (p. 82) for using a concept—in this case the concept of the so-called ‘practical ought’. For it is precisely these rules that make the ‘practical ought’ the concept it is. And as these rules specify that in the deliberation of a rational agent a judgment ‘I ought to φ ’ disposes the agent to intend to φ , a conceptual role semantics for ‘ought’ entails NJI.

In more detail: Wedgwood’s conceptual role semantics for ‘ought’ tells us that ‘for any proposition p , “ $O_{\langle A, t \rangle}(p)$ ” is true just in case there are fully correct plans for A to have about what to do at t that logically entail p , and no such

fully correct plan that logically entails the negation of p ' (p. 102). 'Ought' is here taken to be a propositional operator that indexes a proposition p to an agent A and a time t .

In chapter 6, the argument changes from semantics to metaphysics. Wedgwood aims to give a 'constitutive' (p. 136) account of normative properties, but not a reduction. Such an account must take the form of biconditionals that state necessary and sufficient conditions for the exemplification of normative properties, and these conditions should state the essences of the properties. But it will not be possible, Wedgwood claims, to state these conditions non-circularly, i.e., entirely in non-normative terms.

The reason for this is given in chapters 7 and 8. In chapter 7, Wedgwood argues for the principle that 'the intentional is normative' (p. 158): the essence of intentional states can only be specified in normative terms, while oughts can only be specified in terms of intentional states. Thus, 'the normative and the intentional must each be mentioned in giving an account of the other' (p. 163). Chapter 8 offers additional, but related, arguments against reductive naturalism, which in this case consists in forms of functionalism.

Wedgwood's metaphysical argument takes a route that is different from that of the argument about semantics, but note that there is a close connection. In the case of semantics explaining NJI, i.e., conceptually linking up 'ought' with motivation, was the central issue; in the case of metaphysics, it is the normative essence of the intentional. But as the intentional involves practical reasoning and motivation, we are again looking at the linkage between 'ought' and motivation: formerly a conceptual, now a constitutive linkage.

It was already argued that an ought, indexing a proposition p to an agent A and a time t , consists in there being fully correct plans for A about what to do at t that imply p . The question must now be answered what it means for any 'plans' (which are here understood as intentional states) to be 'correct'. Wedgwood's answer is that it is essential to these states that they are causally regulated by standards of rationality. Just as the standards of rational *belief* must be oriented towards the goal of having correct beliefs, the standards of rational *choice* must be oriented towards the goal of choosing correctly.

The next question is what it is for someone to be capable of intentional states. Wedgwood's answer is that it is someone's disposition to 'conform to the principles of rationality' (p. 164). But could such dispositions not be specified without mentioning normative properties? Wedgwood argues that any rational disposition (from stimulus conditions to a reaction) can be 'defeated' (p. 170), which means that a disposition is rational only when 'defeating conditions' are absent. What unites defeating conditions *in* being defeaters is precisely the rationality that is defeated, which is a normative relation.

Wedgwood's book is much richer than the above can convey. There are excursions into subjects such as deontic logic (chapter 5), mental causation (chapter 7), and modal logic and supervenience (chapter 9). Ideas developed in the book include an account of 'essence' and a corresponding account of nonreductive naturalism in chapters 6 and 8; a dispositionalist account of mental states in chapter 7 and section 8.4, where the notion of a 'rational disposition' is developed; a version of non-S5 modal logic and a corresponding way of understanding supervenience in chapter 9; and finally, an intuitionist account of *a priori* normative knowledge in Part III. The book would not be easy to summarize in its entirety. But I hope that what I have said above gives a fair impression of Wedgwood's main line of argument.

Oughts and Values

Earlier on, I distinguished three aspects usually associated with normativity: deontology, rationality, and value. From the above, we can learn that Wedgwood's account captures deontology and rationality, but not value. The introduction of Normative Judgment Internalism (NJI) on pages 23–34 is a crucial moment: normative judgments are here declared to be ought-judgments, i.e., judgments about what one ought to do, all things considered. Normativity thus becomes 'oughtness', rather than 'to-be-consideredness'. The reasons that enter into our practical thinking to generate ought-judgments are left out of the picture. Because this is the basis of my claim that Wedgwood's account faces an issue of relevance, let me explain it in some detail.

Consider that people act for reasons that can be diverse. We could say that these reasons concern different types of values that generate corresponding types of reasons. They can be moral or prudential, but also hedonic, conventional, legal, aesthetic, religious, and so on. In practical reasoning people consider reasons, and reach, if they can, a verdict on what they ought to do. There is much of a normative character already present before an agent has reached any conclusion about what she, all things considered, ought to do: needs and interests, values and value judgments, *a priori* duties, rules and conventions, and so on. These things will also be present when no such conclusion can be reached—in situations of what Wedgwood calls 'relevant uncertainty' (p. 31).

Together, they are the input of a process of practical reasoning, while an ought-judgment is the output. Ought I to buy this painting? Well, it is beautiful. But it is also stolen. Aesthetic, moral, prudential and legal considerations come to bear, and they all give me reasons for or against buying. Perhaps

I ought to buy the painting, perhaps I ought not to do it, or perhaps no such conclusion is warranted. But *if* there is an ought, it will depend on reasons—while these reasons do not, in turn, depend on this ought. There may not even be any ought. Note that even if I cannot, out of uncertainty, decide what I ought to do, I might still decide to act, and act for a reason. For instance, I might decide to buy the painting because it is beautiful, even if I am not sure that buying it is what, all things considered, I ought to do. So acting for a reason is not the same thing as acting on the basis of an ought-judgment.

Thus, in order to reach an ought-judgment, reasons of different types have to be weighed, typically involving numerous normatively significant facts. On Wedgwood's own view, oughts are a type of entity with the character of all-things-considered-for-an-agent-at-a-time. Mere 'evaluative' properties like 'nice', 'ugly', 'legal', 'just', 'appropriate', 'sacred', etc. are the input, and oughts the output, of practical reasoning. We might say that the former are the raw materials from which oughts are constructed. What Wedgwood's account ignores are these raw materials.

Admitting that it is reasons that generate oughts, one might claim that normativity-proper is a domain of rational oughts, but not of values. But note how implausible this would be. First, all normative concepts are evaluative. Saying that one ought not to steal, or that moving a rook diagonally is not permitted is, among other things, attributing a negative value to stealing, or to going diagonally with a rook. The converse also seems to hold: from positive or negative value follow *prima facie* reasons for behaviour.

Secondly, values are arguably the more basic category within the normative domain. For values generate reasons, which generate oughts. And while values are the basis for ought-judgments, the converse does not seem to hold. Think, for instance, of a useless rule-governed procedure: it will ring hollow to say that the rules ought to be followed because they are rules, hence prescriptive. Even if we do see a type of 'oughtness' here, this will only generate value for those who value the following of rules just because they are rules.

Nowhere in his book does Wedgwood suggest that the domain of value *should* be ignored. Chapter 6, for instance, is not primarily on oughts but on, indeed, value, rightness, and moral wrongness. Wedgwood does acknowledge such properties, be it only in passing. On page 68, he does distinguish evaluative terms, like 'good', 'desirable' and 'valuable', from "normative terms like 'ought'", but he stresses that the two are closely related.

Be this as it may, Wedgwood does restrict his main argument to oughts—to the couple of deontology and rationality. The point is not that the questions about 'ought' that he chooses to deal with cannot be properly addressed without addressing value. It seems clear that they can. But these questions, in

bypassing the domain of value, lead to a theory of normativity that has hardly any connections to the real natural world. Wedgwood is mainly occupied with essences and necessities. His way of theorizing is to forge tight conceptual and constitutive linkages between mind, theoretical and practical rationality, motivation, and ‘ought’. There is no room here for value. True, values are linked, conceptually and constitutively, with oughts by being reasons. But values will only become reasons for an agent in virtue of his being situated in a natural world.

This is why studying value requires engaging with the discourses of common sense and empirical research. And this is clearly not Wedgwood’s priority. When he does address the way his abstract normativity relates to the natural world, what he offers is a purely abstract theory of supervenience. Although he claims that normativity and intentionality are ‘metaphysically irreducible’ (p. 6), he also argues at some length (chapters 8 and 9) that his account is compatible with a moderate form of ‘metaphysical naturalism’ (p. 3). Normative facts are said to be ‘realized in’ (p. 6) natural facts and causally efficacious. But none of this tells us much about the interplay of reasons motivating ordinary agents.

Relevance

I have argued that Wedgwood’s theory of normativity, by focusing on rational oughts, abstracts from the domain of value. But why should this be a problem? If we deplore Wedgwood’s exclusive focus on ‘ought’, he still explores it thoroughly and in many inventive ways. *The Nature of Normativity* is impressive as an ‘intersubdisciplinary’ (p. vii) *tour de force*. Wedgwood points out that ‘the goals of this book require drawing on ideas from many different branches of philosophy’ (p. vii): metaphysics, epistemology, philosophy of language, philosophy of mind, ethics, and logic. The result of all this contributes to different philosophical fields in its turn.

Then, why worry about relevance? Let me mention two general considerations, and substantiate them by discussing four relevance-related themes. My first consideration is that Wedgwood’s main results may be relevant to debates in the philosophical subdisciplines just mentioned, but that they have little to offer to precisely those fields and debates in philosophy where normativity is an issue. Put differently: Wedgwood contributes to the very philosophical subdisciplines that he draws upon, but he does not have much to offer outside of these subdisciplines. In the Introduction, I mentioned as relevant fields: ‘metaphysics, the philosophy of biology, the philosophy of mind and action,

normative ethics, meta-ethics, aesthetics, and all the different branches of philosophy on social, political and cultural issues'. Of these, only metaphysics and philosophy of mind may profit from *The Nature of Normativity*. In the Conclusion, Wedgwood explores the implications of his results for the 'philosophy of mind and language' (p. 267), the 'theory of rational belief and rational choice' (p. 271), 'first-order ethical theory' (p. 273), and 'the philosophy of religion' (p. 277). Note that considering the importance of normativity as a subject this is a rather limited list, and in the first two cases Wedgwood does little more than recapitulate his conclusions. We will also see below that he has in fact very little to offer to ethics.²

My second consideration is that Wedgwood's rationalistic approach generates a conception of agency that has little to do with real-world agents. Illustrative of Wedgwood's approach is one of the two paragraphs in the entire book that refers to actual empirical psychology (p. 270),³ where he points out that there need not be a conflict between his own approach and empirical research about the contingent features of mental states and capacities. There is no hint that his philosophical approach could *inform*, or itself be informed by, any empirical research.

Let me now discuss four themes about which Wedgwood's account has a visible problem with relevance.

Reductionism

In our thinking about normativity, we may be concerned about various forms of reductionism. For instance, are interpersonal relations reducible to the results of cost-benefit analyses, as social exchange theories say? Or is human morality merely a sophisticated form of experiences and patterns of cooperation that we already find among primates?

Reductionism is a big issue in *The Nature of Normativity*, but the nonreductionism that Wedgwood defends in chapters 8 and 9 will in no way help settle issues like the ones just mentioned. Consider what happens in chapter 9. Here, Wedgwood proposes a version of modal logic that allows (a) non-necessary necessity and possibility, thereby allowing (b) the facts of strong supervenience in the actual world to be explainable not just by necessary, but also by contingent facts, thereby allowing (c) a failure of strong supervenience

² To philosophers of religion, he claims to have shown that impersonal normative requirements have the 'unconditional authority' (p. 277) that was long monopolized by traditional theology.

³ The other one is on p. 177.

to imply global supervenience, in order, finally, to allow (d) a nonreductionist notion of ‘realized in’.

This discussion will no doubt be of interest to specialists in metaphysics, but hardly to philosophers working on normativity. Would it really matter, one is tempted to ask, if normative and intentional properties might be ‘reducible’ in *some* sense, especially because the reducibility would at best amount to the coherence of some purely abstract metaphysical (here: functionalistic) framework? It is hard to see what difference Wedgwood’s nonreductionism could make even to those philosophers, cognitive scientists, neuroscientists, cognitive psychologists and sociobiologists who pursue reductively-naturalistic programmes. Their reductionism amounts to the exploration of mental capacities at some explanatory level other than the intentional, something that is quite different from the metaphysically ambitious and abstract forms of reductionism that Wedgwood argues against.

Ethics

Wedgwood claims that his account has implications for ‘first-order ethical theory’ (p. 273), in helping it to answer the question ‘Why should I be moral?’ Wedgwood does not explicitly state his answer, but the answer that *can* be given on the basis of his account is that, given our capacity to ask this question—which is in fact an ought-question—, we are necessarily disposed to be sensitive to the standards of rationality that are oriented towards the goal of choosing correctly. But what makes the question ‘Why should I be moral?’ pertinent to an ethicist is that it generally means: ‘Why should my moral reasons have priority over my other reasons?’ To this question, an answer such as the one just given is obviously inadequate.

Indeed, Wedgwood’s contribution to ethical theory could only be a merry-go-round of concepts such as ‘correct’, ‘rational’ and ‘ought’. Consider:

It may be that virtuous character involves a disposition to engage in praiseworthy sorts of practical reasoning, and such praiseworthy sorts of practical reasoning may themselves involve normative thinking (such as thinking about what one ought to do). Clearly any such account of the virtues, or of the praiseworthy forms of practical reasoning and decision making, will have to be informed by a good account of the nature of normative judgments, such as the account that I have developed here (p. 274).

What it is that makes practical reasoning ‘praiseworthy’, or choices and plans ‘correct’, however, is not part of the theory—it abstracts from value. But especially in ethics, it is vital to know how to navigate in this domain, i.e., among the many different kinds of reasons real-life situations give us.

Are there plausible principles available to rank them (such as: ‘Moral reasons must have priority over prudential reasons’)? Can such principles be justified? Would all rational creatures choose the same ranking? And so on. These issues do not become vivid in Wedgwood’s account of normativity because, as we have seen, this account starts at the point where all the different ingredients of practical reason have already been processed into an ought.

The Actual Role of Oughts

Wedgwood’s account of normativity is all about oughts. Here, we need not question the assumption that it is our essence that we are rational beings, nor that we are capable of forming ought-judgments and of acting on them. But how much of our agency *is* actually governed by ought-judgments—at least such as are the products of practical reasoning? We have already seen that acting for a reason need not coincide with acting on the basis of an ought-judgment. And it seems that we engage in practical reasonings that issue in an ought-judgment mainly in those rare situations when we face a really difficult decision.

There is some reason to believe that the space for oughts is even narrower than this. First, if social psychologists are correct in thinking that the source of much of our behaviour is our social environment, rather than our individual personality (hence the deliberations of us as individuals), the number of our actions that can be traced back to oughts may shrink considerably.

Secondly, psychologists such as Haidt argue that the moral reasons we offer for our actions are (often inadequate) rationalizations after the fact of decisions that have been made unconsciously.⁴ Because in Wedgwood’s account ought-judgments are the all-things-considered result of practical reasoning, rather than the outcome of unconscious processing, they are here sidelined. Not that this proves moral deliberation to be irrelevant. The point is just that real-world agents may turn out to differ significantly from what a rationalist philosophical model assumes about them. And if such a model is merely concerned with abstract deontology and rationality, but not with real agents, irrelevancy threatens.

Explaining NJI

Wedgwood starts off by presenting NJI as an important basic principle, and then argues that any account of normative judgments is supposed to explain it.

⁴ J. Haidt, ‘The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment’, *Psychological Review* 108 (2001), pp. 814–34.

But is NJI really so important when we wish to understand real agents? It is hard not to see NJI as a formal principle of rationality, rather than a perspicuous matter of fact about real agents. NJI, as Wedgwood presents it, is an *a priori* principle about the structure of rationality and agency. There exist many types of explanation, but NJI, it seems, could only be explained by another formal principle that entails it. This, indeed, is what Wedgwood's conceptual role semantics of 'ought' seems to do.

My worry is that such a type of explanation does not tell us much about real agents. True, there is something NJI-like about such agents, namely the fact that they do sometimes seem to be motivated by their ought-judgments. But note that they are motivated by their simpler value judgments just as well.

Conclusion

Above, I have summarized what I consider the main line of argument in Ralph Wedgwood's book *The Nature of Normativity*. I have also developed an objection, not to the arguments and conclusions as such, but to the philosophical approach to normativity of which these are the result. The basis of the objection is that the book concentrates on the concept of 'ought'. The objection itself is that the argument developed in the book faces an issue of relevance. *The Nature of Normativity* lacks intelligible connections with common sense and empirical science.

Whether or not one likes *The Nature of Normativity* will depend on what one expects a philosophical text to do. It can be argued that the issues discussed in the book are just intrinsically interesting, and that this should be sufficient to pre-empt complaints about relevance. It can also be argued that even if the relevance of a technical philosophical work is not clearly visible at present, the work might become highly relevant later on. And it can be argued that in philosophy, it is one task to invent theories, and another to apply them. I cannot prove considerations like these to be mistaken. But philosophers should be wary of them even so: history has just seen too many philosophical debates turn scholastic and sterile. My worry is that the kind of analytic philosophy found in *The Nature of Normativity* also heads in that direction. Stating my worry has been my way of asking what doing philosophy is about—or ought to be.