

A linguistic turn in the philosophy of normativity?

1. Introduction

The discipline of linguistics contains an account of the semantics of ‘ought’ that is well enough established to count to as the ‘standard’ or ‘classical’ account. It stems largely from the work of Angelica Kratzer.¹ It is sophisticated and intricate. It has foundations in modal logic, which gives it a strong philosophical basis.

We philosophers of normativity also have a keen interest in the meaning of ‘ought’. Ours is different from the linguists’ interest. We are not much concerned with the idiosyncrasies of ‘ought’ in natural language, and we are much more concerned with the underlying metaphysics of normativity. Nevertheless, two different theories about the meaning of ‘ought’ – both with philosophical implications – cannot ignore each other. On the face of it, too, they have very different implications. Philosophers of normativity take ‘ought’ to denote one of their prime normative properties, whereas linguists treat it as a not particularly special modal verb.

Yet it is only within the last decade or so that the linguistic theory of ‘ought’ has begun to influence the philosophy of normativity.² This paper considers what its influence should be. My conclusion will be rather ungracious. I shall argue that we philosophers of normativity need not change our thinking much in response to the work of linguists. And I shall be bold enough to suggest that we can even contribute to improving the linguistic theory.

2. The place of ‘ought’ in linguistics

First I shall outline the linguists’ standard account of the meaning of ‘ought’. ‘Ought’ is treated as one out of a range of modal verbs. Others are ‘must’, ‘may’, ‘can’, ‘could’, ‘have to’, ‘need’ and more. There is a broad semantic theory of modality that applies to all of them, and extends further to a wider modal vocabulary that includes ‘possibly’, ‘requires’, ‘necessary’ and much more.

‘Ought’ has no very special place in the theory. The modal verbs are classified as either ‘necessity modals’ or ‘possibility modals’. ‘Ought’ is one of several necessity modals along with ‘must’ and ‘have to’. It is generally taken to express a weaker necessity than those two.

It is also unspecial in another way. According to the linguists there are different sorts of modality – technically known as ‘flavours’. For instance, there is epistemic modality, which has to do with what must or may be so, given what is known. There is bouletic modality, which has to do with achieving what is wanted, and there is deontic modality, which is concerned with rules or duties. There are many other flavours too. In principle, modal verbs are portable between flavours. For example:

You may leave the room after fifteen minutes. (Deontic.)

You may find a surprise waiting for you. (Epistemic.)

Drivers must look in the mirror before turning right. (Deontic.)

There must be life somewhere else in the universe. (Epistemic.)

These examples show that ‘ought’ is one of the portable verbs:

A doctor ought to be tolerant towards her patients. (Deontic.)

As this rate, we ought to arrive around two o’clock. (Epistemic.)

In view of your dislike of early rising, you ought not to be a dairy farmer. (Bouletic.)

It is largely the context that determines what flavour of modality is meant in a particular utterance. The modal verbs are not ambiguous between flavours. They retain their meaning in shifting from one to another, but their meaning has parameters that are set by the context. Among the parameters are the flavour of the modality. They are like pronouns: ‘she’ is not ambiguous, but what this word refers to depends on the context.

In real life, some verbs are more portable than others. Their portability varies between

languages. French ‘pouvoir’ is said to be a fully portable possibility modal. On the other hand, English-speaking children are taught not to use ‘can’ in ‘Can I leave now?’, because ‘can’ is not supposed to serve as a deontic modal. The correct deontic modal is supposed to be ‘may’; children should say ‘May I leave now?’. So the flavour is not set entirely by the context; it can come partly from the meaning of the specific modal word.

But at any rate ‘ought’ is portable between epistemic, deontic and bouletic modalities. Most importantly, it is not particularly a deontic word in the linguistic theory. Conversely, deontic modality is not particularly associated with ought; ‘must’ and ‘may’ are other deontic modals.

Why is ‘ought’ so closely assimilated in the standard linguistic theory to other modal verbs, and so closely integrated into the theory of modality in general? One reason comes from syntax. ‘ought’ is an auxiliary verb. This means that, in a sentence, it always joins with the infinitive of another verb – known as a ‘lexical’ or ‘main’ verb – to constitute a single compound verb. In

Julia ought to catch the evening flight.

‘Ought to catch’ is the compound verb. Simple tests distinguish an auxiliary from a lexical verb. Compare the sentence above with one that has a superficially similar structure:

Julia hurried to catch the evening flight.

The negations of these two sentences have plainly different structures:

Julia ought not to catch the evening flight.

Julia did not hurry to catch the evening flight.

So do the corresponding questions:

Ought Julia to catch the evening flight?

Did Julia hurry to catch the evening flight?

The rule for negation in English is that ‘not’ is placed between the auxiliary and the main verb. The rule for forming a question is to invert the order of the subject and the auxiliary. To make this possible in a sentence that does not have an explicit auxiliary, the all-purpose auxiliary ‘do’ or one of its inflexions has to be inserted. The ‘ought’ and ‘hurry’ sentences behave differently because ‘ought’ is an auxiliary whereas ‘hurry’ is a lexical verb.

As an auxiliary, ‘ought’ works as a modifier on the main verb of a sentence, which means that in effect it modifies the whole sentence. This is a mark of modality. Modality is a feature of a sentence that allows it to refer to something other than what is actually the case. The underlying nonmodal sentence (called the ‘prejacent’ by linguists) ‘Julia catches the evening flight’ says something about the plain, unflavoured actual world. ‘Julia may catch...’, ‘Julia could catch...’, ‘Julia ought to catch...’ all add some flavour by referring implicitly to possibilities other than what actually happens. ‘May’, ‘could’, ‘ought’ and various other verbs are called modal auxiliaries.

But the syntax of modal auxiliaries is not the important thing for the semantic theory of modality. The important thing is that linguists have a unified account of modality that extends far beyond modal auxiliaries. For example, it covers conditional sentences, and sentences that ascribe propositional attitudes to people, since these constructions also refer implicitly to possibilities other than what actually happens. This account is thought to cover ‘ought’ well, and that is the real reason for including ‘ought’ within the semantic theory of modality in general.

3. The standard semantic theory

I shall give a brief outline of the account itself. It is based on modal logic and takes over the standard possible-world semantics of modal logic. It postulates a domain of possible worlds W .³ Each world w is associated with a subset of W made up of worlds that are said to be ‘accessible’ from w . These accessible worlds are in some sense possible from the point of

view of w . This possibility can be understood in various different ways. For example, the accessible worlds might be the ones that are possible given what is known at w . Or they might be the ones that are metaphysically possible. Or they might be the ones that are still possible at some particular time; in this case, the worlds accessible from w will have the same history as w up to that time. The relation between w and worlds that are accessible from w is known as the ‘accessibility relation’. Different interpretations of possibility will give the accessibility relation different formal structures. Fortunately we need not trouble ourselves with these structures apart from one feature of them: we assume that every world is possible from its own point of view, so w is accessible from w . This means the accessibility relation is reflexive.

Logicians initially applied this structure to give truth conditions for the two modal operators ‘Necessarily’ and ‘Possibly’, which operate on sentences ‘ p ’ to form modal sentences ‘Necessarily p ’ and ‘Possibly p ’. ‘Necessarily p ’ is taken to be true at a world w if and only if ‘ p ’ is true at every world that is accessible from w . ‘Possibly p ’ is taken to be true at w if and only if ‘ p ’ is true at one or more of the worlds that are accessible from w .

This semantics, together with the structure of the accessibility relation, entails that various formulae are valid. The valid formulae constitute a ‘modal logic’. Different structures for the accessibility relation yield different logics.

von Wright adapted this same possible-world semantics to cover the deontic operators ‘ought’ and ‘Permissible’.⁴ The trick is to treat the worlds accessible from w , not as those that are possible from the point of view of w , but those that are in some way ideal from the point of view of w . Then we can say that ‘Ought p ’ is true at w if and only if ‘ p ’ is true at all worlds accessible from w , which is to say at all the worlds that are ideal from the point of view of w . ‘Permissible p ’ is true at w if and only if ‘ p ’ is true at one or more of these ideal worlds. ‘ought’ is like ‘Necessarily’ in this respect, which is why it counts as a ‘necessity modal’. The collection of formulae that are made valid by this semantics constitute a ‘deontic logic’.

What mainly distinguishes deontic logic from the original modal logic (which we can now label with the name ‘alethic modal logic’) is that we do not assume that the accessibility relation is reflexive. That is to say, we do not assume that every world is accessible from itself. That would be to assume that each world is ideal from its own point of view, which is false. This different structure for the accessibility relation produces a different logic. Nevertheless, it is a version of modal logic in general.

That is a brief account of the semantics of modal logic. Within linguistics this semantic theory has been made broader and more elaborate. It has been made to cover a wider range of modal operators and a wider range of applications. In a simple modal sentence, some modal operator ‘ M ’ is attached to a prejacent sentence ‘ p ’ to form a modal sentence ‘ Mp ’. The semantics for ‘ M ’ is built from a set of possible worlds and an accessibility relation among them. This structure of this relation is set partly by the operator itself and partly by the context.

One innovation within linguistics is that the accessibility relation is broken into two factors. To each world is assigned a set of worlds known as the ‘modal base’ and also an ordering of those worlds. The worlds treated as accessible from a world are those within the modal base that are highest ranked in the ordering. When ‘ M ’ is a necessity modal ‘ Mp ’ is true at a world w if and only if ‘ p ’ is true at all those highest-ranked worlds. The modal base is the set of worlds that are possible in some sense at the given world. Some sort of an aim is supposed to be implicit in the modal, and worlds are ordered according to how well they satisfy that aim. For a deontic modal, the ordering is given by some rules or requirements that are relevant in the context. For an epistemic modal, it is given by how likely the worlds are in

some sense. The relevant ordering is set partly by the context.

For an example, take the sentence

‘You ought to thank the cashier for the receipt she has just given you’.

This is a deontic modal sentence. For a deontic modal, the modal base is generally given by the circumstances. In this case, it is the set of all worlds compatible with the circumstances you are in, having just been given a receipt. The ordering is given by whatever rules are implied by the context – in this case rules of politeness. The relevant rule in this case is that you thank a person who gives you something. Presumably, any of the highest-ranked worlds compatible with the circumstances will be ones where you thank the cashier. So the sentence comes out true.

The idea of inserting an ordering into the semantic theory came from Angelica Kratzer,⁵ taking her lead from David Lewis.⁶ Its main purpose is to deal properly with conditional modal sentences. It has the extra benefit of opening up a way of understanding how ‘ought’ can be a necessity modal that is weaker than ‘must’. Here is one way this can be done.⁷ ‘Must *p*’ is true if and only if ‘*p*’ is true at all worlds in the modal base that are sufficiently highly ranked, whereas ‘Ought *p*’ is true if and only if ‘*p*’ is true at all worlds in the modal base that are highest ranked. Suppose, say, that you must put a grade on each student’s exam, and you ought also to add useful comments. Then worlds where you put a grade on each student’s exam are ranked higher than those where you do not, and are ranked sufficiently highly. Those worlds where you put a grade on each student’s exam and also add useful comments are ranked most highly.

4. *The deontic and the normative*

That will do as a sketch of modality in linguistics. It leads to the conclusion that ‘ought’ is not a particularly special word, but one among several modal verbs that can be used for various purposes. It is not even confined to modality.

We need to attend to the terms ‘deontic’ and ‘normative’. I am using ‘deontic’ as linguists use it, to mean ‘having to do with rules or duties’. I use ‘normative’ to mean ‘having to do with reasons’. I admit this definition of ‘normative’ is unsatisfactory, since it is meant to refer specifically to normative reasons rather than motivating or explanatory reasons. The definition is therefore circular. I am really just telling you that I use ‘normative’ in the sense that appears in ‘normative reason’, and assuming you already understand that sense. I hope you do. This is the standard meaning of ‘normative’ in the philosophy of normativity.

There are things you have reason to do that are nothing to do with rules or duties. For example, in a restaurant you have reason to choose the food you like, but you are not under any rule or duty to do so. There are also rules you have no reason to conform to, such as the rule not to end a sentence with a preposition. So ‘normative’ and ‘deontic’ refer to different things. So far as I know, linguists have not recognized normativity as a flavour of modality, but there is no reason why they should not. If they did, they would undoubtedly conclude that ‘ought’ is not a particularly normative word, any more than it is particularly a deontic word.

On the other hand, philosophers of normativity generally think of ‘ought’ as referring to a very special, central feature of normativity. Can these viewpoints be reconciled?

5. *Specialized vocabulary*

Philosophy, like many disciplines, has to specialize its vocabulary. Mechanics uses ‘force’ in a very much more precise way than English does, and there is no harm in that. For the sake of philosophical analysis, we must expect to have to give ‘ought’ a more precise meaning than it has in common English. At the very least, we may exclude some ordinary uses of ‘ought’ in order to avoid ambiguity.

True, in the philosophy of normativity we do need to stay in close touch with ordinary language. We aim to talk about things that people ordinarily refer to. For example, suppose we conclude that there are no normative dilemmas, by which I mean that it cannot be the case both that you ought to do something and ought not to do it. For this conclusion to be useful, it needs to be one that people ordinarily understand. It is pointless to arrive at such a conclusion by using 'ought' in an unrecognizable way. Furthermore, our philosophical methods often call on intuitions that arise from our ordinary lives, and for those to be reliable we cannot stray too far from ordinary terminology. Still, we are entitled to make our language more precise, and amend it to an extent, provided we keep open an easy route back to our ordinary language.

Indeed, for philosophical purposes, we even need to modify the grammar of 'ought'. Its grammar is plainly defective for our purposes. A trivial example is that 'ought' has no inflections. It cannot be put into the past tense. When we say 'You ought to have been more careful', we probably mean to put 'ought' into the past as well as 'be', but we cannot grammatically do that. Moreover, 'ought' has no gerund; indeed it has no nominalization of any sort. Yet in philosophy, we very often need a corresponding noun. For this reason, some philosophers use 'obligation' as a nominalization of 'ought', but that is misleading as well as incorrect. 'Obligation' is a moral or legal word. For the sake of accuracy, many philosophers nowadays boldly use 'ought' as a noun as well as verb. This is a clear violation of English grammar, but these days it passes without comment. I do it myself.

Another grammatical problem arises because 'ought' is an auxiliary verb. Because it is attached to the main verb of a sentence, it is impossible in English to make 'ought' take scope over a conditional sentence. In the sentence

You ought, if you spill wine, to spill white wine,
the condition 'if you spill wine' is not grammatically within the scope of 'ought'. It is a parenthesis within the sentence as a whole, so it can be moved to another position. The sentence means exactly the same as

If you spill wine, you ought to spill white wine.

Kratzer's theory of modals confirms this feature of English.⁸ Kratzer takes the antecedent of a conditional sentence to place a restriction on the modal base of whatever is the relevant modal. The position of the conditional clause makes no difference. Wherever in the sentence it appears, it restricts the modal base. This means it is not within the scope of the modal.

In the wine example, the modal is 'ought'. Its modal base is the set of worlds that are compatible with present circumstances. The conditional clause 'if you spill wine' restricts its modal base to those worlds, compatible with the circumstances, in which you spill wine. To evaluate the sentence, let us suppose that the best worlds compatible with the circumstances are ones where you spill no wine, the next best are those where you spill white wine, and the worst are those where you spill wine of another colour. Among worlds where you spill wine, the best are those where you spill white wine. The sentences 'You ought, if you spill wine, to spill white wine' and 'If you spill wine, you ought to spill white wine' have the same meaning and are both true.

But for philosophical purposes we sometimes need a conditional to be within the scope of 'ought'. Suppose you ought not to spill wine at all, but actually you do. It surely might still be true that you ought, if you spill wine, to spill white wine. If the condition is outside the scope of 'ought', it follows by modus ponens (which is supported by Kratzer's analysis) that you ought to spill white wine. But this is an undesirable conclusion, since you ought not to spill wine at all. We can avoid this undesirable conclusion by taking the condition to be within the scope of 'ought'.

Since English offers no means of placing the condition within the scope of 'ought', we need an ungrammatical device. We might say 'You ought (if you spill wine, to spill white wine)' or 'You ought that, if you spill wine, you spill white wine'. True, if the conditional is a material conditional, we could alternatively substitute a logical equivalent, and write 'You ought either not spill wine or else spill white wine'.⁹ This gets us a sentence in grammatical English, but it has the cost of assuming the conditional is a material conditional. This might be an assumption worth avoiding. So philosophers of normativity will need to depart some distance from the use of 'ought' in common English.

6. *The central ought of normativity*

What about the contextuality of 'ought', which linguists stress? Undoubtedly, the meaning of 'ought' in common English has parameters that can vary and are set by the context. It has epistemic uses, and several uses that are not epistemic. It has 'bouletic' uses that are given by a context of something desired or aimed at. An example is

'You ought to try this one'

spoken at a cheese-tasting. Here the context makes 'ought' relative to your desires. Deontic uses give 'ought' a meaning that is relative to particular sets of rules or requirements. 'Catholics ought not to eat meat on Fridays' is an example. According to the standard semantics, the modal base for non-epistemic 'ought' is generally the set of worlds that are compatible with the circumstances of utterance, but the ordering can vary widely with the context.

The philosophy of normativity can readily accept that many oughts are relative to the circumstances. You ought to drive on the left in the circumstances of Australia, but not in the circumstances of Sweden. It can also accept that there is a use of 'ought' that makes it relative to some particular ordering. However, it must insist that 'ought' also has a meaning that is not relative in this second way. This is a normative meaning, which can fairly be considered the core meaning of 'ought'. Meanings relative to different orderings are peripheral by comparison. I shall next explain why.

I do not wish to deal with the epistemic use of 'ought'. It raises too many complications. In any case, it is rare and historically recent. There are some examples recorded from the late 17th century, whereas there are clear examples of non-epistemic 'ought' from the 11th century.¹⁰ In the corpus of English, it is still vastly outnumbered by non-epistemic 'ought's'.¹¹ In identifying a central meaning, I think it is safe to set epistemic uses aside.

In thinking about non-epistemic uses, start with this conversation at a cheese-tasting:

Me: 'You ought to try this one. You'd love it.'

You: 'I'm sure I would, but do you think I really ought to?'

Me: 'Well, as far as the state of your health is concerned, of course you ought not to. Eating hard cheese is always dangerous to your health.'

You: 'That's not what I asked. I know I'd love the cheese and I know it might damage my health. My question is whether or not, recognizing these things, I ought to try it.'

Here I have presented you with two oughts that are each relative to a particular consideration: what you like and what is good for your health. You are not content with those. You ask for an absolute, nonrelative ought: for what is sometimes called an

‘all-things-considered’ ought or a ‘final ought’. Yours is a perfectly acceptable question; it makes good sense and you are entitled to ask it.

If you ask what you ought to do, and you receive a relativist answer, you are entitled to be discontented. You are presumably looking for an answer that you may rationally act on, and a relativist answer does not give you the rational entitlement you want. It would be irrational to try the cheese just because you will love it, when you know it may affect your health. It would also be irrational not to try it just because of your health, when you know you will love it. Relative oughts are not enough to make an action rational.

True, you cannot always expect an answer to your request for a final ought. So far as the cheese is concerned, the position might be that it is not the case that you ought finally to try it, nor the case that you ought finally not to try it. In general when you have a decision to make, you may not be able to form a belief that you ought finally to act one way, nor a belief that you ought finally to act another way. Even so, you may act, and you may be rational in doing so. Sadly, Buridan’s ass had not learnt this lesson. It was determined to be rational, and it thought an act can be rational only if it is based on a belief that you ought finally to do it. There it was mistaken. But even in cases where there is no final ought, acting on the basis of a belief in a mere relative ought is not rational. You need also to believe that no final ought applies in the case.

7. *Enkrasia*

The sense of ‘ought’ it is rational to act on is the final, nonrelative one. This is the sense that is referred to in the requirement of rationality I call ‘Enkrasia’: the requirement that you intend to do what you believe you ought to do. More precisely:

Enkrasia. Rationality requires of *N* that, if
N believes at *t* that she herself ought that *p*, and if
N believes at *t* that it is up to her herself then whether or not *p*, then
N intends at *t* that *p*.

The technical details do not matter here; they are examined thoroughly in my *Rationality Through Reasoning*.¹² Ralph Wedgwood takes this principle as actually defining a meaning for ‘ought’ by means of its conceptual role.¹³ I do not take it to define ‘ought’, but I do take it to be picking out an important meaning of ‘ought’ that we understand very well. It can fairly be called central because it is the one that is practical. Enkrasia specifies exactly the sense in which it is practical.¹⁴

The existence of this central ought need not be inconsistent with the standard semantics for ‘ought’. It may be that there is a final ordering of worlds that takes everything into account, and that you ought (centrally) to do whatever is true in all worlds that are highest in this final ordering. If so, the central ought is just one example of the semantics, in which the context picks out the final ordering. Since the meaning would then be relative to this ordering, it would not then be strictly accurate to call it ‘nonrelative’.¹⁵ However, sections 8 and 9 below cast doubt on this possibility.

Even if the central ought is consistent with the standard semantics, it has an important practical status that other oughts lack. This is the status specified in Enkrasia. For that reason, this is the ought that philosophers of normativity are chiefly interested in it. Consequently, they do not have to be much bothered by the fluid and contextual nature of ‘ought’ in common English. We have our one central meaning, and we should concentrate on that. Moreover, this is not just a central meaning of ‘ought’; it is at the centre of normativity itself. The philosophy of normativity is principally about this ought: what its metaphysical nature is, what determines when it obtains, and so on.

Indeed, I recommend that, in the philosophy of normativity, we should try to use ‘ought’ in

this central sense only. We should eschew uses such as mine in the cheese conversation above. We should also eschew such expressions as ‘morally ought’ and ‘prudentially ought’, which are explicitly relativist. Instead of these ‘adverbial oughts’, I recommend the word ‘requires’. We should say, for example: ‘Morality requires of you that you F’ rather than ‘You morally ought to F’. Sticking to the central ‘ought’ and using ‘requires’ in other contexts would make the philosophy of normativity less confusing.

Several commentators have pointed out to me one important exception to this recommendation of mine: the so-called objective ought. The objective ought is defined relative to what actually happens in the future, rather than to expectations of what will happen. The objective ought is not the central ought,¹⁶ but nevertheless it plays an important role in the philosophy of normativity. It cannot easily be replaced with the ‘requires’ terminology.

8. A problem with deontic logic

So far I have raised no objection to the standard semantics of ‘ought’ itself, but now I come to an objection. I think the semantic theory made a mistake in nailing its colours to the mast of deontic logic. Alethic modal logic with its possible-worlds semantics is a highly successful theory, and it was natural to hang the linguistic theory of modals in general on that. But the extension of modal logic to the deontic realm has been highly unsuccessful. From the beginning, it has been riven by problems and paradoxes.¹⁷ This has given a stimulus to logicians to try and create a deontic logic that is more independent of alethic modal logic, but most philosophers of normativity gave up on deontic logic long ago because of its difficulties. Now it is interesting to see many of the old arguments about deontic logic being re-run within linguistics.¹⁸

The ordering semantics that Kratzer and Lewis introduced overcomes some of the difficulties. However, a fundamental problem remains. The most disturbing of the paradoxes are counterexamples to modal logic’s most fundamental axiom, known as the distribution axiom or axiom K. Applied to the deontic operator ‘ought’ it may be called ‘deontic detachment’, and be written:

$$(\text{Ought } (p \rightarrow q) \ \& \ \text{Ought } p) \rightarrow \text{Ought } q.$$

The arrows in this formula stand for material implication, because modal logic is built on ordinary propositional calculus. Since we are told that material implication is not correctly expressed in English by ‘if ... then’, it is safest to avoid expressing it that way. Since this is an axiom, it is supposed to be necessary. So it may be written:

$$\text{Ought not } (p \text{ and not } q) \text{ and Ought } p \text{ together entail Ought } q$$

Nearly all systems of modal logic – the ones known as ‘normal modal logic’, which includes the deontic logic I have described – are constructed by adding further axioms to K. K is a consequence of the standard semantics of modal logic, including the ordering semantics.

Let us check it is a consequence of the semantics. Suppose that, in some context, ‘Ought not $(p \text{ and not } q)$ ’ and ‘Ought p ’ are both true. This means that, given the modal base and ordering determined by the context, both ‘ p ’ and ‘not $(p \text{ and not } q)$ ’ are true in all highest-ranked worlds. That is to say, in all these worlds ‘ p ’ is true, and moreover ‘ p ’ is not true unless ‘ q ’ is true. Therefore ‘ q ’ is true in all top-ranked worlds. So ‘Ought q ’ is true.

Still, K is questionable. There are many counterexamples to it.¹⁹ I shall mention two. I do not mean to suggest that a couple of counterexamples are enough to cut through the long debate there has been about the validity of K. Of course the counterexamples can themselves be questioned. But they will at least reveal the nature of the problem.

First, suppose you ought to visit your mother-in-law, and that you ought not to visit her

without telling her you are coming. According to K, it follows that you ought to tell your mother-in-law you are coming. Does that really follow? It does not. It might be that you are not going to visit your mother-in-law; you do not necessarily do what you ought to do. If you are not going to visit her, it cannot be the case that you ought to tell her you are coming.

This means that, in a context where you do not visit your mother-in-law, it is not the case that you ought to tell her you are coming. But that is not the context I am concerned with. I am concerned with the context at a time where it is still open to you whether or not you visit your mother-in-law. I hold that context constant. In that context it is true that you ought to visit your mother-in-law, and also true that you ought not to visit her without telling her you are coming. But it may not be true in that context that you ought to tell her you are coming. There are some worlds compatible with that context where it is not true; these are worlds where you do not visit her. This conclusion is therefore not entailed by the premises, since there are worlds where the premises are true and the conclusion false.

My other example is about a relativist ought. It is relative to the International Regulations for the Prevention of Collisions at Sea. Suppose you are the master of a power-driven vessel that is meeting another head-on. According to article 4 of the Colregs, you ought to turn to starboard. ('When two power-driven vessels are meeting on reciprocal or nearly reciprocal courses so as to involve risk of collision each shall alter her course to starboard'.) According to article 34a of the Colregs, you ought not to turn to starboard without emitting one short blast on the ship's whistle. ('A power-driven vessel underway, when manoeuvring as authorized or required by these Rules, shall indicate that manoeuvre by the following signals on her whistle: one short blast to mean "AI am altering my course to starboard"...').

Axiom K tells us it follows that you ought to emit one short blast on the whistle. But this does not follow. You might hold your course rather than turn. For example, you might behave as the masters of many big ships behave when meeting a smaller vessel: you might hold your course and expect the smaller vessel to get out of your way. If you held your course, it would be extremely dangerous to signal that you are turning to starboard, and the Colregs do not require you to.

9. *An alternative semantics*

My diagnosis of the problem is this. The standard semantics for deontic logic implies that you ought to do anything you do at all worlds that are ideal relative to the actual world – that is to say: anything that is necessary for doing everything you ought to do. In the mother-in-law example, if you are to do everything you ought to do, you visit your mother-in-law and you do not visit her without telling her you are coming. So you tell her you are coming. In the Colregs example, if you are to do everything you ought to do, you turn to starboard and you do not turn to starboard without emitting one short blast on the whistle. So you emit one short blast on the whistle. Necessarily, if you do not tell your mother-in-law you are coming, or you do not emit one short blast on the whistle, you fail to do something you ought to do. But why should it follow that you ought to do either of these particular things? Why ought you to do everything that is necessary for doing everything you ought to do?

I do not know. I think this is just a presumption that hung over from alethic modal logic when it was transmuted into deontic logic. It is most implausible. It means, for one thing, that every tautology is something you ought to do: you ought either to walk to the shops or not walk to the shops, for example. Since this is necessarily true, it is necessary for doing everything you ought.

This conclusion for tautologies is often thought to be a harmless consequence of

formalizing deontic logic. But more substantive things are necessary if you are to do everything you ought. It is necessary that the Earth is not destroyed by an asteroid, for instance. If the Earth is destroyed by an asteroid, you will not repay your debts, which you ought to do. Does it follow that this is something you ought to achieve? I think not.

I think this failure of axiom K for deontic logic gives us a good reason to separate the semantics of normative modality from the semantics of alethic modality. Ordinary modal logic is not able to give a satisfactory unified account of all modals including 'ought'.

'Ought' behaves differently from others.

There is an alternative semantics available for 'ought' that does not imply axiom K. It is known as a 'neighbourhood semantics'.²⁰ In the place of an accessibility relation, which is a relation between worlds and worlds, there is a relation between worlds and propositions. Each proposition related to a world specifies something you ought to do at that world. In effect, for each world, there is simply a list of all the things you ought to do at that world. It is not implied by this semantics that you ought to do whatever is necessary for doing everything you ought to do.

K is not implied by a neighbourhood semantics, and nor is much else. This semantics gives us a very weak – almost empty – deontic logic. This is as it should be. We should not expect semantics to tell us much about the structure of ought. This is a matter for substantive normative theory. Even if there is no ultimate distinction to be drawn between the analytic and the synthetic, there is a useful division of labour to be made between semantic theory and normative theory.

Normative theory might issue in some principles that can be expressed in a logical form. For example, if we conclude that there are no normative conflicts, which means it cannot be the case that you both ought to do something and ought not to do it, that would be equivalent to axiom D, one of the standard axioms of modal logic. And indeed, it could be argued that this is genuinely an analytic principle stemming from the meaning of 'ought', so this could fairly be called a principle of logic. But many principles will not take such a simple form. Suppose, say, that what you ought to do is given by decision theory. Richard Jeffrey called decision theory 'the logic of decision',²¹ but it is hard to think of it as just a working-out of the meaning of ought.

In any case, to learn about how the things we ought to do are connected to each other, we need to investigate normativity, not language.

References

- Broome, John, *Rationality Through Reasoning*, Wiley-Blackwell, 2013.
- Chisholm, Roderick M., 'Contrary-to-Duty Imperatives and Deontic Logic', *Analysis*, 24 (1963), pp. 33-6.
- Forrester, James William, *Being Good and Being Logical: Philosophical Groundwork for a New Deontic Logic*, M. E. Sharpe, 1996.
- Hale, Bob, and Crispin Wright, *A Companion to the Philosophy of Language*, Blackwell, 1997.
- Jeffrey, Richard C., *The Logic of Decision*, Second Edition, University of Chicago Press, 1983.
- Kratzer, Angelika, 'Conditionals', in *Semantics: An International Handbook of Contemporary Research*, edited by A. von Stechow and D. Wunderlich, de Gruyter, 1991, pp. 651-6. Reprinted in her *Modals and Conditionals*, pp. 86-107.
- Kratzer, Angelika, *Modals and Conditionals*, Oxford University Press, 2012.
- Kratzer, Angelika, 'The notional category of modality', in *Words, Worlds and Contexts*,

edited by H. J. Eikmeyer and H. Rieser, de Gruyter, 1981, pp. 39-74. Reprinted in her *Modals and Conditionals*, pp. 27-69.

Lewis, David, 'Semantic analyses for dyadic deontic logic', in *Logical Theory and Semantic Analysis*, edited by S. Stendlund, Reidel, Dordrecht, 1974, pp. 1-14.

Nordlinger, Rachel, and Elizabeth Traugott, 'Scope and the Development of Epistemic Modality: Evidence from *ought to*', *English Language and Linguistics*, 1 (1997), pp. 295-317.

Pacuit, Eric, *Neighbourhood Semantics for Modal Logic: an Introduction*.

von Fintel, Kai, 'The best we can (expect to) get?: challenges to the classic semantics for-deontic modals'.

von Fintel, Kai, and Sabine Iatridou, 'How to say *ought* in foreign: the composition of weak necessity modals', <http://web.mit.edu/fintel/fintel-iatridou-2008-ought.pdf>.

Von Wright, G. H., 'Deontic logic', *Mind*, 60 (1951), pp. 1-15.

Wedgwood, Ralph, *The Nature of Normativity*, Oxford University Press, 2007.

Comment [J1]: Something wrong with the formatting here. I don't know how to fix it.

Notes

Presented at the *Analytic Philosophy* symposium at Austin in December 2014. An earlier version was the Jack Smart Lecture I gave in Canberra in July 2014. I have received a great deal of help in the course of writing this paper. Dan Fogal in particular has given me a great deal of his time, and his advice has been invaluable. I also particularly want to thank Janice Dowell, Julien Dutant and Steve Yablo for their excellent advice. Research for this paper was supported by ARC Discovery Grant DP140102468.

1. Kratzer, *Modals and Conditionals*.
2. It had little influence even in the philosophy of language until fairly recently. There is no reference to Kratzer in the *Blackwell Companion to the Philosophy of Language*, published in 1997. I ignored it in my own book *Rationality Through Reasoning*; this paper tries to make up for that omission.
3. The linguistic theory does not actually require these possibilities to be whole worlds. They could be possible situations, possible choices, possible choices or something else. I assume they are worlds simply to make the language more familiar within philosophy.
4. Initially in his 'Deontic logic'.
5. 'The notional category of modality'.
6. 'Semantic analyses for dyadic deontic logic'.
7. It resembles the one in von Fintel and Iatridou, 'How to say *ought* in foreign'.
8. 'Conditionals'.
9. Thanks here to Dan Fogal.
10. Nordlinger and Traugott, 'Scope and the Development of Epistemic Modality', pp. 312-3.
11. Nordlinger and Traugott, 'Scope and the Development of Epistemic Modality', p. 299.
12. pp. 170-4.
13. *The Nature of Normativity*, chapter 1.
14. The central *ought* is the subject of chapters 2 and 3 of my *Rationality Through Reasoning*.
15. Thanks to Dan Fogal.
16. See my *Rationality Through Reasoning*, pp. 40-1.
17. There is a survey in James Forrester's *Being Good and Being Logical*.
18. See Kai von Fintel, 'The best we can (expect to) get?'
19. They are broadly modelled on Chisholm's Paradox, in 'Contrary-to-Duty Imperatives and

Deontic Logic’.

20. Pacuit, *Neighbourhood Semantics for Modal Logic*. There is an independent account and justification for it in my *Rationality Through Reasoning*, chapter 7.

21. See Jeffrey, *The Logic of Decision*.