7-1-2014

Consequences of reasoning with conflicting obligations

Gopal Shyam NAIR
Lingnan University, Hong Kong

Recommended Citation
Nair, S. (2014). Consequences of reasoning with conflicting obligations. Mind, 123(491), 753-790. doi: 10.1093/mind/fzu085

Consequences of Reasoning with Conflicting Obligations

SHYAM NAIR

University of Southern California

gnair@usc.edu

Since at least the 1960s, deontic logicians and ethicists have worried about whether there can be normative systems that allow conflicting obligations. Surprisingly however, little direct attention has been paid to questions about how we may reason with conflicting obligations. In this paper, I present a problem for making sense of reasoning with conflicting obligations and argue that no deontic logic can solve this problem. I then develop an account of reasoning based on the popular idea in ethics that reasons explain obligations and show that it solves this problem.

0. Introduction

Since at least the 1960s, deontic logicians and ethicists have worried about whether there can be normative systems that allow conflicting obligations (Brink 1994, Chellas 1980, Gowans 1987b, Lemmon 1962, Marcus 1980, McConnell 2010, Pietroski 1993, Sinnott-Armstrong 1988, van Fraassen 1973, Williams 1965). Many of these worries stem from problems that show a tension between a normative system allowing conflicting obligations and plausible logical principles for such a system. This has led to a rich literature in deontic logic concerned with developing logics that allow conflicting obligations while still validating a number of desirable inferences (Goble 2009, Hansen 2005, Horty 2003, McNamara 2004, van der Torre and Tan; Goble forthcoming provides a survey).

On the other hand, issues concerning how we may reason with conflicting obligations have received little direct attention. This is somewhat surprising because cases that illustrate compelling patterns of reasoning have been an important source of evidence about what the correct principles of deontic logic are. And while these cases certainly are a good source of evidence about this, they provide even more direct evidence for what the principles of good reasoning are.

For example, consider the case of Smith, a citizen of a just country. Smith's country requires him to fight in the army or perform alternative public service. So Smith ought to fight or serve. Smith is also deeply religious. Smith's religion is committed to pacifism so it requires him not to fight. So Smith ought not to fight. Given that Smith ought to fight or serve and that Smith ought not to fight, we may conclude that Smith ought to serve (cf. van Fraassen 1973, p. 18; Horty 1993, p. 73; Goble 2004, p. 80).[1]

In Smith's case, it is good reasoning to conclude 'Smith ought to serve' from 'Smith ought to fight or serve' and 'Smith ought not to fight'. This means (perhaps among other things) that if we permissibly believe that Smith ought to fight or serve, permissibly believe that Smith ought not to fight, and come to believe that Smith ought to serve by competently reasoning from these beliefs, we permissibly believe that Smith ought to serve.[2] More generally, Smith's case illustrates that the pattern of reasoning from 'it ought to be that *a* or *b*' and 'it ought to be that not

[1] In this case the obligation to fight or serve is based on the law and the obligation to not fight is based on religious commitments while the resulting obligation to serve is not solely based on either of these sources. So while some cases in which obligations derive from different sources involve 'ought's with different meanings, Smith's case shows that other cases involve 'ought's with the same meaning and that there are patterns of good reasoning concerning these kinds of 'ought's (cf. Goble forthcoming, pp. 112–13). My discussion in this paper only applies to this second kind of case involving 'ought's with the same meaning. Thanks to the referee who pressed me to be clear about this.

[2] So, as I am thinking of it, the reasoning we may do in Smith's case is a mental process that starts with two beliefs and concludes in the formation of another belief. In this paper, I focus exclusively on reasoning with qualitative beliefs understood in this way.

a’ to ‘it ought to be that *b*’ is a pattern of good reasoning.[3] This pattern of good reasoning involves performing disjunctive syllogism ‘under several “ought”s’.

We can illustrate another pattern of good reasoning by considering the following case concerning a speeding law. The law sets the speed limit on major streets at fifty miles per hour. So drivers ought to drive less than fifty miles per hour. Given that drivers ought to drive less than fifty miles per hour, we may conclude that drivers ought to drive less than one hundred miles per hour (cf. Cariani forthcoming, n. 1).

In this speeding law case, it is good reasoning to conclude ‘Drivers ought to drive less than one hundred miles per hour’ from ‘Drivers ought to drive less than fifty miles per hour’. More generally, the speeding law case illustrates that the pattern of reasoning from ‘it ought to be that *a*’ to ‘it ought to be that *b*’ when *a* entails *b* is a pattern of good reasoning.

Having used these two cases to identify two principles of good reasoning, we are now in a position to see why it is worthwhile to pay attention to issues concerning reasoning about conflicting obligations. It is worth paying attention because these two principles suggest that we cannot fruitfully reason with conflicting obligations.

To see this, let us begin by supposing that there are conflicting obligations. In particular, let us assume that it ought to be that *a* and that it ought to be that not *a*. According to the second principle of reasoning that we identified, we may conclude ‘it ought to be that *a* or *b*’ from ‘it ought to be that *a*’ because *a* or *b* follows from *a*. Next according to the first principle of reasoning that we identified, we may perform disjunctive syllogism ‘under several “ought”s’. So we may conclude ‘it ought to be that *b*’ from ‘it ought to be that not *a*’ and ‘it ought to be that *a* or *b*’. Thus, our principles of reasoning taken together look to entail that we may reason from

[3] For simplicity, I assume ‘ought’ and ‘obligation’ are roughly synonymous and I assume ‘ought’ is an unary operator on sentences. While both of these assumptions are controversial, the issues involved in these controversies are orthogonal to the central issues of this paper.

conflicting obligations—'it ought to be that *a*' and 'it ought to be that not *a*'—to an explosion of obligations—'it ought to be that *b*' for any *b*.

This result suggests that we cannot fruitfully reason with conflicting obligations. Reasoning that leads to explosion would be fruitless because agents who engage in such reasoning would be committed to thinking that there is no important distinction between claims of which it is true that they ought to hold and those of which it is false. And since no interesting normative system fails to treat this distinction as important, the fact that reasoning with conflicting obligations commits us to treating this distinction as unimportant should make us reluctant to believe that there are conflicting obligations. More conservatively, even if we would like to allow for the abstract possibility that there can be conflicting obligations of some sort, the fact that we cannot fruitfully reason with them suggests that these kinds of obligations cannot play an important role in our lives.

In short, it is worth paying attention to issues concerning reasoning because advocates of conflicting obligations have a problem making sense of how we may reason with such obligations. They face the problem of telling us how to reason with such obligations without reasoning explosively.

More generally, even those who do not believe in conflicting obligations have reason to be interested in this problem. This is because there are analogs of it for other important normative notions as well. For example, consider the notion of a reason. We can use Smith's case and the speeding law case to motivate analogous principles of reasoning for reasons by simply substituting 'there is a reason to do' for 'it ought to be that' in these cases and noticing that the resulting reasoning is intuitively acceptable. Since it is a platitude that reasons can conflict, this shows that everyone faces the problem of understanding how to reason with normative notions that allow for conflicts (cf. Goble 2009, pp. 454–8). And the solution that I will develop in this paper can be generalized so that it applies to the notion of a reason. That said, this generalization will not be obvious and I will not have the space to present it here.[4] Instead, I will focus exclusively on the issue of conflicting obligations because it, unlike the issue of conflicting reasons, has been subject to considerable discussion.

In particular, I will present an account of reasoning with obligations on behalf of the advocate of conflicting obligations. This account not only explains why explosive reasoning is bad reasoning but also explains why the reasoning in Smith's case and the speeding law case are pieces of good reasoning that are neither enthymematic (that is, reliant on a suppressed premiss) nor equivocating (that is, starting with an 'ought' that means one thing and ending with an 'ought' that means something else). The account is closely connected to the popular idea in ethics that reasons explain obligations and not only solves this problem, but also looks like a promising general theory of reasoning with conflicting obligations.

In order to make the case for my account, I begin in section 1 by considering whether we can make use of the rich literature on conflicting obligations in deontic logic to solve our problem. I argue that no deontic logic can solve our problem based on two structural features of the logical consequence relation (cut and monotonicity). Having used this argument to reject the most obvious potential solutions to our problem, I turn to developing my own account. I begin by presenting the popular idea in ethics that reasons explain obligations in section 2. While this idea does not directly give us an account of reasoning, I show in section 3 how to develop a plausible account of reasoning with conflicting obligations that is closely connected to this popular idea. This account of reasoning solves our problem. And it does so, in part, by entailing that the 'reasoning consequence relation' does not share the structural features of the logical consequence

[4] The key first component of this generalization is independently motivated in Nair MS-a.

relation. These are the main results of this paper. The remainder of the paper is dedicated to extending these ideas by looking at other problems about reasoning (Sect. 4), considering alternatives to my account (Sect. 5), and drawing an important moral about the bearing of my account of reasoning on issues in deontic logic (Sect. 6).

1. Deontic logic and the problem

I began the paper by noting that substantial attention has been paid to the task of providing an adequate logic of conflicting obligation. I also noted that logicians have used cases that illustrate patterns of good reasoning like Smith's case and the speeding law case as an important source of evidence about what the correct principles of deontic logic are. It is unsurprising, then, that a problem analogous to the reasoning problem that I just presented has been discussed in this literature (Goble 2009, p. 459).

This suggests that a promising strategy to solve our problem about reasoning is to mine the resources of this rich literature in deontic logic. We might use one of the logics to provide an account of reasoning by claiming that we may reason in ways that correspond to valid inferences according to the logic of our choice. For this strategy to be successful, we would need to find a deontic logic that claims that the reasoning in Smith's case and the speeding law case correspond to valid inferences, but explosive reasoning does not.

In this section, I argue that this seemingly promising strategy is actually hopeless: no deontic logic could say that the reasoning in Smith's case and the speeding law case correspond to valid inferences while the explosive reasoning does not.

1.1 'Logic'

In order to make this argument, I must begin by clarifying what I mean by 'logic' when I say that no deontic logic could get us the results that we want. As I will use 'logic', a logic allows us to define a logical consequence relation. More precisely, if we use upper case letters like A, B, C etc. to stand for sets of arbitrary sentences and lower case Greek letters like α, β, χ etc. to stand for arbitrary sentences, a logic tells us under what conditions φ is a logical consequence of A. Typically logicians symbolize the logical consequence relation with a turnstile, $\vdash$. So in symbols, according to my usage of 'logic', a logic tells us under what conditions A $\vdash \varphi$ (where this is read as 'φ is a logical consequence of A').[5]

And as I will use 'logical consequence', a logical consequence relation is intended to capture when φ *deductively* follows from A. That is, we should say that A $\vdash \varphi$ only if the truth of the sentences in A *guarantees* the truth of φ. In other words, the claim that A $\vdash \varphi$ tells us that an argument with the sentences in A as premisses and φ as a conclusion is a valid argument in the sense of 'valid' that we are all taught in Logic 101.

Of course, this is not the only reasonable way to use 'logic'. John Burgess notes another important usage:

> Among the more technically oriented 'logic' no longer means a theory about which forms of argument are valid, but rather means any formalism, regardless of intended application, that resembles a logic in the original sense enough to allow it to be usefully studied by similar methods. (2009, p. viii)

[5] I will also occasionally use these symbols as names for themselves and trust that context makes it clear when I am doing this.

The sense of 'logic' Burgess identifies is wider than the sense that I have identified.[6] And in fact, the view I will go on to defend could be fairly called a deontic logic in this wider sense of 'logic'. While this way of discussing my view is unobjectionable to my ears, I have found that it often obscures the import of the arguments of this paper as well as the central features of the account of reasoning that I will go on to develop. For this (purely pragmatic) reason, I will adopt the more narrow usage of 'logic' that I identified in the previous paragraphs.

1.2 The argument

Having clarified what I mean by 'logic', I am in a position to state the main claim of this section more precisely. The claim is that no deontic logic in this narrower sense of 'logic' just identified can solve our problem. In this subsection, I will argue for this claim.

1.2.1 Properties of logical consequence My argument begins by isolating two claims about logical consequence that any deontic logic must accept. The first claim is that if φ is a logical consequence of A, then φ is a logical consequence of A together with any other claim, β. We can see why any deontic logic must accept this by noting that this claim just amounts to the obviously true claim that if A guarantees the truth of φ this must mean that A and β together also guarantee the truth of φ. This feature of logical consequence is often called *monotonicity*:

Monotonicity: if $A \vdash \varphi$, then $A, \beta \vdash \varphi$

where we understand A, β as shorthand for $A \cup \{\beta\}$.

The second claim is that if β is a logical consequence of A and φ is a logical consequence of A and β together, then φ is a logical consequence of A. While this claim is more complicated than the claim that logical consequence is monotonic, it is also a claim that any deontic logic

[6] It would, for example, consider the system of default reasoning presented in Reiter 1980 to be a logic.

must accept. After all, this claim just amounts to saying that if A guarantees the truth of β and A and β together guarantee the truth of φ, then A guarantees the truth of φ. This feature of logical consequence is often called *cut*:

Cut: if A $\vdash \beta$ and A, $\beta \vdash \varphi$, then A $\vdash \varphi$

Thus, any deontic logic must define a logical consequence relation that is both cut and monotonic.

1.2.2 Why deontic logic cannot solve the problem I now want to make use of this result to show that no deontic logic can solve our problem. Our problem, recall, was to develop an account that says the reasoning in Smith's case and the speeding law case correspond to valid inferences but explosive reasoning does not.

The good reasoning in Smith's case was reasoning from 'it ought to be that *a* or *b*' and 'it ought to be that not *a*' to 'it ought to be that *b*'. And the idea of using deontic logic to solve our problem says that it can explain this reasoning by showing that it corresponds to a valid inference. In symbols, the idea that the reasoning in Smith's case corresponds to a valid inference is the idea that the following claim holds:

Logical Ought Disjunctive Syllogism ($\text{ODS}_{\vdash}$): $O(a \vee b), O(\neg a) \vdash O(b)$

where O is read as 'it ought to be the case that', the connectives $\vee$, $\neg$ are understood as they usually are, and lower case italicized letters such as *a*, *b*, *c* etc. are for atomic sentences.

Analogously, we said that the speeding law case illustrated that it is good reasoning to conclude 'it ought to be that β' from 'it ought to be that α' when α entails β. For the purposes of my argument, I will only need to discuss the special case of this inference that involves concluding 'it ought to be that *a* or *b*' from 'it ought to be that *a*'.[7] In symbols, the idea that this reasoning corresponds to a valid inference is the idea that the following claim holds:

Logical Ought Disjunction Introduction (ODI$_\vdash$): $\mathrm{O}(a) \vdash \mathrm{O}(a \vee b)$

Thus, *ODS*$_\vdash$ and *ODI*$_\vdash$ are formal statements of the idea that the reasoning in Smith's case and the speeding law case correspond to valid inferences.

To use deontic logic to solve our problem, we must find a deontic logic that accepts *ODS*$_\vdash$ and *ODI*$_\vdash$ but rejects explosive reasoning. Unfortunately, as I will now argue, there is no such deontic logic.

As I have already explained, every deontic logic defines a cut and monotonic consequence relation. I now will make use of this fact to show that if we assume *ODS*$_\vdash$ and *ODI*$_\vdash$ hold we can derive that explosive reasoning corresponds to a valid inference. So begin by assuming that *ODS*$_\vdash$ and *ODI*$_\vdash$, here numbered (1$_\vdash$) and (2$_\vdash$) respectively, hold:

(1$_\vdash$) $\mathrm{O}(\neg a), \mathrm{O}(a \vee b) \vdash \mathrm{O}(b)$
(2$_\vdash$) $\mathrm{O}(a) \vdash \mathrm{O}(a \vee b)$

[7] Some reject *ODI*$_\vdash$ based, for example, on a generalization of Ross 1941's puzzle concerning imperatives. Others find it to be undeniable, see Hilpinen and Føllesdal 1971, p. 22 and Nute and Yu 1997, p. 5. While I do not have the space to consider this issue in the detail it deserves here, three points can be made. First a popular solution to Ross's paradox is to accept *ODI*$_\vdash$ and explain the paradox away on pragmatic grounds (Castañeda 1981). Second though not clearly impossible, it is not obvious that the reasoning in the speeding law case can correspond to a valid inference if *ODI*$_\vdash$ does not hold. In fact, this reasoning can be seen as an instance of *ODI*$_\vdash$ going from 'it ought to be that drivers drive forty-nine mph or forty-eight mph or…' to 'it ought to be that drive ninety-nine mph or ninety-eight mph or … or forty-nine mph or forty-eight mph or….'. Third, *ODI*$_\vdash$ is entailed by other plausible principle so those who reject it face the task of explaining which of these principles they reject. Two examples: *ODI*$_\vdash$ follows from *ODS*$_\vdash$ and the assumption that O(*t*) where *t* can be any tautology; *ODI*$_\vdash$ follows from the principle that 'it ought to be that α' follows from 'it ought to be that α and β' together with the principle that 'it ought to be that β' follows from 'it ought to be that α' when α and β are logically equivalent. Thanks to the Thomas Baldwin for pressing me to address this issue.

Monotonicity, recall, tells us that we may add a sentence to the left and keep what is on the right. So we can get (3$\vdash$) by adding O(a) to the left of (1$\vdash$) and similarly get (4$\vdash$) by adding O($\neg a$) to the left of (2$\vdash$):

(3$\vdash$) $\mathbf{O(\boldsymbol{a})}$,O($\neg a$), O($a \vee b$) $\vdash$ O(b)
(4$\vdash$) $\mathbf{O(\neg \boldsymbol{a})}$,O($a$) $\vdash$ O($a \vee b$)

Cut then allows us to derive the result that explosive reasoning corresponds to a valid inference from (3$\vdash$) and (4$\vdash$). To see this, notice that (4$\vdash$) is of the form A $\vdash$ β:

$$\underbrace{\mathrm{O}(\neg a),\ \mathrm{O}(a)}_{\mathrm{A}} \ \vdash\ \underbrace{\mathrm{O}(a \vee b)}_{\beta}$$

and (3$\vdash$) is of the form A, β $\vdash$ φ:

$$\underbrace{\mathrm{O}(a),\ \mathrm{O}(\neg a),}_{\mathrm{A}} \quad \underbrace{\mathrm{O}(a \vee b)}_{\beta} \ \vdash\ \underbrace{\mathrm{O}(b)}_{\varphi}$$

So cut tells us that A $\vdash$ φ:

(5$\vdash$) O(a), O($\neg a$) $\vdash$ O(b)

(5$\vdash$) says that inferring ‘it ought to be that b’ from ‘it ought to be that a’ and ‘it ought to be that not a’ is a valid inference.

This shows why the strategy of mining the resources of deontic logic to solve our problem is hopeless. After all, for that strategy to work it must show that the reasoning in Smith’s case and the speeding law case correspond to valid inferences even though explosive reasoning does not. But this derivation shows that it is impossible for the reasoning in Smith’s case and the speeding law case to correspond to valid inferences without explosive reasoning corresponding to a valid inference.

And in fact this result tells us something more general than this. It tells us that if we are to solve our problem, the relation of good reasoning or the 'reasoning consequence relation' must not be both cut and monotonic. For if it were both cut and monotonic, then a parallel argument would show that it cannot be that the reasoning in Smith's case and the speeding law case is good reasoning without it also being true that explosive reasoning is good reasoning.

Of course, cut and monotonicity are non-negotiable for a logical consequence relation. For a reasoning consequence relation, however, they are negotiable. Monotonicity about reasoning says that, if you may reason to φ from A, then you may reason to φ from A plus anything. In essence, this means that learning something new cannot make it so you have to take back any of your old conclusions. Monotonicity for reasoning is negotiable simply because it is sensible to wonder whether learning something new might make it so you have to retract your old conclusions.

Similarly cut about reasoning says that if you may reason to β with the premisses A and you may reason to φ with the premisses A and β, then you may reason to φ with only the premisses A. This basically means that you can draw the same conclusions whether you arrive at β as a conclusion or have it from the start as a premiss. Much like monotonicity, cut is negotiable because it is at least sensible to ask whether premisses really do play the same role in our reasoning as conclusions.[8] Thus, while solving our problem with deontic logic is impossible because the logical consequence relation is both cut and monotonic, it may yet be possible to

[8]Some suggest reasoning that does not satisfy cut is unstable (see for example, Makinson 1994, pp. 43–4). Though I do not have space to consider this issue here, Nair MS-b responds to this charge of instability from a general perspective.

solve our problem about reasoning. What we need to do to solve this problem is develop an account that entails that the reasoning consequence relation is not both cut and monotonic. So what I will do is develop such an account.

1.3 Reconsidering the problem

But before we turn to my account, we may wish to try out a different strategy that allows us to stay closer to the idea of using deontic logic to solve our problem. The strategy is to solve our problem by claiming that the reasoning in either Smith's case or the speeding law case is enthymematic or equivocating reasoning.

Of course, these are simple examples where we have clear judgements that we are dealing with non-enthymematic and non-equivocating reasoning. So rejecting these judgements without some good independent reason would be an *ad hoc* solution to our problem (cf. Goble 2009, pp. 466–7). But we have just seen that the reasoning in Smith's case and the speeding law case cannot correspond to valid inferences if we are to avoid saying that explosive reasoning corresponds to a valid inference. And it is tempting to think that this is a good independent reason for rejecting our judgements that our reasoning in these cases is non-enthymematic and non-equivocating.

As tempting as this line of thought might be, it is too quick. What makes the thought too quick is that it is not generally true that we should be willing to reconstrue some apparently good piece of reasoning as enthymematic or equivocating if we find that it does not correspond to a valid inference. Remember we are using 'valid inference' to pick out an inference from A to φ such that the truth of the sentences in A guarantees the truth of φ. But there are many forms of reasoning that do not involve only inferring what is guaranteed to be true. For instance, we engage in inductive reasoning and common sense reasoning (for example, concluding 'John smokes' from 'A reliable witness told me that John smokes'). Such reasoning does not correspond to some valid inference but this is not a good reason to reject our judgement that these types of reasoning are forms of good non-equivocating and non-enthymematic reasoning.

To give one example that displays just how implausible the suggestion that we are currently discussing is, consider the idea that the reasoning from 'A reliable witness told me that John smokes' to 'John smokes' is enthymematic. This idea would say the reasoning from 'A reliable witness told me that John smokes' to 'John smokes' is good reasoning because it relies on a suppressed premiss, *p*, and the inference to 'John smokes' from *p* and 'A reliable witness told me that John smokes' is a valid inference.

But consider what *p* would have to be in order for this inference to be valid. One thing it could be is 'Reliable witnesses always tell the truth'. But this claim is obviously false and rational agents do not generally accept it. So if the reasoning from 'A reliable witness told me that John smokes' to 'John smokes' implicitly relies on a premiss like this, then, far from explaining why this is good reasoning, the strategy of treating this reasoning as enthymematic would conclusively show this reasoning is bad. Of course, we can consider other candidate premisses. But all of the examples that I can think of are unacceptable for reasons similar to the one just given or end up making the explicitly given premiss essentially irrelevant (for example, an implicit premiss like 'The thing the reliable witness said is true'). For this reason, it is implausible to explain all forms of good reasoning as equivocating or enthymematic for reasoning that corresponds to a valid inference.

That said, since these forms of reasoning do not correspond to valid inferences, we may wonder what explains why they are forms of good reasoning. And in fact, one way of understanding what philosophers are up to when they are tackling the (new and old) riddle of induction is to understand them to be trying to answer a question like this (Hume 1739, Bk. 1, Pt. 3, Sect. 6; Goodman 1955, Ch. 3). But the example of induction also teaches us that taking questions about what explains why a certain form of reasoning is good reasoning seriously and not yet having fully worked out answers to such questions should not undermine our confidence that this form of reasoning really is good.

In light of this, I conclude that we should continue to take Smith's case and the speeding law case to provide us with strong evidence that certain patterns of reasoning with obligations are patterns of good reasoning. What I will do in the next two sections is develop an account of reasoning that entails that the reasoning in these cases is good reasoning. And while my account does not provide a fully worked out story of what explains why this reasoning is good, it will at least shed some light on this question.

2. The popular idea from ethics

In order to develop my account of reasoning with obligations, it will help to first say something about the nature of obligation itself. And for that, I turn to the popular idea in ethics that reasons explain obligations (Dancy 2004, Nagel 1970, Parfit 2011, Raz 1975, Scanlon 1998).

2.1 Versions of the popular idea

I am not going to be arguing for the popular idea in ethics in this paper. Instead, I will simply be assuming that it is true. In fact, I will be assuming that a particular version of it is true and the primary task of this section is to explain the version of this view that I will be assuming is true.

We get different versions of the idea that reasons explain obligations by considering different accounts of *how* reasons explain obligations. Unfortunately, while the idea that reasons explain obligations is popular, no particular version of this idea enjoys the same celebrity. To illustrate, one account that has been discussed is that we explain why you ought to do *a* by showing that you have better reasons to do *a* than to do any alternative to *a*.

Unfortunately, I cannot adopt this account because I am interested in understanding reasoning with conflicting obligations on behalf of the advocates of such obligations, and this account looks to rule out conflicting obligations. According to this account, we ought to do *a* and ought to do *b* when *a* and *b* conflict only if the reasons to do *a* are better than the reasons to do any alternative and the reasons to do *b* are better than the reasons to do any alternative. But given natural assumptions about what alternatives are, *a* and *b* are alternatives to one another. And given natural assumptions about what it is for one reason to be better than another, a reason to do *a* cannot be better than a reason to do *b* while that same reason to do *b* is better than that reason to do *a*. But since this would have to be true in order for there to be conflicting obligations, this account rules out conflicting obligations.

So this version of the popular idea in ethics will not work for us because it is not compatible with the existence of conflicting obligations. Luckily, there are other versions of the popular idea in ethics and, in fact, one of the most sophisticated versions of the popular idea in ethics does allow for the existence of conflicting obligations. This account is due to John Horty (2012).

Though one of the virtues of Horty's account is how detailed it is, I will only introduce a simplified version of it in this paper.[9] While I can develop my account of reasoning using the full details of Horty's system, such details are distracting for our purposes. We are not trying to develop the best version of the popular idea in ethics in this paper. I simply want to make use of this idea from ethics to develop an account of reasoning. So all we require are those details that we will need to develop our account of reasoning.

2.2 $\mid\!\sim_{Horty}$

We can start to get a feel for this simplified version of Horty's account by noticing that it is not just any reasons that explain obligations. I might, for example, have a reason to watch television all day. But it is usually not true that I ought to do this. Instead, I usually ought to do something else like teach my students, work on a paper, or clean my house. And this is generally because my reasons to do these things are much better than my reasons to sit at home and watch television all day. Now Horty's account can tell us how reasons of different strength interact with one another. But since we are interested in understanding reasoning in this paper and not the strength of reasons, we will not worry about how we determine whether one reason is better than another. Instead, we will take it for granted that there is a class of *good* or undefeated reasons where an *undefeated* reason is understood to be a reason that is not worse than any reason that is incompatible with it.

Since we will later use these ideas from ethics in developing our account of reasoning and since I will want to be able to compare this account with the ideas from deontic logic that we discussed earlier, I will need to introduce a bit of formalism. So far all I have done is point out

[9] In fact, the simplified version of this system is equivalent to a theory first proposed in van Fraassen 1973. I will none the less continue to call it 'Horty's system' because I intend my account of reasoning to ultimately embrace the much richer system developed in Horty 2012.

that the notion of a good reason plays a role in Horty's system, so let me introduce a formal device for discussing good reasons. I will write $!(\alpha)$ for 'there is a good reason for it to be the case that α'. And for simplicity once again, we will not bother with formally representing what the reason is.

Since we will often be interested in what we ought to do in cases where we have more than one reason, it will also be helpful to have a formal device for a set of good reasons. I will use $\mathfrak{R}$ for such a set. Finally, it will be helpful to have a way of picking out the thing that we have a good reason to do (for example, 'I go to the store' in 'there is a good reason for it to be the case that I go to the store'). So let me define a function *Consequent* that allows us to do this:

$$\textit{Consequent}[!(\alpha)]=\alpha$$

It is also useful to generalize this so that we can take a set of claims about what there is a good reason to do and return a set of those things that there is a good reason to do:

$$\textit{Consequent}[\mathfrak{R}]=\{x \mid \text{there is some } y \text{ such that } y\in\mathfrak{R} \text{ and } \textit{Consequent}[y]=x\}$$

So far then we have introduced a formal way of talking about good reasons ($!(\alpha)$) and about sets of good reasons ($\mathfrak{R}$) along with a function to pick out the things that there are good reasons to do (*Consequent*).

Now that we have identified the kinds of reasons that we are going to use to explain obligations, we need to say *how* these reasons explain obligations. Sometimes this is clear enough. For example, suppose you promise to meet Sally for lunch and you do not have anything else pressing to do. In this situation, you ought to meet Sally for lunch and this is explained by the fact that there is a good reason for you to meet her—namely the promise. So here you have a good reason to meet Sally for lunch and you also ought to meet her.

But for other claims about what you ought to do in this case, things are less clear. For example, many people accept that if you ought to do α and α entails β, then you to ought to do β. Assuming that this is correct, this means that you also ought to meet Sally for lunch or meet Bill for lunch because this follows from meeting Sally for lunch. It is not clear however that you have good reason to meet Sally or Bill for lunch. After all, we usually do not think promising Sally is a reason to meet Sally or Bill.[10] So while it is true that you ought to meet Sally or Bill, it is not obviously true that there is a good reason for you to meet Sally or Bill. If this is right, then it may be that you ought to do something even though you do not have good reason to do that thing.[11]

This is not a counterexample to the idea that reasons explain obligations. It just tells us that the connection between reasons and obligations is not one that allows us to say that every time you ought to do something there is a reason to do that very thing. Instead of saying this, what we can do is say that what explains why you ought to meet Sally or Bill is that this *follows* from what you have good reason to do. That is, you ought to meet Sally or Bill because you have a good reason to meet Sally and because meeting Sally entails meeting Sally or Bill. Horty's account essentially embraces this idea about how reasons explain obligations. Roughly, obligations are the things that follow from what you have good reasons to do.

[10] It may seem plausible that if there is a reason to do α and α entails β, then there is a reason to do β. And it may also seem plausible that if an agent ought to do α, then there is a reason for an agent to do α. Though it is not obvious given what I have said in this paper, I show in Nair MS-a that, suitably interpreted, these ideas are compatible with and explicable by the framework presented here. Seeing this is also the key to seeing how to generalize the present framework so it can accommodate reasoning with reasons.

[11] This is not intended as a conclusive argument for this claim. I am merely presenting a line of thought that can lead someone to this conclusion as a helpful way of introducing the reader to Horty's framework. As I mention in n. 4 and n. 10, there is a more general and independently motivated way of developing the framework that allows us to cling to the claim that every time you ought to do something there is a reason to do it.

It is important that this is only a rough idea and not the whole story. After all, everyone should accept that there can be situations in which you have two equally good reasons. So consider such a case (for example, you made equally important promises to Sally and to Bill) and suppose further that these reasons conflict (for example you cannot fulfil both promises) (cf. Marcus 1980, p. 125). In this kind of case, we do not want to identify what you ought to do with what follows from what you have good reasons to do. After all, you have a good reason to do each of a pair of inconsistent things. And since everything follows from an inconsistent set of claims and since everyone should agree that it is not true that you ought to do everything in this case, it cannot be true that you ought to do what follows from what you have good reasons to do.

We avoid this result by focusing not on what follows from the whole set of reasons but rather on what follows from some 'most inclusive' consistent subset of reasons. Since this idea can be hard to get your head around when it is stated this abstractly, it may be useful to have a heuristic that lets us think about this idea in less abstract terms. So think of it like this: A rational agent treats the things she thinks she has good reasons to do as goals and what such a rational agent ought to do are the steps in a plan that allows her to achieve those goals. In cases where she has incompatible good reasons, she has incompatible goals. So there is no single plan that can accomplish all of them. A 'most inclusive' consistent set of reasons is a collection of goals that is consistent and such that you cannot add more goals to that collection without making it inconsistent. This is the most inclusive set of goals that can be accomplished by a single plan. And so what follows from some such set can be thought of as a step in a plan that helps achieve the most inclusive set of goals that she can achieve with a single plan.

Formally, we say a set of good reasons, $\mathfrak{R}$, includes incompatible reasons when the things we have good reason to do are inconsistent—that is, when *Consequent*[$\mathfrak{R}$] is inconsistent. And we will treat the idea of a 'most inclusive' consistent subset of a set of reasons as a *maximal consistent subset* of *Consequent*[$\mathfrak{R}$]:

> A is *maximal consistent subset* of B iff (1) A $\subseteq$ B, (2) A is consistent, and (3) it is not the case that there is a C such that C is consistent and A $\subset$ C $\subseteq$ B.

What we have said is that what you ought to do is what is *entailed* by some maximal consistent subset. We will use the logical consequence relation of ordinary propositional logic $\vdash_{PL}$ as our formal model of entailment. This allows us to define a relation $\mid\!\sim_{Horty}$ between a set of good reasons, $\mathfrak{R}$, and a claim about what we ought to do, $O(\alpha)$, that holds just in case the set of good reasons explains the obligation:

> $\mathfrak{R} \mid\!\sim_{Horty} O(\alpha)$ iff M $\vdash_{PL}$ α for some maximal consistent subset, M, of *Consequent*[$\mathfrak{R}$]

This then is the idea of how reasons explain obligations that we will be using. And I will refer to this particular version of the popular idea in ethics by the name *Reasons Explain Obligations* or *REO* for short.

To summarize, the basic idea is that what we ought to do is what follows from what we have good reason to do. We add the complication for cases where we have equally good reason to do two incompatible things that what we ought to do is what follows from some maximal consistent subset of what we have good reason to do. And it is precisely in these cases that this system allows for conflicting obligations.[12]

[12] Sinnott-Armstrong (1988) and Nagel (1979) argue that these are the cases where we face conflicting obligations.

Thus, I will be assuming *REO* is correct in what follows.[13] And in the next section I will build my account of reasoning on this idea about the nature of obligation.

2.3 Comparison to deontic logic

Before I go on to do develop the account of reasoning that will solve our problem, I want to take a moment to spell out the relationship between *REO* and the results about deontic logic in section 1. *REO* is a metaphysical thesis about the relationship between reasons and obligations. As such, it does not directly define a logical consequence relation between obligations in the way a deontic logic does.

However, we can use the resources of *REO* to define a relation that tells us when a collection of obligations guarantees the truth of another obligation: one obligation will guarantee the truth of another just in case any collection of reasons that can explain why the first obligation holds suffices to explain why the second obligation holds. This leads to the following definition of a relation $\vdash_{\text{Truth}}$ that tells us when a collection of obligations guarantees the truth of another obligation:

$$A \vdash_{\text{Truth}} \varphi \text{ iff for all } \mathfrak{R} \text{ such that } \mathfrak{R} \mathrel{|\!\sim}_{\text{Horty}} \alpha \text{ for all } \alpha \in A,\ \mathfrak{R} \mathrel{|\!\sim}_{\text{Horty}} \varphi$$

Thus, $\vdash_{\text{Truth}}$ is a relation that tells us when one obligation guarantees the truth of another.

Given the way that we have defined $\vdash_{\text{Truth}}$, it does not take much work to prove that it must be both monotonic and cut. And given our results from section 1.2, this means that we cannot say that the reasoning in Smith's case and the speeding law case correspond to truth guaranteeing inferences without saying that the explosive reasoning corresponds to a truth

[13] This is not to say that this is the only view that I could have assumed in order to develop my account. For another view see Hansen 2007.

guaranteeing inference. Let us consider then which of these inferences corresponds to a truth guaranteeing inference and which of these does not.

The reasoning in the speeding law case does correspond to a truth guaranteeing inference. Formally, this means that the following claim is true:

$$O(\alpha) \vdash_{Truth} O(\beta) \text{ if } \{\alpha\} \vdash_{PL} \beta$$

And it is not hard to see why it is. After all, this claim will be true if any collection of reasons that explains why $O(\alpha)$ is true also explains why $O(\beta)$ is true. For a collection of reasons to explain why $O(\alpha)$ is true is just for α to follow from a maximal consistent subset of that set of reasons. Since we are understanding *following from* in terms of $\vdash_{PL}$ and since $\vdash_{PL}$ is a transitive relation, we know that if α follows from this set of reasons and $\{\alpha\} \vdash_{PL} \beta$ then β must follow from this set of reasons as well. Thus, any set of reasons that explains why $O(\alpha)$ is true must also explain why $O(\beta)$ is true.

Now that we have seen that the reasoning in the speeding law case corresponds to a truth guaranteeing inference, we know that the reasoning in Smith's case must not correspond to a truth guaranteeing inference on pain of saying that the explosive reasoning corresponds to a truth guaranteeing inference. And in fact we can demonstrate with a single abstract example that neither the reasoning in Smith's case nor the explosive reasoning corresponds to a truth guaranteeing inference.

The example is a case where we have a good reason to do a and a good reason to do $\neg a$, $\mathfrak{R}=\{!(a), !(\neg a)\}$. Evidently, this set of reasons has two maximal consistent subsets, $\{a\}$ and $\{\neg a\}$. Thus, this set of reasons explains why $O(a)$ and $O(a \vee b)$ are true—they follow from a maximal consistent subset, $\{a\}$—and explains why $O(\neg a)$ is true—it follows from a maximal consistent subset $\{\neg a\}$. But importantly this set of reasons does not suffice to make O(b) true—it follows from neither $\{a\}$ nor $\{\neg a\}$. Thus, this set of reasons shows that it is possible for O(a), O($a \vee b$), and O($\neg a$) to be true while O(b) is false. So this shows that the reasoning we do in Smith's case from 'it ought to be that *a* or *b*' and 'it ought to be that not *a*' to 'it ought to be that *b*' does not correspond to a truth guaranteeing inference. And it shows that the explosive reasoning from 'it ought to be that *a*' and 'it ought to be that not *a*' to 'it ought to be that *b*' doesn't either.

This means that *REO* and $\vdash_{\text{Truth}}$ alone do not allow us to provide an account that solves our problem about reasoning. But what I will show in the next section is that the ideas that I have introduced in this section together with certain other assumptions do allow us to develop an account of reasoning with conflicting obligations.

3. How to reason with conflicting obligations

My account of reasoning is that we reason well when we reason as if we have good reasons to do what our premisses say we ought to do. Let me explain what this mouthful means.

Begin by noticing that there are two respects in which *REO* does not directly tell us anything about the reasoning that we are interested in. First *REO* is a metaphysical theory about the nature of obligation—it tells us how reasons *explain* obligation—and not a broadly epistemological theory of *reasoning*. Second, *REO* concerns the connection between reasons and obligations, but the reasoning that we are interested in concerns the connection between obligations and other obligations.

My account bridges this gap between *REO* and the reasoning that we are interested in by taking on board two further claims. First, I claim that it is good reasoning to conclude 'it ought to be that *a*' from a set of premisses about what you have good reason to do just in case *REO* says those reasons must make it true that it ought to be that *a*. Second, I claim that when ‘it ought to be that *a*’ is your premiss, the conclusions you can draw about what you ought to do are the conclusions you can draw about what you ought to do from ‘there is a good reason to do *a*’.

This does not mean that when your premiss is ‘it ought to be that *a*’ you reason enthymematically by relying on the suppressed premiss ‘there is a good reason to do *a*.’ And it does not mean that when your premiss is ‘it ought to be that *a*’ you equivocate on its meaning by taking it to mean the same thing as ‘there is a good reason to do *a*’. My proposal instead is just that when you accept ‘it ought to be that *a*’ as a premiss, the conclusions you can draw about what you ought to do are the same as the conclusions you can draw from ‘there is a good reason to do *a*’. In short, my idea is that the premiss ‘it ought to be that *a*’ plays the same functional role in reasoning about what you ought to do as ‘there is a good reason to do *a*’.

Of course, this is just a statement of the account. So in the remainder of this section, I motivate it (Sect. 3.1) and demonstrate that it solves the problem (Sect. 3.2).

3.1 Motivating the account

There are two reasons why this account is plausible. The first reason is that it jointly satisfies two desiderata. The second is that it provides a plausible picture of reasoning.

3.1.1 First reason We want to understand Smith’s case and the speeding law case as pieces of non-equivocating and non-enthymematic good reasoning. This means that these are cases of reasoning directly from claims about what we ought to do to other claims about what we ought to do. Call reasoning like this *deontic reasoning*. What is important to notice about deontic reasoning is that it involves reasoning from obligations to obligations directly and does not involve appealing, for example, to premisses about our reasons. So the first, perhaps obvious, desideratum is just this: the correct account treats the reasoning that we are interested in as deontic reasoning.

And the account that I have proposed does this. It is an account of how to reason from claims about obligations to other claims about obligations. Of course, I said that when you have ‘it ought to be that *a*’ as a premiss you reason as if you have a good reason to do *a*. But as I insisted above, this does not mean you implicitly rely on ‘there is a good reason to do *a*’ in your reasoning and it doesn’t mean that when ‘it ought to be that *a*’ is your premiss it means the same thing as ‘there is a good reason to do *a*’. Rather it is just that the conclusions that you can draw from ‘it ought to be that *a*’ are the same as the conclusions that you can draw from ‘there is a good reason to do *a*’.

The second desideratum for an account of the reasoning in Smith’s case and the speeding law case is that it should fit well with *REO*. Of course, as I have emphasized, *REO* is, in the first instance, a metaphysical theory about the nature of obligation and not a broadly epistemology theory about deontic reasoning. That said, it is still natural to want your metaphysics and your epistemology to fit well together.

And my account of deontic reasoning does fit closely with *REO* because it builds on *REO* with the help of the two further claims that I took on board. The fact that my account jointly satisfies these desiderata adds to its plausibility. And since my account rests on *REO* and two further claims and since we are taking *REO* for granted, the fact that my account of deontic reasoning satisfies these desiderata lends plausibility to the two further claims that I am making.

3.1.2 Second reason My second reason provides a direct illustration of the plausibility of my account. The way to see this is to return to our heuristic about planning.

Consider then two ways an account of deontic reasoning might tell you to treat your premisses. An account of deontic reasoning could tell you to treat the things your premisses say that you ought to do as goals of a plan or as steps in a plan. In particular, if your premiss in deontic reasoning is that you ought to do *a*, an account could tell you to treat *a* as a mere step in your plan or treat *a* as a goal. But if all you know is that you ought to do *a*, then it seems that you should structure your planning around *a* like you would structure your planning around a goal. If all you know is that you ought to do *a*, then *a* seems like it should have this kind of guiding role in your planning. This suggests that an account of deontic reasoning should tell you to treat the things your premisses say that you ought to do as goals.

If we think of deontic reasoning in this way and plausibly assume, as we did in section 2.2, that a rational agent treats the things she thinks she has good reason to do as goals, we can see that treating what our premisses say that we ought to do as goals amounts to reasoning as if we have good reasons to do what our premisses say we ought to do. This means that an account of deontic reasoning should tell us to reason as if we have good reasons to do what our premisses say we ought to do.

This discussion of planning is not intended as a description of what deontic reasoning is rather it is intended as a heuristic to bring out what is underlyingly plausible about my account of such reasoning. The idea in more abstract terms is this: As agents, we sometimes must reason with obligations even though we do not know which reasons explain those obligations. In this environment, it is reasonable to think that we treat the things our premisses say we ought to do as guiding our reasoning about what else we ought to do in the way claims about reasons guide our reasoning about what we ought to do.[14] The idea is that this is a sensible ampliative reasoning strategy.[15]

For this reason, I think there is some plausibility to the idea that rational agents reason well when they to treat the premiss 'it ought to be that *a*' as playing the same role in their reasoning as 'there is a good reason to do *a*'. This is not, of course, a full account of what justifies this reasoning strategy. But this is analogous to the situation that we find ourselves in with regard to other forms of reasoning such as inductive reasoning and common sense reasoning. For these forms of reasoning, we do not have a fully worked out story of why they are forms of good reasoning either.

That said, my account does shed at least some light on why the deontic reasoning we engage in is good reasoning by satisfying the two desiderata and by presenting a plausible picture of what such reasoning consists in. Because of this, I believe that the account of deontic reasoning that I have offered has some independent plausibility. And as I will now demonstrate, it solves our problem.

3.2 Solving the problem

To do this, we will need to develop the account that I have just presented informally in enough formal detail (Sect. 3.2.1) to verify that it entails that the reasoning in Smith's case and the speeding law case is good reasoning but explosive reasoning is not (Sect. 3.2.2). And we will

[14] So the conclusions that you can draw about what you ought to do when you accept 'it ought to be that *a*' as a premiss are the same as the conclusions you can draw from 'there is a good reason to do *a*' not solely because of the meaning of 'it ought to be that *a*' but also because of the fact that you accept it as a premiss in the context of deontic reasoning. This is why deontic reasoning with 'it ought to be that *a*' is not enthymematic or equivocating for reasoning with 'there is a good reason to do *a*'—it is a feature of acceptance as a premiss in deontic reasoning not a feature of meaning that makes the difference.

[15] By *ampliative* reasoning I mean reasoning that allows us to conclude more than what is guaranteed to be true by our premisss.

need to verify that this account avoids the argument in Section 1 that showed that our problem cannot be solved if the reasoning consequence relation is both cut and monotonic (Sect. 3.2.3).

3.2.1 The formal implementation So let us begin our discussion by formally implementing the idea that we have been informally discussing so far. I said that this account takes *REO* for granted and builds on it with two further claims. Since *REO* is already presented in formal terms, all we need to do in order to formally implement this idea is formalize the two further claims.

The first claim I added to *REO* was that we can reason from some claims about what we have good reason to do to the conclusion 'it ought to be that *a*' just in case according to *REO*, any set of reasons that includes these claims about what we have good reason to do explains why 'it ought to be that *a*' holds. So in symbols, you can conclude, for example, 'it ought to be that *b*' from 'there is a good reason to do *a*' just in case for every $\mathfrak{R}$ such that $\{!(a)\} \subseteq \mathfrak{R}$, $\mathfrak{R} \mid\sim_{\text{Horty}} \mathrm{O}(b)$.

The second claim that I added to *REO* is that when your premiss is 'it ought to be that *a*', the conclusions that you can draw from it are the conclusions that you can draw from 'there is a good reason to do *a*'. To see how this works formally, recall that we just said that you can conclude 'it ought to be that *b*' from 'there is a good reason to do *a*' just in case for every $\mathfrak{R}$ such that $\{!(a)\} \subseteq \mathfrak{R}$, $\mathfrak{R} \mid\sim_{\text{Horty}} \mathrm{O}(b)$. Since my account says that when your premiss is 'it ought to be that *a*', the conclusions that you can draw about what you ought to do are the conclusions that you can draw from 'there is a good reason to do *a*', this means that you can conclude 'it ought to be that *b*' from 'it ought to be that *a*' just in case for every $\mathfrak{R}$ such that $\{!(a)\} \subseteq \mathfrak{R}$, $\mathfrak{R} \mid\sim_{\text{Horty}} \mathrm{O}(b)$.

This tells us how to formalize my account of deontic reasoning for the special case of reasoning from 'it ought to be that *a*' to 'it ought to be that *b*'. Of course, we will want to know how to reason with sets of obligations rather than just a single obligation so we will need to generalize this idea to sets. This leads us to the following definition of a reasoning consequence relation that I will call for reasons that will become apparent later $\Vdash_{\text{NonCut}}$:

$$
\begin{aligned}
&O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n) \Vdash_{\text{NonCut}} O(\beta) \text{ iff, for all } \Re \text{ such that:}\\
&\qquad\qquad \text{(i) } \Re \mid\!\sim_{\text{Horty}} O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n)\\
&\qquad\qquad \text{and}\\
&\qquad\qquad \text{(ii) } \{!(\alpha_i) \mid 1 \leq i \leq n\} \subseteq \Re\\
&\qquad \Re \mid\!\sim_{\text{Horty}} O(\beta)
\end{aligned}
$$
[16]

This definition tells us that we can conclude $O(\beta)$ from a set of obligations just in case $\Re \mid\!\sim_{\text{Horty}} O(\beta)$ where $\Re$ is a set of reasons that we have placed conditions (i) and (ii) on. Conditon (i) captures the idea that you take your premisses to be true in reasoning.[17] Condition (ii) captures the idea that the premiss 'it ought to be that *a*' plays the same functional role in reasoning as 'there is a good reason to do *a*'. And the rest of the definition captures the idea that we may reason from some claims about what we have good reason to do to the conclusion 'it ought to be that *b*' just in case according to *REO*, any set of reasons that includes these claims about what we

[16] $\Vdash_{\text{NonCut}}$ is formally similar to a relation defined in Horty 1993 (and other systems inspired by this one). More generally, this paper owes many of its ideas to previous work by John Horty. In essence, the present account newly interprets, further develops, and connects various strands of Horty's work. Let me briefly explain. Horty 1993 and 1997 define a consequence relation that is formally similar to the one defined here. At this time however, Horty did not distinguish between reasons and obligations. So this relation though formally similar to mine was not given the interpretation that I give it—that when you have 'it ought to be that *a*' as a premiss you reasons as if you have good reason to do *a*. For the first time in Horty 1994 and more elaborately in Horty 2003, 2012, Horty distinguished between reasons and obligations. This distinction allowed him to recharacterize the relation defined in Horty 1993 and 1997 as a relation between reasons and obligations—namely, $\mid\!\sim_{\text{Horty}}$. But as we saw this relation is one that does not directly tell us about deontic reasoning. So my account shows how we can accept the theory developed in Horty's later work and build on top of it an account of deontic reasoning that is formally similar to the one given in his earlier work.

[17] This condition is independent of our core idea that we reason well when we reason if we have good reason to do what our premisses say we ought to do. While I do find this further condition plausible, I am officially neutral about it because there are a number of interesting issues that I cannot adequately discuss here that turn on whether we reject (i): rejecting (i) would make a difference to how we treat obligations to do contradictory things and would change certain formal properties of the relation (it would make it so the relation is not reflexive and so that not all logically valid inferences are pieces of good reasoning). I will however assume (i) holds from now on because it allows me to state some of my main ideas in a more straight forward way (for example, it simplifies the proof of $SOC_{\Vdash}$ below)

have good reason to do explains why 'it ought to be that *b*' holds. So this is just a generalization and formalization of the idea that we reason well when we reason as if we have good reasons to do what our premisses say we ought to do.[18]

3.2.2 The cases Having formally implemented this account of deontic reasoning, I am now in a position to demonstrate in a precise way that it solves our problem. To solve our problem is to provide an account of deontic reasoning that entails that the reasoning in Smith's case and the speeding law case is good reasoning but explosive reasoning is not. Formally, this means we want to show that the following claims hold:

Reasoning Ought Disjunctive Syllogism ($ODS_{\Vdash}$): $O(a \vee b), O(\neg a) \Vdash_{\text{NonCut}} O(b)$
Reasoning Single Ought Closure ($SOC_{\Vdash}$) $O(\alpha) \Vdash_{\text{NonCut}} O(\beta)$ when $\{\alpha\} \vdash_{PL} \beta$

even though the following one does not:

Reasoning Ought Explosion ($OE_{\Vdash}$) $O(a), O(\neg a) \Vdash_{\text{NonCut}} O(b)$

Let us verify that my account gets these results.

Begin with $SOC_{\Vdash}$. According to the definition of $\Vdash_{\text{NonCut}}$, you may conclude $O(\beta)$ from $O(\alpha)$ when $\{\alpha\} \vdash_{PL} \beta$ just in case every set of reasons of a certain kind make $O(\beta)$ true. The relevant sets of reasons according to our definition are ones where (i) $O(\alpha)$ is true and where (ii) you have a good reason to do α. As we showed in section 2.3, $O(\beta)$ where $\{\alpha\} \vdash_{PL} \beta$ holds in every set of reasons where $O(\alpha)$ holds. Thus, $O(\beta)$ must also hold in sets of reasons where (i)

[18] While this tells us how to reason with obligations that have prejacents of arbitrary logical complexity, it does not tell us how to reason with arbitrarily complex sentences about obligation. I do not present an account of this here because it is not needed to solve the main problems about reasoning with conflicting obligation and introduces unnecessary complications. In work in progress, I present a fully general account.

$O(\alpha)$ holds and (ii) you have a good reason to do α. So $SOC_{\Vdash}$ holds. This shows that our account of reasoning entails that the reasoning in the speeding law case is good reasoning.

Next let us consider $ODS_{\Vdash}$. To show that this holds according to our definition, we need to look at sets of reasons where (i) $O(\neg a)$, $O(a \vee b)$ hold and where (ii) you have a good reason to do $\neg a$ and a good reason to do $a \vee b$. Since $\neg a$ and $a \vee b$ are consistent, they must be a subset of some maximal consistent subset, M, of your good reasons.[19] Obviously $\{\neg a, a \vee b\} \vdash_{PL} b$. So since $\{\neg a, a \vee b\} \subseteq M$, it follows that $M \vdash_{PL} b$.[20] Since we have shown that b must follow from a maximal consistent subset of reasons, $O(b)$ holds in every such set of good reasons. This suffices to show $ODS_{\Vdash}$ and therefore shows our account entails that the reasoning in Smith's case is good.

But importantly our account does not say that $OE_{\Vdash}$ holds. To see why, we need to show that there is a set of reasons that does not make $O(b)$ true even though it includes a good reason to do a and a good reason to do $\neg a$. And this is easy to do: Consider a set of just these reasons, $\mathfrak{R}=\{!(a), !(\neg a)\}$. There are two maximal consistent subsets $\{a\}$ and $\{\neg a\}$ neither of which entails b. So $O(b)$ doesn't hold. Because of this our account of reasoning entails that reasoning explosively is not good reasoning.

Thus, our account of reasoning gets us exactly what we want: it entails that the reasoning in Smith's case and the speeding law case is good reasoning but the explosive reasoning is not.

3.2.3 The reasoning consequence relation I argued in section 1 that we can get what we want only if the reasoning consequence relation is not both cut and monotonic. That argument appealed to the following derivation to establish this point. Replacing $\Vdash$ for $\vdash$, the derivation

[19] Any consistent set, C, must be a subset of some maximal consistent subset, M. For suppose, C was not a proper subset of any M, then C it itself is a maximal consistent subset.
[20] This follows by the monotonicity of $\vdash_{PL}$.

began with ($1_{\Vdash}$) which is $ODS_{\Vdash}$ and ($2_{\Vdash}$) which is a special case of the reasoning in the speeding law case that we have called $ODI_{\Vdash}$:

($1_{\Vdash}$) $O(\neg a), O(a \vee b) \Vdash O(b)$
($2_{\Vdash}$) $O(a) \Vdash O(a \vee b)$

By monotonicity, we got ($3_{\Vdash}$) and ($4_{\Vdash}$):

($3_{\Vdash}$) $O(a), O(\neg a), O(a \vee b) \Vdash O(b)$
($4_{\Vdash}$) $O(\neg a), O(a) \Vdash O(a \vee b)$

And then by cut, we are able to derive ($5_{\Vdash}$) from ($3_{\Vdash}$) and ($4_{\Vdash}$):

($5_{\Vdash}$) $O(a), O(\neg a) \Vdash O(b)$

This derivation taught us that if the reasoning consequence relation is both cut and monotonic then our problem cannot because solved.

We know that my account solves our problem because, as we just saw, it accepts ($1_{\Vdash}$) and ($2_{\Vdash}$) but rejects ($5_{\Vdash}$), but do not yet know how it solves the problem. One way it could solve the problem is by rejecting monotonicty in a way that blocks the derivation of ($3_{\Vdash}$) or ($4_{\Vdash}$). Another other way it could solve the problem is by rejecting cut in a way that blocks the derivation of ($5_{\Vdash}$). Nothing in our discussion so far has suggested that one way of solving the problem is preferable to another.

The account I have developed, as the name suggests, solves the problem by rejecting cut. And I will show in section 5 of this paper that we can define a variant of it that solves the problem by denying monotonicity instead. While this variant shares in many of the advantages of the account that I am developing here, it will be easiest to understand this non-monotonic account after having appreciated the non-cut account.

So for now let us focus on $\Vdash_{\text{NonCut}}$ and see how it blocks this derivation. Recall that intuitively cut told us that there is no important difference between premisses and conclusions. Cut fails according to the current approach because premisses have a special role. You reason as if you have good reason to do what your *premisses* say that you ought to do. You are not committed to doing this for your *conclusions*. To see that this thought does lead to a denial of cut, let us consider ($3_{\Vdash}$), ($4_{\Vdash}$) and ($5_{\Vdash}$) more carefully (cf. Horty 1997, p. 29).

To see that ($3_{\Vdash}$) holds, the sets of good reasons to consider are the ones where you have a good reason to do a, a good reason to do $\neg a$, and a good reason to do $a \vee b$. Since the good reason to do $\neg a$ and the good reason to do $a \vee b$ are consistent, they will be elements of a maximal consistent subset. As we just saw in showing that the reasoning in Smith's case is good reasoning, O(b) holds when we have such a maximal consistent subset. So ($3_{\Vdash}$) holds.

To see that ($4_{\Vdash}$) holds, the sets of reasons to look at are ones where you have a good reason to do $\neg a$ and a good reason to do a. Since a is consistent, it must be a subset of some maximal consistent subset M. Obviously, $a \vdash_{\text{PL}} a \vee b$ and M $\vdash_{\text{PL}} a$, so M $\vdash_{\text{PL}} a \vee b$. Thus, O($a \vee b$) must hold. Hence, ($4_{\Vdash}$) holds.

Notice what is happening here. In ($4_{\Vdash}$) we do not reason as if we have a good reason to do $a \vee b$ even though in ($3_{\Vdash}$) we do reason as if we have a good reason to do $a \vee b$. This is because O($a \vee b$) is a premiss in ($3_{\Vdash}$) while it is a conclusion in ($4_{\Vdash}$). And notice that in showing why ($3_{\Vdash}$) held—in showing why we can conclude O(b)—we crucially relied on the fact that we reason as if we have good reason to do $a \vee b$.

As we have already seen, ($5_{\Vdash}$), which is the formal statement of the claim that we may reason explosively, does not hold. Since ($4_{\Vdash}$) and ($5_{\Vdash}$) have the same premisses and since our premisses determine what we reason as if we have good reasons to do, we know that we will not reason as if we have a good reason to do $a \vee b$ when we have O(a) and O($\neg a$) as our premisses. And as I just said in discussing (3$_{\Vdash}$) reasoning as if we have a good reason to do $a \vee b$ is crucial to allowing us to conclude O(b). This is why we are not able to conclude O(b) from O(a) and O($\neg a$) alone.

This shows that the solution to our problem flows directly from the features of the account that shed light into why the deontic reasoning that we are interested in understanding is good reasoning. A feature of that account was that we reason as if we have good reasons to do what our premisses say we ought to do. It is this special role played by premisses that leads to the denial of cut.

3.3 Taking stock

What we have done in this section is build an account of deontic reasoning that fits closely with *REO*. The basic idea is that you reason well when you reason as if you have good reasons to do what your premisses say you ought to do. We showed that this idea provides an *explanatory* solution to our problem because it vindicates the reasoning in Smith's case and the speeding law case and condemns explosive reasoning while shedding some light on why such reasoning is good and why such reasoning does not satisfy cut.

That said, it is worth highlighting two limitations of this account and the arguments given for it. First, I do not believe that the independent motivation for my account decisively rules out every alternative to it. There are other accounts such as $\vdash_{\text{Truth}}$ that also could be used to give an account of deontic reasoning that fits with *REO*. However as we have seen, $\vdash_{\text{Truth}}$ does not solve our problem. So $\Vdash_{\text{NonCut}}$ has the advantage of solving our problem while still capturing this independent motivation.

A second limitation of my account is that it is built around *REO*, and *REO* was designed to allow conflicting obligations. In this respect, my solution does not add anything novel about what explains the existence of conflicting obligations. What my account contributes instead is a theory of *deontic reasoning*. This has been the central concern of this paper and it is to the distinctive problems about reasoning with conflicting obligations that my account provides an explanatory solution.

Of course, part of why my solution is explanatory is that it is built around *REO*. But in my opinion, it is a virtue of my view that it allows for this kind of tight fit between our account of deontic reasoning and the popular idea in ethics. For these reasons then, I conclude that, my account of deontic reasoning provides an explanatory solution to the problem about reasoning that I introduced at the outset of this paper.

We have now accomplished the main task of this paper. We began the paper with a problem about how to reason with conflicting obligations (Sect. 0). We saw that this problem cannot be solved by deontic logic and, more generally, cannot be solved by any account that entails that the reasoning consequence relation is both cut and monotonic (Sect. 1). We then took a detour to introduce a particular version of the popular idea in ethics, *REO* (Sect. 2). And in this section we used *REO* along with two further claims to develop an independently plausible account of deontic reasoning (Sect. 3.1) that solves our problem and entails that the reasoning consequence relation is not cut (Sect. 3.2).

Having established our main results, I want to close the paper by considering some other problems (Sect. 4), discussing alternatives to my account (Sect. 5), and drawing a moral about the bearing of my account of reasoning on issues in deontic logic (Sect. 6).

4. Other problems

As we have seen from our earlier discussion, there is a close albeit imperfect connection between deontic logic and reasoning with conflicting obligations. Since there are a host of problems about conflicting obligations in deontic logic, it is natural to wonder about the corresponding problems for reasoning with conflicting obligations. While I believe that my account can solve the reasoning analogs of the problems about conflicting obligations from deontic logic that I am aware of, it would not be worthwhile for me to discuss each of these problems here. Instead, in this section I will provide a more abstract perspective on my account that will enable the interested reader to easily apply it to different problems concerning reasoning and verify for herself that it solves these problems.[21]

To get this more abstract perspective, consider the following principle:

The General Reasoning Principle ($GP_{\Vdash}$): $O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n) \Vdash O(\beta)$ if M $\vdash_{PL}$ β for some M such that M is a maximally consistent subset of $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$.[22]

$GP_{\Vdash}$ succinctly encapsulates the ideas about deontic reasoning that are central to my account.

Remember my idea about deontic reasoning says that we reason as if we have good reasons to do what our premisses say we ought to do (Sect. 3). And recall that good reasons explain obligations by showing that they follow from a maximal consistent subset of such reasons (Sect. 2). Putting these together, we get that you may conclude 'it ought to be that β' if β follows from a maximal consistent subset of the things your premisses say that you ought to do. And this is what $GP_{\Vdash}$ says.

[21] In making these remarks, I set aside issues concerning permission and conflicting obligations (see, for example, Brink 1994, pp. 235–6). While these issues can be addressed within a generalization of my system that I develop in work in progress, that generalization does not follow from anything that is said below. Thanks to Errol Lord for pressing me about permissions.

[22] Cases where a premiss says that a contradiction ought to be prevent this from being a biconditional. To strengthen this to a biconditional I would need to drop clause (i) in my definition of $\Vdash_{NonCut}$ or add the condition that none of your premisses say that it ought to be the case that a contradiction holds.

GP$_{\Vdash}$ also makes it is easy to see how my account avoids the problem that we began the paper with. To appreciate this, let us consider our original problem in a slightly different light. That problem relies on *ODS*$_{\Vdash}$ and *SOC*$_{\Vdash}$. A natural thought is that we can subsume both of these principles under a more general one:

> *General Ought Closure* (GOC$_{\Vdash}$): $O(\alpha_1), O(\alpha_2), \ldots , O(\alpha_n) \Vdash O(\beta)$ if $\{\alpha_1,\alpha_2, \ldots , \alpha_n\} \vdash_{PL} \beta$

Evidently, if we accept *GOC*$_{\Vdash}$, we are immediately committed to explosive reasoning because $\{a, \neg a\} \vdash_{PL} b$. So the problem that we began the paper with can be thought of as a problem about how to restrict *GOC*$_{\Vdash}$.[23]

What is nice about *GP*$_{\Vdash}$ is that it tells us what restriction my account puts on *GOC*$_{\Vdash}$: it restricts *GOC*$_{\Vdash}$ by only considering what follows from maximal *consistent* subsets of your premisses. And this restriction to maximal consistent subsets seems like the most natural restriction of *GOC*$_{\Vdash}$ that an advocate of conflicting obligations can adopt. What's more, with a little effort it is easy to prove that *GP*$_{\Vdash}$ entails that the reasoning consequence relation has the structural properties that my account says it has (for example, it is not cut and not closed under substitution).[24] And similarly it is not hard to prove that *GP*$_{\Vdash}$ is equivalent to *GOC*$_{\Vdash}$ if your premiss set does not contain conflicting obligations.[25]

---

[23] I owe this way of thinking of our original problem to Jeff Horty. Cf. Schotch and Jennings 1981, pp. 160–161.

[24] We say:

> *R* is *closed under substitution (in $\mathcal{L}$)* just in case if
>
> (1) A *R* φ (where A is a set of sentences of $\mathcal{L}$ and φ is sentence of $\mathcal{L}$) and
> (2) A* and φ* are the result of uniformly substituting some arbitrary sentence (in $\mathcal{L}$) for an atomic sentence in A and φ,
>
> then A* *R* φ*.

To see that $\Vdash_{NonCut}$ is not closed under substitution, we need only use *GP*$_{\Vdash}$ to check that while we accept *ODS*$_{\Vdash}$ we do not accept the following clam:

> (6$_{\Vdash}$) $O(\neg a), O(a \vee (a \,\&\, b)) \Vdash O(a \,\&\, b)$.

In this way, $GP_{\Vdash}$ is a succinct general statement of the main idea behind my account. And since this principle is simple, it will be easy for the interested reader to apply it for herself to various problems concerning reasoning.

5. Alternatives

As we know, my account solves the problem that I began the paper with by defining a non-cut reasoning consequence relation. But our result from section 1 suggested that our problem depends on both cut and monotonicity. And since nonmonotonic accounts of reasoning in other domains are familiar, this raises the question of whether there is a nonmonotonic account of deontic reasoning that can lead to an explanatory solution to our problem. So I want to show in this section how to develop a nonmonontic account of deontic reasoning that is a slight variant of $\Vdash_{\text{NonCut}}$.[26]

5.1 $\Vdash_{NonMon}$

The easiest way to think of this nonmonotonic account of deontic reasoning is as a relaxation of our original account of deontic reasoning. Recall that our original account says that for *every* premiss you reason as if you have good reasons to do what that premiss says you ought

---

This is the right result. Assuming that $O(a \vee (a \,\&\, b))$ is logically equivalent to $O(a)$ because a is logically equivalent to $a \vee (a \wedge b)$, ($6_{\Vdash}$) is equivalent to

$$(7_{\Vdash})\ O(\neg a), O(a) \Vdash O(a \,\&\, b)$$

which is a kind of explosion result that we should not accept (cf. McNamara 2004, p. 20; Goble 2009, p. 468).

[25] These comments require either one of the strengthenings to a biconditional mentioned in n. 22.

[26] What about other formal theories of reasoning that are not logics in the narrow sense? While I cannot discuss the details of these views here, it is worth pointing out that one such view known as adaptive deontic logic is strikingly similar to the views developed in this paper and especially this section. Adaptive logic can be thought of as a general theory of 'reasoning as if' (cf. Batens 2008). And adaptive deontic logics are a kind of non-monotonic formalism that has been used to approach our problem (Beirlaen, Straßer, Meheus 2013, Straßer, Meheus, Beirlaen 2012). What's more, this formalism treats the problems discussed here in a similar way. Unfortunately, I must leave it to further work to compare in detail the differences and similarities between my proposal and the interesting proposals made by adaptive logicians.

to do. The alternative view says that you only reason as if you have good reasons to do what *some privileged subset* of your premisses say that you ought to do.

To get a sense of why you might do this, consider the following abstract example. You have as premisses O(a) and O($a \vee b$). Now O($a \vee b$) is not generally sufficient to ensure the truth of O(a). But we may suppose O(a) is sufficient for the truth of O($a \vee b$). If that is right, it might seem like you should only reason as if you have a good reason to do *a*. After all, O(a) is sufficient to ensure that O($a \vee b$) holds. So you really only need to reason as if you have a good reason to do *a* for O(a) because O($a \vee b$) can 'come along for the ride'.

More generally: when you have some premisses, there will be some (perhaps improper) subset of these premisses, D, such that (i) D ensures the truth of the whole premiss set and (ii) no subset of D entails the truth of the whole premiss set—call this a *minimal deontically sufficient* subset. The idea is that you should only reason as if you have good reasons to do the things in a minimal deontically sufficient subset.

Formally, this idea just amounts to tweaking $\Vdash_{\text{NonCut}}$ so that it only applies to minimal deontically sufficient subsets:

$O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n) \Vdash_{\text{NonMon}} O(\beta)$ iff, for some minimal deontically sufficient subset D of $\{O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n)\}$, and for every $\mathfrak{R}$ such that:
(i) $\mathfrak{R} \mid\!\sim_{\text{Horty}} \delta$ for all $\delta \in D$
and
(ii) $\{!(x) \mid O(x) \in D\} \subseteq \mathfrak{R}$
$\mathfrak{R} \mid\!\sim_{\text{Horty}} O(\beta)$[27]

[27] I have chosen to write 'some minimal deontically sufficient subset' rather than writing 'every minimal … ' only because it is simplest to introduce the account in this way. All of the points made in the section also apply to the account that we get by writing 'every' rather than 'some'.

There are however two related problems that support adopting the 'every' proposal over the 'some' proposal (these remarks rely on the definition given in n. 28). First the 'some' proposal leads to an intuitively strange account of under we may conclude it O(a & b) from premise set that contains O(a) and O(b). We would have:

$O(a), O(b) \Vdash O(a \,\&\, b)$

As is easy to see, this account is exactly the same as $\Vdash_{\text{NonCut}}$ except that it says to reason as if you have good reasons to do what is in each minimal deontically sufficient subset. Of course, I have not presented a formal definition of a minimal deontically sufficient subset here. While this can be done, the intuitive idea of such a set is sufficient for our purposes so I will not bother to define the notion in the main text.[28]

Though I will leave the proof of this to the reader who is interested in working through all of the formal details, as promised this account solves our problem because it denies monotonicity. Specifically, while this account like $\Vdash_{\text{NonCut}}$ accepts ($1_{\Vdash}$), ($2_{\Vdash}$), and ($4_{\Vdash}$), it rejects:

($3_{\Vdash}$) $O(a)$ $O(\neg a)$, $O(a \vee b) \Vdash O(b)$

Since ($3_{\Vdash}$) follows from ($1_{\Vdash}$) assuming monotonicity, this shows that $\Vdash_{\text{NonMon}}$ is not monotonic. Thus, $\Vdash_{\text{NonCut}}$ and $\Vdash_{\text{NonMon}}$ alike solve the problem that we have been considering.

5.2 Two reasons to prefer $\Vdash_{\mathit{NonCut}}$

And we would lack:

$O(\neg a \ \& \ b)$, $O(a \ \& \ \neg b)$, $O(a)$, $O(b)$, $\Vdash O(a \ \& \ b)$

But none the less, if we add $O(\neg b)$ to the premisses we would strangely get $O(a \ \& \ b)$. That is, we have the following:

(*) $O(\neg b)$, $O(\neg a \ \& \ b)$, $O(a \ \& \ \neg b)$, $O(a)$, $O(b) \Vdash O(a \ \& \ b)$

This oddity is related to a second strange result. While we would not have:

($3_{\Vdash}$) $O(a)$,$O(\neg a)$, $O(a \vee b) \Vdash O(b)$

we would have:

(**) $O(\neg b)$, $O(a)$,$O(\neg a)$, $O(a \vee b) \Vdash O(b)$

That is, if we add $O(\neg b)$ to our premisses, we can suddenly conclude $O(b)$.

Though it takes work to show this, we only get (*) and (**) because the conclusions follows from one but not all of the minimal deontically sufficient subsets. So the 'every' proposal avoids these problems. I thank the anonymous referee who brought these problems to my attention.

[28] Here is how to do it:

A is a *deontically sufficient* subset of B iff (1) $A \subseteq B$ and (2) $A \vdash_{\text{Truth}} \beta$ for all $\beta \in B$.

A is a *minimal deontically sufficient subset* of B iff (1) A is a deontically sufficient subset of B and (2) there is no C such that $C \subset A$ and C is a deontically sufficient subset of B.

That said, I think that there are two reasons to prefer $\Vdash_{\text{NonCut}}$ to $\Vdash_{\text{NonMon}}$. The first reason concerns ($3_{\Vdash}$):

$$\mathrm{O}(a), \mathrm{O}(\neg a), \mathrm{O}(a \vee b) \Vdash \mathrm{O}(b)$$

($3_{\Vdash}$) is accepted by $\Vdash_{\text{NonCut}}$ but rejected by $\Vdash_{\text{NonMon}}$. My first reason for preferring $\Vdash_{\text{NonCut}}$ is that reasoning in accordance with ($3_{\Vdash}$) is good reasoning.

To see why I believe that this is good reasoning, consider the following variant of Smith's case. As before, Smith is a citizen of a just country. The laws of this country require Smith to either fight or perform alternative public service. So Smith ought to fight or serve, $\mathrm{O}(f \vee s)$. Also as before, Smith is a member of a pacifist religion. So Smith ought not to fight, $\mathrm{O}(\neg f)$. Now let us add that Smith comes from a family with a strong tradition of military service. So out of filial duty, Smith ought to fight, $\mathrm{O}(f)$.

In this case, it seems to me that we may conclude that Smith ought to serve, $\mathrm{O}(s)$. If Smith is deciding what to do, one of the steps in one of his (two incompatible) plans should be serving. Since this reasoning from $\mathrm{O}(f)$, $\mathrm{O}(\neg f)$, $\mathrm{O}(f \vee s)$ to $\mathrm{O}(s)$ is an instance of ($3_{\Vdash}$), I believe that we should accept ($3_{\Vdash}$). This is my first reason for preferring $\Vdash_{\text{NonCut}}$.

My second reason to prefer $\Vdash_{\text{NonCut}}$ is that it gives a more elegant account of deontic reasoning. To see this, consider the question of how we reason with some arbitrary premiss, $\mathrm{O}(a)$. The answer according to $\Vdash_{\text{NonCut}}$ is that we reason as if we have a good reason to do a. The answer according to $\Vdash_{\text{NonMon}}$, however, is not straightforward.

According to $\Vdash_{\text{NonMon}}$, a premiss first plays a holistic role in determining a minimal deontically sufficient subset. Then if the premiss is not in a minimal deontically sufficient subset, it plays no further role. If the premiss, O(*a*), is in a minimal deontically sufficient subset, then you reason as if you have a good reason to do *a*.

Obviously this is more complicated than the answer provided by $\Vdash_{\text{NonCut}}$. And one vivid way to see the undesirable results of this complication is to consider the restriction $\Vdash_{\text{NonMon}}$ places on *GOC*$_{\Vdash}$:

> $O(\alpha_1), O(\alpha_2), ..., O(\alpha_n) \Vdash O(\beta)$ if, for some D such that D is a minimal deontically sufficient subset of $\{O(\alpha_1), O(\alpha_2), \ldots, O(\alpha_n)\}$ and for some M such that M is a maximal consistent subset of $\{x \mid O(x) \in D\}$, M $\vdash_{PL} \beta$.

Evidently, this restriction lacks *GP*$_{\Vdash}$'s straightforward elegance. So this my second reason to prefer $\Vdash_{\text{NonCut}}$.

For these two reasons then, $\Vdash_{\text{NonCut}}$ may be the better account of deontic reasoning.[29] None the less, we should remember that both of these accounts of deontic reasoning provide explanatory solutions to our problems concerning reasoning with conflicting obligations.

6. Conclusion

Interestingly, having such an account of reasoning with conflicting obligations not only solves the problems that we set out to solve in this paper but also has consequences for issues in deontic logic. As I hinted at earlier, the rich literature in deontic logic has been concerned with developing logics that validate many inferences and allow for conflicting obligations without leading to inconsistency or explosion. It is important to know that since at least Chellas 1974 there have been logics that allow conflicting obligations without leading to inconsistency and

[29] Another potential problem with $\Vdash_{\text{NonMon}}$ is that it does not have the property of *cautious monotonicity*:

> *Cautious Monotonicity*: if A $\Vdash \varphi$ and A $\Vdash \gamma$, then A, $\varphi \Vdash \gamma$.

To see that $\Vdash_{\text{NonMon}}$ is not cautious monotonic, the interested reader may work through an example where A is {O(*a*), O(¬*a*), O(*b*)}, φ is O(*a* & *b*), and γ is O(¬*a* & *b*). Thus both $\Vdash_{\text{NonCut}}$ and $\Vdash_{\text{NonMon}}$ fail to be *cumulative relations* (relations that are both cut and cautious monotonic). In this respect, they are both subject to the worries discussed in n. 8. I thank Lou Goble for asking me whether $\Vdash_{\text{NonMon}}$ is cautious monotonic.

explosion. The problem that continues to occupy the attention of theorists is the problem of developing a logic that not only does this but also validates lots of inferences. In particular, deontic logicians have tried to develop logics that have enough valid inferences to explain the reasoning in cases like Smith's case and the speeding law case.

My account calls the idea that we have good reason to search for a logic that validates inferences corresponding to those cases in doubt. Because my account provides a direct explanation of the reasoning in these cases, there is little reason for us to insist on using the resources of deontic logic to explain them.[30] This means that it is mistaken to think that we should be unsatisfied with a deontic logic that does not validate enough inferences to explain these cases.

The upshot is that the account of reasoning with conflicting obligations that I have provided alleviates the need to move beyond the logics of conflicting obligations that we have had since the mid-1970s. Thus, one of the consequences of the account developed here is that we can be satisfied with these relatively weak logics.

Let us take stock: What we saw in this paper is that those who believe that there are conflicting obligations face a problem about understanding reasoning that cannot be adequately resolved by appeal to deontic logic. We then developed an account of reasoning with conflicting obligations that provides an explanatory solution to this problem. And we saw that this account looks to be promising general theory of deontic reasoning. Ultimately, it is this ability to provide

[30] So I disagree with Lou Goble (2005) who worries about using a consequence relation that is not a logical consequence relation to account for Smith's case: 'It does not give an account of the ordinary validity of the original argument regarding Smith's service. Instead it offers a weaker substitute, perhaps to explain the appearance of (ordinary) validity' (p. 407). My view is that in the first instances our firmest judgement about Smith's case is that it is a piece of good reasoning. As I see it, deontic logicians have taken this as evidence that the relevant inference is valid. But since my account directly explains why it is good reasoning, there is no reason to try to account for this in terms of validity.

a general explanatory solution to our problems concerning deontic reasoning that I think is the main piece of evidence in favour my account.

And while I admit more work remains to be done to provide a fully worked out picture of deontic reasoning, the work I have done here is enough to equip the advocate of conflicting obligations with an account that they can use to solve the main problems about reasoning for their approach. More generally, I hope to have shown that there are distinctive problems concerning reasoning with conflicting obligations and that addressing these problems requires us to think hard about issues in both ethics and deontic logic.[31]

[31] Some of the ideas in this paper have been presented at the 2011 central APA session on deontic logic, the USC Speculative Society, Formal Ethics 2012, and the USC Deontic Modals Workship. I thank the organizers, audience, and participants of these conferences. For helpful comments and discussion on this ideas of this paper, I thank Justin Dallman, Steve Finlay, Ben Lennertz, Errol Lord, Matt Lutz, Ryan Millsap, Kenny Pearce, Doug Portmore, Indrek Reiland, Jacob Ross, Barry Schein, Johannes Schmitt, Sam Shpall, Justin Snedegar, Scott Soames, Gabriel Uzquiano-Cruz, Aness Webster, Ralph Wedgwood, and especially Lou Goble and Jeff Horty. I also thank the anonymous referees at *Mind* as well as Thomas Baldwin for detailed feedback. I would most of all like to thank Mark Schroeder who has provided me with invaluable advice and criticism on every issue at every stage of this project. Finally, I thank the USC Provost's PhD Fellowship and the Russell Fellowship for support.

References

Batens, Diderik 2007. 'A Universal Logic Approach to Adaptive Logics'. *Logica Universalis*, 1, pp. 221–42.

Beirlaen, M., C. Straßer, and J. Meheus 2013: 'An Inconsistency-Adaptive Deontic Logic for Normative Conflicts'. *Journal of Philosophical Logic*, 42, pp. 285–315.

Brink, David 1994: 'Moral Conflict and Its Structure'. *Philosophical Review*, 103, pp. 215–47.

Burgess, John 2009: *Philosophical Logic*. Princeton: Princeton University Press.

Cariani, Fabrizio forthcoming: '*Ought* and Resolution Semantics'. Forthcoming in *Nous*.

Castañeda, Henri Neri. 1981: 'The Paradoxes of Deontic Logic'. In Hilpinen 1981, pp. 37–86.

Chellas, Brian 1974: 'Conditional Obligation'. In Stenlund 1974, pp. 23–33.

—1980: *Modal Logic*. Cambridge: Cambridge University Press.

Dancy, Jonathan 2004: *Ethics without Principle*. Oxford: Oxford University Press.

Gabbay, D., C. Hodger, and J. Robinson 1994: *Handbook of Artificial Intelligence and Logic Programming*, volume 3. Oxford: Oxford University Press.

Gabbay, D., J. Horty, R. van der Meyden, and L. van der Torre (eds) forthcoming: *Handbook of Deontic Logic and Normative Systems*. London: College Publications

Goble, Lou 2005: 'A Logic for Deontic Dilemmas'. *Journal of Applied Logic*, 3, pp. 461–83.

—2009: 'Normative Conflicts and the Logic of 'Ought''. *Nous*, 43, pp. 450–89.

—forthcoming: 'Prima Facie Norms, Normative Conflicts, and Dilemmas'. Forthcoming in Gabbay, Horty, van der Meyden, and van der Torre forthcoming.

Goodman, Nelson 1955: *Fact, Fiction, and Forecast*. Cambridge: Harvard University Press.

Gowans, Christopher (ed) 1987a: *Moral Dilemmas.* Oxford: Oxford University Press.

—1987b: 'Introduction'. In his 1987a, pp. 3–33.

Hansen, Jörg 2005: 'Conflicting Imperatives and Dyadic Deontic Logic'. *Journal of Applied Logic*, 3, pp. 484–511.

—2007: 'Prioritized Conditional Imperatives'. *Autonomous Agents and Multi-Agent Systems*, 17, pp. 11–35.

Hilpinen, Risto (ed) 1971: *Deontic Logic: Introductory and Systematic Readings.* Dordrecht: D. Reidel.

—(ed) 1981: *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*. Dordrecht: D. Reidel.

— and Dagfinn Føllesdal 1971: 'Deontic Logic: An Introduction'. In Hilpinen 1971, pp. 1–35

Horty, John 1993: 'Deontic Logic as Founded on Nonmonotonic Logic'. *Annals of Mathematics and Artificial Intelligence*, 9, pp. 69–91.

—1994: 'Moral Dilemmas and Nonmonotonic Logic'. *Journal of Philosophical Logic*, 23, pp. 35–65.

—1997: 'Nonmonotonic Foundations for Deontic Logic'. In Nute 1997, pp. 17–44.

—2003: 'Reasoning with Moral Conflicts'. *Nous*, 37, pp. 557–605.

—2012: *Reasons as Defaults*. Oxford: Oxford University Press.

Hume, David 1739: *A Treatise of Human Nature*, P.H. Nidditch (ed), 2nd edition. Oxford: Oxford University Press. 1978.

Lemmon, E.J 1962: 'Moral Dilemmas'. *The Philosophical Review*, 70, pp. 139–158.

Makinson, David 1994: 'General Patterns in Nonmonotonic Reasoning'. In Gabbay, Hodger, and Robinson 1994, pp. 35–110.

Marcus, Ruth Barcan 1980: 'Moral Dilemmas and Consistency'*Journal of Philosophy*, 77, pp. 121–36.

McConnell, Terrance 2010: 'Moral Dilemmas'. *Stanford Encyclopedia of Philosophy* [online encyclopedia], Summer 2010 Edition, <http://plato.stanford.edu/archives/sum2010/entries/moral-dilemmas/>, accessed 23 July 2010.

McNamara, Paul. 2004: 'Agential Obligation as Non-Agential Personal Obligation Plus Agency'. *Journal of Applied Logic*, 2, pp. 117–52.

Nagel, Thomas 1970: *The Possibility of Altruism*. Princeton: Princeton University Press.

—1979: *Mortal Questions*. Cambridge: Cambridge University Press.

—1979: 'The Fragmentation of Value'. In Gowans 1987a, pp. 174–87. Originally published in Nagel 1979.

Nair, Shyam MS-a: 'Conflicting Reasons, Unconflicting Obligations'

Nair, Shyam MS-b: 'Must Good Reasoning be Cumulative Transitive?'

Nute, Donald (ed) 1997: *Defeasible Deontic Logic*. Dordrecht: Kluwer.

— and Xiaochang Yu 1997: 'Introduction'. In Nute 1997, pp. 1–18.

Parfit, Derek 2011: *On What Matters*. Oxford: Oxford University Press.

Pietroski, Paul 1993: 'Prima Facie Obligation, Ceteris Paribus Laws in Moral Theory'. *Ethics*, 103, pp. 489–515.

Raz, Joseph 1975: *Practical Reasoning and Norms*. Oxford: Oxford University Press. 2002. Originally published by in 1975 by Hutchinson & Co.
Reiter, Raymond 1980: 'A Logic for Default Reasoning'. *Artificial Intelligence*, 13, pp. 81–132.
Ross, Alf 1941: 'Imperatives and Logic'. *Theoria*, 7, pp. 53–71.
Scanlon, T.M. 1998: *What We Owe to Each Other.* Cambridge: Harvard University Press.
Schotch, Peter and Raymond Jennings 1981: 'Non-Kripkean Deontic Logic'. In Hilpinen 1981, pp. 149–162.
Sinnott-Armstrong, Walter 1988: *Moral Dilemmas*. Oxford: Basil Blackwell.
Stenlund, Sören (ed) 1974: *Logical Theory and Semantic Analysis: Essays Dedicated to Stig Kanger on His Fiftieth Birthday*. Boston: D. Reidel
Straßer, C., J. Mehus, and M. Beirlaen 2012: 'Tolerating Deontic Conflict by Adaptively Restricting Inheritance'. *Logique et Analyse*, 219, pp. 477–506.
van der Torre, Leendert and Yao Hua Tan 2000: 'Two Phase Deontic Logic'. *Logique et Analyse,* 171-2, pp. 411–56.
van Fraassen, Bas 1973: 'Values and the Heart's Command'. *Journal of Philosophy*, 70, pp. 5–19.
Williams, Bernard 1965: 'Ethical Consistency'. *Proceedings of Aristotelian Society, Supplementary Volumes*, 39, pp. 103–138.