

The Irreducibility of Personal Obligation¹

How are claims about what people *ought to do* related to claims about what *ought to be the case*? That is, how are claims about of personal obligation, of the form *s ought to ϕ* , related to claims about impersonal obligation, of the form *it ought to be the case that p* ?² Many philosophers and logicians have held that the former type of claim can be reduced to the latter. In particular, they have held a view known as the Meinong-Chisholm Reduction, which is sometimes stated as follows:

MCR: ‘*s ought to ϕ* ’ is logically equivalent to ‘it ought to be the case that *s ϕ* .’

and which is sometimes stated, more restrictedly, as follows:

MCR_{stit}: ‘*s ought to see to it that p* ’ is logically equivalent to ‘it ought to be the case that *s sees to it that p* .’³

In the next two sections, I will present counterexamples to MCR and to MCR_{stit}, respectively, counterexamples that involve multiple agents. In the third section, I will generalize the results of the first two sections, and argue that no reduction of personal obligation to impersonal obligation can succeed. And in the conclusion I will briefly discuss the significance of these results.

1. A Counterexample to MCR

Consider the following case:

Campus Visit: Professor Wagstaff is visiting Huxley University where he has been offered a professorship. Baravelli is the chair of the hiring committee, and it is his job to accompany Wagstaff throughout his visit. Wagstaff has promised that he will go to departmental colloquium at noon, and he has no good reason to do otherwise. But instead he goes to the saloon on the other side of town, and Baravelli is unable to prevent this course of action.

In this case, the following claims all appear to be true at noon:

- (1) Wagstaff goes to the saloon.
- (2) Baravelli ought to accompany Wagstaff.
- (3) Under Baravelli's circumstances, Baravelli's going to the saloon is a necessary means to his accompanying Wagstaff.
- (4) Wagstaff ought to go to the colloquium.
- (5) Necessarily, if Wagstaff goes the colloquium, and Baravelli accompanies him, then Baravelli does not go to the saloon.

These five claims are jointly inconsistent with the conjunction of MCR and the following four principles:

- O1. For any two propositions, p and q , if it ought to be the case that p , and it ought to be the case that q , then it ought to be the case that (p and q).
- O2. For any two propositions, p and q , if it ought to be the case that p , and if (necessarily, if p then q), then it ought to be the case that q .
- O3. There is no proposition, p , such that it ought to be the case that (p and not- p).
- O4. For any agent, s , and any two action types, ϕ and ψ , if s ought to ϕ , and if, under s 's circumstances, her ψ -ing is a necessary means to her ϕ -ing, then s ought to ψ .

This inconsistency can be shown by the following derivation:

- (6) Baravelli ought to go to the saloon. [2, 3, O4]
- (7) It ought to be the case that Baravelli goes to the saloon. [6, MCR]
- (8) It ought to be the case that Wagstaff goes to the colloquium. [4, MCR]
- (9) It ought to be the case that Baravelli accompanies Wagstaff. [2, MCR]

(10) It ought to be the case that Wagstaff goes to the colloquium and Baravelli accompanies him. [8, 9, O1]

(11) It ought to be the case that Baravelli does not go to the saloon. [5, 10, O2]

(12) It ought to be the case that Baravelli goes to the saloon and Baravelli does not go to the saloon. [7, 11, O1]

Clearly, (12) is incompatible with O3. And so if we assume claims (1) through (5), as well as principles O1 through O4 and MCR, we can derive a contradiction. Thus, if we accept these claims and principles, we must reject MCR.

The defender of MCR might respond by denying that claims (1) through (5) are all true in *Campus Visit*. Now claim (1) is true by stipulation. And claim (3) follows from the description of the case. For Baravelli's circumstances include the fact that Wagstaff goes to the saloon, and that he cannot prevent Wagstaff from doing so. And under these circumstances, a necessary means to Baravelli's accompanying Wagstaff is that he goes to the saloon. Claim (4) also follows from the description of the case, since it has been stipulated that Wagstaff has promised to go to the colloquium (which provides at least a *pro tanto* reason for him to do so), and that he has no reason not to do so. And claim (5), we may assume, follows from the geographical separation of the colloquium and the saloon, which implies that Baravelli cannot now go both to the colloquium and to the saloon. And so if any of these five claims can be contested, it is claim (2). One might object to this claim as follows:

In *Campus Visit*, we are concerned with the actions Baravelli and Wagstaff perform at noon, and so we are concerned with the simultaneous actions of two agents. But if these actions are genuinely simultaneous, then whether Baravelli accompanies Wagstaff will depend not only on the choices Baravelli makes, but also on the choices Wagstaff makes. Consequently, whether Baravelli accompanies Wagstaff isn't up to Baravelli. And if it isn't up to Baravelli whether he accompanies Wagstaff, then it cannot be the case that Baravelli ought to accompany Wagstaff. And so (2) is false.

This objection, however, cannot coherently be made by anyone who accepts MCR in conjunction with O1 and O2. For consider a case in which Mary ought to brush her teeth at noon, and Peter ought to brush his teeth at noon. In this case, by MCR, it ought to be the case that Mary brushes her teeth at noon, and it ought to be the case that Peter brushes his teeth at noon. And so, by O1, it ought to be the case that Mary brushes her teeth at noon and Peter brushes his teeth at noon. And necessarily, if Mary and Peter both brush their teeth at noon, then Mary brushes her teeth at the same time as Peter brushes his. And so, by O2, it ought to be the case that Mary brushes her teeth at the same time as Peter brushes his. Hence, by MCR, Mary ought to brush her teeth at the same time as Peter brushes his. But whether Mary brushes her teeth at the same time as Peter brushes his depends not only on Mary's choices, but also on Peter's choices. Therefore, anyone who accepts O1, O2 and MCR must allow that it can be the case that an agent ought to do something that requires the cooperation of other agents. And so anyone who accepts these principles cannot object to (2) on the grounds that it isn't up to Baravelli whether he accompanies Wagstaff.

It seems, therefore, that the defender of MCR cannot plausibly deny that claims (1) through (5) are all true in Campus Visit. The only other way she can resist the rejection of MCR is by rejecting one or more of principles O1 through O4. She might appeal to the standard paradoxes of deontic logic in order to provide independent grounds for rejecting some of these principles. However, I argue in the appendix that this response is unsuccessful. For several of the standard paradoxes rely on the implicit assumption of MCR, and so they show only that my principles cannot be held in conjunction with MCR, which is of course a conclusion I accept. And the remaining paradoxes can be resolved by restricting the scope of these principles in ways that have no effect on my argument.

Note that, while MCR is a formula of logical equivalence, the argument given above relies only on the entailment from left to right. Thus, it serves equally as an argument against the following one-way entailment:

MCR_{L→R}: 's ought to ϕ ' logically entails 'it ought to be the case that s ϕ s.'

However, the Campus Visit case also appears to provide a counterexample to the converse entailment, namely:

MCR_{R→L}: ‘it ought to be the case that $s \phi$ ’ logically entails ‘ s ought to ϕ .’

To argue for falsity of this converse entailment, we will need some principle that enables us to derive claims about impersonal obligation from claims about personal obligation. It will be useful to begin with a definition. Let us say that a possibility (or a way that it is practically possible for the world now to be), w , is *wrongless* just in case everyone in w does what she ought to do in w . We may now state a very weak principle connecting personal and impersonal obligation, as follows:

O5: If, in actuality, s ought to ϕ , and if $s \phi$ s in every wrongless possibility, then it ought to be the case that $s \phi$ s.

This principle is clearly much weaker than MCR_{L→R}. And it is immune to the problem we raised for the latter principle. This problem arose because MCR_{L→R} licenses the inference from (6) Baravelli ought to go to the saloon to (7) it ought to be the case that Baravelli goes to the saloon. But O5 does not license this inference. For while, in actuality, Baravelli ought to go to the saloon, there is a wrongless possibility in which Baravelli does not go to the saloon, namely, a possibility in which Baravelli and Wagstaff both go to the colloquium. O5 does, however, allow us to infer (8) from (4) together with the fact that Wagstaff goes to the colloquium in every wrongless possibility. Similarly, O5 allows us to infer (9) from (2) together with the fact that Baravelli accompanies Wagstaff in every wrongless possibility. Hence, we can make the inference from (8) and (9) to (11): *it ought to be the case that Baravelli does not go to the saloon*. And MCR_{R→L} licenses the inference from (11) to the false conclusion that Baravelli ought not to go to the saloon. Therefore, if we accept O5, along with principles O1 through O4, we should reject MCR_{R→L}. Hence, Campus Visit motivates the rejection of both directions of entailment implied by MCR.

2. A Counterexample to MCR_{stit}

Perhaps we can solve this problem for MCR by restricting it in the following way:

MCR_{stit} : ‘ s ought to see to it that p ’ is logically equivalent to ‘it ought to be the case that s sees to it that p .’

This is a specialization of MCR in which ϕ ranges not over any way of acting, but only over ways of acting that consist in seeing to it that some proposition obtains. John Horty has persuasively argued that MCR_{stit} has a number of advantages over the more general MCR. He has also clarified MCR_{stit} by offering an account of what it is for an agent to see to it that p , for an arbitrary proposition p . On his account, for s to see to it that p is for s to act in such a way as to ensure that p obtains regardless of how other agents may act.⁴

When MCR_{stit} is read in this way, it may be able to avoid the problem for MCR that we saw in the previous section. For this problem arises because it seems that in *Campus Visit*, Baravelli ought to accompany Wagstaff. And hence, it follows, by MCR, that it ought to be the case that he accompanies Wagstaff. But no analogous conclusion can be drawn from MCR_{stit} . Of course, if we assumed that Baravelli ought to *see to it* that he accompanies Wagstaff, then we could indeed infer, from MCR_{stit} , that it ought to be the case that Baravelli sees to it that he accompanies Wagstaff. But Baravelli cannot see to it that he accompanies Wagstaff, since there is nothing he can do that would ensure that he accompanies Wagstaff regardless of how Wagstaff acts. And so it is false that Baravelli ought to see to it that he accompanies Wagstaff. And so we cannot infer, from MCR_{stit} , that it ought to be the case that he sees to it that he accompanies Wagstaff.

One objection that could be raised to MCR_{stit} is that it provides a reduction of only a very special class of claims about personal obligation, namely, claims about what agents ought to *see to*, or act in such a way as to ensure. And yet we often say things of the form “ X ought to dance with Y ,” or “ X ought to marry Y ,” or “ X ought to buy a loaf of bread from Y .” And in such cases, what we are saying X ought to do is some action that she is not in a

position to guarantee, or to see to it that she performs, since it requires the cooperation of *Y*. Thus, there is a wide variety of claims about personal obligation for which MCR_{stit} fails to provide a reduction. But while MCR_{stit} does not by itself provide a complete reduction of personal obligation claims, it might perhaps do so when conjoined with some appropriate bridge principle connecting what an agent ought to do with what she ought to see to. Perhaps, for example, an agent ought ϕ just in case, in her actual circumstances, if she were to fulfill all her obligations to see to it that propositions obtain, she would thereby ϕ .

A more serious problem for MCR_{stit} is that it appears to have unacceptable consequences. Horty himself has argued that MCR_{stit} has unacceptable consequences in cases involving gambles. His argument, however, rests on a particular analysis of impersonal obligation that the defender of MCR_{stit} needn't accept.⁵ Here I will present what I take to be a more theory-neutral objection to MCR_{stit} . My argument is based on the following case:

Coordination. Wagstaff has been offered a professorship at Huxley University, and he has promised to meet Baravelli at 4 p.m. to sign the contract. Baravelli has likewise promised to meet Wagstaff at 4 p.m. If they fail to meet at this time, then the deal will fall through, which will be worse for everyone. Unfortunately, they didn't specify where they would meet, and there are two possible places where they could meet, at opposite ends of campus: Baravelli's office, and the main office.

Here there are two options available to Baravelli: he can either go to his own office, or he can go to the main office. And which office he goes to does not depend on the actions of any other agent. Thus, if he were to go to his own office, he would thereby see to it that he goes to his own office, and if he were to go to the main office, he would thereby see to it that he goes to the main office. And so we can represent Baravelli's choice situation as one in which his two options are *seeing to it that he goes to Baravelli's office* and *seeing to it that he goes to the main office*. And we can likewise represent Wagstaff's choice situation as one in which he has these same two options.

On its own, MCR_{stit} , does not allow us to derive any normative consequences concerning *Coordination*, but it does allow us to derive some very peculiar normative consequences if

we adopt any of a range of very standard normative theories. Consider the following normative theory:

Objective Act Consequentialism: Necessarily, in any given choice situation, an agent ought to perform whatever action would have the best consequences under her circumstances.

where these circumstances include all those facts over which she has no control, including the independent actions of other agents. It follows from Objective Act Consequentialism that

(13) Baravelli ought to see to it that he goes to the main office.

For, since we are assuming that Wagstaff goes to the main office, it follows that, under Baravelli's circumstances, his going to the main office is the action that would have the best consequences. Hence, by MCR_{stit} ,

(14) It ought to be the case that Baravelli sees to it that he goes to the main office.

Similarly, it follows from Objective Act Consequentialism that

(15) Wagstaff ought to see to it that he goes to Baravelli's office.

For, since we are assuming that Baravelli goes to Baravelli's office, it follows that, under Wagstaff's circumstances, his going to Baravelli's office is the action that would have the best consequences. Hence, by MCR_{stit} ,

(16) It ought to be the case that Wagstaff sees to it that he goes to Baravelli's office.

And so, by O1,

(17) It ought to be the case that Baravelli sees to it that he goes to the main office and Wagstaff sees to it that he goes to Baravelli's office.

However, in *Coordination*, necessarily, if the two agents act in this manner, then the deal will fall through, which is the worst possible outcome. That is,

- (18) Necessarily, if Baravelli sees to it that he goes to the main office and Wagstaff sees to it that he goes to Baravelli's office, then the worst possible outcome will obtain.

And so we can infer, by O2,

- (19) It ought to be the case that the worst possible outcome obtains.

Thus, if we accept O1, O2, and MCR_{stit} , then we must maintain that it follows logically from Objective Act Consequentialism that, in *Coordination*, it ought to be case that the worst possible outcome obtains. And this is a rather surprising result: while it is generally agreed that Objective Act Consequentialism has counterintuitive implications, the implication that the worst possible outcome ought to obtain is not generally thought to be among them. Indeed, this implication would seem to undermine Objective Act Consequentialism, for if, in *Coordination*, it ought to be the case that the worst possible outcome obtains, then it's hard to see how each agent could have a duty to act in such a way as to make things go as well as possible.

Furthermore, if we accept MCR_{stit} , along with (O1) and (O2), then we must conclude that Objective Act Consequentialism has even more surprising implications. Let q be the proposition that Baravelli sees to it that he goes to the main office and Wagstaff sees to it that he goes to Baravelli's office. Thus, (17) is equivalent to

- (17*) It ought to be the case that q .

Now the following two claims are clearly true:

- (20) Necessarily, if q , then Baravelli sees to it that he goes to the main office.
- (21) Necessarily, if Baravelli sees to it that he goes to the main office, then under Wagstaff's circumstances, the option available to Wagstaff with the best consequences is to see to it that he goes to the main office.

And from Objective Act Consequentialism, we may infer

- (22) Necessarily, if, under Wagstaff's circumstances, the option available to Wagstaff with the best consequences is to see to it that he goes to the main office, then Wagstaff ought to see to it that he goes to the main office.

But assuming MCR_{stit} ,

- (23) Necessarily, if Wagstaff ought to see to it that he goes to the main office, then it ought to be the case that Wagstaff sees to it that he goes to the main office.

And so it follows, from (20), (21), (22) and (23) that

- (24) Necessarily, if q , then it ought to be the case that Wagstaff sees to it that he goes to the main office.

Moreover,

- (25) Necessarily, if Wagstaff sees to it that he goes to the main office, then not- q .

And from O2, by the necessitation rule,

- (26) Necessarily, (if it ought to be the case that Wagstaff sees to it that he goes to the main office, and if, necessarily (if Wagstaff sees to it that he goes to the main office, then not- q), then it ought to be the case that not- q)

And so we can infer, from (25) and (26),

- (27) Necessarily, if it ought to be the case that Wagstaff sees to it that he goes to the main office, then it ought to be the case that not- q .

Hence, from (24) and (27),

- (28) Necessarily, if q , then it ought to be the case that not- q .

Thus, q is a proposition that ought to obtain, and yet, necessarily, if it were to obtain, it would be impermissible. And so, assuming O1, O2, and MCR_{stit} , it follows from Objective Act Consequentialism that there are obligations that cannot be permissibly satisfied.

Nor is this a special problem for Objective Act Consequentialism. Assuming (O1), (O2), and MCR_{stit} , any normative theory that allows for the possibility of coordination problems of the form just considered will entail that there are obligations that cannot be permissibly satisfied. Consider, for example, a deontological theory that includes the following principle:

Fidelity Principle: Necessarily, in any given situation, an agent ought to do what is required in order to break the fewest promises she can under her circumstances.

where these circumstances include all those facts over which she has no control, including the independent actions of other agents. The Fidelity Principle, like Objective Act Consequentialism, entails that, in *Coordination*, since Baravelli goes to his own office and Wagstaff goes to the main office, Baravelli ought to go to the main office and Wagstaff ought to go to Baravelli's office. And the Fidelity Principle likewise entails that, in *Coordination*, necessarily, if Baravelli goes to the main office and Wagstaff goes to Baravelli's office, then Baravelli ought to go to his own office and Wagstaff ought to go to the main office. Therefore, assuming O1, O2, and MCR_{stit} , the Fidelity Principle, like Objective Act Consequentialism, entails that there are circumstances in which it ought to be the case that q , and yet, necessarily, if q were true, then it ought to be the case that not- q .

Now there is a clear sense in which a prescription that cannot be permissibly satisfied is an unacceptable prescription. Thus, if we accept O1, O2, and MCR_{stit} , we must conclude that in situations like *Coordination*, a wide range of standard normative theories make prescriptions which, by their own lights, are unacceptable. And so if we want to avoid this result, and we don't want to reject O1 or O2, then we must reject MCR_{stit} . And, since MCR is stronger than MCR_{stit} , if we reject MCR_{stit} we must also reject MCR.

3. Why Personal Obligation is Irreducible to Unconditional Impersonal Obligation

I have argued against two ways of reducing personal to impersonal obligation, namely MCR and MCR_{stit} . I will now argue, more generally, that no attempt to analyze personal

obligation in terms of impersonal obligation can succeed, since two theories can differ in their implications concerning what agents ought to do without differing in their implications concerning what ought to be the case.

Let T1 be a normative theory that can be expressed thus:

T1: At any given time, if a job candidate is visiting a department where a colloquium is about to begin,

- (a) the job candidate ought to go to the colloquium;
- (b) the chair of the hiring committee ought to go to the colloquium;
- (c) it ought to be the case that the obligations stated in (a) and (b) are satisfied.

And let T2 be a normative theory that can be expressed thus:

T2: At any given time, if a job candidate is visiting a department where a colloquium is about to begin,

- (a*) the job candidate ought to go to the colloquium;
- (b*) the chair of the hiring committee ought to go wherever the job candidate goes;
- (c*) it ought to be the case that the obligations stated in (a*) and (b*) are satisfied.

Theories T1 and T2 have the exact same implications concerning what ought to be the case. For each of these theories applies only in contexts in which a job candidate is visiting a department where a colloquium is about to begin. And in such contexts, what ought to be the case, according to T1, is whatever is required for the joint satisfaction of the obligations stated in (a) and (b). And what ought to be the case, according to T2, is whatever is required for the joint satisfaction of the obligations stated in (a*) and (b*). But the joint satisfaction of the former pair of obligations is logically equivalent to the joint satisfaction of the latter pair.

And so what ought to be the case, according to T1, coincides with what ought to be the case, according to T2.

However, these two theories disagree concerning what agents ought to do. In *Campus Visit*, where Wagstaff goes to the saloon instead of the colloquium, T1 entails that Baravelli ought to go to the colloquium, whereas T2 entails that Baravelli ought to go to the saloon. Therefore, since T1 and T2 agree in all their implications concerning what ought to be the case, and differ in their implications concerning what agents ought to do, it follows that claims about what agents ought to do are not logically equivalent to claims about what ought to be the case.

Of course, if we were to accept MCR, along with (O1) and (O2), then we would have to maintain that T1 and T2 are logically equivalent theories. It seems clear, however, that T1 and T2 are not logically equivalent, as they differ in their implications concerning what Baravelli ought to do in *Campus Visit*. And so this gives us further reason to reject MCR.

So far I have been focusing on the view that claims about personal obligation are *logically equivalent* to claims about impersonal obligation. I have done so because the proposed reductions of personal to impersonal obligation generally take the form of claims about logical equivalence. But one might propose a reduction that is not meant to be analytic. One might hold that while claims about personal obligation may not be logically equivalent to claims about impersonal obligation, these claims have the same truth conditions. Or one might hold the still weaker view that the truth values of claims about personal obligation supervene on the truth values of claims about impersonal obligation. Might these weaker claims be true?

I believe the answer is no, for the following reason. As *Campus Visit* and *Coordination* both illustrate, it appears that what choice a given agent ought to make, at a time, t , may depend on what choices other agents actually make at t . It would seem, however, that what *ought to be the case* at t does not depend on what choices agents actually make at t . But if

what an agent ought to do depends on the simultaneous choices of other agents, while what ought to be the case does not so depend, then it follows that what agents ought to do cannot supervene on what ought to be the case. And so it follows that personal obligation cannot be reduced to impersonal obligation.⁶

4. Conclusion

I have argued that statements of the form ‘ s ought to ϕ ’ cannot be reduced to statements of the form ‘it ought to be the case that p .’ One might suppose, however, that the former statements can be reduced to claims of the form ‘conditional on q , it ought to be the case that p .’ In other words, one might suppose that claims about personal obligation can be reduced to claims about *conditional* impersonal obligation. In particular, one might propose a reduction according to which ‘ s ought to ϕ ’ is true just in case, *conditional on all those facts that are outside s ’s control*, it ought to be the case that s ϕ s (or, alternately, s ought to ϕ just in case, conditional on all those facts that are outside s ’s control, it ought to be the case that s *sees to it* that s ϕ s). Such an account would allow that what a given agent, s , ought to do can depend on what other agents in fact do, so long as the actions of these other agents are outside s ’s control. Thus, it could allow that, while it ought to be the case that Baravelli goes to the colloquium, he nonetheless ought to go to the saloon, since, conditional on all the facts outside his control—including the fact that Wagstaff goes to the colloquium—it ought to be the case that Baravelli goes to the colloquium as well. Hence, such an account would avoid many of the problems we have discussed.

However, I argue elsewhere [Ross, manuscript] that no such reduction of personal obligation to conditional impersonal obligation can succeed, at least given standard accounts of conditional impersonal obligation. The basic problem is that, on standard accounts of conditional impersonal obligation, it is understood in terms of some ordering relation that obtains among propositions. The relation in question has been described in many ways: in [Hansson, 1969], it is described as the *more ideal than* relation; in [Lewis, 1973], it is

described as the *better than* relation; in [Zimmerman, 1996] it is described as the *more deontically valuable than* relation; and in [Hilpinen, 2001] it is described as the *normatively less objectionable than* relation. But in each case it is understood as a single relation by which deontically evaluable propositions are ordered. The difficulty for the proposed reduction of personal obligation, however, is that according to non-consequentialist normative theories, facts about personal obligation do not supervene on any single ordering relation: to represent such facts in terms of ordering relations, we would need a different ordering relation for each agent. Hence, if we accept any of these standard accounts of conditional impersonal obligation, and we want to allow for the logical coherence of non-consequentialist normative theories, then we must conclude that claims about personal obligation are irreducible to claims about conditional impersonal obligation, just as they are irreducible to claims about unconditional impersonal obligation.

If I am right that the *ought* of personal obligation is irreducible to the *ought* of impersonal obligation, then there is reason to regard the former *ought* as the more central of the two. For only the former *ought* plays a direct role in the guidance of action. This can be seen if we consider a case in which personal and impersonal obligation come apart. Suppose that in *Campus Visit*, Baravelli knows that Wagstaff is going to the saloon and that he ought to accompany Wagstaff, and so he knows that he ought to go to the saloon. This knowledge should guide him to go to the saloon. But suppose he also knows that it ought to be the case that Wagstaff goes to the colloquium and that he accompanies Wagstaff, and so he knows that it ought to be the case that he goes to the colloquium. This knowledge should not guide him to go to the colloquium. In such cases in which the two kinds of obligation come apart, it is clear that the agent's actions should be directly guided by his beliefs about what he ought to do, not by his beliefs about what it ought to be the case that she does. Hence, the normative question that guides deliberation must be a question about personal obligation: it is the question 'what ought I to do.' And so *practical reasons*, or reasons that can figure in rational practical deliberation, must be reasons that bear on this question of personal

obligation. Hence, the *ought* of personal obligation can plausibly be regarded as the fundamentally normative *ought*—as both irreducible to, and more central than, the *ought* of impersonal obligation.⁷

The methodological implications of this conclusion are obvious. For if the *ought* of personal obligation is the fundamentally normative *ought*, then it should be of central concern to deontic logicians. And if it is irreducible to the *ought* of impersonal obligation, then it cannot be represented by a propositional operator, but only by a two-place predicate one of whose arguments denotes an agent. Moreover, the conclusion that personal obligation is irreducible to impersonal obligation may have much broader implications as well. For it suggests that the concept of the agent or person or self may be indispensable to anyone engaged in deliberation. There is a venerable tradition according to which the concept of the self is an illusion that we should strive to overcome. But the very possibility of such an overcoming is called into question by the above arguments. For if these arguments are sound, then in deciding what to do, our deliberative question cannot be settled by a judgment about what it ought to be the case that we do, but only by a judgment about what *we ought to do*. Here, the concept of ourselves figures both as the subject of the action to be decided upon, and as the subject of the obligation to perform this action. Hence, if personal obligation is irreducibly personal, then even if it were possible to adopt a conceptual scheme in which we could understand *our doing* of an action in impersonal terms that invoked no concept of ourselves as the agent of the action, such a scheme would nonetheless be inadequate to represent the *ought* that is relevant to deliberation. As rational agents, therefore, we could accept no such scheme.⁸

Appendix: Why the Standard Deontic Paradoxes do not Threaten my Argument

Three of the premises on which my arguments are based, namely O1 through O3, are very similar to axioms of standard deontic logic, and may appear to be vulnerable to the same counterexamples. I will argue, however, that the apparent counterexamples to these principles do nothing to undermine my argument. For all they show is that these principles cannot be accepted

in conjunction with MCR. And so the difficulties these cases present can be resolved by rejecting MCR. And the difficulties that cannot be resolved in this manner can be resolved by restricting the principles in ways that have no effect on my arguments.

Here is a general recipe for generating a deontic paradox:

- (a) Begin with a set of apparently consistent claims, at least one of which is a statement of personal obligation.
- (b) Restate the claims about personal obligation as claims about impersonal obligation, in a manner that presupposes MCR or some similar thesis.
- (c) Show that, so reformulated, these claims are jointly incompatible with the standard principles of deontic logic.
- (d) Conclude that these principles should not all be accepted.

Here is one example. Aqvist begins with the following pair of intuitively consistent sentences:⁹

- (1) It ought to be the case that Smith refrains from robbing Jones
- (2) I ought to know that Smith robs Jones

He then formalizes these claims as follows:

- (3) $O\neg p$
- (4) OKp

Where ‘O’ represents the impersonal obligation operator (‘it ought to be the case that’), ‘K’ represents a knowledge operator (‘I know that’), and ‘ p ’ represents the proposition that Smith robs Jones. Since, necessarily, if Kp then p , O2 licenses the inference from (4) to Op . And O1 licenses the inference from this conclusion, along with (3), to $O(p \ \& \ \neg p)$, which is incompatible with O3. And so this case seems to show that principles O1 through O3 cannot all be accepted.

An alternative response, however, is to reject MCR, and hence to reject the inference from (2) to (4). For one might plausibly claim that while I ought to know that Smith robs Jones, it is not true that it ought to be the case that I know that Smith robs Jones. I ought to believe that Smith robs Jones, one might argue, because, *under my actual circumstances*, that’s the correct belief for me to have. However, one might argue, what ought to be the case does not depend, in the same manner, on my actual circumstances, but is rather what would be the case if everything were as it ought to be. And if everything were as it ought to be, then Smith would not rob Jones, and I would know that Smith does not rob Jones.

As a second illustration, consider cases involving moral dilemmas. These are often said to create problems for the conjunction of O1 and O3. As one typical example, Paske (1990) considers a case in which John encounters brigands who tell him that unless he kills one innocent person, they will kill ten. In this case, Paske says that it is plausible that John is under the following conflicting moral obligations:

- (5) John ought to kill the one innocent person (because he ought to save the ten).
- (6) John ought not to kill the one innocent person (because killing innocent persons is a no-no).

Paske then formalizes these two claims as follows:

- (7) Oa
- (8) $O\neg a$

where ‘O’ represents the impersonal obligation operator, and ‘*a*’ to the proposition that John kills the one innocent person. By O1, these two claims entail $O(a \ \& \ \neg a)$, which is incompatible with O3. Hence, Paske concludes that “an axiom in standard systems of deontic logic must be abandoned” (p. 316).

But, once more, an alternative solution is to reject MCR, and hence to reject the inference from (5) to (7). Even if one grants that, given John’s actual circumstances, he is under an obligation to kill the one innocent person, one might hold that, if things were as they ought to be, the brigands would release all the prisoners, and no one would kill anyone. Hence, one might conclude, it is not true that it ought to be the case that John kills one innocent person. More generally, all the standard cases that are offered as examples of moral conflicts involve conflicts in what an agent ought to do, and insofar as these cases appear to involve a conflict in what ought to be the case, this is because something like MCR is implicitly assumed.¹⁰ Apart from such a supposition, there is little reason to suppose, in any of these cases, that two contradictory propositions each ought to be the case.

As a final illustration, consider Alf Ross’s paradox. While Ross (1941) states his problem in terms of imperatives, it is frequently reformulated in terms of declarative sentences about obligation. Thus, for example, Danielson (2005) says that

- (9) You ought to post the letter.

does not entail

- (10) You ought to post the letter or burn it.

He takes this to be “an objection to the Standard System of deontic logic,” according to which “if B follows from A , then OB follows from OA .” But even if we grant, for the sake of argument, that (9) does not entail (10),¹¹ this will provide an objection to SDL only if (9) and (10) are equivalent to the following two claims, respectively.

(11) It ought to be the case that you post the letter.

(12) It ought to be the case that you post the letter or burn it.

For if this is so, then (11) will fail to entail (12), which is inconsistent with SDL, and, more specifically, with O2. But if we reject MCR, then we can deny this equivalence. Hence, even if we reject the inference from (9) to (10), we can still endorse the inference from (11) to (12). And such a move may be independently motivated, for there are accounts of semantics of personal obligation claims that imply that while the latter inference is valid, the former is not.¹²

There are, admittedly, some deontic paradoxes that do not rely on MCR or any similar assumption. Consider the following case, which appears to be a counterexample to O2:

The Good Samaritan: Smith the Samaritan comes across Jones who, hours ago, was beaten, robbed, and left half dead. The assault on Jones is now an established fact. But while Smith cannot change the past, he can restore Jones to health.

In this case, since it is part of the historical record that Jones has been assaulted in the manner described, it seems that the best thing that could happen at this point is that Smith restores him to health. And so it seems that

(13) It ought to be the case that Smith restores Jones to health.

Moreover,

(14) Necessarily, if Smith restores Jones to Health, then Jones is initially in poor health.

And so O2 licenses an inference to the false conclusion:

(15) It ought to be the case that Jones is initially in poor health.

We can solve this problem, however, if we restrict the scope of the propositional variables in O2 to propositions that are historically contingent, or to propositions such that neither they nor their negations are inevitable. This is not an *ad hoc* restriction, since it is plausibly only to such historically contingent propositions that deontic operators apply.¹³ Alternately, rather than restricting the range over which the variables range, we can simply strengthen the antecedent of O2, as follows:

O2*. For any two propositions, p and q , if it ought to be the case that p , and if (necessarily, if p then q), and if q is not historically necessary, then it ought to be the case that q .¹⁴

If O2 is weakened in either of these ways, then it will no longer license the inference from (13) and (14) to (15). For Jones's being initially in poor health is historically necessary in the example. Moreover, neither of these restrictions has any effect on any of my arguments in this paper. For none of these arguments involve premises or conclusions of the form "it ought to be the case that p ," where q is some historically necessary proposition.

Closely analogous resolutions are possible in the following case, which may appear to be a counterexample to O4:

Professor Procrastinate [Jackson and Pargetter, 1986]: Procrastinate is invited to write a book review. He can write the review only if he received the book, and he will receive the book only if he accepts the invitation. What would be best would be for him to accept the invitation and then write the review. However, regardless of what Procrastinate now intends, if he were to accept the invitation, he would not in fact get around to writing the review. And it would be better for him to decline the invitation than to accept it but fail to write the review.

In this case, one might hold that the following two claims are true:

- (16) Professor Procrastinate ought to write the book review
- (17) A necessary means to Professor Procrastinate's writing the book review is that he accepts the invitation.

And yet one might hold that the following claim is false:

- (18) Procrastinate ought to accept the invitation.

Hence, since O4 licenses the inference from (16) and (17) to (18), one might argue that we should reject O4.

But just as we can plausibly restrict O2 to propositions that are historically contingent, so we can plausibly restrict O4 to action types that are *practically* contingent for the agent. Let us say that ϕ -ing is *practically necessary* for s at t just in case, given s 's actual circumstances, no matter what s intends at t , s will ϕ . And let us say that ϕ -ing is *practically contingent* for s at t just in case neither ϕ -ing nor its negation is practically necessary for s at t . Restricting the range of ϕ and ψ to action types that are practically contingent for s is not *ad hoc*, since, plausibly, deontic predicates such as *obligatory* and *permissible* apply to s 's ϕ -ing only where the latter is practically contingent. Alternately, rather than restricting the range of the variables, we can simply strengthen the antecedent of O4, as follows:

O4*. For any agent, s , and any two action types, ϕ and ψ , if s ought to ϕ , and if s 's ψ -ing is a necessary means to her ϕ -ing, and if ϕ -ing is practically contingent for s , then s ought to ψ .

If O4 is weakened in either of these ways, then it will no longer license the inference from (16) and (17) to (18). For, *ex hypothesi*, regardless of what Procrastinate now intends, he will not write the book review.¹⁵ Hence, writing the book review is not historically contingent for Procrastinate, and so it falls outside of the scope of the restricted principles. Moreover, neither of these restrictions has any effect on any of the arguments of this paper. For the only application of O4 is in deriving the conclusion that Baravelli ought to go to the saloon from the premises that Baravelli ought to accompany Wagstaff and that Baravelli's going to the saloon is a necessary means to his to his accompanying Wagstaff. And in Campus Visit, accompanying Wagstaff and going to the saloon are both practically contingent for Baravelli. And so even if we restrict O4 in either of the manners suggested above, it will still license the inference in question.

References

- Anderson, A. (1962) Logic, Norms and Roles. *Ratio* 4: 32-49.
- Åqvist, L. (1967) Good Samaritans, Contrary-to-Duty Imperatives, and Epistemic Obligations. *Noûs* 1: 361-379.
- Belnap, N., and M. Perloff and M. Xu. (2001) *Facing the Future: Agents and Choices in Our Indeterminist World*. New York: Oxford University Press.
- Broome, J. (unpublished) *Reasoning* Unpublished lectures given at Brown University and the University of Stockholm
- Brown, M. & J. Carmo (1996) *Deontic Logic, Agency and Normative Systems*. New York: Springer Verlag.
- Danielsson, S. (2000) What Shall We Do with Deontic Logic? *Theoria* 66: 97-114.
- Danielsson, S. (2005) Taking Ross's Paradox Seriously: A note on the original problems of deontic logic. *Theoria* 71: 20-28.
- Casteneda, H. (1981) The Paradoxes of Deontic Logic: The simplest solution to all of them in one fell swoop. In Hilpinen, 1981, 37-85.
- Chisholm, R. (1964) The Ethics of Requirement. *American Philosophical Quarterly* 1: 147-153.

- Føllesdal, D. and R. Hilpinen (1971) Deontic Logic: An introduction. In Hilpinen, 1971, 1-35.
- Goble, L. (2001) *Blackwell Guide to Philosophical Logic*. Oxford: Blackwell Publishers.
- Hansson, B. (1969) An Analysis of Some Deontic Logics *Noûs* 3: 373-398.
- Hilpinen, R. (1971) *Deontic Logic: Introductory and Systematic Readings*. Dordrecht: D. Reidel.
- Hilpinen, R. (1981) *New Studies in Deontic Logic*. Dordrecht: Reidel.
- Hilpinen, R. (2001) Deontic Logic. In Goble, 2001, 159-182.
- Horty, J. (1996) Agency and Obligation. *Synthese* 108: 269-307.
- Horty, J. (2001) *Agency and Deontic Logic*. Oxford: Oxford University Press.
- Jones, A. and M. Sergot (1996) A Formal Characterization of Institutionalized Power. *Journal of the IGPL* 4: 429-445.
- Jackson, F. and R. Pargetter (1986) Oughts, Options and Actualism. *The Philosophical Review* 95: 233-255.
- Kanger, S. (1957/1971) New Foundations for Ethical Theory. Reprinted in Hilpinen, 1971, 36-58.
- Krogh, C. and H. Herrestad (1996) Getting Personal: Some Notes on the Relationship between Personal and Impersonal Obligation. In Brown & Carmo, 1996, 134-153.
- Lewis, D. (1973) *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Meinong, A. (1917/1972) *On Emotional Presentation*, tr. M. L. S. Kalsi. Evanston: Northwestern University Press.
- Parfit, D. (1999) Experiences, Subjects, and Conceptual Schemes. *Philosophical Topics* 26: 217-70.
- Paske, G. (1990) Genuine Moral Dilemmas and the Containment of Incoherence. *Journal of Value Inquiry* 24: 315-323.
- Ross, A. (1941) Imperatives and Logic. *Theoria* 7: 53-71.
- Ross, J. (2006) The Scope of Obligation. In his *Acceptance and Practical Reason*, Dissertation, Rutgers University.

Ross, J. (manuscript) Conditional Analyses of Personal Obligation. <http://www-rcf.usc.edu/~jacobmro/ppr/conditional-ross.pdf>

Santos, F. and J. Carmo (1996) Indirect Action, Influence and Responsibility. In Brown & Carmo, 1996, 194-215.

Schroeder, M. (manuscript) Oughts, Agents and Actions.

Williams, B. (1981) *Ought* and Moral Obligation. In his *Moral Luck*, 114-123. Cambridge: Cambridge University Press.

Zimmerman, J. (1996) *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.

¹ In writing this paper, I benefited greatly from discussions with, and comments from Paul Bartha, Stephen Finlay, John Horty, Johannes Schmitt, and two anonymous referees. My greatest debt is to Mark Schroeder for many indispensable discussions.

² I am using “obligation” as a nominalization of the verb “ought.” Thus, I am using it in a sense that is in one respect broader than its sense in ordinary language. For there are things we ought to do but that we have no obligation, in the ordinary language sense, to do, such as brushing our teeth. And, arguably, I am using “obligation” in a sense that is in another respect *narrower* than its sense in ordinary language. For, arguably, there are cases in which we have an obligation, in the ordinary language sense, to do something, and yet it is not the case that we ought to do it, e.g., fulfilling our promise to return a weapon to someone who is bent on mayhem.

³ A version similar to MCR is first presented in [Meinong, 1917/1972]. And a version similar to MCR_{stit} is first presented in [Chisholm, 1964]. For other early works containing versions the Meinong-Chisholm Reduction, see [Kanger, 1957/1971] and [Anderson, 1962]. For more recent works in which this reduction figures, see [Jones and Sergot, 1996], [Santos and Carmo, 1996], and [Belnap, Perloff and Xu, 2001]. For a defense of the Meinong-Chisholm Reduction, see [Williams, 1981]. And for critical discussions, see [Krogh and Herrestad, 1996], [Horty 1996], [Horty, 2001], and [Schroeder, manuscript].

⁴ For Horty’s presentation of MCR_{stit} , see [Horty, 2001] 44-47. And for his discussion of the advantages of MCR_{stit} over MCR, see *ibid.* 50-53.

⁵ Horty, 2001, 53-58. Horty’s problem arises within a particular utilitarian framework in which what ought to obtain is whatever obtains in the best possible outcomes. If one adopts this framework, and one also accepts MCR_{stit} , then one must conclude that an agent ought to take any gamble, no matter how risky, so long as the best possible outcome is one in which she take the gamble and wins. The defender of MCR_{stit} might, however, hold that what ought to obtain is not what is true in the best possible outcomes *simpliciter*, but rather what is true in the outcomes that are *deontically best*, or best with respect to the fulfillment of obligations. Hence she

might say that since no one is under any obligation to win gambles, it is not true that the deontically best outcomes are all ones in which one takes risky gambles and wins.

⁶ I have argued that facts about personal obligation do not supervene on, and hence cannot be reduced to, facts about unconditional impersonal obligation *alone*. It is compatible with this argument that facts about personal obligation may be reducible to some set of facts that *includes* facts about impersonal obligation.

⁷ See [Broome, unpublished] where he argues that the fundamental *ought* must be directly action-guiding.

⁸ We must therefore disagree with Parfit's claim that 'the concepts of *reason* and *person* need not go together. An impersonal conceptual scheme could use the concept of a reason' [Parfit, 1999]. Or at least, we must disagree with this claim so long as the reasons under consideration are practical.

⁹ [Aqvist, 1967] p. 367.

¹⁰ As another illustration of an argument with this structure, see [Danielsson, 2000] p. 107.

¹¹ And there is reason to be skeptical of this claim. See [Føllesdal and Hilpinen, 1971] p. 22 and [Castañeda, 1981] p. 65.

¹² See [Belnap, Perloff and Xu, 2001] pp. 84-85.

¹³ This claim is defended in [Ross, 2006]. For a related response, see [Castañeda, 1981].

¹⁴ See [Danielsson, 2005]. p. 24 for a very similar proposal.

¹⁵ Note that if this were not so, and hence Procrastinate's now forming an appropriate intention would result in his writing the book review, then (27) would be true, and so the problem for O2 would not arise.