

Conflicting Reasons, Unconflicting ‘Ought’s

Shyam Nair

0. Introduction

One of the popular albeit controversial ideas in the last century of moral philosophy is that what we ought to do is explained by our reasons.¹ I will call this idea as REASONS EXPLAIN ‘OUGHT’S.’² One of the central features of reasons that accounts for their popularity among normative theorists is that they can conflict. For example, the fact that reasons conflict promises to give a neat explanation of (i) why in standard choice situations, an agent faces the task of choosing among a variety of incompatible options each of which has something to be said for it and (ii) how the competition of these conflicting considerations determines what such an agent ought to do. But in this paper, I argue that the fact that reasons conflict actually poses a problem for those who accept REASONS EXPLAIN ‘OUGHT’S.’

In fact, I argue that there are two closely related problems (or, if you prefer, two different ways of arriving at the same problem) that illustrate how difficult it is to make sense of cases involving conflicting reasons if we accept REASONS EXPLAIN ‘OUGHT’S.’ The first problem is a generalization of a problem in deontic logic concerning the existence of conflicting ‘ought’s’ (§1). The second problem arises from a tension between three ethical principles (§2-3). Having presented each of these problems, I develop a unified solution to them that is informed by results in both ethics and deontic logic. An important implication of this solution is that we must distinguish between derivative and non-derivative reasons and reconsider what the commitments of REASONS EXPLAIN ‘OUGHT’S’ really are (§4-5).

The approach that I will take to these questions will be mildly formal. I do not assume that this methodology is to be privileged over other approaches. Rather, what I will do is present a number of simple cases that need to be explained and a number of plausible ethical principles that need to be accommodated. And I will show that in order to do this, we must spell out in more detail the connection between reasons and ‘ought’s’ that REASONS EXPLAIN ‘OUGHT’S’ claims. I will then consider certain simple and tempting principles describing formal connections among reasons and ‘ought’s’ that might be used to explain these cases and accommodate these principles. I show that these initially tempting principles are inadequate and then go on to describe my preferred set of principles for solving these problems. The ability to handle these cases and accommodate these principles will be one piece of evidence in favor of my preferred set of principles and approach to the problem. That said, those who prefer an approach that eschews formal connections and formal modelling are free to develop alternative solutions to these problems. What I wish to do is only put

¹ Though he uses different terminology, I take this idea to originate in Ross 1930. Other classic discussions include Dancy 2004a, Nagel 1970, Parfit 2011, Raz 2002, and Scanlon 1998.

² For simplicity, I adopt the assumption that ‘ought’ and ‘should’ are roughly synonymous.

the problems on the table and develop one systematic solution to these problems that spells out the details of the connection between reasons and ‘ought’s. Comparing this solution to alternatives is a task that must be postponed until the time when alternative solutions are developed.

1. The First Problem

To begin to see the first problem, consider the following pair of cases:

Speeding Law. The laws of the country set the speed limit at fifty miles per hour. We may suppose then that drivers have a reason to drive slower than fifty miles per hour. Given that drivers have a reason to drive slower than fifty miles per hour, drivers also have a reason to drive slower than one hundred miles per hour.³

Fighting or Serving. Smith’s country requires him to fight in the army or perform alternative public service. We may suppose then that Smith has a reason to fight in the army or perform alternative public service. Smith is also deeply committed to a pacifist religion that requires him not fight in the army. We may suppose then that Smith has a reason to not fight in the army. Given that Smith has a reason to fight or serve and given that Smith has a reason to not to fight, Smith also has a reason to serve.⁴

In *Speeding Law*, there is a reason to drive slower than one hundred miles per hour given that there is a reason to drive slower than fifty miles per hour. In *Fighting or Serving*, there is a reason to perform alternative public service given that there is a reason to fight or serve and that there is a reason to not fight. The first problem arises when we try to develop a systematic account why there are these reasons in cases that have the structure of *Speeding Law* and *Fighting or Serving*.

Before we look at what some systematic accounts might look like, it is worth discussing which of the many senses of ‘reason’ and ‘ought’ we are focusing on in this paper. While I believe that the main ideas of this paper will apply to any sense of ‘reason’ and ‘ought’ for which REASONS EXPLAIN ‘OUGHT’S is plausible, I will for concreteness fix on one. I will be discussing what can be called *objective* normative reasons. To borrow the now standard terminology of Scanlon 1998, an objective normative reason for an agent to do some act is a fact or true proposition that “counts in favor” of the agent doing that act. What we have objective normative reason to do depends, in the first instance, on what the facts are and not what our beliefs are. The ‘ought’s that I will be discussing can be called *objective all-things-considered* ‘ought’s. These objective ‘ought’s similarly depend on the facts and not our beliefs about the facts. What makes these ‘ought’s *all-things-considered* is that they hold in light of all the relevant normative considerations rather than only considerations from one domain such as, e.g., morality or prudence.

³ Cf. Cariani 2013: n. 1.

⁴ Cf. van Fraassen 1973: 18, Horty 2003: 578, Goble 2009: 459. It may also be worth noting that this case assumes a kind of pluralism about the sources of our reasons (in this case, one reason is due to the laws and the other is due to personal commitments). While the details of our reasons’ sources are controversial, pluralism in general has been a commitment in the spirit of the popular idea in ethics since Ross. In any case, this pluralism is inessential to the structure of this case and is included only out of deference to the history of the case.

With this in mind, let us return to considering how to give a systematic account of the reasons involved our opening pair of cases. While there are many possible answers to this question, formal principles describe entailments among reasons perhaps provide the simplest general answer to this question: these principles provide natural explanations of why there are certain reasons *given* the presence of other reasons. In particular, the following two principles seem like the simplest and most obvious ones that do the job:

Single Reason Closure: if there is a reason for S to do *a* and doing *a* involves doing *b*, then there is a reason for S to do *b*⁵

Consistent Reasons Agglomeration: if there is a reason for S to do *a*, there is a reason for S to do *b*, and S can do *a* and *b*, then there is a reason for S to do *a* and *b*⁶

Single Reason Closure allows us to explain *Speeding Law*. Why is there a reason to drive less than one hundred miles per hour in that case? Given that there is a reason to drive less than fifty miles per hour and given that driving less than fifty miles per hour involves driving less than one hundred miles per hour, *Single Reason Closure* tells us that there must be a reason to drive less than one hundred miles per hour.

Single Reason Closure alone cannot however show that there is a reason to perform alternative public service in *Fighting or Serving*. But *Single Reason Closure* would give us this reason if somehow we could get a reason to [(fight or serve) and not fight]. Here *Consistent Reasons Agglomeration* helps. Given that there is a reason to fight or serve, that there is a reason to not fight, and that Smith can [(fight or serve) and not fight], *Consistent Reasons Agglomeration* tells us there must be a reason to [(fight or serve) and not fight]. This then allows *Single Reason Closure* to apply and complete our explanation of why there is a reason to serve in *Fighting or Serving*: since [(fighting or serving) and not fighting] involves serving, there is a reason to serve. So together *Single Reason Closure* and *Consistent Reasons Agglomeration* provide a tidy treatment of our opening cases.

But despite its tidiness, this treatment of these cases cannot be correct. This is because *Single Reason Closure* and *Consistent Reasons Agglomeration* cannot both be true. The crucial premise in the argument that demonstrates this is the platitude that we began the paper with, reasons can conflict. Though the argument

⁵ Throughout the main text, I will rely on our pretheoretical grasp of the notion of doing one act involving doing another. Appendix two provides a simple formal model of this notion of involvement in terms of logical entailment in propositional logic (other more complicated formal models are also possible though I do not have the space in this paper to develop them).

⁶ Three clarifications may be in order. First, for it to be the case that S can do *a* is for doing *a* to be under S's intentional control in the sense that both (a) if S intends to do *a*, S does *a* and (b) if S intends to not do *a*, then S does not do *a*.

Second, claims about reasons are naturally understood to be doubly-tensed in much the way other constructions such as 'want'-constructions are. For example, 'John wants to eat dinner' is tensed both on the time of the wanting and the time of the thing wanted (this can be made vivid by comparing the meaning of 'Tomorrow it will be the case that John wants to eat dinner' and 'John wants to eat dinner tomorrow'. In most natural contexts, the time of the wanting is tomorrow in the first but not the second. And the time of the thing wanted is tomorrow in the second). I intend *Consistent Reasons Agglomeration* to concern cases where the time of the reason is the same but the time of the thing that there is a reason to do may be different (though of course whether an agent can do two acts may depend on the time of the things that there is a reason to do).

Third, throughout this paper, I will concentrate on what there is a reason to do. An interesting further question is what the reason is to do it. A generalization of the formal model that I describe in appendix two answers this question, but I do not have the space here to present this theory (though the theory entails that the reason do the conjunction is the a conjunction of the reasons to do each conjunct). In any case, it is not strictly speaking required to present or solve the problems that I describe in this paper. Thanks to an anonymous referee for pressing me to clarify these issues.

doesn't require this, it is easiest to illustrate how it works by assuming there is a particular instance of conflicting reasons, one to do an act and one to not do that act. For example:

- (1) There is a reason to meet Sam for a drink
- (2) There is a reason to not meet Sam for a drink

By *Single Reason Closure* applied to (1), we may infer:

- (3) There is a reason to [meet Sam for a drink or torture a kitten]

Since you can [[meet Sam for a drink or torture a kitten] and not meet Sam for a drink], we may apply *Consistent Reasons Agglomeration* to (2) and (3) and infer:

- (4) There is a reason to [[meet Sam for a drink or torture a kitten] and not meet Sam for a drink]

Finally by *Single Reason Closure* again, we may infer:

- (5) There is a reason to torture a kitten

Thus, the existence of conflicting reasons concerning meeting Sam for a drink together with *Single Reason Closure* and *Consistent Reasons Agglomeration* entail the intuitively unacceptable existence of reason to torture a kitten. In fact, this is an instance of a more general result that does not require the assumption that we have reasons to do an act and its negation—it is enough that we merely have any conflicting reasons—and does not rely on the assumption that we have no reason to torture a kitten—it is enough that there is at least one act that we can do that we have no reason to do. Though I leave the proof to a note, the result is this: *Single Reason Closure* and *Consistent Reasons Agglomeration* entail that if there is a reason for S to do *a*, there is a reason for S to do *b*, and S can't do *a* and *b*, then there is a reason for S to do *c* for any *c* that S can do.⁷ In other words, this result says that if reasons can conflict, then we have an *explosion* of reasons: for any act an agent can do, that agent has a reason to do that act.

This shows that the initially tempting account of the cases is incorrect and some other account is needed. And there are many possible accounts of these cases. But before embarking on the task of

⁷ Suppose there is a reason for S to do *a* and there is a reason for S to do *b* and S cannot do *a* and *b*. Now doing *b* involves not doing *a* because S cannot do *a* and *b*, so by *Single Reason Closure* there is a reason for S to do not-*a*. And by another two application of *Single Reason Closure*, we get a reason to do *a* or *c* and a reason to do not-*a* or *c* for any *c* such that S can do *c*. Next notice that for any such *c* it must be that either S can do *a* and *c* or S can do not-*a* and *c*.

Suppose S can do *a* and *c*. This means S can do [[not-*a* or *c*] and *a*]. So by *Consistent Reasons Agglomeration*, we can generate a reason to do [[not-*a* or *c*] and *a*]. Finally since doing [[not-*a* or *c*] and *a*] involves doing *c*, we get the results by *Single Reason Closure* there is a reason to do *c*.

If we suppose instead that S can do not-*a* and *c* and start this time with the reason to do *a* or *c*, we get by analogous reasoning the result that there is a reason to do *c*. So either way there is a reason to do *c* for any *c* that S can do. This completes the proof.

considering all of these accounts, it is worth knowing that the problem that I just presented involving conflicting reasons has the same structure as a problem in deontic logic concerning conflicting ‘ought’s.⁸ The problem arises when trying to explain the above cases except replacing each occurrence of ‘reason’ with an occurrence ‘ought’ while also maintaining that ‘ought’s can conflict.

This structural similarity teaches us two important lessons. First, the problem in deontic logic is well-studied. And deontic logicians have shown that the problem that we have just described is robust under a choice of a large variety of alternatives to the particular principles that we just looked at. That is, deontic logicians have looked at many principles other than the ‘ought’ analogs of *Single Reason Closure* and *Consistent Reasons Agglomeration* and shown that striking a balance between explaining our cases and avoiding something like the explosion problem is quite difficult. Indeed, in a recent survey article that is over one hundred pages long Lou Goble (2013) meticulously considers over twenty different packages of principles and illustrates the trade-offs they have between explaining our cases and avoiding explosion.

For this reason, I cannot catalog every twist and turn one could take to try to explain these cases (but n. 22 provides a rough guide to the different packages of principles that bear on our problem). I will simply leave it as a challenge to resolve this tension between explaining these cases and the platitude that reasons conflict. And a promising idea that I will end up making good on is that we can resolve it by mining the resources of the rich literature in deontic logic about the problem of conflicting ‘ought’s. Ultimately, this will mean that I will resolve this tension by continuing to accept *Single Reason Closure*, by rejecting *Consistent Reasons Agglomeration*, and by developing an alternative to it that still allows us to explain the cases that we need to explain. This may come as a surprise to some readers who find *Single Reason Closure* independently implausible and so see it as the obvious culprit in our problem. Appendix one addresses these readers’ worries.

The second thing to be learned by comparing our problem to the problem in deontic logic is that our problem is strictly more difficult than the problem in deontic logic. This is because one solution to the problem in deontic logic is to simply accept the ‘ought’ analogs of *Single Reason Closure* and *Consistent Reasons Agglomeration* and reject the assumption that ‘ought’s can conflict. If ‘ought’s never conflict, there is no harm in the trivially true claim that if ‘ought’s conflict, we get an explosion of ‘ought’s. Indeed, I will take it as a working assumption in this paper that ‘ought’s never conflict. But what is important to see is that the analogous solution to the problem concerning reasons is simply a non-starter for our problem. After all, it is a platitude that reasons can conflict.⁹

⁸ For general discussion of the problem of conflicting obligations, see Brink 1994, Chellas 1980, Foot 1983, Gowans 1987b, Lemon 1962, Marcus 1980, McConnell 2010, Pietroski 1993, Sinnott-Armstrong 1988, van Fraassen 1973, and Williams 1965. And see the helpful collection Gowans 1987a. For more recent discussion that bears particularly on the issues discussed in this section see Hansen 2004, McNamara 2004, van der Torre and Tan 2000, and especially Goble 2009, 2013 and Horty 2003.

⁹ Given my working assumption that ‘ought’s never conflict, we can, in fact, raise the first problem in a slightly different way. Almost everyone who accepts the popular idea in ethics accepts the following intuitively plausible connection between reasons and what we ought to do:

‘Ought’s Entail Reasons: if S ought to do *a*, then there is a reason for S to do *a*

This then is the first problem: there is a tension between giving a systematic account of certain simple cases and the platitude that reasons conflict. As I said, my solution will be to accept *Single Reason Closure* and develop an alternative to *Consistent Reasons Agglomeration* that explains our cases and is informed by work in deontic logic.

2. Conflicting Reasons and How Reasons Explain What We Ought to Do

But before I present my solution, we should consider a second problem. This problem is similar to one that John Horty and Lou Goble have discussed in earlier work.¹⁰ But the contribution of the problem that I will present is that it highlights a tension between three plausible ethical principles. We have already encountered the first two principles in passing:

Reasons Allow Conflicts: there can be situations in which S has a reason to do *a* and a reason to do *b* but cannot do *a* and *b*

'Ought's Cannot Conflicts: there cannot be situations in which S ought to do *a* and ought to do *a* but cannot do *a* and *b*

As I said earlier, *Reasons Allow Conflict* is a platitude and we will be granting *'Ought's Cannot Conflicts* for the sake of argument.

The third principle that is involved in this tension is a claim about *how* reasons explain what we ought to do:

If we accept *'Ought's Entail Reasons*, this means that there is a reason corresponding to each thing we ought to do. Now since we are assuming that 'ought's never conflict and since given this assumption, the 'ought' analogs of *Single Reason Closure* and *Consistent Reasons Agglomeration*:

Single 'Ought' Closure: if S ought to do *a* and doing *a* involves doing *b*, then S ought to do *b*

Consistent 'Ought' Agglomeration: if S ought to do *a*, S ought to do *b*, and S can do *a* and *b*, then S ought to do *a* and *b*

provide a simple explanation of the 'ought' analogs of *Speeding Law* and *Fighting or Serving*, we may accept these two claims describing entailments among 'ought's.

If we accept these entailments and accept *'Ought's Entail Reasons*, we would need to ensure that there are reasons corresponding to each thing we ought to do that is generated by *Single 'Ought' Closure* and *'Consistent 'Ought' Agglomeration*. A natural question to ask the advocate of the popular idea in ethics who accepts *'Ought's Entail Reasons* is what guarantees that there are reasons corresponding to each thing we ought to do that is generated by these entailments: what guarantees that reasons and what we ought to do walk in lock-step in the way that *'Ought's Entail Reasons* requires? An elegant answer would be that the analogous entailments hold among reasons. That is, *Single Reason Closure* and *Consistent Reasons Agglomeration* hold. This then shows that there is a tension between providing a systematic the 'ought' analogs of *Speeding Law* and *Fighting or Serving* while preserving natural connections between 'ought's and reasons and the platitude that reasons conflict.

¹⁰ Horty 2003: 572-573 was the first to notice the issue discussed in §3.1. But, as he presents it, the difficulty is narrowly tailored to the view in Brink 1994. Horty does not isolate the problem as one that arises from accepting the three principles discussed below. Horty also does not consider, as we will below, whether the view can be saved by supplementing it with principles concerning entailments among reasons.

Goble 2013: 22 is the first to recognize the issue in §3.2 The contribution of my development of this difficulty is that it shows that the problem arises as a tension between the three ethical principle discussed below. Goble's presentation focuses on validating so-called Standard Deontic Logic, delivering the disjunctive response, and explaining certain cases. My work is heavily indebted to Horty and Goble's discussions.

Reasons Explain 'Ought's Directly: if S ought to do *a*, then this is in part (essentially) explained by the fact that there is a reason for S to do *a*¹¹

Philosophers who accept REASONS EXPLAIN 'OUGHT'S are almost unanimous in their acceptance of *Reasons Explain 'Ought's Directly*.^{12, 13} While philosophers often discuss why they believe 'ought's are explained in terms of reasons, they rarely discuss why that explanation must validate *Reasons Explain 'Ought's Directly*. So unfortunately there is, to my knowledge, no argument in print in favor of *Reasons Explain 'Ought's Directly*. Perhaps this is because it seems obvious that accepting REASONS EXPLAIN 'OUGHT'S requires accepting *Reasons Explain 'Ought's Directly*. Or perhaps it is because it is hard to see what alternative there could be to accepting *Reasons Explain 'Ought's Directly*.

But whatever the case maybe, it will be worthwhile to take a moment here to show that there is a certain kind of argument or motivation for *Reasons Explain 'Ought's Directly* that can be given. This will be worthwhile because the solution that I will offer involves rejecting *Reasons Explain 'Ought's Directly*. So it is important to have clearly in view ahead of time what can be said in favor of this principle that I will be committed to rejecting.

The argument or motivation for *Reasons Explain 'Ought's Directly* can be introduced by comparing it to the following plausible theoretical principle which almost everyone who accepts REASONS EXPLAIN 'OUGHT'S accepts:

'Ought's Entail Reasons: if S ought to do *a*, then there is a reason for S to do *a*

The difference between these two claims is this: *'Ought's Entail Reasons* says that if you ought to do something, there is a reason to do that very thing.¹⁴ But it does not make any claims about explanation. *Reasons Explain 'Ought's Directly*, on the other hand, makes the stronger claim that the reason is part of what explains why you ought to do it.

Because *Reasons Explain 'Ought's Directly* is logically stronger than *'Ought's Entail Reasons*, an advocate of REASONS EXPLAIN 'OUGHT'S can explain why *'Ought's Entail Reasons* holds by citing *Reasons Explain 'Ought's*

¹¹ The explanation may be *partial* because facts about other reasons and the strength of reasons may need to be mentioned. I The explanation is supposed to be *essential* in the sense that there is no complete explanation of why S ought to do *a* in terms of reasons that fails to mention the reason for S to do *a*.

¹² According to my usage, REASONS EXPLAIN 'OUGHT'S and therefore *Reasons Explain 'Ought's Directly* assume that facts about reasons are prior to and explanatory of facts about what we ought to do So according to this usage, Broome 2004 does not accept the popular idea in ethics. This is because Broome explains the notion of a reason in terms of the prior notions of what we ought to do and a weighing explanation.

¹³ For our purposes, certain variants of *Reasons Explain 'Ought's Directly* and *'Ought's Entail Reasons* will do just as well. While it is difficult to provide a general statement of what the weakest commitment is that would still be sufficient for my argument that is not hopelessly abstract, it is important to know that those who favor revisions to simplistic formulation of REASONS EXPLAIN 'OUGHT'S such as Bedke 2011, Dancy 2004b, Gert 2007, and Greenspan 2007 will be committed to variants of *Reasons Explain 'Ought's Directly* that will suffice for our problem. John Horty, Lou Goble, and Douglas Portmore are the only theorists that I know of who do not accept some such variant of *Reasons Explain 'Ought's Directly*. Horty also rejects *'Ought's Entail Reasons*. Appendix two discusses their work.

¹⁴ Of course, many philosophers who reject the popular idea in ethics also reject *'Ought's Entail Reasons*. For example, Foot 1972 famously rejects an analog of *'Ought's Entail Reasons* that concerns what we morally ought to do (rather than what we ought to do all things considered; Foot's stance on our version is not transparent). Similar comments apply to *Reasons Explain 'Ought's Directly*.

Directly. This is some evidence in favor of *Reasons Explain 'Ought's Directly*. What's more, *Reasons Explain 'Ought's Directly* can look like the only plausible explanation of *'Ought's Entail Reasons* that an advocate of REASONS EXPLAIN 'OUGHT'S can adopt as the following suggestive line of reasoning shows.¹⁵

Consider how else an advocate of REASONS EXPLAIN 'OUGHT'S explain the truth of *'Ought's Entail Reasons*. One explanation of *'Ought's Entail Reasons* is that anytime an agent ought to do *a*, the fact that she ought to do *a* explains why there is a reason for her to do *a*. If the fact that she ought to do *a* explains the reason to do *a*, then *'Ought's Entail Reasons* would follow as a corollary. Unfortunately, this explanation looks to give up on REASONS EXPLAIN 'OUGHT'S because it is an instance of an 'ought' explaining a reason. So someone who accepts REASONS EXPLAIN 'OUGHT'S would not be able to explain *'Ought's Entail Reasons* in this way.

Another possible explanation might go like this: What explains why an agent ought to do something is that it leads to the best outcome. What explains why an agent has a reason to do something is that it leads to an outcome that is better than another outcome in some respect. So if you ought to do *a*, then this must be because doing *a* has the best outcome. And if doing *a* has the best outcome, then doing *a* has an outcome that is better than another outcome in some respect. Thus, if you ought to do *a*, there must be a reason to do *a*. So *'Ought's Entail Reasons* follows as a corollary of this view about the connection between what we ought to do, reasons, and goodness.

This explanation of the truth of *'Ought's Entail Reasons* is problematic in two ways. First it is incompatible with REASONS EXPLAIN 'OUGHT'S because it denies that reasons explain 'ought's and instead says both reasons and what we ought to are explained by a third thing, values.

Second even if we could modify this explanation of *'Ought's Entail Reasons* so that it was consistent with REASONS EXPLAIN 'OUGHT'S and only claimed that reasons (and not 'ought's) are explain in terms something else, e.g., values, it would nonetheless be a less than ideal explanation of *'Ought's Entail Reasons* because it requires not just the resources of REASONS EXPLAIN 'OUGHT'S but also a further theory of the nature of reasons. This would be less than ideal is not because such theories are false. Rather it is less than ideal because those who accept REASONS EXPLAIN 'OUGHT'S have diverse commitments about which if any theory about the nature of reasons is correct. So I will set aside this approach in hopes of finding a more ecumenical solution to our problem.

So far then we have considered two alternative explanations of *'Ought's Entail Reasons* and found them unsatisfying. And it can seem like all the other explanations must be similarly unsatisfying as well. It can, at least at first glance, seem like the only alternative to explaining *'Ought's Entail Reasons* in terms of *Reasons Explain 'Ought's Directly* is to either claim that what we ought to do explains what we have reason to do or rely

¹⁵ I say this line of reasoning is *suggestive* because it makes a number of unstated assumptions about the nature of explanation and the space of options. Indeed, I end up rejecting one such assumption in giving my solution and so I am presenting this line of thought as foil for the solution that I will eventually offer.

on some further theory about the nature of what we ought to do and have reason to do. This leave us in the following situation: We can explain *'Ought's Entail Reasons* with the help of *Reasons Explain 'Ought's Directly*. The only alternative explanations seem to involve giving up on REASONS EXPLAIN 'OUGHT'S or relying on a further theory about the nature of reasons. Thus, *Reasons Explain 'Ought's Directly* looks to be the most plausible and ecumenical way for the advocate of REASONS EXPLAIN 'OUGHT'S to explain *'Ought's Entail Reasons*.

This, to my mind, is a powerful line of thought in *Reasons Explain 'Ought's Directly's* favor. And it is one that was worth developing because my solution will require us to give up on *Reasons Explain 'Ought's Directly* and therefore face up to this argument in its favor.

We have, then, the three elements of our second problem *Reasons Can Conflict*, *'Ought's Cannot Conflict*, and *Reasons Explain 'Ought's Directly*. In the next section, I will show how there is a tension between these three theses.

3. The Tension

We can illustrate the tension between these principles by considering two cases.

3.1 Two Breakfasts

The first case to consider is the following one:

Two Breakfasts

Mary promised Jeff that she would meet him downtown for breakfast. Mary also promised Scott that she would meet him by the beach for breakfast. These promises are equally important. Even though Mary can make it to either of the breakfasts, she can't make it to both because downtown and the beach are too far apart.

In *Two Breakfasts*, Mary's promise to Jeff gives her a reason to meet Jeff downtown for breakfast and her promise to Scott gives her a reason to meet Scott by the beach for breakfast. Since she cannot do both, this means that her reasons conflict. Given *Reasons Can Conflict*, we should expect there to be cases involving conflicting reasons like Mary's. And while *Reasons Can Conflict* doesn't entail that there are cases involving equally good conflicting reasons, *Two Breakfasts* is plausibly such a case: since the promises that Mary made are equally important, plausibly the reasons that stem from them are equally important as well. Finally, we may also stipulate that Mary does not have any other reasons that conflict with keeping either promise.

Given that Mary's reasons are equally good, what should she do? There seem to be only two plausible answers to this question. One answer is that Mary ought to meet Jeff for breakfast and Mary ought to meet Scott for breakfast. Unfortunately, this answer is ruled out by *'Ought's Cannot Conflict*. Since Mary cannot go to both breakfasts, *'Ought's Cannot Conflict* entails that it cannot be that she ought to go to each breakfast.

The only other plausible answer to our question about Mary's case is that she ought to meet Scott for

breakfast *or* meet Jeff for breakfast.¹⁶ What Mary ought to do is this *disjunctive* act. To support this judgment, we can notice that we do not want to say that there is nothing Mary ought to do. If, for example, Mary chose to eat breakfast at home she would not be doing what she ought to be doing. But if Mary ought to do something and if it isn't true that Mary ought to meet Jeff and ought to meet Scott, it looks like the only other candidate for what Mary ought to do is meet Jeff or Scott.¹⁷

Thus, we have seen that *Reasons Can Conflict* and *'Ought's Cannot Conflict* commit us to thinking that in cases involving equally good conflicting reasons like *Two Breakfasts*, what agents like Mary ought to do is a disjunctive act like either meeting Jeff or meeting Scott. Now that we know this, we can ask how REASONS EXPLAIN 'OUGHT'S can explain why we ought to do a disjunctive act like meeting Jeff or Scott. And *Reasons Explain 'Ought's Directly* tells us that in order to explain why we ought to do *a* we need to at least show that there is a reason to do that very act, *a*. So if we are to explain why Mary ought to meet Jeff or Scott, we need to show that there is a reason for Mary to meet Jeff or Scott.

And maybe there is a reason to meet Jeff or Scott in *Two Breakfasts*. Maybe the promise to Jeff is a reason to do that. Or maybe her promises taken together are the reason. It is not important for our purposes to decide which of these claims is correct.

What is important is to see is the general point that *Two Breakfasts* illustrates. It illustrates that in cases involving conflicting reasons, we need to provide some general principle that entails that there is a reason to do disjunctive acts like meeting Jeff or Scott. *Two Breakfasts* illustrates this general point because it was the structural feature of this case—that it involved conflicting equally good reasons—combined with our three theses—*Reasons Can Conflict*, *'Ought's Cannot Conflict*, and *Reasons Explain 'Ought's Directly*—that led us to the result that we need to have a reason to do this disjunctive act. What this means is that unless we accept a general principle that entails that there are reasons to do disjunctive acts in cases involving conflicting reasons like *Two Breakfasts*, our three theses will be incompatible.

So what principle might we use to get the right result in cases like Mary's? We have already encountered one that looks promising, *Single Reason Closure* (which, recall, says that if there is a reason for *S* to do *a* and doing *a* involves doing *b*, then there is a reason to do *b*). As we saw, *Single Reason Closure* allows us to

¹⁶ According to Horty 2003: 570-571, this so-called "disjunctive response" to cases like Mary's was first explicitly stated in Donagan 1984. It has since been endorsed by Brink 1994, Horty 2003 and Goble 2013.

¹⁷ I have claimed that we do not want to say that Mary ought to do nothing. But one might object that it would be unfair to keep one promise rather than the other so Mary ought to keep neither and do nothing. I have three responses to this objection. First, though here is not the place to argue about first order issues in the theory of promises and fairness, I will simply report that I do not think it is unfair to keep one of the promises in this case. Second, insofar as we accept that it is unfair, the most natural way to capture this is to claim that there is a strong reason not to keep exactly one promise in this case. If we were to do that, then the case at hand would not be one of the structure that I intended it to be (one where there are only two relevant reasons and they are equally good but incompatible), but the objection would also would not undermine the structural claim that I am making. Third, consider a different case where the competition is not between two promises but between a promise and, e.g., harm to others, the agent's own rational aims, or pleasure. Though people's judgments about exactly which cases involving these kinds of reasons are ones where the reasons that are not worse than one another, people generally will agree that there are such cases. In such cases, I claim that the agent ought to perform the disjunctive act. And in such cases issues of fairness have no obvious role to play. Thanks to Nick Laskowski and an anonymous referee for pressing me on this point.

explain *Speeding Law*. And it also can similarly help explain *Two Breakfasts*. Since there is a reason to meet Jeff in this case and meeting Jeff involves meeting Jeff or Scott, *Single Reason Closure* tells us that there is the desired reason to meet Jeff or Scott.¹⁸

3.2 Lunch-Coffee-Dinner

Unfortunately, *Single Reason Closure* alone won't suffice to explain the full range of cases involving conflicting reasons. Consider for example the case of Sally:

Lunch-Coffee-Dinner

Sally has made three promises. She promised Tom that she will meet him for lunch downtown. She promised Jack that she will meet him for coffee by the beach. She promised Ann that she will meet her for dinner in Santa Barbara. While Sally can make it to lunch and coffee, can make it to lunch and dinner, and can make it to coffee and dinner, she cannot make it to lunch, coffee, and dinner. There just isn't enough time for all that driving.

In this case, Sally has a reason to meet Tom for lunch downtown, a reason to meet Jack for coffee by the beach, and a reason to meet Ann for dinner in Santa Barbara. While Sally can do any two of these things, she cannot do all three. As before, *Reasons Can Conflict* should lead us to expect that there are cases like *Lunch-Coffee-Dinner*. And it is independently plausible that *Lunch-Coffee-Dinner* is one involving equally good conflicting reasons.

Given that Sally's reasons are equally good, what should she do? One plausible answer is that Sally ought to meet Tom, ought to meet Jack, and ought to meet Ann. But presumably the spirit of '*Ought's Cannot Conflict*' tells us that just as pairs of 'ought's cannot conflict, sets of 'ought's cannot conflict either.

The only other plausible thing to say about this case is that Sally ought to perform the disjunctive act of meeting Tom and Jack *or* meeting Tom and Ann *or* meeting Jack and Ann. That is, Sally ought to perform this act that is a disjunction of conjunctions.¹⁹ And, recall, according to *Reasons Explain 'Ought's Directly*, to explain why we ought to do *a* we need a reason to do *a* itself. So to explain why Sally ought to do this disjunction of conjunctions we need to show that there is a reason to do this disjunction of conjunctions.

This means that we need our general principles to entail that such a reason exists in cases like *Lunch-Coffee-Dinner*. Unfortunately, *Single Reason Closure* alone does not do this. If we apply *Single Reason Closure* to any of the individual reasons, the best we can do is generate a reason to meet Tom or meet Jack or meet Ann. But this is not the disjunction that we want. We want a disjunction of *conjunctions*. This is important because we want to be able to say that Sally ought to do at least two of the three acts and the simple disjunctions that *Single Reason Closure* gets us do not entail that there is a reason to do at least two of the three acts. Thus, *Single*

¹⁸ This is of course not the only principle that could be used to tackle this case. One could instead adopt the principle that if there is a reason to do *a* and a reason to do *b*, then there is a reason to do *a* or *b*. As I said, I will not have the space here to discuss all the alternative principles that one might try because there are too many to discuss (but see n. 22 for a guide to what the issues with different families of principles are). My aim is to only present one solution to the problems developed in this paper. That said, it is worth noting that adopting this principle and rejecting *Single Reason Closure* would leave *Speeding Law* unexplained.

¹⁹ Goble 2013: 22 concurs with this judgment about the case.

Reason Closure alone is not sufficient to solve our problem.

Now if we were to have a reason to do two of the acts, *Single Reason Closure* would be enough to generate the reason to do the disjunction that we want. So the natural thing to do is adopt a principle that will allow us to generate a reason to do a pair of acts from the reasons to do each act. And we have already encountered such a principle, *Consistent Reasons Agglomeration* (which, recall, says that if there is a reason for S to do *a*, there is a reason for S to do *b*, and S can do *a* and *b*, then there is a reason for S to do *a* and *b*.) With *Consistent Reasons Agglomeration* supplementing *Single Reason Closure*, we are in a position to explain *Lunch-Coffee-Dinner*. Since there is a reason to meet Tom and a reason to meet Jack and since Sally can meet Tom and Jack, *Consistent Reasons Agglomeration* entails that there is a reason to meet Tom and Jack. We may then apply *Single Reason Closure* to the reason to meet Tom and Jack to generate the desired reason to meet Tom and Jack *or* meet Tom and Ann *or* meet Jack and Ann. Thus, *Consistent Reasons Agglomeration* and *Single Reason Closure* look to get us what we want.

But of course, as we know from the first problem, we cannot accept both *Single Reason Closure* and *Consistent Reasons Agglomeration*. They entail an explosion of reasons. So this is the second problem: In order to resolve the tension between *Reasons Can Conflict*, *'Ought's Cannot Conflict*, and *Reasons Explain 'Ought's Directly*, we are led into the project of developing principles concerning the entailments among reasons. And once again, *Single Reason Closure* and *Consistent Reasons Agglomeration* initially look like promising candidates. But since they lead to an explosion of reasons, they are unacceptable. In order to solve this problem we need to develop alternatives to *Single Reason Closure* and *Consistent Reasons Agglomeration* that suffice to explain *Two Breakfasts* and *Lunch-Coffee-Dinner* but do not lead to explosion.

In light of these two problems concerning conflicting reasons, I conclude that there is work to be done to make sense of cases involving conflicting reasons if we accept REASONS EXPLAIN 'OUGHT'S. Having illustrated this, I turn to developing my own solution.

4. Toward a Solution

As I have already said, there are many different principles that we could use to try to solve our two problems. But I will not pursue the project of canvassing the full range of these principles. Instead, since the problem of conflicting 'ought's is relatively well-known (in deontic logic at least), I begin by taking up the conjecture for solving the first problem that I mentioned earlier. The conjecture was that the solution to the first problem will have the same structure as solutions to the problem of conflicting 'ought's that allow for the existence of such 'ought's. I take up this conjecture in two stages. First I describe a certain structure shared by many solutions to the problem of conflicting 'ought's (§4.1). Second I turn to some familiar ideas in moral philosophy in order to adapt this structural feature to cases involving reasons (§4.2).

While this only directly tells us how to solve the first problem, I observe that it also indirectly suggests that the solution to the second problem is to deny *Reasons Explain 'Ought's Directly* (§4.3). §5 develops

these suggestions in greater detail and shows that they solve our problem.

4.1 The Structure of Agglomeration

The conjecture that we are pursuing is this: since our problem has the same structure as the problem of conflicting obligations in deontic logic, the literature in deontic logic might hold a solution.

Though I cannot discuss all of the details of this literature here, I can draw out, I believe for the first time, a certain structural similarity shared by many accounts developed in it. These accounts avoid explosion by denying the ‘ought’ analog of *Consistent Reasons Agglomeration* (*Consistent ‘Ought’s Agglomeration*: if S ought to do *a*, S ought to do *b*, and it is possible to do *a* and *b* then S ought to do *a* and *b*).²⁰ But since *Consistent ‘Ought’s Agglomeration* had the virtue of explaining cases like *Fighting or Serving*, these accounts owe us some alternative treatment of this case if they are genuinely to solve the problem of conflicting ‘ought’s. And each of these accounts has its own distinct treatment of *Fighting or Serving*.

For example, Paul McNamara (in his 2004) defines what we can call a ‘basically ought’ operator. The details of McNamara’s definition are too complex for us to discuss here, but these details will not be needed for us to appreciate the *structure* of his proposal. So let’s use the expression ‘basically ought’ to discuss these special ‘ought’s and continue to use ‘ought’ for ‘ought’s whether they are basic or not. McNamara’s idea then is that though *Consistent ‘Ought’s Agglomeration* does not hold, the following principle does:

Consistent ‘Basically Ought’s Agglomeration: if S basically ought to do *a*, basically ought to do *b*, and S can do *a* and *b*, then S ought to do *a* and *b*.

This allows McNamara to explain *Fighting or Serving* by claiming that in this case Smith basically ought to fight or serve and Smith basically ought to not fight. So by *Consistent ‘Basically Ought’s Agglomeration*, Smith ought to fight or serve and not fight. And then by ‘ought’ analog of *Single Reason Closure* (*Single ‘Ought’ Closure*) it follows that Smith ought to serve.

Importantly accepting *Consistent ‘Basically Ought’s Agglomeration* rather than *Consistent ‘Ought’s Agglomeration* allows McNamara to avoid the explosion result that we discussed earlier. This is because while McNamara accepts *Single ‘Ought’ Closure* and accepts that ‘basically ought’s entail ‘ought’s, he does not accept that a ‘basically ought’ to do *a* entails a ‘basically ought’ to do *b* if doing *a* involves doing *b*. To see how this avoids the explosion problem, suppose that ‘basically ought’s conflict:

- (1) You basically ought to meet Sam for a drink
- (2) You basically ought not to meet Sam for a drink

Since ‘basically ought’s are a kind of ‘ought’, we may apply *Single ‘Ought’ Closure* to (1) and generate:

²⁰ See Brown 1999; Hansen 2004; Horty 2003 and 2012; McNamara 2004, van der Torre and Tan 2000.

(3) You ought to [meet Sam for a drink or you torture a kitten]

The next step in our derivation would be to use *Consistent 'Ought's Agglomeration* to derive the following claim from (2) and (3):

(4) You ought to [[meet Sam for a drink or torture a kitten] *and* not meet Sam for a drink]

But since we reject *Consistent 'Ought's Agglomeration*, this step is blocked. Now *Consistent 'Basically Ought's Agglomeration* would apply if (2) and (3) were claims about 'basically ought's. But since (3) is not a claim about 'basically ought's and since we said that McNamara rejects the idea that a 'basically ought' to do *a* entails a 'basically ought' to do *b* if doing *a* involves doing *b*, we cannot generate some 'basically ought' version of (3) from (1). Thus, McNamara is able to avoid this explosion result by blocking the move to (4).

It is a significant fact that McNamara's solution is just one of a family of solutions that have this structure. Other theorists solve this problem in the same way except that they do not focus on 'basically ought's. Instead they discuss good reasons, phase-1 obligations, prima facie obligations, etc.²¹ All of these solutions to the problem do the same thing. Instead of accepting *Consistent 'Ought's Agglomeration* they accept something like *Consistent 'Basically Ought's Agglomeration* but replace 'basically ought's with good reasons, phase-1 obligations, prima facie obligations, etc.

It is the *structure* of these proposals that allows them to solve the problem. They can explain (perhaps, with unequal plausibility) Smith's case by saying that in that case Smith basically ought or has good reason or phase-1 ought or prima facie ought, etc. to fight or serve and Smith basically ought or has good reason or phase-1 ought or prima facie ought, etc. to not fight and show that it follows from this and their principles that Smith ought to serve. They can avoid explosion because in order to apply *Consistent 'Basically ought's Agglomeration* or a version of it that talks about good reasons, phase-1 obligations, prima facie obligations, etc. to get (4) we would need (3) to be a claim about 'basically ought's, good reasons, phase-1 obligations, prima facie obligations, etc. But each of these views denies that we can use *Single 'Ought' Closure* on (1) to generate such a claim. Thus, it is the shared *structure* of these views that allows them explain to Smith's case while avoiding explosion. It is no coincidence, then, that so many deontic logicians have converged on this structure despite the fact that these logicians have approached the problem of conflicting 'ought's from very different theoretical perspectives.²²

²¹ Horty 2012 discusses good reason, van der Torre and Tan 2000 discusses phase-1 obligations, Horty 2003 discusses prima facie obligations. Brown 1999 and Hansen 2004 make similar distinctions.

²² There are four kinds of views about conflicting 'ought's that do not have this structure. First Goble 2009 solves the problem of conflicting 'ought's by modifying *Single 'Ought' Closure*. We should not go in for this as a solution to our problem because Goble's modification would not allow us to generate the required reasons in *Two Breakfasts*. Second others accept something like *Single 'Ought' Closure* and *Consistent 'Ought' Agglomeration* but go on to deny the derivation of explosion (see, e.g., Beirlaen, Straßer, Meheus 2013). Though these solutions are plausible as theories of reasoning, I believe they are not adequate for dealing with (necessary) entailments among reasons because they give up on certain structural properties of logical consequence (see Nair 2014 for discussion). Third some views either fail to avoid explosion or fail to give any account of the cases that motivate our principles, for a comprehensive survey see Goble 2013. Thus, these three kinds of views cannot solve our problem. Finally, I have recently learned of a fourth kind of view

Since numerous deontic logicians have converged on this structure as a solution to the problem of conflicting obligations and since our problem has the same structure as this problem, a promising idea is that we can get a solution to our problem by taking advantage of this structure. What this means is that we should hold on to *Single Reason Closure* and reject *Consistent Reasons Agglomeration*. And we should develop some reasons-analog to *Consistent 'Basically Ought's Agglomeration*.

4.2 Non-Derivative CRA

This insight will help us solve our problem only if we can determine what to use as an analog of 'basically ought's in our principle. To answer this question, I turn to a familiar distinction in ethics.

The familiar distinction is the distinction between things that are intrinsically, or in my preferred terminology, non-derivatively good and things that are relationally or, in my terminology, derivatively good. Things that are derivatively good are good in virtue of standing in some important relation to other things that are good. Things that are non-derivatively good, on the other hand, are good but not in virtue of standing in some important relation to other things that are good. So for example, we might think that pleasure is non-derivatively good. And since pleasure is non-derivatively good, it may be that eating ice cream is derivatively good because it stands in an important relation to the non-derivative good of pleasure; it causes it. So causal relations count as one of the important relations that something can stand to something else that is good in order to count as being derivatively good.

But we should leave it open whether there are other important relations.²³ For example, a Kantian might believe that the only thing that is non-derivatively good is the good will but claim other things are derivatively good in virtue of standing in intentional relations to the good will. More generally, we should leave it open exactly which relations are the ones that are capable of making something derivatively good.²⁴

This familiar distinction in moral philosophy is important for two reasons. First, it allows theorists can settle questions about what is good in two stages. First they decide what the non-derivative goods are. Then they give an account of relations to non-derivative goods that explains the existence of derivative goods. Together, these claims provide an account of what all the good things are in the ordinary non-theoretical sense; the primary data theories of goodness attempt to explain are what the good things are in this non-theoretical sense of 'good'.

Second this distinction allows theorists to capture certain subtle properties of good things that would otherwise be missed. For example, pleasure is good and eating ice cream is good but these goods have different modal profiles. Pleasure is necessarily good but eating ice cream is contingently good. This distinction is explained by the fact that pleasure is a non-derivative good, the fact that eating ice cream is a

being developed in work in progress by Lou Goble. This view involves a radical departure from *Single 'Ought' Closure*, but has other features that may allow it to solve our problem. Unfortunately, I do not have the space here to discuss the prospects of this proposal.

²³ Cf. Kagan 1998, Korsgaard 1983.

²⁴ To say that something is non-derivatively good is to say that its goodness is not explained by something else being good. This does not rule out that it can be explained in terms of some other normative or non-normative notion.

derivative good, and the relation it needs to stand to pleasure to be good is a contingent one. Thus, this distinction allows us to make perspicuous the explanatory structure of our theories and explain certain features of the things that we want to explain that we would otherwise miss.

Now that we have introduced this distinction between derivative and non-derivative and seen what is important about it by talking about goodness, we can apply it to other normative notions. We can then help ourselves to a distinction between derivative and non-derivative reasons as well as a distinction between derivative and non-derivative ‘ought’s.²⁵ So from now on, let’s explicitly write ‘non-derivative reason’ and ‘derivative reason’ to discuss non-derivative reasons and derivative reasons respectively and use ‘reason’ without a modifier before it to discuss reasons without making a claim about whether they are derivative or not. Similarly, for ‘ought’s.

And let’s use this distinction in our reasons-analog of *Consistent ‘Basically Ought’s Agglomeration*. In particular, the principle is the following:

Consistent Non-Derivative Reasons Agglomeration: if there is a non-derivative reason for S to do *a*, a non-derivative reason for S to do *b*, and S can do *a* and *b*, then there is a derivative reason to do *a* and *b*²⁶

Consistent Non-Derivative Reasons Agglomerations says that there is a derivative reason to do *a* and *b* because doing *a* and *b* stands in an important relation to what there are non-derivative reasons to do: Just as the distinction between derivative and non-derivative goods allows us to explain the different features of goods like pleasure and ice cream (their different modal profiles), the distinction between derivative and non-derivative reasons allows us to explain different features of certain reasons. In particular, the distinction together with *Consistent Non-Derivative Reasons Agglomeration* tells us exactly which reasons give rises to reasons to do a conjunction and which reasons do not.

In short, the proposal is that we can explain why there are certain derivative reasons by showing that these derivative reasons stand in an important relation to non-derivative reasons. Since *Consistent Non-Derivative Reasons Agglomeration* is structurally analogous to *Consistent ‘Basically Ought’s Agglomeration* and since this structure is one that deontic logicians have used to approach the problem, this suggests that adopting this principle will solve our first problem.²⁷ We will look at this in detail in §5. But before that, I want to make a

²⁵ Cf. Parfit 2011: vol. 1 39, Väyrynen 2011: 190 and n. 20.

²⁶ Two comments may be in order. First, the exact analog of *Consistent ‘Basically Ought’s Agglomeration* would not say ‘there is a derivative reason’ in the consequent but would merely say ‘there is a reason (perhaps derivative, perhaps non-derivative)’. So the principle given in the text is strictly stronger than the exact analog of *Consistent ‘Basically Ought’s Agglomeration*. But this stronger principle is plausible and also solves our problem. Indeed, I accept the even stronger claim that says the non-derivative reasons explain the derivative reasons.

Second, there are complications that arise from the fact that I am only focusing on there being reasons to do acts in this paper and not discussing what those reasons are. In particular, there can be cases in which there is a derivative reasons as well as a non-derivative reason to do an act where these reasons are provided by different facts. While the ideas I develop below are compatible with this, they do not provide as clear verdicts about these cases. However, a generalization of the model in appendix two that I develop in work in progress can provide a clear treatment of these cases.

²⁷ The derivative/non-derivative distinction also allows us to notice a more restricted version of *Reasons Explain ‘Ought’s Directly* that

final observation that will suggest a promising solution to the second problem as well.

4.3 Rejecting Reasons Explain ‘Ought’s Directly

The observation concerns how we might be able to accept REASONS EXPLAIN ‘OUGHT’S without accepting *Reasons Explain ‘Ought’s Directly*: just as we can explain derivative reasons by showing that what we have a derivative reason to do is that which stands in an important relation to our non-derivative reasons, we might analogously explain what we ought to do by showing that what we ought to do is that which stands in an important relation to our non-derivative reasons. Of course, we have not yet identified the exact relationship involved in explaining ‘ought’s. Nonetheless, this is enough to see, abstractly at least, that *Reasons Explain ‘Ought’s Directly* might not be the only way that reasons could explain ‘ought’s. This suggests that we can straightforwardly resolve the tension between *Reasons Can Conflict* and ‘*Ought’s Cannot Conflict* on one hand and *Reasons Explain ‘Ought’s Directly* on the other by siding with *Reasons Can Conflict* and ‘*Ought’s Cannot Conflict* and rejecting *Reasons Explain ‘Ought’s Directly*.

Of course, we also thought *Reasons Explain ‘Ought’s Directly* was appealing because it looked like the most promising explanation of ‘*Ought’s Entail Reasons* (which, recall, says that if you ought to do *a*, then there is a reason for you to do *a*). But what makes the idea that we are considering so promising is that it shows that *Reasons Explain ‘Ought’s Directly* might not be the only ecumenical explanation of ‘*Ought’s Entail Reasons* that is compatible with REASONS EXPLAIN ‘OUGHT’S. Instead, if non-derivative reasons explain ‘ought’s as well as derivative reasons, we might be able to explain ‘*Ought’s Entail Reasons* without accepting *Reasons Explain ‘Ought’s Directly*. ‘*Ought’s Entail Reasons* would be true because ‘ought’s and derivative reasons have a “common cause”: non-derivative reasons.

In the next section, I describe the solution suggested by our discussion in this section in greater detail and show how it solves our problem.

5. The Solution

The idea that is suggested by the observations of the last section is that distinguishing between derivative reasons and non-derivative reasons will be useful in two ways. First, the distinction will allow us to develop a principle of agglomeration (namely, *Consistent Non-Derivative Reasons Agglomeration*) that explains our

would suffice for our second problem:

Reasons Explain ‘Non-Derivative Ought’s Directly: If an agent non-derivatively ought to do *a*, then this is in part because the agent has a non-derivative reason to do *a*.

We can see that the problem arises even if we assume only *Reasons Explain ‘Non-Derivative Ought’s Directly* rather than *Reasons Explain ‘Ought’s Directly*. This is because what Mary and Sally non-derivatively ought to do are those disjunctive acts. This ‘ought’ is not explained by some other ‘ought’.

Though I do not have the space to discuss the details here, the fact that our problem arises even only assuming *Reasons Explain ‘Non-Derivative Ought’s Directly* distinguishes it from superficially similar problems concerning instrumental reasons and ‘ought’s. See Bedke 2009, Raz 2005, Schroeder 2009, and especially Kolodny Forthcoming and Millsap ms for discussion of instrumental reasons.

cases and avoids explosion. Second the distinction will allow us to explain both derivative reasons and what we ought to do in terms of non-derivative reasons. This second point is important because it opens the door to there being cases where you ought to do so something and have a reason to do it, but that reason does not partially explain why you ought to do it. Instead, there can be cases where you ought to do some act and merely have a derivative reason to do that act and in these cases, both what you ought to do and what you have (derivative) reason to do is explain by having a non-derivative reason to do some (other) act.

In this section, we will develop these ideas in enough detail to show that they really do solve our problem. Appendix two provides a simple formal model of these ideas that allows us to formally check this solution.

5.1 The Theory

Let me now describe how non-derivative reasons explain derivative reasons as well as what we ought to do.

Roughly, my idea is that given a collection of non-derivative reasons, we have a derivative reason to do anything that is involved in doing what there are non-derivative reasons to do. For example, suppose I have non-derivative reason to give comments on my friend's paper (e.g., because I promised her that I would) and suppose I have non-derivative reason to meet a different friend for lunch (e.g., because I promised this other friend that I would). In this case doing what I have non-derivative reasons to do involves both giving comments and going to lunch. So in this case, I have a derivative reason to both give comments and go to lunch—there is a derivative reason to do this conjunction in virtue of the fact that there is a non-derivative reason to do each conjunct and doing conjuncts involves doing the conjunction.

Now this is only *roughly* my idea because non-derivative reasons can conflict. For example, it may be that there is a non-derivative reason to go to the store and a non-derivative reason to not go to the store. Perhaps, there are these reasons because I have made conflicting promises. In this case, I do not wish to claim that there is a derivative reason to do the impossible act of both going and not going to the store.

In order to avoid this problem, my proposal is this: what we have a derivative reason to do is anything that is involved in doing some most inclusive collection of acts such that (i) we have non-derivative reason to do each act in that collection and (ii) it is possible to do all the acts in the collection together. For short, I will say that we have a derivative reason to do what is involved in doing some most inclusive compossible set of acts that we have non-derivative reasons to do.

Let me illustrate how this account works. Suppose I have a non-derivative reason to go to the store, a non-derivative reason to not go to the store, and a non-derivative reason to go the dry cleaners. In this case, there are three acts that I have a non-derivative reason to do: go to the store, not go to the store, and go to the dry cleaner. Suppose that I can [go to the store and go to the dry cleaner] and can [refrain from going to the store and go to the dry cleaner]. Of course, I cannot [go to the store and not go to the store] and I

cannot [go to the store, not go to the store, and go to the dry cleaner]. So these collections of acts are not compossible.

In this case then, there are two most inclusive collections of compossible acts that I have non-derivative reasons to do. One consists of the act of going to the store and the act of going to the dry cleaner; the other of the act of not going to the store and the act of going to the dry cleaner. The (singleton) collection consisting of just going to the store is not a *most inclusive* collection because there are other acts (e.g., going to the dry cleaner) that we could add to that collection and still have a collection of compossible acts that there are non-derivative reasons to do. Similar remarks apply to the collection just consisting of not going to the store and the collection just consisting of going to the dry cleaners.

So according to my account, you have a derivative reason to do anything that is involved in doing the collection of acts consisting of going to the store and going to the dry cleaner and you have a derivative reason to do anything that is involved in doing the collection of acts consisting of not going to the store and going to the dry cleaner. For example, then, you have a derivative reason to go to the store and the dry cleaner and you have a derivative reason to not go to the store and the dry cleaner.

This same basic idea can be extended to deal with what we ought to do with two tweaks. First, what we ought to do does not depend on any old non-derivative reasons. For example, sometimes you might have much stronger non-derivative reason to do *a* than to do *b* where you can't do *a* and *b*. In this case, you ought to do *a*. So the first tweak will be to focus not on all the non-derivative reasons but the subset of those non-derivative reasons that are *undefeated* in the sense that you have an undefeated non-derivative reason to do *a* just in case there is no act incompatible with doing *a* that you have a better non-derivative reason to do.²⁸ So far then, the tweak to our account of derivative reasons would give us this: what we ought to do is anything that is involved in doing some most inclusive collection of compossible acts that we have *undefeated* non-derivative reasons to do.

Another tweak that is needed is that we should focus on what is involved in doing *any* collection of acts of the relevant sort rather than what is involved in doing *some* collection. The reason why we should do this is that we are assuming '*Ought's Cannot Conflict*' holds. Let me explain.

Suppose I have made two equally important promises one to go to the store, the other to meet Tom for lunch. And suppose in this case I cannot both go to the store and meet Tom for lunch. In this case each of these reason is undefeated because, we may suppose, you have no relevant reason to do anything else and the reasons are equally strong because the promises are equally important. If we were to say that what we ought to do is what follows from *some* most inclusive set of compossible acts that we have undefeated non-derivative reasons to do, we would get the result that I ought to go to the store and that I ought to meet Tom for lunch even though I cannot do both of these. This is because the collection consisting solely of going to

²⁸ See appendix two for some details about how the collection of undefeated non-derivative reasons is determined and some complications associated with this issue.

the store is a most inclusive compossible collection of acts that I have non-derivative reasons to do; similarly for the collection consisting solely of meeting Tom for lunch.

To avoid this result, we should say that what we ought to do is any act that is involved in doing *any* most inclusive collection of acts of the relevant sort. In this case, going to the store is not involved in doing any most inclusive collection of acts of the relevant sort and neither is meeting Tom for lunch. However, either meeting Tom for lunch or going to the store *is* involved in doing any most inclusive set of act. Thus, we get the correct result that you ought to meet Tom for lunch or go to the store.

To summarize then, we have an account of derivative reasons and what we ought to do in terms of non-derivative reasons. The account of derivative reasons is this:

There is a derivative reason for S to do *a* just in case and because doing *a* is involved in doing some most inclusive collection of acts such that (i) there are non-derivative reasons for S to do each act in that collection and (ii) it is possible for S to do all the acts in the collection together.

The account of what we ought to do is this:

S ought to do *a* just in case and because doing *a* is involved in doing any most inclusive collection of acts such that (i) there are undefeated non-derivative reasons for S to do each act in that collection and (ii) it is possible for S to do all the acts in the collection together.

To adopt one way of thinking about the kinds of explanations involved in normative theorizing, we can say that this account says facts about derivative reasons are complex facts about non-derivative reasons and relations of involvement among acts. And similarly, facts about what we ought to do are complex facts about (undefeated) non-derivative reasons and relations of involvement among acts.

The explanatory structure, then, of the theory has non-derivative reasons as basic and uses the fact that certain acts involve doing other acts together with non-derivative reasons to explain derivative reasons and what we ought to do. In this sense, non-derivative reasons function as a “common cause” of both derivative reasons and what we ought to do.

There is however a different explanatory structure that may have become apparent at this point. This is a structure in which the non-derivative reasons explain the derivative reasons and then derivative reasons explain what we ought to do. There are two comments to make about this alternative.

First, this alternative explanatory structure is strictly stronger than the one that I have offered. This is because while it is generally thought that if A explain B and B explains C, then A explains C, it is not generally thought that if A explains B and A explains C, then A explains B and B explains C. So the alternative explanatory structure entails my explanatory structure but not vice-versa. It is interesting fact that we need not adopt the more committal explanatory structure to solve our problem.

Second, it is not obvious absent further explanation why the stronger structure is correct for two

reasons. First, the stronger explanation requires us to spell out the details of how what we ought to do can be explained in terms of derivative reasons. And this turns out to be surprisingly difficult for reasons that I sketch in a note.²⁹ Second, and more importantly, even if this can be done the explanation of what we ought to do can be given without referencing these complications. I have just stated how it may be done. For this reason, the appeal to derivative reasons is in no way *essential* to explaining what we ought to do.

Because of this, my preferred way of viewing things is this: non-derivative reasons explain both derivative reasons and what we ought to do. Necessary connections between derivative reasons and what we ought to do are epiphenomenal in the sense that they do not arise directly because of explanatory connections between derivative reasons and what we ought to do. Rather they arise because of the fact that derivative reasons and what we ought to do are both explained by what we have non-derivative reason to do and involvement relations among actions.

With these clarifications in mind, let me now illustrate how this account solves our problem.

5.2. The Solution

We will begin by verifying that the account on offer captures the principles that we discussed and then turn to how each of the cases will be treated.

5.2.1 The Principles

I said that the package of principles that would avoid explosion are *Single Reason Closure* and *Consistent Non-Derivative Reasons Agglomeration*. Let us first verify that we avoid explosion and then see that this is precisely because we accept these principles.

The explosion result said that we have a reason to do *a* and a reason to do *b* and can't do both, then there is a reason to do any *c* such that we can do *c*. Here is a simple counterexample that shows that we do not get this result whether we are discussing derivative or non-derivative reasons. Suppose there is a reason to go to the store (perhaps a derivative reason or perhaps a non-derivative reason) and a reason to not go to the store (perhaps a derivative reason or perhaps not a derivative reason). Now one thing that I can do is throw my laptop on the floor right now. Let us consider then whether we get the result that I must have a reason to

²⁹ In order to explain why we ought to do *a* in terms of reasons to do *a*, we must be able to say something about how the strength of derivative reasons is determined. It is in fact easy enough to say when there is decisive derivative reason to do *a*: there is a decisive derivative reason to do *a* just in case doing *a* is involved in doing any most inclusive compossible collection of acts that there is undefeated non-derivative reason to do. That said, whether a reason is derivative is determined by whether the reasons for doing something are better than the reasons against doing this. So to actually vindicate this definition, I take it, we must be able to say in general when one derivative reason is stronger than the other and show that the definition of a decisive reason can be recovered as the special case of the reasons to do *a* being better than the reasons to do any alternative. This task is one that no one to my knowledge has accomplished. In work in progress, I show how this can be done with some considerable complication. But while the possibility of doing this is interesting, I do not believe that it shows that 'ought's are explained in the way the alternative explanatory structure suggests. This is because the explanation directly in terms of non-derivative reasons is considerably simpler and more elegant.

Now the alternative explanatory structure may seem more attractive because it vindicated *Reasons Explain 'Ought's Directly*. But there are two points to be made about this. First, I do not believe that *Reasons Explain 'Ought's Directly* is a principle that there is independent reason to capture even though many theorists accept it. This is why I developed a motivation for it based on *'Ought's Entail Reasons*. Second, as I am understanding *Reasons Explain 'Ought's Directly*, it claims that the reason to do *a* is an *essential* component of an explanation of why we ought to *a* in terms of reasons. This proposal does not vindicate that idea as I explain in the main text. Thanks to an anonymous referee for pressing me on this issue.

throw my laptop on the floor now.

We can look at a few different cases. First, suppose I have a non-derivative reason to go to the store and a non-derivative reason to not go to the store. In this case, I do not have a reason (derivative or otherwise) to throw my laptop on the floor because throwing my laptop on the floor is not involved in going to the store and is not involved in not going to the store. Second suppose I have derivative reason to go to the store because it is involved in buying vegetables and suppose I have derivative reason to not go the store because it is involved in caring for my sick child. In this case too, there is no reason (derivative or otherwise) to throw my laptop on the floor because throwing my laptop on the floor is not involved in buying vegetables and it is not involved in caring for my sick child.

Of course, it can be that I have a reason to throw my laptop on the ground when I also have a reason to go to the store and a reason to not go to the store. But this will only be true in cases where throwing my laptop on the ground is involved in doing some compossible collection of things that I have non-derivative reasons to do. And in these cases, there is no problem with this result. The explosion result is problematic because it gives rise to things that there is a reason for you to do that have no bearing whatsoever on what other things that you have reasons to do. My account does not do this.

And as I said, the account does this by accepting *Single Reason Closure* and *Consistent Non-Derivative Reasons Agglomeration*. *Single Reason Closure* says this:

if there is a reason for S to do a and doing a involves doing b , then there is a reason for S to do b

To see that my account validates this principle, suppose that there is a reason for S to do a . Now this reason may be derivative or not. If it is non-derivative, then there must be a reason to do b because you have a derivative reason to do anything that is involved in doing some most inclusive compossible set of actions that you have non-derivative reasons to do. Since a will be part of some such collection and since doing a involves doing b , it follows that there is a derivative reason and so a reason to do b . Suppose instead that there is a derivative reason to do a , then there must be a reason to do b as well because, by the transitivity of involving, if a is involved in doing some most inclusive compossible collection of acts that you have non-derivative reasons to do and if doing a involves doing b , b is involved in doing that most inclusive compossible collection of acts.

Next consider *Consistent Non-Derivative Reasons Agglomeration*:

if there is a non-derivative reason for S to do a and there is a non-derivative reason for S to do b , and S can do a and b , then there is a derivative reason for S to do a and b .

To see that this holds suppose there is a non-derivative reason to do a , a non-derivative reason to do b , and you can do a and b . In this case, a and b both will be part of a particular most inclusive collection of

compossible acts that you have non-derivative reasons to do. Since doing any such collection involves doing *a* and *b*, it follows that there is a derivative reason to do *a* and *b*.

Notice here that *Consistent Reasons Agglomeration* does not hold. There can be cases where there is a reason to do *a* and a reason to do *b* and you can do *a* and *b* even though there is no reason to do *a* and *b*. Suppose for example you have a non-derivative reason to do [*a* and *c*] and a non-derivative reason to do [*b* and not-*c*]. In this case, there are two most inclusive compossible collections of acts that you have non-derivative reason to do. The first contains the conjunctive act of doing *a* and *c*; the second the conjunctive act of doing *b* and not-*c*. In this case, there is a reason to do *a* because it is involved in doing the first collection and there is a reason to do *b* because it is involved in doing the second collection, but there is no reason to do *a* and *b* because it is not involved in doing either collection.

Next notice that *Reasons Can Conflict* and *Ought's Cannot Conflict* both hold. *Reasons Can Conflict* holds because we allow non-derivative reasons to conflict and set things up to allow derivative reasons to conflict as well. We also ensured that *Ought's Cannot Conflict* holds by focusing on what follows from *any* most inclusive compossible collection of acts that there are undefeated non-derivative reasons to do.

Finally turn to *Reasons Explain Ought's Directly*. To see this need not hold, consider *Two Breakfasts*. There Mary promised to meet Jeff downtown and Scott by the beach. Plausible, there is a non-derivative reason to meet Jeff and a non-derivative reason to meet Scott and they are each undefeated. In this case, what explains why Mary ought to either meet Jeff or Scott on my view is that this follows from *each* collection of most inclusive compossible acts that there are undefeated non-derivative reasons to do. Now since neither collection contains a non-derivative reason to [meet Jeff or Scott], *Reasons Explain Ought's Directly* does not hold in this case.

It is still nonetheless true that there is such a reason to meet Jeff or Scott. And this is because this reason is derivative has those same non-derivative reasons as a “common cause”. The reason to meet Jeff or Scott is derivative and arises from the fact that it follows from some most inclusive compossible set of non-derivative reasons.

More generally, even though *Reasons Explains Ought's Direct* does not hold *Ought's Entail Reasons* does hold. This is because if doing *a* is involved in doing *any* most inclusive collection of compossible acts that there are undefeated non-derivative reasons to do, doing *a* is also involved in doing *some* most inclusive collection of compossible acts that there are non-derivative reasons to do.

This shows then that the account developed here accepts the exact set of principles that I have claimed will avoid our problem.

5.2.2 Cases

I now turn to the cases that we have discussed. There are four.

First begin by considering *Speeding Law*. We want to show that there must be a reason to drive less than one hundred miles per hour given that there is a reason to drive less than one hundred miler per hour.

And as we know, *Single Reason Closure* is enough to explain this case and my account accepts *Single Reason Closure*.

Two Breakfasts was already discussed in §5.2.1. I showed that we get the desired result that I ought to meet Jeff or Scott and that there is a reason for me to meet Jeff or Scott.

Third consider *Lunch-Coffee-Dinner*:

Sally has made three promises. She promised Tom that she will meet him for lunch downtown. She promised Jack that she will meet him for coffee by the beach. She promised Ann that she will meet her for dinner in Santa Barbara. While Sally can make it to lunch and coffee, can make it to lunch and dinner, and can make it to coffee and dinner, she cannot make it to lunch, coffee, and dinner. There just isn't enough time for all that driving.

Here we have three non-derivative reasons that are equally good.³⁰ In this case, the acts are pairwise compossible but setwise impossible. So there are three most inclusive compossible collections of acts that there are undefeated non-derivative reasons to do. The first collection consists of the act of going to lunch and the act of going to coffee; the second the act of going to lunch and the act of going to dinner; third the act of going to coffee and the act of going to dinner. Doing any of these collections of acts involves doing the disjunction of conjunctions: going to lunch and coffee or going to lunch and dinner or going to coffee and dinner. Thus, in this case there is a derivative reason to do this and this ought to be done.

Fourth and finally consider *Fighting our Serving*:

Smith's country requires him to fight in the army or perform alternative public service. We may suppose then that Smith has a reason to fight in the army or perform alternative public service. Smith is also deeply committed to a pacifist religion that requires him not fight in the army. We may suppose then that Smith has a reason to not to fight in the army. Given that Smith has a reason to fight or serve and given that Smith has a reason to not to fight, Smith also has a reason to serve.

If we assume there is non-derivative reason to fight or serve and non-derivative reasons to not fight, we get the right result in this case. Since these are the only reasons in this case and you can do both of these things, they form a most inclusive compossible collection of acts that there is undefeated reason to do. Since serving is involved in fighting or serving and not fighting, it follows by our account that there is a reason to serve and that you ought to serve.

³⁰ Since it is an open question in moral philosophy what the non-derivative reasons are, what entitles me to make the assumption each promise provides a non-derivative reason? I make these assumptions for simplicity. Many other assumptions but not all assumptions would work. All that can be done to answer suspicion about my solution on this score is to consider alternative proposals on a case-by-case basis.

For example, the alternative view that we have non-derivative reason to keep as many promises as we can or that we have non-derivative reason to minimize promise-breaking would work for my purposes. But the idea that we *only* have non-derivative reason to keep *all* of our promises would not.

This is the right result. If the *only* thing that we have non-derivative reason to do is keep *all* of our promises and we can't keep all of our promises, then it is not the case that we ought to keep two of three promises. Indeed, this consequence of my view shows just what is so implausible about this view of the normative significance of promises.

Similar comments apply to *Fighting or Serving*. Thus, though I cannot prove this, I believe my view will get the right results in these cases.

6. Conclusion

In this paper, I have argued that two related problems show that it is difficult to make sense of cases involving conflicting reasons if we accept REASONS EXPLAIN 'OUGHT'S. I then developed a unified solution to these problems. An important upshot of this solution is that the distinction from moral philosophy between derivative and non-derivative reasons and work in deontic logic on conflicting 'ought's are important to the project of understanding how reasons explain 'ought's. It was these ideas that allowed us to clearly appreciate the structure of our problem. And it was these ideas that led to the central thought behind my solution: what there is a derivative reason to do and what we are ought to do is that which stands in an important relation to what there are non-derivative reasons to do. If reasons explain 'ought's, then this is, at least to first approximation, *how* they explain them.³¹

³¹ For helpful comments thanks to the audience of the USC speculative society, Tony Anderson, Josh Crabill, Steve Finlay, Lou Goble, Keith Hall, Matthew Hanser, Ben Lennertz, Alida Liberman, Nick Laskowski, Errol Lord, Doug Portmore, Indrek Reiland, Henry Richardson, Jacob Ross, Barry Schein, Sam Shpall, Justin Snedegar, Julia Staffel, Sigrún Svavarsdóttir, Gabriel Uzquiano-Cruz, Ryan Walsh, Aness Webster, Ralph Wedgwood, Aaron Zimmerman, and especially two anonymous referees at *Philosophical Studies*. Thanks most of all to Mark Schroeder for advice and criticism on every issue at every stage of this project. Finally, I thank the USC Provost's PhD Fellowship and Russell Fellowship for support.

A.1 The Life of *Single Reason Closure*

In my experience, *Single Reason Closure* is a principle that some people find highly suspect. This appendix is dedicated to briefly discussing the status of this principle.³²

I should start by admitting that it is unlikely that what I can say in this paper that is not directly dedicated to considering all aspects of *Single Reason Closure* will satisfy everyone. This is because this principle, or more accurately a version of it concerning ‘ought’s, has a rather complicated status in ethics and in deontic logic: some think it is obviously true (Hilpinen and Føllesdal 1971: 22 and Nute and Yu 1997: 5); other obviously false. But, I submit, this fact—that people’s assessments differ—tells us that a thorough and even handed assessment of *Single Reason Closure* requires careful consideration of both its merits and its problems. This paper will show some positive applications of *Single Reason Closure*. This contributes one important piece of evidence to the large body of evidence that we must consider in order to responsibly evaluate *Single Reason Closure*. What this piece of evidence tells us is that those who wish to reject *Single Reason Closure* must identify an alternative to it that is capable of resolving our two problems. So the main contribution of this paper with regard to *Single Reason Closure* will not be to defend it from problems others have raised but to highlight some underappreciated virtues of this principle.

That said, let me make two points about *Single Reason Closure* and the cases that motivate it. First let me mention one weakening of *Single Reason Closure* that avoids many problems for this principle that people raise (e.g., the problem that *Single Reason Closure* entails that we have reasons to do “tautological actions” such as murder or not murder). The weakening is this:

if there is a reason for S to do *a*, doing *a* involves doing *b*, and doing *b* is under S’s intentional control, then there is a reason for S to do *b*

where we say that doing an act is under an agent’s intentional control just in case both (1) if she intends to do it, she does it and (2) if she intends to not do it, she does not do it. It will be easy for the interested reader to check that even this restricted version of *Single Reason Closure* gives rise to the problem that I have described when combined with *Consistent Reasons Agglomeration*. And though I did not complicate the presentation of my solution so that it only validates this restricted principle, it will not be hard to see that the solution could be modified in this way.

Next, recall that we considered *Single Reason Closure* in the first place because we wanted to explain why there is a reason to drive less than one hundred miles per hour given that there is a reason to drive less than fifty miles per hour in *Speeding Law*. The second issue I would like to consider is an objection to the idea that there really is a reason to drive less than one hundred miles per hour in that case.

The objection begins with the observation that one can drive less than one hundred miles per hour by driving exactly eighty miles per hour. But, plausibly, there is no reason to drive exactly eighty miles per

³² Thanks to an anonymous referee for pressing me on this issue.

hour in this case. On the basis of these two premises, the objection concludes that there is no reason to drive less than one hundred miles per hour.

Before responding to what I take the heart of the objection, let me make two preliminary points. First *Single Reason Closure* alone does not imply that there is a reason to drive exactly eighty miles per hour so it gets this fact right. Second, not only is there not a reason to drive exactly eighty miles per hour, there is a reason not to drive exactly eighty miles per hour in this case. And *Single Reason Closure* is the most natural explanation of why there is this reason. Since driving less than fifty miles per hour involves failing to drive exactly eighty miles per hour, *Single Reason Closure* predicts that there is a reason not to drive exactly eighty miles per hour. These are two correct results about the case and the second result follows from *Single Reason Closure* and is a piece of evidence in its favor.

However, the heart of the objection, I take it, does not focus on whether *Single Reason Closure* alone predicts that driving exactly eighty miles an hour is something that you have a reason to do (as I said, it does not). Rather the heart of the objection is that *Single Reason Closure* predicts that there is a reason to do something (namely, drive less than one hundred miles per hour) and one way of complying with this reason (namely driving exactly eighty miles per hour) is not something you have a reason to do and indeed have a reason against doing. This observation seems to be that it cannot be the case that there is no reason to do an act which is a way of complying with something you have a reason to do.

So the bridge premise this objection relies on is this:

if there is a reason to do a and b is a way of doing a , then there must be a reason to do b

or something similar to this perhaps restricted to cases in which you ought to do a as well. Since there is no reason to drive exactly eighty miles per hour, the objector concludes by modus tollens that there is no reason to drive less than one hundred miles per hour. My response to the objection is that this bridge premise is false.

My response rests on three claims. First, at least in some cases (though perhaps not all) we have a reason to accomplish an end E that gives us a reason to take a (perhaps necessary, perhaps not) means M to E that is not a sufficient means to E . For example, I may have an excellent reason to drive my friend to the airport and that gives me a reason to clear the snow off of my car.

Next consider that there will always be a sufficient means M' to accomplishing M that precludes accomplishing E .³³ To return to the example, a sufficient means to clearing the snow off of my car would be to murder my friend, steal his shovel, and use it to shovel the snow off of my car.

Finally I submit there is no reason to do M' (at least only given the reason to do E and the reason to do M that it generates). To return to the example, there is no reason to murder my friend, steal his shovel, and shovel the snow off of my car.

These three general claims then provide a counterexample to the bridge premise because there is a

³³ There will always be such an M' because we can always let M' be $[M$ and not- $E]$.

reason to do M, M' is a way of doing M, and nonetheless there is no reason to do M'. More generally, this argument suggest that the bridge principle is false when applied to reasons to do acts that are generated by other reasons. It should be no surprise then that there is no reason to drive exactly eighty miles per hour in *Speeding Law*.^{34, 35}

This conclude my response to concerns about *Single Reason Closure*. That said, these comments do not suffice to address all of the problems that have been offered for *Single Reason Closure*. But as I said, my aim is not to solve all the problem for *Single Reason Closure* here. It is to show that there are certain problems understanding conflicting reasons and their connection to what we ought to do and that *Single Reason Closure* can be a part of a solution to these problems. That said, for those interested in these problems, I briefly discuss them in a note.³⁶

A.2 A Simple Formal Model

The formal theory that I will develop is an adaptation of a formal system developed by John Horty (2012). I adapt the system in two ways. First I considerably simplify the system. Though my ideas could be

³⁴ It is important to be clear here about why this response is not question begging against a reasonable objector. A reasonable objector may reject the motivation for *Single Reason Closure* for the reason provided, but she should not reject the first claim of my argument (the claim that sometimes we have reasons to take non-sufficient means to our ends). For those who wish to reject the first assumption of my argument, I do not offer any response. I simply note that this is a highly theoretical commitment that, to my mind, hardly provides a starting point for an objection to simple judgements about cases unless the commitment can be bolstered by strong independent argument.

³⁵ Counterexamples to the 'ought' version of the bridge premise can also be had with the help of the simple observation that for almost anything we ought to do there are sufficient means for doing it that are impermissible. Some of these counterexamples may also falsify the bridge premise stated in the text but judgments about these examples involving reasons are less clear.

³⁶ I discuss only two puzzles here. Begin with, a generalization of Ross's paradox (Ross 1941). Though I do not have the space here to engage with the literature on this topic in the level of detail it deserves, let me make four points. First a popular solution to Ross's paradox explains the paradox away on pragmatic grounds (Castañeda 1981). Second though not clearly impossible, it is not obvious that the speeding law case can be explained without allowing the Rossian inference: that case can be thought of as going from 'there is a reason to drive forty-nine mph or forty-eight mph or...' to 'there is a reason to drive ninety-nine mph or ninety-eight mph or ... or forty-nine mph or forty-eight mph or...'. Third, *Single Reason Closure* is entailed by other plausible principle so those who reject it face the task of explaining which of these principles they reject. For example *Single Reason Closure* follows from the principle that if there is a reason to do *a* and *b*, then there is a reason to do *a* together with the principle that there is a reason to do *a* just in case there is a reason to do *b* when doing *a* involves doing *b* and doing *b* involves doing *a*. Fourth and finally, the Rossian intuition is suspect. Apply the same intuition pump to neither mailing nor burning the letter. Neither mailing nor burning the letter is an act that you do not have a reason to do and indeed an act you have a reason to not do. So this intuition pump suggests that you have strong reasons to not [neither sending nor burning the letter]. Now if we make the extremely weak assumption that you have a reason to do *a* then you also have a reason to do *b* where *b* is the result of performing a de Morgans transformation on *a*, it follows that you do indeed have a reason to mail or burn the letter. So the Rossian intuition itself pulls in both directions and those who wish to draw the conclusion that *Single Reason Closure* fails based on it must face up to denying much more than *Single Reason Closure*.

Next consider a generalization of Professor Procrastinate type cases (Jackson and Pargetter 1986). Again, there is a rich literature on this topic that I cannot do justice to, but let me say two things in response. First and most importantly let me note that one prominent family of accounts of this case and solution to the so-called actualism/possibilism debate it is involved in preserves *Single Reason Closure*. This family of accounts is motivated by consideration independent of the ones here and I myself find these solutions plausible (see Portmore 2013 and Ross 2012 for a recent defense). Second, one of the original arguments used to suggest that we have a failure of *Single Reason Closure* in the Professor Procrastinate case involved appealing to a certain simple consequentialist semantics for 'ought'. But this semantics is controversial not only for its implication in first order ethics but also because (i) it is not the only consequentialist approach possible and there are alternatives that validate *Single Reason Closure* (Portmore 2011), (ii) non-consequentialist semantics abound that do not have this result such as the usual modal semantics given for the so-called Standard Deontic Logic and such as a more recent semantic treatment due to Wedgwood 2006 that integrates insights from logic, linguistics, and moral philosophy, and (iii) the approach that I develop is an alternative account of the nature of what we ought to do so the argument from Jackson's theory simply begs the question against it.

developed using the full resources of Horty’s system, I am simplifying here in order to introduce my main ideas in the most approachable form. Second as will become clearer as we go on, I interpret some of the objects in my system differently than Horty interprets the analogous objects in his system. As we will see, this difference in interpretation is important because it is what will allow me to do something Horty cannot do: reject *Reasons Explain ‘Ought’s Directly* while accepting *‘Ought’s Entail Reasons* and thereby solve our problem.³⁷

Let’s develop the system. To begin, instead of discussing action and involvements among actions, I will simplify things by discussing sentences and logical entailments among them. Let lowercase Greek letters α , β , γ , etc. be sentences. Accordingly, we will also officially have to treat ‘reason’ and ‘ought’ as operators on sentences.

Next since we know that the notion of a non-derivative reason is going to play an important role in this system, we should introduce a way of representing them. So let $!(\alpha)$ represent a non-derivative reason to do α . While this formalism does not allow us to represent what the reason is to do α , we do not need to add such details to our system because our problem does not turn on exactly what the reason is.

Since we will often be interested in discussing not just individual non-derivative reasons but also collections of non-derivative reasons, we should introduce a device for representing these as well. We therefore use \mathfrak{R} for a set of non-derivative reasons.

It will also turn out to be useful when we are giving some definitions below to have a function, *Consequent*, that takes a non-derivative reason and returns the thing that we have a non-derivative reason to do and similarly for sets of non-derivative reasons. So we define:

$$\begin{aligned} \text{Consequent}[!(\alpha)] &= \alpha \\ \text{Consequent}[\mathfrak{R}] &= \{x \mid \exists y \in \mathfrak{R} \text{ and } \text{Consequent}[y]=x\} \end{aligned}$$

Next, we must distinguish the undefeated non-derivative reasons from the non-derivative reasons.

To do this, let’s help ourselves to an ordering on non-derivative reasons \leq . We read $!(\alpha) \leq !(\beta)$ as ‘there is a least as good of a non-derivative reason to α as there is to do β ’. We assume \leq is a reflexive relation in the sense that the following claim holds:

$$!(\alpha) \leq !(\alpha)$$

and assume that \leq is a transitive relation in the sense that the following claim holds:

$$\text{if } !(\alpha) \leq !(\beta) \text{ and } !(\beta) \leq !(\gamma), \text{ then } !(\alpha) \leq !(\gamma)$$

Finally, since we have seen that it plausible to think that there are equally good conflicting reasons, we know

³⁷ Horty notices that his system cannot accept *‘Ought’s Entail Reasons* and, for this reason, says he has an “austere” theory of reasons (2012: ch. 2, §1). By my lights, such austerity is a cost and the interpretation provided here can be thought of as an amendment to his system that does not have this cost.

that there can be situations where α is inconsistent with β , $!(\alpha) \leq !(\beta)$, and $!(\beta) \leq !(\alpha)$.

We can now define an operator *Undefeated* that takes us from a set of non-derivative reason to the subset of it that contains undefeated non-derivative reasons:

$$Undefeated_{\leq}(\mathfrak{R}) = \{x \in \mathfrak{R} \mid \text{there is no } y \in \mathfrak{R} \text{ such that (i) } x < y \text{ and (ii) } Consequent[x] \text{ is inconsistent with } Consequent[y]\}$$

In other words, a reason to do some act is defeated if there is a better reason to do some act that conflicts with it. And an undefeated reason is just a reason that is not defeated.^{38, 39}

There are many simplifications involved in this formalism (see n. 36), but it is enough for our purpose. And our purpose, recall, is to develop an account on which non-derivative reasons explain derivative reasons and ‘ought’s by showing that they stand in an important logical relation to non-derivative reasons. What we still need to do is describe this logical relation.

To do this, we need the notion of a most inclusive collection of non-derivative reasons. We define this in terms a maximal consistent subset of a certain set as follows:

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

With this in hand, we can now define a relation $|\sim_{\text{HNRE}}$ that we interpret as telling us *how non-derivative reasons explain* ‘ought’s.

This relation will hold between a collection of reasons of a certain weight and an ‘ought’ just in case the reasons explain the ‘ought’. Recall that a collection of reasons of a certain weight is formally represented by a pair $\langle \mathfrak{R}, \leq \rangle$ and we can represent an ‘ought’ with the operator O . So we may define this relation as follows:

$$\langle \mathfrak{R}, \leq \rangle |\sim_{\text{HNRE}} O(\alpha) \text{ iff } M \vdash \alpha \text{ for every maximal consistent subset } M \text{ of } Consequent[Undefeated_{\leq}(\mathfrak{R})]$$

where \vdash is the logical consequence relation of ordinary propositional logic. And we can also define when when $|\sim_{\text{HNRE}}$ holds between a pair $\langle \mathfrak{R}, \leq \rangle$ and the claim ‘there is a reason for it to be the case that α ’ which

³⁸ Strictly speaking, what I have defined is the notion of a reason being not worse than any other reasons. A reason being not worse and a reason being at least as good come apart in cases where reasons are incomparable with one another. But in order to present my ideas as simply as possible, I will ignore incomparability among reasons. Indeed, there are a whole host of phenomena concerning the weight of reasons that I will be ignoring in order to present my system simply (e.g., undercutting defeat, attenuation, how multiple reasons can “add up” to provide more support for an act, reinstatement). Luckily, as I said before, my system can be developed using the full resources of Horty’s system and these extra resources were developed precisely to understand these phenomena (see Hansen 2008 and Horty 2012 for further discussion).

³⁹ We do not focus on non-derivative reasons that are *better than* reasons that conflict with them because the set of such non-derivative reasons is empty in cases like *Two Breakfasts* and *Lunch-Coffee-Dinner* where we have equally good non-derivative reasons (cf. Horty 2003: 572-573).

we write as $R(\alpha)$:

$$\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha) \text{ iff } M \vdash \alpha \text{ for some maximal consistent subset } M \text{ of } \textit{Consequent}[\mathfrak{R}]$$

As I have defined it, $R(\alpha)$ is a claim about reasons derivative or otherwise. Having defined this notion, we may define a derivative reason as a reason that is not non-derivative. Formally this ends up looking like this:

$$\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} \text{DR}(\alpha) \text{ iff } \langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha) \text{ and } !(\alpha) \notin \mathfrak{R}.$$

where $\text{DR}(\alpha)$ is read as ‘there is a derivative reason for it to be the case that α ’.

Thus, my formal framework adapts Horty’s formal framework by thinking of the elements of \mathfrak{R} as non-derivative reasons.⁴⁰ It is easy to use this framework to now formally verify each of the claims I made about the principles.

Here I will verify *Single Reason Closure*, *Consistent Non-Derivative Reasons Agglomeration*, and ‘*Ought’s Entail Reasons*’. I leave it to the reader to perform the simple verification of other claims that she may be interested in:

- (1) Formally, *Single Reason Closure* says that if $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$ and $\{\alpha\} \vdash \beta$, then $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\beta)$. Suppose then $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$ and $\{\alpha\} \vdash \beta$. Since $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$, there must be some maximal consistent subset M of $\textit{Consequent}[\mathfrak{R}]$ such that $M \vdash \alpha$. Since \vdash is a transitive relation and since $\{\alpha\} \vdash \beta$, $M \vdash \beta$. Thus, $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\beta)$.
- (2) Formally, *Consistent Non-Derivative Reasons Agglomeration* says that if $!(\alpha) \in \mathfrak{R}$, $!(\beta) \in \mathfrak{R}$, and $\{\alpha, \beta\}$ is consistent, $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha \wedge \beta)$. Assume $!(\alpha) \in \mathfrak{R}$, $!(\beta) \in \mathfrak{R}$, and $\{\alpha, \beta\}$ is consistent. There is then some maximal consistent subset M of $\textit{Consequent}[\mathfrak{R}]$ such that $\{\alpha, \beta\} \subseteq M$. Since $\{\alpha, \beta\} \vdash \alpha \wedge \beta$ and since \vdash is monotonic, $M \vdash \alpha \wedge \beta$. Thus, $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha \wedge \beta)$.
- (3) Formally ‘*Ought’s Entail Reasons*’ says that if $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} O(\alpha)$, then $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$. So suppose $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} O(\alpha)$. Consider some maximal consistent subset of $\textit{Consequent}[\textit{Undefeated}_{\leq}[\mathfrak{R}]]$, M . By our supposition, $M \vdash \alpha$. By the definition of an *Undefeated*_≤, $M \subseteq \textit{Consequent}[\mathfrak{R}]$ and M is consistent. So M is either a maximal consistent subset of $\textit{Consequent}[\mathfrak{R}]$ or it isn’t. If it is, then we have $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$ because $M \vdash \alpha$. If it is not, then by definition there is some maximal consistent subset of $\textit{Consequent}[\mathfrak{R}]$, M' , such that $M \subseteq M'$. By the monotonicity of \vdash , it follows that $M' \vdash \alpha$. So $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$. Thus, $\langle \mathfrak{R}, \leq \rangle \mid \sim_{\text{HNRE}} R(\alpha)$.

⁴⁰ It may be worth noting how this system is related to two other systems. First, the system that can be found in Goble 2013: §4.4 is very similar to mine. When I came up with my system, Goble’s paper did not contain the system that is now found in his §4.4. At that time, the system that was closest to mine did not validate ‘*Ought’s Entail Reasons*’. Goble since has, perhaps independently, developed the system now found in §4.4. Second Portmore 2013 argues on very different grounds for a special case of my view (see n. 36 for further discussion). According to Portmore, if an agent has a non-derivative reason to do α , a non-derivative reason to do β , and $\alpha \neq \beta$, then α and β are inconsistent. In effect, Portmore thinks that there are only non-derivative reasons to do maximal consistent acts.

References

- Beirlaen, Mathieu, Christian Straßer, and Joke Meheus. 2013. “An Inconsistency-Adaptive Deontic Logic for Normative Conflicts” in *Journal of Philosophical Logic* 42: 285-315.
- Bedke, Matthew. 2009. “The Iffiest Oughts” in *Ethics* 119: 672-698.
- . 2011. “Passing the Deontic Buck” in *Oxford Studies in Metaethics* 6: 128-153.
- Brink, David. 1994. “Moral Conflict and Its Structure” in *Philosophical Review* 103: 215-247.
- Broome, John. 2004. “Reasons” in R Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith eds. *Reason and Value*. Oxford: Oxford University Press. 28-55
- Brown, Mark. 1999 “Agents with Changing and Conflicting Commitments” in Paul McNamara and Henry Prakken eds. *Norms, Logics, and Information Systems*. Amsterdam: IOS Press. 109-126.
- Cariani, Fabrizio. 2013. “Ought and Resolution Semantics” in *Nous* 47: 534-558.
- Castañeda, Henri Neri. 1981. “The Paradoxes of Deontic Logic” in Risto Hilpinen ed. *New Studies in Deontic Logic*. Dordrecht: D. Reidel. 37–86.
- Chellas, Brian. 1980. *Modal Logic*. Cambridge: Cambridge University Press.
- Dany, Jonathan. 2004a. *Ethics without Principles*. Oxford: Oxford University Press.
- . 2004b. “Enticing Reasons” in R. Jay Wallace *et al.* eds. *Reason and Value*. Oxford: Oxford University Press. 91-118.
- Donagan, Alan. 1984 “Consistency in Rationalist Moral Systems” in *Journal of Philosophy* 81: 291-309.
- Foot, Philippa. 1972. “Morality as a System of Hypothetical Imperatives” in *Philosophical Review* 81: 305-316.
- . 1983. “Moral Realism and Moral Dilemma” reprinted in Gowans 1987a. 250-270.
- Gert, Joshua. 2007. “Normative Strength and the Balance of Reasons” in *Philosophical Review* 116: 533-562.
- Goble, Lou. 2009. “Normative Conflicts and the Logic of ‘Ought’” in *Nous* 43: 450-489.
- . 2013. “Prima Facie Norms, Normative Conflicts, and Dilemmas” in Leendert van der Torre, John Horty, Dov Gabbay, and Ron van der Meyden eds. *Handbook of Deontic Logic and Normative Systems*. Milton Keynes: College Publications. 241-352.
- Gowans, Christopher eds. 1987a. *Moral Dilemmas*. Oxford: Oxford University Press.
- . 1987b. “Introduction” in Gowans 1987a. 3-33.
- Greenspan, Patricia. 2007. “Practical Reasons and Moral ‘Ought’” in Russ Shafer-Landau eds. *Oxford Studies in Metaethics* 2: 172-194.
- Hansen, Jörg. 2004. “Problems and Results for Logics about Imperatives” in *Journal of Applied Logic* 2: 39-61.
- . 2008. “Prioritized Conditional Imperatives” in *Autonomous Agents and Multi-Agent Systems*. 11-35
- Hilpinen, Risto and Dagfinn Føllesdal. 1971. “Deontic Logic: An Introduction” Risto hilpinen ed. *Deontic Logic: Introductory and Systematic Readings*. Dordrecht: D. Reidel. 1-35
- Horty, John. 2003. “Reasoning in Moral with Moral Conflicts” in *Nous* 37: 557-605.
- . 2012. *Reasons as Defaults*. Oxford: Oxford University Press.
- Jackson, Frank and Robert Pargetter. 1986. “Oughts, Options, and Actualism” in *Philosophical Review* 95: 233-255.
- Kagan, Shelly. 1998. “Rethinking Intrinsic Value” in *Journal of Ethics* 2: 277-297.
- Kolodny, Niko. Forthcoming. “Instrumental Reasons” in Daniel Star ed. *Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Korsgaard, Christine. 1983. “Two Distinctions in Goodness” in *Philosophical Review* 92: 169-195.
- Lemmon, E.J. 1962. “Moral Dilemmas”, in *The Philosophical Review* 70: 139-158.
- Marcus, Ruth Barcan. 1980. “Moral Dilemmas and Consistency” in *Journal of Philosophy* 77: 121-136.
- Millsap, Ryan. ms. “The Balancing Theory of Ought and Reasons Transmission”
- McConnell, Terrance. 2010. “Moral Dilemmas”, *Stanford Encyclopedia of Philosophy* Summer 2010 Edition. <<http://plato.stanford.edu/archives/sum2010/entries/moral-dilemmas/>>.
- McNamara, Paul. 2004. “Agential Obligation as Non-Agential Personal Obligation Plus Agency” in *Journal of Applied Logic* 2: 117-152.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Princeton: Princeton University Press.
- Nair, Shyam. 2014. “Consequences of Reasoning with Conflicting Obligations” in *Mind* 123: 753-790.
- Nute, Donald and Xiaochang Yu. 1997. “Introduction” in Donald Nute ed. *Defeasible Deontic Logic*. Dordrecht: Kluwer. 1–18.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- Portmore, Douglas. 2013. “Perform Your Best Option” in *Journal of Philosophy* 110: 436-459.
- Pietroski, Paul. 1993. “Prima Facie Obligation, Ceteris Paribus Laws in Moral Theory” in *Ethics* 103: 489-515.
- Ross, Alf. 1941. “Imperatives and Logic” in *Theoria* 7: 53-71.
- Ross, Jacob. 2012. “Actualism, Possibilism, and Beyond” in Mark Timmons ed. *Oxford Studies in Normative Ethics* 2. New York: Oxford University Press. 74-96.
- Ross, W.D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Raz, Joseph. 2002. *Practical Reasoning and Norms*. Oxford: Oxford University Press.
- . 2005. “The Myth of Instrumental Rationality” in *Journal of Ethics and Social Philosophy* 1.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Schroeder, Mark. 2009. “Means-End Coherence, Stringency, and Subjective Reasons” in *Philosophical Studies* 143: 223-248.
- Sinnott-Armstrong, Walter. 1988. *Moral Dilemmas*. Oxford: Basil Blackwell.
- van der Torre, Leendert and Yao Hua Tan. 2000. “Two Phase Deontic Logic” in *Logique et Analyse* 171-172: 411-456.
- van Fraassen, Bas. 1973. “Values and the Heart’s Command” in *Journal of Philosophy* 70: 5-19.
- Väyrynen, Pekka. 2011. “A Wrong Turn to Reasons?” in Michael Brady ed. *New Waves in Metaethics*. New York: Palgrave. 185-207.

- Wedgwood, Ralph. 2006. "The Meaning of 'Ought'" in Russ Shafer-Landau ed. *Oxford Studies in Metaethics* 1. New York: Oxford University Press. 127-160.
- Williams, Bernard. 1965. "Ethical Consistency" reprinted in Geoffrey Sayre-McCord ed. *Essays on Moral Realism*. 1988. Ithaca: Cornell University Press.