philosophies

MDPI

Article

Some Worries About Deontic Closure

Kevin Kimble

Department of Philosophy, National Chung Cheng University, Chiayi 621301, Taiwan; kekimble@ccu.edu.tw

Abstract: The Deontic Principle of Closure (DCL) appears initially to be a highly plausible principle. The DCL is commonly assumed in practical ethical reasoning, as when we make certain inferences about what we (morally) ought to do in particular situations. For example, if I am standing beside a burning house with several victims trapped inside and I have an obligation to rescue them, then if it is necessary for me to open the front door in order for me to lead them out, then it seems that I am morally bound in this situation to open the door. Similarly, if it is the case that I ought to keep my Friday morning 8:00 appointment with my student and it is a necessary condition for keeping the appointment that I wake up some time before 8:00, then I ought to wake up before 8:00. In spite of its attractiveness, however, various worries have been raised about the plausibility of this closure principle. In what follows, I shall critically examine DCL and discuss its plausibility in the face of various objections that have been raised against it. I shall argue that a slightly modified version of the principle circumvents the main objections and holds under several of the more or less standard interpretations of the "ought" operator.

Keywords: deontic closure; ought; moral dilemma; moral reason

Citation: Kimble, K. Some Worries About Deontic Closure. *Philosophies* **2024**, *9*, 182. https://doi.org/10.3390/philosophies9060182

Received: 2 October 2024
Revised: 23 November 2024
Accepted: 25 November 2024
Published: 5 December 2024

1. Introduction

The Deontic Principle of Closure (DCL) appears initially to be a highly plausible principle. DCL is commonly assumed in practical ethical reasoning, as when we make certain inferences about what we (morally) ought to do in particular situations. For example, if I am standing beside a burning house with several victims trapped inside and I have an obligation to rescue them, then if it is necessary for me to open the front door in order for me to lead them out, then it seems that I am morally bound in this situation to open the door. Similarly, if it is the case that I ought to keep my Friday morning 8:00 appointment with my student and it is a necessary condition for keeping the appointment that I wake up some time before 8:00, then I ought to wake up before 8:00. In spite of its attractiveness, however, various worries have been raised about the plausibility of this closure principle. In what follows, I shall critically examine DCL and discuss its plausibility in the face of various objections that have been raised against it. I shall argue that a slightly modified version of the principle circumvents the main objections and holds under several of the more or less standard interpretations of the "ought" operator.

2. Slote's Objection

Let's begin by examining an objection to closure raised by Michael Slote [1] (p. 24).[1] Slote observes that there are a whole host of related closure principles, corresponding to various kinds of necessity that are involved. There are closure or transfer principles for deontic necessity, epistemic necessity, and unavoidability or power necessity. The general form which captures the logical structure of these principles is what Slote calls the *main modal principle*

$$[\mathrm{N}\ (\mathrm{p}) \bullet \mathrm{N}\ (\mathrm{p} \supset \mathrm{q})] \supset \mathrm{N}\ (\mathrm{q})$$

where 'N' is an operator that stands for whatever type of necessity the principle employs. (The '⊃' is used here because I am reproducing Slote's principle verbatim; I use '→'

throughout the paper when expressing material or logical implication.) Slote's objection is centered around two main claims. The first of these is that any valid modal principle entails closure with respect to both conjunction introduction and logical implication. The first of these properties, sometimes referred to as 'agglomerativity', says that for any type of necessity you take (that is, for any 'N'), N (p) • N (q) implies N (p • q). The second of these properties claims that one can validly move from N (p) and □(p → q) to N (q). According to Slote, then, any valid closure principle entails both agglomerativity and closure under logical entailment. The failure of either of these properties for a given type of necessity undermines the relevant closure principle involving that particular kind of necessity. Slote's second important claim is that if these properties fail for a certain kind(s) of necessity (or if it turns out that the relevant closure principle(s) is simply invalid), and there are other forms of necessity that have certain salient features in common with the first kind, then there is prima facie reason for thinking that the modal (closure) principle fails as well with respect to those other forms of necessity. While Slote concedes that the known failure of some kinds of necessity to obey the main modal principle is not a *decisive* reason for inferring a like failure with respect to other relevantly similar forms of necessity, nevertheless it would still give us *some* reason for thinking that such closure principles stand or fall together.

With regard to his second main claim, Slote offers several examples which purport to show that agglomerativity and closure under logical implication fail for epistemic and deontic necessity, and then goes on to offer an explanation for their failure. That feature which explains the failure of both the deontic and epistemic closure principles is also present in power necessity and other alethic modalities; hence, closure fails for power necessity as well.[2] In other words, Slote sees a crucial similarity between the deontic version of the main modal principle, the epistemic version, and the 'power necessity' version; all have an important factor in common, and if that factor renders invalid one of the closure principles, then there are good grounds for thinking that the others fail for the same reason. The relevant feature which Slote isolates is what he terms *selectivity*:

> Some forms of relational necessity can arise only in narrowly circumscribed ways; and when restrictions on the way a given kind of relational necessity can come into being unhinge it from agglomerativity (or closure or our main principle), we may say that such necessity is selective. Thus if obligation is nonagglomerative, that is, as we have seen, because of limitations on the way (relational) obligations can arise, it is because obligations to do specific things typically derive from undertakings (to individuals) to do those very things. So, obligation is not only relational, but selective. And both these factors enter into the logic of the notion [1] (p. 13).[3]

When Slote says that a form of necessity is "relational" and "selective", he means first that the necessity exists in relation to certain "factors" present for that kind of necessity (e.g., the specific person(s) or undertaking(s) involved when a particular obligation is incurred), and second that the ways in which such relationality arise may prevent that necessity from being "transferred" via certain forms of inference. Slote argues that the logical properties of agglomerativity and closure under logical entailment are lacking for a certain form of deontic necessity (obligation), owing to the fact that such necessity is (he thinks) selective; hence, this deontic version of the main modal principle is invalid. And, Slote continues, since this feature of selectivity is present with respect to unavoidability or power necessity, the way is paved for arguing that this latter version of closure also fails. It is clear from his discussion that Slote also intends his argument to work the other way around—if power necessity could be shown to fail the main modal principle, then there would be good grounds for drawing the same conclusion with respect to deontic necessity.

Even if it turned out to be true that power necessity is selective in the way that Slote suggests and that the transfer principle associated with it fails [2,3],[4] nothing of any significance would automatically follow for deontic necessity and its respective closure principle under certain standard interpretations of ought, for there are many interpretations of the deontic operator which arguably are not selective. The examples Slote offers for moral

necessity in which closure and its logical properties fail due to selectivity are restricted to cases in which an agent puts himself under a specific obligation to some other person(s). If *S* is under an obligation to return a book to a friend *and* has an obligation to another friend to meet him at a certain time, it does not follow that *S* has *one* obligation to do both of those things. In other words, there is no single obligation *S* incurs to perform some joint act of returning-the-book-and-meeting-a-friend. Obligation thus fails to be agglomerative, and the reason why is that the necessity involved is selective—obligation is a moral boundness that arises in virtue of certain specific kinds of performances or undertakings to specific persons, and exists only in relation to those particular undertakings. Such selectivity "unhinges" this necessity from logical properties such as agglomerativity and implication. But while this insight may be valid with respect to *obligation*, a problem arises in that deontic necessity in general does not seem to be selective in this way when the 'ought' operator is interpreted in some of the standard ways, measured in terms of the strength or weight of moral reasons:

O_p *(PMR) There is a prima facie moral reason(s) for S to do A in circumstances C*
O_o *(OMR) There is an overriding or all-things-considered moral reason(s) for S to do A in C*
O_n *(NMR) There is an undefeated or non-overridden moral reason(s) for S to do A in C*
O_d *(DMR) There is a definitive moral reason(s) for S to do A in C* [4][5]

For example, perhaps *S* ought *all-things-considered* to return the book and ought *all-things-considered* to meet his friend, from which it plausibly follows that *S* ought *all-things-considered* to do both. This inference might be valid even though there is no one specific "obligation" that demands that *S* perform the conjunction of both actions. So, Slote's examples do not undermine the plausibility of the DCL on the OMR construal of 'ought'.

This same point applies to Slote's discussion of deontic necessity and closure under logical entailment. According to Slote, if *Y* has promised to meet her friend *Z* at 3:00 tomorrow, then *Y* is under an obligation to do so. Moreover, *Y's* meeting *Z* at that time entails that *Y* will be alive tomorrow. But Slote claims that it is wrong to say that *Y* has incurred some specific obligation to be alive tomorrow. Hence, he concludes that closure under logical implication fails for obligation, and once again the explanation of the failure is to be found in the feature of selectivity that such necessity carries with it. But as in the case of agglomerativity, if we interpret 'obligation' as, say, a PMR all the way through the chain of reasoning, then this example does not evidently fail our given closure principle. It is not readily apparent that deontic necessity construed in any of the ways suggested above is selective in the way Slote alleges it is for other types of necessity. Such being the case, there is no compelling reason to think that such forms of deontic necessity are analogous to obligation, knowledge, and power necessity in the way required for one to be justified in evaluating the validity of the deontic closure principle on the basis of the success or failure of those other closure principles.[6]

Let us now return to the first major claim of Slote's article, the contention that any valid modal principle (that follows the form of his main modal principle) entails closure with respect to both conjunction introduction and logical implication. Commenting on the main modal principle, for any kind of necessity one wishes to substitute, Slote argues

> Anyone who assumes the validity of arguing from 'N (p)' and 'N (p ⊃ q)' to 'N (q)' would seem to be tacitly assuming that the necessity expressed in the operator 'N' is both agglomerative (closed with respect to conjunction introduction) and closed under logical implication, so that one can, e.g., validly move from 'N (p)' and 'N (p ⊃ q)' to 'N (p • p ⊃ q)' and from the latter to 'N (q)' [1] (p. 10).[7]

I deny that closure principles for certain forms of necessity logically require or presuppose agglomerativity and closure under logical implication. In particular, I deny that certain plausible deontic closure principles *by themselves* entail these properties. It is far from clear, for example, why reasoning from the main modal principle logically requires that 'N (p · (p ⊃ q))' also be true in order to be able to infer 'N (q)'—especially if the

necessity involved is non-alethic [6].[8] Slote merely asserts this without giving any supporting argument.

To illustrate, closure under conjunction introduction and logical implication seem not to be (logically) necessary for certain forms of the deontic closure principle. Take the following highly plausible principle, *closure of 'ought' under 'ought'* (COO):

$$\text{(COO) } [O(p) \bullet O(p \rightarrow q)] \rightarrow O(q)$$

Ethical reasoning according to COO seems to be quite commonplace. Suppose that a judge ought to sentence a certain convicted criminal to prison, and that it is also morally incumbent upon the judge that *if* she sentences the criminal, *then* the punishment not exceed the seriousness of the crime. We can then conclude that the judge morally ought not hand down an excessive sentence.[9] Agglomerativity (AG) cannot be derived from COO alone or in conjunction with other standard *axioms* of deontic logic. Nor does COO entail closure of 'ought' under logical implication (COL):

$$\text{(COL) } [O(p) \bullet \Box(p \rightarrow q)] \rightarrow O(q)$$

In order to be able to derive COL from COO, we would need to have available some standard inference rule that allows us to go from $\Box(\varphi)$ to $O(\varphi)$. But such a rule is clearly invalid. And there is no way to derive AG from COO without further assuming another closure-like deontic principle, namely

$$(DCL)\ [O(p) \bullet (p \Box\!\!\rightarrow q)] \rightarrow O(q)$$

(Here, '$\Box\!\!\rightarrow$' stands for 'nomologically implies', that is, the implication holds in all possible worlds in which the actual laws of nature are held constant. I use '$\Box\!\!\rightarrow$' to designate nomological necessity throughout the rest of this paper.) Thus, the claim that all closure principles entail agglomerativity and closure under logical implication for their respective types of necessity seems clearly mistaken.

At the same time, I am willing to grant that certain closure principles, sometimes in conjunction with other plausible principles, do in fact logically presuppose agglomerativity and closure under logical implication. I will now show that DCL in conjunction with COO and one other intuitively plausible principle together entail AG, and that DCL plus this second principle alone together entail COL. This "other principle" is the uncontroversial truth that logical necessity implies nomological necessity, or

$$\text{(LIN) } \Box(p \rightarrow q) \rightarrow (p \Box\!\!\rightarrow q).$$

Proof. Proof that DCL $\models$ COL:

1. $[O(p) \bullet (p \Box\!\!\rightarrow q)] \rightarrow O(q)$	(DCL)
2. $O(p)$	Assumption
3. $\Box(p \rightarrow q)$	Assumption
4. $(p \Box\!\!\rightarrow q)$	3, (LIN)
5. $O(p) \bullet (p \Box\!\!\rightarrow q)$	2, 4, $\bullet$ I
6. $O(q)$	1, 5, $\supset$E
7. $[O(p) \bullet \Box(p \rightarrow q)] \rightarrow O(q)$ (COL)	2–6, $\supset$I

Proof. Proof that DCL (together with COO) ⊧ AG

1. [O (φ) • (φ □→ λ)] → O (λ)	(DCL)
2. [O (φ) • O (φ → λ)] → O (λ)	(COO)
3. O (p)	Assumption
4. O (q)	Assumption
5. □{p → [q → (p • q)]}	Necessity of a logical truth
6. {p □→ [q → (p • q)]}	5, (LIN)
7. O [q → (p • q)]	1,3,6, ⊃E, substitution
8. O (p • q)	2,4,7, ⊃E, substitution
9. [O (p) • O (q)] → O (p • q) (AG)	3–8, ⊃I

If these proofs are correct, then DCL does indeed logically presuppose COL; and given the highly plausible assumption of COO, DCL does entail AG.[10] One implication of this is that the 'ought' operator in DCL possesses these logical properties and therefore must obey both agglomerativity and closure under logical entailment. If either fails to hold for a given sense of the operator 'O' employed in DCL, then DCL itself is also rendered invalid on that reading and so is any reasoning based upon it. Since COL is a weaker principle than DCL, if the former is untenable then so is the latter. So, although these derivations by themselves "prove" nothing about what can be derived from other kinds of principles employing quite different sorts of necessity that instantiate Slote's main modal principle (e.g., epistemic closure, unavoidability), there is perhaps some reason to think that at least certain readings of other necessity operators might generate inference principles that are logically related in this way. Thus, although I deny his second claim, I am in partial agreement with the first central claim in Slote's overall argument. But as I have shown, that is not sufficient to undermine the plausibility of deontic closure.

3. Some Alleged Paradoxes for COL

Recall that COL states that 'ought' is closed under strict (necessary) logical implication, where 'p → q' stands for a material conditional, and '□' is the operator for logical necessity:

$$(1)\ (COL)\ [O\ (p) \bullet \Box(p \rightarrow q)] \rightarrow O\ (q)$$

Some philosophers have attacked COL on the grounds that the principle generates unacceptable paradoxes [7,8].[11] I will briefly review two of these that have been discussed by Walter Sinnott-Armstrong and Donald Davidson, and then consider two other types of cases that they do not discuss.[12] First, there is the so-called *paradox of disjunctive obligation* [9,10].[13] If Alf ought to mail a letter, then since it is necessarily true that mailing the letter entails either mailing or burning it, Alf ought to either mail the letter or burn it. Yet it seems wrong to say this when in fact what he ought to do is mail it. Sinnott-Armstrong's reply is to the point: while it may seem *strange* to affirm this disjunction, it is nevertheless still *true* in virtue of the fact that Alf ought to perform the action which is the first disjunct.[14] It is essentially the same kind of oddity that is involved in claiming that if Alf raises his hand, then he makes it the case that, say, either he raises his hand or Superman leaps over the Sears Tower. It may also *seem* misleading to say that Alf does what he ought to do by merely burning the letter. But the "ought" Alf would fulfil in this case is *disjunctive* (made true by the fact that he ought to do the first disjunct), and not the singular act of merely burning the letter. Since O (φ v λ) does not imply O (λ), one cannot plausibly claim that he morally ought to perform the latter action.

The second type of problem, a version of which was set forth in an article by James W. Forrester, has been dubbed "The Good Samaritan Paradox" [11].[15] Let us suppose that (1) If Smith murders Jones, he ought to murder him gently, and that in fact (2) Smith does murder Jones. (1) and (2) jointly seem to imply that (3) Smith ought to murder Jones gently. But it is necessarily true that (4) If Smith murders Jones gently, then Smith murders Jones. So, if 'ought' is closed under logical implication, then we can infer that (5) Smith ought to murder Jones. But this consequence seems absurd; the fact that Smith *does* murder Jones should not

make it true that he morally *ought* to murder Jones. Sinnott-Armstrong follows Davidson's suggestion that (3) and the consequent of (1) can be analyzed either as (3′) O(∃x)(Mxsj & Gx) or (3″) (∃x)(Mxsj & O (Gx)) [12].[16] (3″) is clearly the correct reading, but since (3″) includes only the adverb 'gently' within the scope of the 'ought' and '(∃x)Gx' alone does not logically imply '(∃x)Mxsj', (3″) does not yield the paradox if COL is true. Thus, by attending carefully to the logical form of action sentences and the scope of the 'ought' operator, Forrester's paradox can be avoided without giving up the closure principle.

Although I concur with the Davidson/Sinnott-Armstrong analysis, there seems to be a more natural reading of (1) and (3) that is equally effective in dissolving the paradox. Given (3″), Davidson/Sinnott-Armstrong would read (1) as

$$(1'')\ (\forall x)[Mxsj \supset (Mxsj\ \&\ O(Gx))]$$

But (1″) is logically equivalent to

$$(1''')\ (\forall x)(Mxsj \supset O(Gx))$$

So (3) can be symbolized as

$$(3''')\ O(\exists x)Gx$$

Since the sentence within the scope of the 'ought' in (3‴) is not equivalent to the antecedent of (4), we do not have a genuine instance of COL in (3‴)-(5). To produce such an instance, we would have to rewrite (4) as

$$(4^*)\ (\forall x)(Gx \supset Mxsj),$$

which is clearly false. And (4) together with (1‴), (2), (3‴), and COL do not jointly entail (5). Thus, the Good Samaritan Paradox is unsuccessful in undermining COL.

Although COL appears to be immune to the above types of counter-examples, might not there be others that force us to revise COL or abandon it altogether? The next type of case I would like to consider is a kind of "paradox of created obligation". In many situations, the fact that one is morally bound to do something is at least partially a result of some prior action on her part. Perhaps the kind of example that most readily comes to mind is the duty to keep one's promises. Suppose that I am obligated to attend a friend's wedding because I promised him that I would. According to COL, since

(1) I ought to fulfil my promise to attend my friend's wedding, and
(2) Fulfilling my promise to attend my friend's wedding entails that I make a promise to attend his wedding,

it therefore follows that

(3) I ought to make the promise to attend his wedding.

It also seems to follow that any promise that I ought to fulfil is a promise that I was obligated to make; and this is clearly absurd. This case can also be generalized to include any kind of created obligation whatsoever, such as paying a debt, returning stolen money, etc.

This example can be analyzed in a manner similar to that of Forrester's paradox, by paying close attention to the exact nature of the obligation in (1). Recall that COL says [O (p) & □(p → q)] → O (q). The antecedent of (2) in our exampl[17] can be symbolized as

$$(2a)\ (\exists xy)(Pyif\ \&\ Fxyif),$$

and so (2) would be read as

$$(2')\ (\exists xy)(Pyif\ \&\ Fxyif) \rightarrow (\exists y)Pyif$$

Now COL requires that the entire sentence expressed by (2a) falls under the scope of the deontic operator in 'O (p)'. So, (1) would be symbolized as

$$(1') \; O(\exists xy)(Pyif \; \& \; Fxyif)$$

(1′) and (2′) together with COL entail

$$(3') \; O(\exists y)(Pyif)$$

The problem is that (1′) is an incorrect reading of (1), which should instead be read as

$$(1'') \; (\exists xy)(Pyif \; \& \; O(Fxyif))$$

(1) by itself does not say that I ought to make a promise. But since the scope of the operator in (1″) includes only part of 'p', (1″) along with (2′) and (3′) does not constitute a genuine substitution instance of COL. Thus, this alleged counter-example to COL fails because if (1) is analyzed correctly, (1)–(3) is not a true instance of COL, and the only way to generate a true instance is to read (1) in a way that does not capture its true meaning and in fact makes (1) turn out false. Hence, I conclude that the "paradox of created obligation" fails to refute COL.

The final type of example I would like to consider is a case in which 'q' is a logical truth. Since any logical truth is entailed by any sentence whatsoever, any action sentence will entail every logical truth. Substituting for COL, for any action sentence D v ~D that is a logical truth, then [O(W) & (W → (D v ~D))] → O(D v ~D). Now suppose that I ought to get up and go to work today. Since getting up and going to work today entails that I will either drink a glass of milk this evening or not drink a glass of milk this evening, it follows that I ought to either drink or not drink that glass of milk. But clearly I am not under any sort of obligation to do either; thus, COL is falsified. Some think that COL can be salvaged at very little cost by simply placing on it the restriction that 'q' is not any action sentence that is a logical truth, or just by restricting 'q' to *simple* action sentences (none of which are logical truths). With this sensible and non-ad hoc restriction, it is argued, COL turns out to be a very plausible principle.

However, I find this move both unnecessary and undesirable. First, it just does not follow, contrary to what the objection assumes, that if I am obligated to do the disjunction, then I am obligated to do one or the other of the disjuncts. So, neither the fact that it is not true that I ought to drink a glass of milk, nor the fact that it is not true that I ought not drink a glass of milk provide any support for the claim that it is not the case that I ought to either drink or not drink the glass of milk. It is widely recognized that O (φ v λ) does not imply O (φ) v O (λ). Second, why not just accept that I ought to perform the above disjunction, or for that matter that I ought to perform any (disjunctive) action which is a logical truth, since all logical truths are entailed by any 'ought' statement whatsoever? Surely one fairly compelling moral reason for my performing Dv~D is that if (per impossible) I were not to perform that action, then the world would cease to be intelligible, including the moral order—for it would then follow that I ought to get up and go to work *and* that it is false that I ought to get up and go to work, and so on for every conceivable 'ought' claim.[18] So, perhaps moral agents are under an obligation after all to perform any actions which are logical or necessary truths. But this may seem a bit far-fetched. More to the point, there does not seem to me to be anything especially problematic about an agent's having moral duties or obligations that she cannot fail to perform or fulfil. For surely it is true that I ought not blow up the universe tomorrow, even though doing so is not within my power and hence is a mandate that I cannot avoid fulfilling. Similarly, when I travel this weekend, I ought to drive well below 500 mph on the freeway, even though there is no way for me to fail to do this. (These examples also show that the fact that we do not ordinarily talk about trivial disjunctive obligations does not mean that they do not exist—for before writing this chapter, I never before heard anyone mention an obligation to drive below 500 mph on the

highway, or an obligation to drive below 499 mph, etc.) Thus, I conclude that COL is valid as it stands and that there is no need to qualify it in the manner suggested above.

Even if COL is true, other ways of interpreting the closure principle still might turn out to be false. Recall a second version of deontic closure, mentioned above, namely

$$(DCL)\ [O\ (p) \bullet (p \mathrel{\Box\!\!\rightarrow} q)] \rightarrow O\ (q)$$

This closure principle implies that 'ought' is closed under 'can', where 'p $\Box\!\!\rightarrow$ q' is equivalent to '~◊(p • ~q)' on a counterfactual or nomological reading of the conditional. Here, '~◊(p • ~q)' means roughly that it is not within the agent's power to perform *p* without also performing *q*—the sentence holding true in all worlds with the same natural laws as the actual world—and hence asserts a nomological connection between *p* and *q*. (My use of "◊" in this section and the next designates nomological possibility and not strict logical possibility). DCL is assumed in many common moral arguments. The scenario I offered earlier about the person who has an obligation to rescue certain people trapped in a burning building, and who is therefore morally required to open the door in order to get them out, is an example of the application of DCL. Notice that cases such as this cannot be justified simply by appealing to closure under logical implication, for it is *logically* possible, for instance, that the victims in the burning building be rescued without opening the door. In the following sections, I shall focus exclusively on the question of whether there is a good objection to DCL, by examining two very different sorts of worries about the principle.

4. Foot's Objection to DCL

The first objection involves a case suggested by Philippa Foot, in which a necessary condition of what ought to be done is a morally wrong action:

> Suppose, for instance, that some person has an obligation to support a dependent relative: an aged parent perhaps. Then it may be that he *ought* to take a job to get some money. The obligation produces an *ought* related to means. So far so good. But what if the only means of getting money is by killing someone?... Should we nevertheless say that although the agent ought not to kill he also ought to kill? Surely this *ought* is destroyed by the superior injunction against taking life? [13] (p. 255).[19]

The agent has an obligation to support an aged parent, but the only means he has of obtaining the money to do this is by killing someone. If the DCL is true, then the person ought to kill someone in order to obtain the money. Foot claims that the latter is clearly false. Other examples raising the same objection are found elsewhere [14]. In commenting on Foot's example, Sinnott-Armstrong points out that whether or not the child ought to kill depends upon how 'ought' is interpreted. On the reading of 'ought' which says that the agent has overriding reasons for performing the action (OMRs), Foot is right to say that the child ought not to kill, since the reasons involved in the action are not *overriding*. But on this interpretation, neither is it true that the child "ought" (in the overriding sense) to obtain the money to support his parent [10] (p. 150).[20] The same result follows if 'ought' means *non-overridden* moral reasons (OMRs), because the child's reason for obtaining the money is overridden by the moral considerations which prohibit taking a life. The only way to make the counter-example seem plausible is to apply differing senses of 'ought' to the two actions involved in the scenario, but this would employ a version of the DCL in which the deontic operator is not interpreted consistently throughout.

However, Sinnott-Armstrong does not consider what would happen if the counter-example were modified so as to make the weight of both 'oughts' equal. Suppose that the child has equal obligations to support two dependent elderly relatives, but unfortunately only has the money and resources to support one. The 'ought' operator in such a scenario is interpreted in the sense of NMRs. Using 'O_n' to designate a non-overridden or undefeated moral requirement or "ought", and where A is 'the agent supports relative A' and B is 'the agent supports relative B', we have O_n (A), and [O_n (A) & ~◊(A & B)] → O_n (~B);

therefore, O_n (~B). But given the above scenario, it is also the case that O_n (B). So, can it still be true that O_n (~B)? There seems to be no good reason to deny this. Just because there is a non-overridden moral reason to take care of relative B does not mean there is *no* non-overridden moral reason of equal (but not greater) weight *not* to take care of B. Hence, this scenario is not a counter-example to the DCL on the NMR reading of 'ought'.

What happens if we interpret 'O' as a *definitive* moral requirement (DMR)? Employing O_d as the operator, our scenario becomes O_d (A) and [O_d (A) & ~◊(A & B)] → O_d (~B); therefore, O_d (~B). But it is also true that O_d (B), and this implies that ~B would be morally wrong. So, can it still be true that O_d (~B)? Anyone who claims that O_d (~B) is false and thus denies closure does so because he believes that an action cannot be *both* morally binding *and* morally wrong; that is, he holds that ~(O_d (B) & O_d (~B)), thus endorsing a consistency thesis with respect to DMRs. But whether or not such a consistency thesis is true is just as much an open question as is the plausibility of DCL. So, appealing to the consistency thesis as the basis for an objection against DCL is inconclusive. I conclude that Foot's argument fails to refute closure under 'can' and hence DCL remains intact.

5. The Problem of Morally Neutral Actions

Another type of objection that is sometimes raised is the worry that DCL implies that we ought to do actions that seem morally *neutral*. Sinnott-Armstrong, David O. Brink, and Ruth Barcan Marcus all discuss this kind of criticism [15–18].[21] Sinnott-Armstrong invites us to consider the following example. Suppose that I ought to pay my taxes but that I cannot do so without opening the drawer where my tax records are kept, and I cannot do that without moving certain air molecules with my hand. Closure under 'can' warrants the conclusion that I morally ought to move some air molecules. But it seems at best odd to say that I ought to move some air molecules, since such acts do not seem to be the kind of thing which one morally ought or ought not do. That is, they are morally neutral. Thus, although it is true that I ought to open the drawer where my tax records are kept, it is false that I ought to move some air molecules. So, DCL fails in this example. Sinnott-Armstrong concedes that such consequences of closure are odd but insists that they are still true. If I pay my taxes, which is what I ought to do, then my act will move some air molecules; thus, it is still the case that I ought (derivatively) to move some air molecules.

What if I do not know that the air molecules exist? Would it still be appropriate to say that I ought to move the air molecules? Sinnott-Armstrong defends DCL by distinguishing there *being* a (moral) reason to do a certain act from an agent *having* a reason to do the act [10] (pp. 149–150). Sometimes it is true that an agent ought to do something even though he does not know the reason(s) why he should do it. For instance, if a store is about to close, and it is the only place Joe can buy a certain medicine that his wife urgently needs, then Joe ought to leave soon, even though he does not know the store is about to close. Thus, Joe ought to do something, the reason for the doing of which he is ignorant. While this may be true, part of the reason for denying, in our example above, that I ought to move some air molecules stems not only from the fact that I do not have a reason to move them (since I do not know they exist, let us suppose), but also from the fact that I do not know that I *can* move them. Knowing or having reasons for performing an act is different from knowing that you can perform the act. Perhaps Joe does not know the full reason why he should leave soon, but at least he does know that he *can* leave soon. If for some odd reason Joe did not know that he *could* leave soon, would we still assume that he (morally) ought to leave soon? I will take up this question again shortly.

Sinnott-Armstrong attempts to circumvent this problem by adopting a restriction on closure, that is, by stipulating that the agent knows that he can perform the action that is implied by the original action which he ought to do. On this revision of the DCL, if I do not know that air molecules exist, then it does not follow that I ought to move them. However, it seems to me that the DCL can be adequately defended against the kind of objection outlined above by considering a line of response quite different from this one, and without having to restrict closure. For on closer inspection it seems that the particular

act of my moving my hand and opening the drawer (in circumstances C) *just is* my moving some air molecules (in C), albeit under a different description. (Let us call the former act 'D' and the latter act 'M'.) Now if D is identical to M, then it does follow from closure that O (M), for paying my taxes (T) requires that I perform D, and D in fact is an action that I am aware I can perform. Given closure, however, since O (T) and (T $\Box\rightarrow$ D), it follows that O (D). But since D = M, it is also true that O (M). Thus, the best course of action is to accept both O(D) and O(M). Sinnott-Armstrong's suggested modification of the DCL—to the effect that we restrict its application to actions the agent knows he is performing—is unnecessary; we can (and should) wholeheartedly accept the conclusion that O (M) despite its oddity.

What then accounts for the oddity of saying that I am morally required to do M? Here we can press into service an insight that Donald Davidson offers regarding actions and their descriptions. Specific actions are often referred to or described in terms of some terminal stage, outcome, consequence, or the like in virtue of which they fall under various descriptions [19] (pp. 686–687).[22] For instance, when I flip the switch, turn on the light, illuminate the room, and even alert a prowler to the fact that I am home, I perform only *one* act (as opposed to four) of which four different descriptions have been given. When I speak of opening a desk drawer to retrieve my tax documents, I normally do not refer to such an action under the description of "moving some air molecules with my hand." When viewed in terms of this way of describing the action, it does not seem to have the characteristic of being such that I morally ought to do it. An action that is morally required may seem like something that I ought to do when viewed under one description, while seeming to be morally neutral under some other description. Since the description involving my moving air molecules does not pick out the relevant feature(s) of my action in virtue of which the action is morally required, the action itself seems neutral with respect to moral considerations. There may also be a tendency to regard the action under this description as an action numerically distinct from the action as described in the usual way. In general, actions are morally required, just permissible, or morally forbidden owing to the fact that they possess certain properties or "right-making" characteristics; but if a given description of an action does not refer or allude to such characteristics, then such a description may make the action appear to be *amoral* or outside the sphere to which moral judgments apply.

The insight that actions fall under various descriptions, coupled with the fact that my act describable as "moving some air molecules" is identical with my act of opening a desk drawer, seems to provide us with a plausible explanation of the claim that O (M). However, in Davidson's discussion of the notion of act identity, he is primarily talking about act-*tokens* and not act-*types*. If I am referring to some particular instance (a dated event) of my opening my desk drawer, perhaps I can plausibly redescribe that particular act-token by referring to my moving certain air molecules in such and such a place at such and such a time (once again, a dated event). But can such an account be extended to act-types? How can we make sense of act-identity for act-types? The act-type of "my moving some air molecules" certainly is not identical to the act-type of "my opening the desk drawer where my tax forms are kept", for there are many tokens of the first type that cannot be plausibly identified with any token of the second type (My moving air molecules when swinging a baseball bat is not identical to any act of opening a desk drawer containing tax documents). So, the act-types (to be identified) must be specified in an appropriate manner and to an appropriate degree in order for the identity to hold.

Perhaps something like the following definition of *act-type identity* will work:

> (AI) An act-type A is identical to an act-type A* if and only if each token of A is identical to some token of A*, and each token of A* is identical to some token of A.

In Sinnott-Armstrong's example, each token of the act-type *in specified circumstances C, my opening the drawer (with my hand) where my tax documents are located* is identical to some token of the type *in circumstances C, my moving air molecules in a certain manner along a particular spatial pathway leading from my hand to the drawer where my tax documents are*

kept; and each token of the act-type *in C, my moving air molecules in a certain manner along a particular spatial pathway leading from my hand to the drawer where my tax documents are kept*, is identical to some token of the type *in C, my opening the drawer (with my hand) where my tax documents are located*. The phrase "in specified circumstances C" describes any other relevant conditions or features of the act-type that are needed to give it the appropriate degree of specificity. Event-types and act-types in general can have varying degrees of specificity. In the example we are considering, there will be some level of specification for act-type D, say D′, such that each token of D′ will be identical to some token of act-type M.[23]

Inherent in the act-type of *my moving some air molecules with my hand* (M), as stated in the example, is a sort of vagueness in the way it is described that makes it seem (on the surface at least) not to qualify as an act that I morally ought to do. On the one hand, M might mean something like

> M* my moving just any air molecules with my hand (and for no apparent reason).

Taken this way, M is not obviously an action that I ought to do. But neither is M* nomologically implied by my opening the desk drawer where my tax forms are kept. What is so implied is

> M** my moving air molecules in a certain manner along a particular spatial pathway leading from my hand to the drawer where my tax documents are kept.

In our example, what must be meant as the act implied by D is M**. And whether this is identical to the act of opening the drawer or whether it is simply a condition or consequence of my opening the drawer, I think that M** may be plausibly regarded as something that I derivatively ought to do. If so, then Sinnott-Armstrong's example is not a genuine counter-example to the DCL after all.

However, one problem with this strategy is that it can answer only a very limited range of cases. It does not enable us to handle, for example, a similar kind of example that David O. Brink discusses. Suppose that someone is obligated to give a talk on a college campus, one consequence of which is that she will see and pass by a certain building on the campus on her way home after the talk. It follows that she morally ought to pass by that building, even though the building is not something she has ever thought about or perhaps even *knows* is there. Brink avers that it is quite unnatural to say she ought to do this. In this case, it is not plausible to argue that, construed as act-types, her passing by the building (P) is simply one description of the identical action of her leaving the campus and going home after the talk (G). And while leaving the campus is arguably something she ought to do since it is counterfactually implied by her giving the talk and is something she knows she will be required to do in order to fulfil the obligation, it is not so clear that the same can be said for (P) [15].[24]

There is another related strategy that may be deployed with equal effectiveness in handling Brink's type of case. Instead of trying to identify act-types which appear under different descriptions and claiming that they are identical with one another, we can say that some act-types are *event-constituents* or *act-constituents* of other act-types. The idea here is simply that some act-types are constituted by or are made up of other act-types. For instance, *the lecturer's going home after the talk (in circumstances C)* includes the act-types of *moving her arms and legs*, *opening an exit door*, *leaving the building*, *passing by various landmarks*, and *crossing an intersection*, etc. Act-constituents (for act-types) can be defined in a manner that parallels our definition for act identity.

> (AC) An act-type *C* is an *act-constituent* of another act-type *A* if and only if each token of *A* has some token of *C* as an *act-constituent*.

Now the lecturer's passing by a certain building after the talk (P) is an act-constituent of her going home after the talk (G). Since, given the supposition made in Brink's example, P is nomologically implied by G, each token of G will include some token of P as an act-constituent. (To contrast: suppose that the speaker would cross Kirkwood Avenue only on one of the several routes home from the campus. Then crossing Kirkwood Avenue

(K) would not be an act-constituent of G, because not every token of G would have as an act-constituent some token of K. But neither then would K be nomologically implied by G.) Taking this a step further, I would like to endorse something like the following claim, which I shall call the *Principle of Closure of 'Ought' Under Act-Constituency*:

> (COAC) If some action *A* is morally required of an agent *S* in a set of circumstances *C*, then *S* is also morally required to perform each of the act-constituents of *A*.

In other words, if *S* (morally) ought to do *A*, then *S* ought also to do all of the things which make up or comprise *A*. This seems to me to be an intuitively plausible claim. The supposition that O (G) (which Brink grants), together with the COAC plus the claim that P is an act-constituent of G, jointly imply O (P); thus, it follows that the lecturer ought to pass by the building. Although the action P itself sounds amoral given the way that it is described in the scenario as something that has little relevance to the original obligatory action, such appearances are deceiving. For when understood as an act-type that partially comprises the larger action of *the lecturer's going home after the talk*, P can be seen as a legitimately derived moral claim. These considerations cast substantive doubt on the contention that Brink's type of scenario works as a genuine counter-example to the DCL; hence (in my judgment), the force of this kind of argument against deontic closure is severely undermined.

Moreover, I think that the reply outlined above also takes care of the worry that is sometimes raised in regard to an action's being something that the agent does not *know* that she can or will perform. At the time she incurred the obligation to give the talk, our agent *knew* that leaving the campus and going home is a necessary consequence of her coming to give the talk. On the other hand, at the time in question, she did not know that she would have to pass by a certain building. Could one reasonably conclude that it is not the case that O (P) because the agent did not originally know that she would have to pass by the building? Could it be that such actions do not fall under the scope of "ought"? That would be a bit too hasty; I see no good reason for demanding in general that every derived "ought" be known by the agent at the time of the original "ought", or even that it be known by the agent at all. Recall the scenario in which I ought to rescue the people trapped just inside the entrance to a burning building. Since saving them requires that I open the front door, I ought (derivatively) to open the front door. Suppose that at the time I incur this obligation (e.g., at the time that I happen to pass by the building, become aware of the fire, and assess that I can safely and easily rescue the victims) I am yet unaware that I will have to carry out a young woman, badly injured, who is lying just inside the entrance. It is still the case that I ought to carry the woman out, even though I did not know that she was lying there until I opened the door. So, one's lack of *prior* knowledge of the implied action does not seem to be a relevant factor (at least generally) in determining whether that action is morally required. Imagine further that I am also unaware that this young woman is six weeks pregnant, and that no one informs me of this. Is it not still true that I ought to rescue the fetus she is carrying, even though I have no knowledge of this at any time? Again, the inference that O (φ) $\rightarrow$ K (φ) does not seem to be generally valid.

I have outlined a strategy for defending DCL that renders completely innocuous a significant class of alleged counter-examples which draw on certain epiphenomenal consequences of our actions. As an initial approximation, any action (act-type) which at first glance *appears* to be morally neutral but is nevertheless an act upon which some moral requirement nomologically depends is in actuality either identical to an action that is a morally relevant derivative of the original requirement, or else is an act-constituent of a required action. Of course, there may yet be other worries about deontic closure that I have not addressed here. I conclude (provisionally) that the various counter-examples and arguments employed to refute deontic closure of 'ought' under both our strict implication and nomological readings of the closure relation all fail, and that COL and DCL remain intact.

Funding: This research received no external funding.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflict of interest.

Notes

1 Slote (1982, 24) [1]. Here I preserve Slote's original formalization.

2 "Power necessity" is the kind of necessity involved in the basic argument for the incompatibility of free will and determinism. The necessity operator means, roughly, that some proposition or state of affairs *p* is unavoidable for an agent, or that it is not within the agent's power to bring it about that *p*. The argument claims that if the thesis of determinism is true, then since propositions about the past and the laws of nature are "unavoidable" for an agent, it follows that any action the agent performs is also "unavoidable" for that agent.

3 Slote (1982, 13) [1]. The notion of selectivity employed here is to be distinguished from occurrences of the term in other contexts in moral philosophy.

4 For impressive criticisms of Slote's argument against closure principles associated with power necessity and unavoidability, see Fischer (1986) and O'Connor (1993) [2,3].

5 For an account of the notion of definitive moral reason or requirement, see Kimble (2013) [4].

6 In a footnote Slote (1982, 14, note 12), Slote acknowledges that it is much more difficult to produce obvious counter-examples to the modal (closure) principle for deontic necessity when the operator is construed as 'ought' [1].

7 This point is nicely summarized by Willer, (2012, 1–8) [5].

8 It may be another story when certain alethic modalities are involved, as I shall argue shortly. For a discussion of related issues, see Fine (2018, 634–664) [6].

9 Someone might argue that in my example what the claim about the judge *means* is simply 'p → O (q)', in which case reasoning according to COO might not be so "commonplace" after all. Fair enough; nevertheless, it is still the case that O (p → q) is *true*. For it is certainly true that the judge ought *not* to sentence the criminal *and* hand down a punishment that exceeds the crime. And this (by the rule of Implication) is equivalent to the formula employed in COO, namely 'O (p → q)'.

10 For similar proofs of how the transfer principle for power necessity entails both agglomerativity and closure under logical implication, see O'Connor (1993) [3].

11 The paradoxes of deontic logic commonly refer to absolute deontic modalities modeled by these logics and not to relative obligations which are sometimes used to avoid those paradoxes. See, for example, Avquist (2002) and Pacuit (2017) [7,8].

12 Sinnott-Armstrong also discusses a third type of case which he calls the "paradox of converse agglomeration", but the paradox only arises when the occurrences of 'O' in COL are interpreted in different senses. Since I have no interest in defending such a reading of the principle, I will say nothing more about it here.

13 This problem was first presented byRoss(1941, 53–71), and further discussed by Sinnott-Armstrong (1988, 142) [9,10].

14 Where Sinnott-Armstrong speaks of Alf's (disjunctive) *obligation*, I prefer to speak of action(s) Alf *ought* to perform. This way of taking the example avoids the potential difficulty of selectivity raised by Slote.

15 This version of the paradox originally appeared inForrester (1984, 193–197) [11]. For recent discussions of the problem, see Davidson(1980) and Sinnott-Armstrong (1988) [10,12].

16 Davidson (1980). The relevant material is found in Essay 6, p. 218 [12]. Again, I preserve Davidson's original logical formalization of the claims in my exposition.

17 Where 'x' ranges over actions, '(∃xy)(Pyif and Fxyif)' may be translated as "there is an act y of promising (to attend f's wedding) done by i to f, and there is an act x done by i to f which is a fulfilling of y."

18 For to not perform (B v ~B) would be to perform a contradiction (i.e., ~B • B), and a contradiction logically entails any proposition whatsoever.

19 Foot(1987, 255) [13]. A similar objection is raised by Marra (2019) [14].

20 For a clear explication of different senses of 'ought', see Sinnott-Armstrong (1988) [10].

21 See Brink, (1996, 123–124), Marcus (1996, 29–30), Sinnott-Armstrong (1992, 403–405), and Giordani (2021, 183–200) [15–18].

22 A more complete account of actions and events is found in Davidson (1980) [12].

23 When I speak of *act-tokens*, I mean to include *possible* tokens as well as *actual* tokens (tokens that have or will obtain in the actual world).

24 Brink (1996, 123–124) [15] distinguishes between the *upstream* consequences of an action and the *downstream* consequences of an action. The former refers to actions which are nomologically implied by the obligatory action and which are performed *prior to* that action; the latter refers to actions nomologically implied by the original obligation but which are performed *after* the obligatory action. While Brink maintains that the upstream consequences of our obligations are themselves obligatory, he leaves it an open question as to whether or not it is defensible to claim that the downstream consequences are morally obligatory.

References

1. Slote, M. Selective necessity and the free will problem. *J. Philos.* **1982**, *79*, 5–24. [CrossRef]
2. Fischer, J.M. Power necessity. *Philos. Top.* **1986**, *14*, 82–88. [CrossRef]
3. O'Connor, T. On the transfer of necessity. *Nous* **1993**, *27*, 209–235. [CrossRef]
4. Kimble, K. Dilemmas that matter. *Int. J. Philos.* **2013**, *1*, 29–37. [CrossRef]
5. Willer, M. Deontic paradoxes from a dynamic perspective. *Reason. Syst.* **2012**, *23*, 1–8.
6. Fine, K. Compliance and command II: Imperatives and deontics. *Rev. Symb. Log.* **2018**, *11*, 634–664. [CrossRef]
7. Avquist, L. Deontic logic. In *Handbook of Philosophical Logic*; Gabbay, D., Gunthner, F., Eds.; Springer: Dordrecht, The Netherlands, 2002.
8. Pacuit, E. *Neighborhood Semantics for Modal logic*; Springer: Berlin/Heidelberg, Germany, 2017.
9. Alf, R. Imperatives and logic. *Theoria* **1941**, *7*, 53–71.
10. Sinnott-Armstrong, W. *Moral Dilemmas*; Blackwell: New York, NY, USA, 1988.
11. Forrester, J. Gentle murder or the adverbial Samaritan. *J. Philos.* **1984**, *81*, 192–197. [CrossRef]
12. Davidson, D. *Essays on Action and Events*; Oxford University Press: New York, NY, USA, 1980.
13. Foot, P. Moral realism and moral dilemma. In *Moral Dilemmas*; Gowans, C., Ed.; Oxford University Press: New York, NY, USA, 1987.
14. Marra, A. Deontic Modality in Rationality and Reasoning. Ph.D. Thesis, Tilburg University, Tilburg, The Netherlands, 2019.
15. Brink, D. Moral conflict and its structure. In *Moral Dilemmas and Moral Theory*; Mason, H.E., Ed.; Oxford University Press: New York, NY, USA, 1996.
16. Marcus, R.B. *More about dilemmas. Moral Dilemmas and Moral Theory*; Mason, H.E., Ed.; Oxford University Press: New York, NY, USA, 1996.
17. Sinnott-Armstrong, W. An argument for consequentialism. *Philos. Perspect.* **1992**, *6*, 383–415. [CrossRef]
18. Giordani, A. Reason-based deontic logic. In Proceedings of the DEON Conference on Deontic Logic and Normative Systems, Munich, Germany, 21–24 July 2021; pp. 183–200.
19. Davidson, D. Actions, reasons, and causes. *J. Philos.* **1963**, *60*, 668–697. [CrossRef]