Inability and Permission*

Kyle Blumberg†

Draft of September 28, 2024

Abstract

Consider the following questions:

> **Obligation**: If S is unable to X, is S obligated to X?
>
> **Permission**: If S is unable to X, is S permitted to X?

Obligation has received a great deal of attention in ethics and allied fields. Many are convinced that it should receive a simple answer: always 'no'. In this paper, I focus on **Permission**, a question which has received much less scrutiny. I consider two simple answers to **Permission**—always 'no' and always 'yes'—and show that, granted plausible background assumptions, they have significant implications for the theory of obligation, on the one hand, and the theory of ability, on the other. I then develop a more sophisticated answer on which the meaning of permission claims has a counterfactual aspect. Among other things, this theory clarifies the relationship between permission and ability, and sheds new light on the connection between permission and obligation.

1   Introduction

I am interested in the following question: if an agent S is unable to X, what is the status of Xing in S's moral economy? In other words, what are S's obligations and permissions with respect to Xing? It will be helpful to divide this question into two subquestions:

> **Obligation**: If S is unable to X, is S obligated to X?
>
> **Permission**: If S is unable to X, is S permitted to X?

Obligation has received a great deal of attention in ethics and allied fields. Many are convinced that it should receive a simple answer: always 'no'. This is usually discussed under the guise of the so-called "ought implies can" principle:

---

†kyle.blumberg@unimelb.edu.au

OUGHT IMPLIES CAN (OIC): If S ought to X, then S is able to X.

$\mathrm{O}X \to \mathrm{A}X$

By contraposition, OIC is equivalent to the claim that if an agent is unable to X, then they are not obligated to X:

INABILITY IMPLIES NO OBLIGATION: If S is unable to X, then S is not obligated to X.

$\neg \mathrm{A}X \to \neg \mathrm{O}X$

Thus, proponents of OIC maintain that the answer to **Obligation** is always 'no'. Although OIC has its detractors, it remains a popular principle.[1] For instance, some argue that it best explains our judgments about examples such as the following:

(1)   a. I ought to snap my fingers and thereby end all suffering in the universe (Graham, 2011).
      b. You ought to travel back in time and prevent the crusades (Streumer, 2007).

These claims are infelicitous, and it is argued that this is captured by OIC. For example, even though it would be a good outcome, I am not able to end all suffering by snapping my fingers. Thus, given OIC, (1a) is false, accounting for its unacceptability. Similar remarks apply to (1b).[2]

In this paper, I focus on **Permission**, a question which has received much less scrutiny. I consider two simple answers to **Permission**—always 'no' and always 'yes'—and show that, granted plausible background assumptions, they have significant implications for the theory of obligation, on the one hand, and the theory of ability, on the other. For example, I show that the answer always 'no' implies the principle of "ability excluded middle", which is a controversial claim in the theory of ability modals (Kenny, 1975, 1976):

ABILITY EXCLUDED MIDDLE (AEM): S is able to X or S is able to not X.

$\mathrm{A}X \vee \mathrm{A}\neg X$

And I show that the answer always 'yes' is incompatible with the plausible principle of "weakening" for 'ought' claims (Cariani, 2013, 2016):

---

[1] The literature on OIC is vast. Proponents include Moore (1922), Prichard (1932), Ross (1939), Sapontzis (1991), Griffin (1992), Streumer (2003, 2007, 2010), Vranas (2007), Zimmerman (2008), Smith (2010), Schroeder (2011), Hedden (2012), Mason (2019). Critics include Lemmon (1962), Williams and Atkinson (1965), Brouwer (1969), Trigg (1971), Fraassen (1973), Heintz (1975), Brown (1977), Sinnott-Armstrong (1984, 1988), Rescher (1987), Saka (2000), Fischer (2003), Graham (2011), King (2019).

[2] See, for example, Streumer (2007) for an argument for OIC from examples such as those in (1).

WEAKENING: If S ought to X and S ought to Y, then S ought to $(X$ or $Y)$.

$(\mathrm{O}X \,\&\, \mathrm{O}Y) \to \mathrm{O}(X \vee Y)$

Since we should avoid these implications, we should not give simple answers to **Permission**. Instead we should answer '(only) sometimes'. This entails that there are actions X and Y such that (i) S is unable to perform either, and (ii) S is permitted to X, but S is not permitted to Y. For example, it is plausible that the actions *curing cancer* and *committing genocide* conform to this pattern. We may suppose that you are neither able to cure cancer nor commit genocide, yet intuitively (2a) is true but (2b) is false:

(2)   a. You are permitted to cure cancer.
      b. You are permitted to commit genocide.

This response, in turn, raises a further question: what accounts for the normative difference between actions such as *curing cancer* and *committing genocide* such that the former is permissible but the latter is not? In answering this question, I propose a novel analysis of permission. On this account, the meaning of permission claims has a distinctively counterfactual aspect; an agent's permissions are intimately connected to the moral value of counterfactual situations in which they perform the relevant action. To illustrate, when we evaluate (2a) we take into account counterfactual situations where you cure cancer. Since such situations are at least as good as any situations that you are actually able to bring about, (2a) is predicted to be true. On the other hand, when we evaluate (2b) we take into account counterfactual situations where you commit genocide. Since such situations are worse than any situations that you are actually able to bring about, (2b) is predicted to be false.

To be clear, I am not the first to maintain that certain deontic modal constructions have a counterfactual aspect.[3] However, as I will show, my theory differs from existing accounts in important ways. Apart from making explicit the connection between permission and ability, my account helps to clarify the relationship between permission and obligation. Along the way, we will also shed new light on a number of topics at the intersection of ethics and action theory, for instance the conception of agentive ability relevant for assessing an agent's obligations, and the notion of optionhood appropriate for moral theorizing.

The paper is structured as follows. In §2 I discuss some of my main background assumptions. Then in §§3-4 I consider the always 'no' and always 'yes' answers to **Permission**, respectively. In §5 I put forward a more sophisticated answer, and argue that permission has a counterfactual element. §6 briefly concludes.

2   Preliminaries

In §3 I consider a view on which the answer to **Permission** is always 'no'. But before we get there, in this section I want to make some of my central assumptions clear.

---

[3]See, for example, von Fintel and Iatridou (2008, 2023), Willer (2014), Silk (2021), Stojnić (2024), and the discussion in §5.3.

First, our discussion is confined to 'ought' (and its cognates, e.g. 'have to', 'must', etc.—also see below) and 'may' (and its cognates, e.g. 'permitted', 'allowed', etc.) as they are used to express concepts in normative ethics. These modals can be used to express other types of relations, for example so-called "epistemic oughts", but these won't be relevant in what follows.[4] We will also deploy modals such as 'can' and 'able'. These terms express a notion of agentive ability, i.e. they appear as agentive modals. It is a good question how exactly to interpret these expressions, and we will return to this issue in §3.2. But as a rough first-run gloss: they denote an ability concept that is stronger than physical possibility, and likely involves a notion of agentive control.[5] As for logical forms, I assume that all of the modals of interest denote a relation between an individual and an action (which, for simplicity, I'll model simply as a property of individuals).

Second, it is common for theorists to distinguish between "objective" and "subjective" normative concepts.[6] Roughly speaking, subjective obligations and permissions track a subject's normative commitments given their (possibly limited) state of information; while objective obligations and permissions track a subject's normative commitments given all of the relevant information. As far as I can tell, this distinction doesn't make any essential difference to my arguments. That said, some theorists have argued that OIC is more plausible when 'ought' is interpreted objectively; while others argue that OIC is more plausible when 'ought' is interpreted subjectively.[7] Thus, to the extent that an answer to **Permission** entails OIC, the attractiveness of this answer could depend on whether 'ought'/'may' are taken to denote objective or subjective moral concepts (for instance, see §3).

Finally, there is evidence that necessity modals differ in the strength of normative commitment that they express. 'ought' is taken to express a relatively "weak" normative concept, whereas modals such as 'have to', 'must', and 'obligation' are taken to express a comparatively "stronger" one.[8] To motivate this, note that the examples below are perfectly felicitous:

(3)　　a.　[The following is displayed on a sign in a restaurant bathroom.] Customers ought to wash their hands, but employees have to. (von Fintel and Iatridou, 2008)

　　　　b.　[You are on a flight and the seatbelt light has just been turned off.] You ought to keep your seat-belt fastened, but you don't have to.

---

[4]See Yalcin (2016), Hawthorne (2021), Boylan (2023) for more on epistemic oughts.

[5]See King (2019), Schwarz (2020) for a discussion of how 'able' and 'can' should be interpreted in the context of principles such as OIC, and see Mandelkern et al. (2017), Willer (2020), Boylan (2022), Mandelkern (forthcoming) for recent work on the fine-grained semantics of these expressions. Also see §3.2 for a discussion of Mandelkern's arguments in particular.

[6]See, for example, Prichard (1932), Parfit (1984, 2011), Jackson (1991), Hedden (2012), Zimmerman (2006, 2008, 2014), Muñoz and Spencer (2021) among many others.

[7]For arguments against a subjectivist interpretation of OIC, see Graham (2011, 2021) and King (2019) (though neither Graham nor King endorses OIC even when 'ought' is interpreted objectively). For arguments against an objectivist interpretation of OIC, see Mason (2003), Andrić (2017).

[8]The observation goes back at least to Sloman (1970). Also see Wertheimer (1972), Jones and Pörn (1986), McNamara (1996), von Fintel and Iatridou (2008), Silk (2021), von Fintel and Iatridou (2023) for discussion.

If 'ought' and 'have to' expressed the same moral concept in (3a), then the example would be infelicitous, owing to constraints on the use of 'but' (for instance, 'Ann is happy, but Bill is happy' is infelicitous because there is no obvious contrast between the two conjuncts). And (3b) would be incoherent. I am sympathetic to the idea that there are important differences between 'ought' on the one hand, and 'have to', 'must', etc. on the other. However, in this paper I will follow much of the ethics literature and essentially ignore these differences by assuming that 'ought' expresses a strong moral concept.[9] In particular, I will assume that the relevant reading of 'ought' expresses a notion of moral obligation. Those who find this troubling are welcome to substitute 'have to', 'must', or some other strong necessity modal for 'ought' in my arguments; this shouldn't lead to any dialectical loss. Indeed, in some cases it will be more convenient to consider the behaviour of strong necessity modals rather than 'ought', so I move freely between these expressions in what follows.[10]

3 No

The first answer to **Permission** which we will consider is relatively straightforward: always 'no'. That is, it consists of the following principle:

> INABILITY IMPLIES NO PERMISSION: If S is unable to X, then S is not permitted to X.
>
> $\neg AX \rightarrow \neg MX$

It will be more helpful to focus on the contraposed version of this claim, which is equivalent to "may implies can":

> MAY IMPLIES CAN (MIC): If S may X, then S is able to X.
>
> $MX \rightarrow AX$

MIC has received little attention in the literature.[11] Nevertheless, many ethical theories are committed to this principle. Consider accounts on which an agent's moral economy is constrained by the agent's "options" in context. Roughly speaking, an agent's option space in a context constitutes the range of actions eligible for deontic status, i.e. the actions which an agent can be obligated/permitted to perform. Crucially, it is widely maintained that ability is a

---

[9]For instance, Zimmerman (2008) assumes that on the relevant reading, 'ought' expresses a notion of all-things-considered moral obligation.

[10]For example, I use 'have to' in constructions that feature wide-scope negation, since 'ought' tends to sound awkward in such configurations; compare 'Ann doesn't have to donate $10 000 to charity' with 'It's not the case that Ann ought to donate $10 000 to charity' (von Fintel, 2012, Iatridou and Zeijlstra, 2013).

[11]One exception is Hughes (2018). Hughes examines arguments for OIC from considerations involving the purported "action guidance" properties of obligation. He maintains that in order for such arguments to go through, action guidance would need to provide justification for MIC. However, he argues that action guidance does not support MIC. Hughes leaves it open whether there are other motivations for MIC, and does not directly consider whether or not this principle is valid (fn.12, 10).

necessary condition on optionhood. Indeed, many accounts make option-actions "thin" and easy to achieve exactly because theorists take it as a fixed point that options must be actions which agents are able to perform.[12] Thus, many "option-sensitive" theories validate MIC.[13]

Granted that obligation asymmetrically entails permission, MIC is a straightforward strengthening of OIC. Consequently, those who reject OIC will reject MIC as well (along with the always 'no' answer to **Permission**). It is a good question whether OIC is valid, and I have some sympathy with those who answer negatively.[14] However, I want to remain neutral on OIC as much as I can here. Instead, I want to draw out some consequences of MIC which do not stem from OIC alone. Regardless of one's stance on OIC, I argue that MIC is problematic.

3.1   MIC and natural language

First, natural language provides some evidence against MIC. To make this vivid, consider the following case:

> *Cruise*: Bill and Charlie are on a cruise ship when the weather turns stormy. They get washed out to sea and will drown if they are not rescued. Ann is the only person on board who is able to help them. However, it is common knowledge that she only has one lifesaver available to her. So, Ann can save exactly one of Bill or Charlie, but not both.

(4)   a. Ann is permitted to save both Bill and Charlie.
　　　b. Ann is not permitted to save both Bill and Charlie.
　　　c. It is morally wrong for Ann to save both Bill and Charlie.

To my ears, and those of my informants, (4a) is felicitous, but (4b) is not. This can be made even sharper if we analyse 'It is morally wrong for S to X' as 'S is not permitted to X'. For then (4c) should be true given MIC, and yet this sentence is unacceptable.

More generally, it is often fine to assert 'S is permitted to X' when X would be a good outcome, even if it is common knowledge that it is impossible for S

---

[12]For instance, some philosophers maintain that options should be identified with "basic actions" familiar from action theory (Howard-Snyder, 2005), others maintain that they are minimally realizable movements (Koon, 2020), and yet others argue that they are internally individuated mental states: tryings (Jeffrey, 1983), willings (Joyce, 1999), intendings (Broome, 2013), decisions (Pollock, 1995, Hedden, 2012), or a *sui generis* type of planning state (Schwarz, 2021). But even those who permit options to be more robust types of actions still maintain that optionhood carries an ability constraint, for example see Portmore (2019).

[13]Some theories appeal to options in a more limited fashion. For instance, a fairly popular style of analysis maintains that S ought to X just in case "S *Xes* follows from S's best options". On many existing attempts to fill in the right-hand side of this biconditional, it is predicted that 'You ought to not cure cancer' is true, since it is common knowledge that you cannot cure cancer (Cariani, 2013, Cariani et al., 2013, Hedden, 2015, Blumberg and Hawthorne, 2023). Given DUALITY introduced below, (5a) ('You are permitted to cure cancer') below is predicted to be false. More generally, these accounts predict that actions which are known to be unperformable are never permitted.

[14]For relevant discussion, see the work cited in fn.1 as well as §5.1.

to bring X about. Correspondingly, 'S is not permitted to X' often sounds bad in these circumstances. To repeat an example from §1, it seems safe to assume that you are not able to cure cancer, and while (5a) is felicitous, neither (5b) nor (5c) is:

(5) a. You are permitted to cure cancer.
     b. You are not permitted to cure cancer.
     c. It is morally wrong for you to cure cancer.

For a final data point here, note that $\neg A\neg X \to \neg M\neg X$ is an instance of IN-ABILITY IMPLIES NO PERMISSION. Given DUALITY (introduced below), this is equivalent to $\neg A\neg X \to OX$ which is the claim that if an agent cannot help but X—that is, they are compelled to X—then they are obligated to X. This predicts that examples such as the following should be true and thus acceptable:

(6) a. If Ann is compelled to gamble, then she ought to gamble.
     b. If serial killers are compelled to murder people, then serial killers ought to murder people.

In fact, they are strikingly infelicitous.[15]

3.2   MIC and control

It is widely accepted that the ability concept expressed by agentive modals such as 'able' and 'can' is *control-entailing*: in order for S to be able to X, Xing needs to be under S's control in a substantive sense.[16]

> CONTROL: If S is able to X, then Xing is under S's control.

Given CONTROL, MIC can be shown to be problematic in two ways.

3.2.1   AEM

Our first argument appeals to DUALITY:

> DUALITY: S ought to X iff it's not the case that S may not X.
> $OX \leftrightarrow \neg M\neg X$

---

[15]Thanks to Annina Loets for suggesting that I consider how INABILITY IMPLIES NO PERMISSION interacts with compulsion, and for bringing the examples in (6) to my attention.

[16]A control constraint on ability is endorsed by Kenny (1975, 1976), Fusco (2020), Willer (2020), Boylan (2022), Kikkert (2022) among others. It's a good question how to interpret the control concept that is being employed here. The literature provides several options. But we don't need to get into this, since on almost any construal of 'control', it will be shown that MIC has problematic consequences.

DUALITY says that 'ought' and 'may' are duals, and is widely endorsed in both the ethics literature as well as the literature on deontic modals.[17] Some motivation for it comes from the fact that denials of this principle sound uniformly terrible:

(7) a. # Ann has to save exactly one person but she is permitted to not save exactly one.

  b. # Ann is neither permitted to save exactly two people nor does she have to not save exactly two.

Note that Ann is unable to save exactly two people in the *Cruise* case. Nevertheless, (7b) borders on sounding incoherent. This suggests that the pattern underlying DUALITY holds even when agents are unable to perform the target actions.

Assuming DUALITY, MIC implies "ability excluded middle":

ABILITY EXCLUDED MIDDLE (AEM): S is able to X or S is able to fail to X.

$AX \lor A\neg X$

As noted above, MIC entails OIC. Thus, given that we are assuming MIC, we may use OIC below. The argument goes as follows:

(1)	$OX \lor \neg OX$	TAUTOLOGY
(2)	OX	PREMISE
(3)	AX	2, OIC
(4)	$AX \lor A\neg X$	3
(5)	$\neg OX$	PREMISE
(6)	$M\neg X$	5, DUALITY
(7)	$A\neg X$	6, MIC
(8)	$AX \lor A\neg X$	7
(9)	$AX \lor A\neg X$	1, 2, 5, 8

The tension between CONTROL and AEM can be brought out by cases such as the following:[18]

> *Target*: Ann is given the opportunity to push a button which she knows will launch a rocket. She also knows that if the rocket is launched, it will land within a square field far from her location. The field is divided into two halves: the northern half and the southern

---

[17]Most theorists that I am aware of either implicitly assume DUALITY, or explicitly endorse this principle; see for example Zimmerman (2008). Furthermore, DUALITY is valid on the standard intensional frameworks used to model natural language modals, for example the influential theory of Kratzer (1977, 1981, 1991, 2012) (see §5.2 for a presentation of Kratzer's account).

[18]This general style of example originates with Kenny (1975, 1976). This case differs from Kenny's only in that it is made explicit that Ann will hit the field no matter how she acts, so that her failing to hit one half of the field is equivalent to her hitting the other half.

half. At noon, she pushes the button, and the rocket is launched. Shortly after noon, she is given a joystick whose movements influence the rocket's path. However, she has no idea (i) where exactly the rocket is heading, or (ii) how the joystick's movements impact the flight of the rocket, and has no means of finding out this information.

(8)   a. (Shortly after noon) Ann is able to hit the northern half of the field.
      b. (Shortly after noon) Ann is able to hit the southern half of the field.
      c. (Shortly after noon) Ann is able to hit the northern half of the field or Ann is able to hit the southern half of the field.

CONTROL implies that none of (8a)-(8c) is true in context, since Ann does not have any control over the flight of the rocket. However, AEM implies that at least one of (8a) or (8b) is true (and thus that (8c) is true). For suppose (8a) is false. Then by AEM, Ann is able to not hit the northern half of the field. Given that she pushed the button at noon, either she will hit the northern half of the field, or she will hit the southern half of the field. So, the only way for her to not hit the northern half of the field is if she hits the southern half of the field. That is, if (8a) is false, then (8b) must be true.[19]

3.2.2 OUGHT-TO-ABILITY

The argument to AEM given above appeals to DUALITY. However, assuming CONTROL we can put pressure on MIC even without this principle. Consider the following claim:

OUGHT-TO-MAY: If S ought to (X or Y), then S may X or S may Y.

$O(X \vee Y) \rightarrow (MX \vee MY)$

OUGHT-TO-MAY is very plausible, and I am not aware of any purported counterexamples in the literature. But it is straightforward to show that OUGHT-TO-MAY and MIC imply:

OUGHT-TO-ABILITY: If S ought to (X or Y), then S is able to X or S is able to Y.

$O(X \vee Y) \rightarrow (AX \vee AY)$

But OUGHT-TO-ABILITY is incompatible with CONTROL. For instance, consider a variant of the *Target* case where if Ann chooses to launch the rocket and hit the field, ten innocent lives will be saved. And if she chooses not to launch the rocket, the ten will be killed. Then 'Ann ought to hit the field' (= 'Ann ought to hit the northern half of the field or the southern half of the field') is true. But as discussed, given CONTROL Ann is unable to hit the northern half of the field and she is unable to hit the southern half of the field.

---

[19]We can make this argument a bit more precise if we assume that ability is closed under known conditional entailment. Then given (i) that it is common knowledge that if Ann does not hit the northern half of the field then she will hit the southern half of the field; and (ii) Ann is able to not hit the northern half of the field; it follows that Ann is able to hit the southern half of the field.

3.2.3 CONTROL VS SUCCESS

Although CONTROL is a popular principle, it has recently been challenged by Mandelkern (forthcoming). Mandelkern's key observation is that claims such as (8a) and (8b) can be felicitously embedded under probability and possibility operators. For instance, one can hear (9) as felicitous, and it is acceptable to reply to (10a) with (10b):

(9) Ann might be able to hit the northern half of the field.

(10) a. What is the chance that Ann will be able to hit the northern half of the field?

      b. 0.5.

However, as Mandelkern points out, it is difficult to square this data with a commitment to (8a) not being true. For example, if (8a) is false, then it will be knowably so by competent speakers. But this makes the felicity of the reply in (9) rather surprising. On the view which Mandelkern defends, successful performance of an action is sufficient for the agent to be able to perform the action. That is, he endorses SUCCESS:

> SUCCESS: If S Xes, then S is able to X.

SUCCESS explains why (9) is true: for all we know Ann might end up hitting the northern half of the field, so for all we know she might be able to do so. And although a claim such as (8a) is not known to be false, it is also not known to be true which explains why it is not assertible.

Mandelkern's observations are intriguing, but I don't think that they undermine a control constraint on ability, at least on the relevant sense of 'ability'. I suggest that although not every notion of agentive ability carries such a constraint, the ability concept which interacts with obligations and permissions (in other words, that which is relevant for the evaluation of OIC and MIC) does.[20] To motivate this, consider denials of ability. Suppose that shortly after noon you order Ann to hit the northern/southern half of the field. It would be perfectly acceptable for her to respond by saying 'I can't! (I don't know how)', and not be forced to hedge by saying 'I'm not sure if I can'. Moreover, it is generally taken as a data point that disjunctions such as (8c) are unacceptable (Kenny, 1976, Boylan, 2022). However, it is quite difficult to see why this should be if agentive modals could not express a control-requiring ability concept.[21]

Finally, even those who reject OIC maintain that some obligations are ability entailing (Graham, 2011, King, 2019). For instance, Graham (2011) maintains that so-called "positive duties"—duties 'demanding that certain things be done in certain situations' (377)—such as saving someone's life, are ability constrained.[22] This allows us to argue for a control constraint as follows. Suppose

---

[20] As far as I can tell, the claim that the ability concept undergirding OIC carries a control requirement is widely endorsed in the ethics literature. See, for example, (Portmore, 2019) for a detailed defence of it.

[21] Mandelkern does not discuss disjunctions such as (8c), or the fairly widespread judgment that they are infelicitous.

[22] See §5.1 for further discussion of Graham's account.

we add the following details to the *Target* case (as common knowledge): Bill is suffering from a deadly disease and is trapped in the northern half of the field; the rocket is carrying a cure for the disease. Consider the following claims:

(11) a. Ann has an obligation to/ought to/has to hit the northern half of the field.

 b. Ann has an obligation to/ought to/has to save Bill's life.

(12) a. Ann might have an obligation to hit the northern half of the field.

 b. Ann might have an obligation to save Bill's life.

(13) a. What is the chance that Ann has an obligation to hit the northern half of the field?

 b. What is the chance that Ann has an obligation to save Bill's life?

I submit that (11a)-(12b) are straightforwardly false, and the only reasonable response to (13a)-(13b) is '0'. A natural explanation of these facts is that Ann does not have an obligation to hit the northern half of the field/save Bill's life because she is unable to do these things, in the relevant sense of 'able'. Plausibly, this sense involves a notion of agentive control.[23,24]

Although proponents of CONTROL will find MIC particularly problematic, it is worth noting that even those who endorse SUCCESS should be reluctant to adopt MIC. For as we saw in §3.1, claims such as (4a)('Ann is permitted to save both Bill and Charlie') and (5a) ('You are permitted to cure cancer') are acceptable. But the relevant actions here are not just ones for which the agents lack requisite control, they are more deeply *unperformable*. For instance, it is common knowledge that there is no outcome whereby Ann ends up saving both Bill and Charlie. So, even those who go in for relatively "thin" conceptions of agentive ability, e.g. those sympathetic to SUCCESS, should think that we can know that Ann is not able to save both Bill and Charlie.

---

[23]Our argument here makes contact with the discussion in Schwarz (2020). Schwarz maintains that the ability concept relevant for the interpretation of obligations such as those expressed by (11a)/(11b) is fairly demanding, and requires something like agentive control. His central argument for this claim rests on the principle that if S knew that they were able to X and knew that Xing would prevent some tragedy (e.g. Bill's death) at little to no cost to themselves, then S has an obligation to X (5). However, as mentioned above, approaches such as Mandelkern's do not predict that a claim such as (8a) is knowable by Ann. Still, if the ability concept relevant for assessing Ann's obligation towards Bill did not carry a control constraint, then one would expect modalized claims such as (12a) and (12b) to be felicitous; and one would expect answers to (13a) and (13b) of the form 'x', where $x > 0$, to be acceptable.

[24]Some might be concerned that if the ability concept which undergirds (at least some) obligations/permissions requires a measure of agentive control, then this will undermine DUALITY. For instance, suppose you are able to push one of two buttons: A or B. You know that if you push button A, there is a 50% chance that ten people will be saved, and a 50% chance that the ten people will be killed. You also know that if you push button B, the ten people will be killed. Assume that 'You ought to save the ten people' is false because you are not able to save the ten, in the relevant sense of 'able'. Given DUALITY, this implies that 'You are permitted to fail to save the ten people' is true. However, the latter is infelicitous when uttered out-of-the-blue. In response, I maintain that 'You are permitted to fail to save the ten people' is misleading but nevertheless true. It is misleading because it suggests that you are permitted to push button B. But this suggestion can be canceled: it is perfectly acceptable to assert 'You are permitted to fail to save the ten people, but you are not permitted to push button B'. Thanks to Ben Holguín for helpful discussion here.

3.2.4 Presupposition

Some might be tempted to respond to the arguments in §§3.2.1-3.2.2 by appealing to the phenomenon of presupposition. A popular approach in philosophical semantics sees presuppositions as preconditions for truth and falsity.[25] The thought is that MIC should be construed in terms of presupposition: 'S may X' is neither true nor false—it is "undefined"—if S isn't able to X. With presuppositions in the picture, there are choice points as to how to define the notion of validity. On one popular conception—so-called "Strawson validity" (see §5.2)—it can be shown that the arguments of §§3.2.1-3.2.2 are rendered harmless. For instance, we will only be able to derive $AX \vee A\neg X$ from $OX \vee \neg OX$ if we assume $A\neg X$, which is uncontroversial.

However, there are straightforward problems with this approach. For one thing, we observed that claims such as (4a)('Ann is permitted to save both Bill and Charlie') and (5a) ('You are permitted to cure cancer') are felicitous. But undefined sentences are usually assumed to be unassertable. For example, (14) presupposes that there exists a unique King of France:

(14)   The King of France is happy.

Given that this presupposition isn't satisfied, (14) is unacceptable.[26]

For another, an important feature of presuppositions is that they *project*: even when presupposition-carrying sentences are embedded, the conditions they impose can escape these embedded environments (Chierchia and Mcconnell-Ginet, 2000). For example, (15a)-(15d), just as much as (14), imply that there exists a unique King of France:

(15)   a.   The King of France isn't happy.
       b.   Is the King of France happy?
       c.   If the King of France is happy, then his subjects will be happy too.
       d.   Bill believes that the King of France happy.

By contrast, none of (16a)-(16d) imply that Ann is able to take an additional class (the class could conflict with her job schedule):[27]

(16)   a.   Ann isn't permitted to take an additional class.

---

[25]See, for example, Strawson (1950), Heim (1983). Note, however, that the criticisms below do not hang on adopting a three-valued logic in order to model presuppositions. For instance, our arguments would also go through on pragmatic approaches, for example those defended by Stalnaker (1974), Schlenker (2009).

[26]von Fintel (2008) observes that an example such as 'I didn't have breakfast with the King of France this morning' is felicitous, even though standard theories of presupposition predict that this sentence also presupposes that there exists a unique King of France (see below for further discussion of the projection properties of presuppositions). However, negation appears to be playing a special role here, since 'I had breakfast with the King of France this morning' is clearly unacceptable. Importantly, examples such as (4a) and (5a) do not feature negation.

[27]Our observations here make contact with arguments given by King (2017). King considers a view on which 'ought' claims presuppose ability, but maintains that ability fails to project from obligation, undermining the proposal.

b. Is Ann permitted to take an additional class?

c. If Ann is permitted to take an additional class, then all of the other students will be as well.

d. Bill believes that Ann is permitted to take an additional class.

Overall, projection tests militate against assimilating ability to a presupposition of permission claims.[28] It is worth emphasizing that there is no problem with building presuppositions into the meaning of deontic modals per se. Indeed, my positive approach developed in §5.2 makes crucial use of definedness conditions on permission claims. But as we will see, these are more sophisticated than a simple ability presupposition.

4 Yes

The second answer to **Permission** which we will focus on is also straightforward: always 'yes'. That is, it consists of the following principle:

INABILITY IMPLIES PERMISSION: If S is unable to X, then S is permitted to X.

$\neg AX \rightarrow MX$

Given DUALITY, by contraposition this is equivalent to the principle "ought implies can not":

OUGHT IMPLIES CAN NOT (OICN): If S ought to X, then S is able to fail to X.

$OX \rightarrow A\neg X$

OICN says that whenever S is obligated to do something, they are able to do otherwise.

---

[28]It is also helpful to consider a related diagnostic. It is usually infelicitous to express uncertainty as to whether a given claim holds, but then use an expression that carries that very same claim as a presupposition:

(17) a. # I'm not sure whether there exists a King of France, but I doubt that the King of France is happy.

b. # I'm not sure whether there exists a King of France, but if the King of France is happy, then his subjects will be happy too.

By contrast, ability doesn't pattern this way; the examples below are perfectly acceptable:

(18) a. I'm not sure whether Ann has time to take an additional class, but I doubt that she is permitted to.

b. I'm not sure whether Ann has time to take an additional class, but if she is permitted to, then all of the other students will be as well.

Unlike MIC, OICN has been discussed a fair bit in the literature. It has its supporters, but is also rejected by several theorists.[29] Most arguments against OICN are case-based, and take it to be obvious that subjects can be prohibited from doing things that they are unable to do. For instance, some have argued that a claim such as (19a) is straightforwardly true and (19b) is straightforwardly false given DUALITY, even if you are unable to commit genocide, lacking the means or the psychological profile (Elzein, 2020, 39).[30]

(19)   a. You are obligated to not commit genocide.
       b. You are permitted to commit genocide.

I'm sympathetic towards this criticism, but here I want to draw out some of the structural features of OICN. More precisely, I consider how this principle interacts with the logic of obligation.[31]

The main result is that OICN is incompatible with the following highly plausible principle:

> WEAKENING: If S ought to X and S ought to Y, then S ought to $(X$ or $Y)$.
>
> $(OX \mathbin{\&} OY) \to O(X \vee Y)$

Indeed, even those who are willing to give 'ought' a non-normal modal logic go to great lengths to validate WEAKENING.[32]

The key observation is that OICN and WEAKENING imply COMPULSION:[33]

---

[29] Proponents of OICN include Haji (2002), Zimmerman (2008), Lockie (2018); critics include Hughes (2018), Elzein (2020), Spafford (2020). The debate usually concerns the principle "ought not implies can", but this is equivalent to OICN granted plausible assumptions, e.g. that oughts are closed under logical equivalence.

[30] Elzein (2020, 41) also puts forward a more sophisticated argument from attitude embedding that is worth mentioning. She claims that one can be certain that it would be impermissible for me to "leave the house with a kitchen knife and stab to death the first person I see". However, one might also be unsure whether I am able to perform this action—perhaps this would be impossible for me owing to the strength and speed of the first person I see. But this pattern in attitudes puts pressure on OICN given that it is irrational to assign higher credence to p than q when it is known that p entails q.

[31] That said, I will register a concern with OICN which comes from cases such as the following. You are able to push one of two buttons: A or B. You know that if you push button A, ten people will be saved. You also know that if you push button B, there is a 50% chance that the ten people will be saved, and a 50% chance that the ten people will be killed. Intuitively, 'You ought to save the ten people' is true here. However, assuming that ability requires a measure of agentive control, you are unable to fail to save the ten (after all, if you push button B, there is only a 50% chance that you will fail to save the ten). Thus, given OICN, 'You ought to save the ten people' cannot be true.

[32] For example Cariani (2013). Also see Cariani (2016) who criticizes a range of analyses of 'ought' on the grounds that they do not validate WEAKENING.

[33] Here is the argument:

(1)	$\neg A \neg (X \vee Y)$	PREMISE
(2)	$\neg O(X \vee Y)$	1, OICN
(3)	$\neg OX \vee \neg OY$	2, WEAKENING
(4)	$\neg A \neg (X \vee Y) \to (\neg OX \vee \neg OY)$	1, 3

COMPULSION: If S is compelled to $(X$ or $Y)$, then either S does not have to X or S does not have to Y.

$$\neg A\neg(X \vee Y) \to (\neg OX \vee \neg OY)$$

But COMPULSION is undermined by cases such as the following:

> *Charity*: The town where Charlie lives contains two charitable clubs: Club A provides food for people in need, and Club B cares for rescued animals. When Charlie turns 13, she can choose to join Club A, Club B, or both clubs. If she makes no choice, then she will automatically be admitted to either Club A or Club B, at random (this happens to all 13-year-olds in the town). Joining both clubs would be best.

(20) a. Charlie is compelled to join at least one of the following: Club A or Club B.

  b. Charlie ought to join Club A.

  c. Charlie ought to join Club B.

Charlie will become a member of at least one of the clubs no matter what she does, i.e. (20a) is true.[34] Moreover, I submit that any adequate theory of 'ought' should predict that both (20b) and (20c) are true in context. After all, it is widely maintained that one of the primary theoretical roles of the notion of obligation is that it is central to the practice of giving and receiving (moral) advice.[35] And we would not hesitate to advise Charlie to both join Club A and join Club B. But the pattern in (20) clearly contravenes COMPULSION. This should give those who are tempted to endorse INABILITY IMPLIES PERMISSION as an answer to **Permission** pause before accepting this principle.

5 Sometimes

We have seen that straightforward answers to **Permission**—namely always 'yes' and always 'no'—have problematic consequences. Since we should avoid these consequences, we shouldn't give simple answers to this question. We must answer '(only) sometimes'.

Since INABILITY IMPLIES NO PERMISSION fails, there are actions X such that S is unable to perform X, and S is permitted to X. And since INABILITY IMPLIES PERMISSION fails, there are actions Y such that S is unable to perform Y, and S is not permitted to Y. Putting things together, there are actions X and Y such that (i) S is unable to perform either, and (ii) S is permitted to X, but S is not permitted to Y. Let us call such pairs of actions *inability-permission pairs* for S. It will be helpful to have some concrete inability-permission pairs on the table. One plausible candidate comes from the *Cruise* case discussed earlier:

---

[34]I express 'S X or Y' as 'S does one of the following: X or Y' in order to ensure that disjunction takes narrow scope with respect to the embedding operator.

[35]See, for example, Schroeder (2011) among many others.

(21)   a. Ann is permitted to save both Bill and Charlie.

        b. Ann is permitted to blow up the cruise ship and kill everyone on board.

(21a) is intuitively true, but (21b) is intuitively false, even though Ann is unable to either save both Bill and Charlie or (we may assume) blow up the ship. Another candidate consists of (22a) and (22b), both of which were briefly considered earlier:

(22)   a. You are permitted to cure cancer.

        b. You are permitted to commit genocide.

We noted that (22a) is intuitively true, but (22b) is intuitively false, even though you are unable to either cure cancer or commit genocide.

We have arrived at an answer to **Permission**. But this answer raises further questions. For it is natural to be puzzled as to how the actions which form inability-permission pairs can differ in terms of their normative status. For example, how can *curing cancer* have positive moral valence for you but *committing genocide* have negative moral valence for you, even though you cannot perform either action? More generally, can we say something systematic about what distinguishes actions which form inability-permission pairs? Let us call this question **Difference**:

> **Difference**: Where X and Y are actions which form an inability-permission pair for S, what accounts for the normative difference between X and Y, i.e. why is S is permitted to X, but not permitted to Y?

In §5.2 I show that **Difference** can be given a satisfying answer. In the course of providing this answer, I will develop a novel account of permission.

5.1   Ability and the moral fabric

Before attempting to answer **Difference**, it will be instructive to briefly contrast the status of **Permission** with that of **Obligation** (repeated below):

> **Obligation**: If S is unable to X, is S obligated to X?

Those who reject always 'no' as an answer to **Obligation** reject OIC. For instance, they say that the following claims are true, even though the relevant agents are unable to act otherwise:

(23)   a. [George, an evil neurosurgeon, can manipulate Jones's brain in such a way that will ensure that Jones kills Smith. George observes Jones and will only intervene if Jones is about to decide not to kill Smith (Frankfurt, 1969).]
Jones ought to refrain from killing Smith.

    b. [Kerry is a severe kleptomaniac and cannot stop herself from stealing whenever she is given the opportunity. Dave has left his wallet on his desk.]
    Kerry ought to refrain from stealing the wallet.

But they also reject always 'yes' as an answer. For instance, they do not maintain that the claims below are true (repeated from §1):

(1)   a. I ought to snap my fingers and thereby end all suffering in the universe (Graham, 2011).
       b. You ought to travel back in time and prevent the crusades (Streumer, 2007).

That is, those who reject OIC answer '(only) sometimes' to **Obligation**: there are actions X (e.g. killing Smith) and Y (e.g. preventing the crusades) such that (i) S is unable to perform either, and (ii) S is obligated to X, but S is not obligated to Y. This raises a puzzle analogous to **Difference** from above: what accounts for the normative difference between the actions X and Y such that the former is obligatory but the latter is not? Why are facts about obligation configured like this?

Let us focus on one of the more developed answers to this question that I am aware of, namely that given by Graham (2011, 372-379). On this worldview, morality is ultimately comprised of a collection of fundamental principles, or rules, whose content consists of general obligations, e.g. *you ought to eliminate suffering*, *you ought not harm others*, etc. Particular moral obligations are derived from these principles given the situations that agents find themselves in. Crucially, some fundamental principles are ability-entailing and others are not. More specifically, principles demanding that things be done in certain situations—so-called "positive principles"—are ability-entailing; while principles forbidding that things be done in certain situations—so-called "negative principles"—are not ability-entailing. Thus, obligations which are derived from negative principles hold regardless of whether the agent is able to perform the relevant action. By contrast, obligations which are derived from positive principles only hold if the agent is able to perform the relevant action. This explains the difference between the examples in (23) and (1): (23a) expresses a negative duty, and therefore holds regardless of whether Jones is able to refrain from killing Smith; while the obligation expressed by (1b) holds only if you are able to travel back in time and prevent the crusades (which you are obviously not able to do).[36]

This is not the place to assess the merits of Graham's proposal, or the general moral framework which it assumes. Instead, I just want to point out that this picture of the structure of morality does not obviously explain the sorts of contrasts which we are interested in, for instance that between (22a) ('You are permitted to cure cancer') and (22b) ('You are permitted to commit genocide').

---

[36]King (2019, 68-76) provides a similar explanation for the contrast between (23) and (1). The key distinction she proposes is that between "natural obligations"—which are ability entailing—and "non-natural obligations"—which are not. However, as King herself recognizes, the natural/non-natural distinction is not altogether very sharp.

More precisely: this framework explains why (22b) should be false, but it does not account for why (22a) should be true. Presumably, 'You ought to refrain from committing genocide' expresses a negative duty, and therefore (22b) is false given DUALITY. However, *you are obligated to cure cancer* is presumably a positive duty, and therefore you only have this obligation if you are able to cure cancer. Since you aren't able to do this, you don't have the obligation. Given DUALITY, this means that 'You are permitted to not cure cancer' is true. But this doesn't imply that (22a) should be true. As far as I can see, it would be consistent with this general moral framework that you are obligated to *refrain* from curing cancer. In short: selectively building ability constraints into the fabric of obligation doesn't account for the heterogeneity of permission which follows from a 'sometimes' answer to **Permission**. A different style of explanation is needed.

5.2   Counterfactual permission

A fairly intuitive thought is that the answer to **Difference** is tied to what agents do in counterfactual situations. For instance, if you were to cure cancer, then this would be a good outcome; and if you were to commit genocide, then this would be a (very) bad outcome. These counterfactual scenarios appear to be relevant for determining what your *actual* permissions are.

There are several ways of making the connection between permission and counterfactuality explicit. I will use some of the ideology of contemporary philosophical semantics, namely the popular framework for modelling modal concepts which originates with the work of Angelika Kratzer. As we will see, apart from allowing us to make the central ideas precise, this account provides a way of handling some fairly subtle data points around permission claims. On this analysis of deontic concepts, moral notions are evaluated relative to a background domain of possibilities. I propose that the connection between permission and counterfactuality should be articulated in terms of requirements on the background domain. To put the idea roughly: it only makes sense to ask whether S is permitted to X if the most counterfactually similar possibilities where S Xes are included in the domain. In what follows, I spell this out more concretely.

On the Kratzerian picture, modals operate as quantifiers over a restricted domain of possibilities.[37] This domain is fixed by two contextually determined parameters: (i) the *modal base* $\mathcal{B}$, which is the broad set of possibilities relevant for evaluating the modal, and (ii) the *ordering source*, which provides an ordering $\geq$ of the possibilities in the modal base.[38] Where M is a modal, and Q_{M} expresses its associated quantificational force, the general idea is that 'Mϕ' is true just in case ϕ is true at Q_{M} of the top-ranked worlds in the modal base. To make this more explicit, theorists make use of a function BEST$(\cdot, \cdot)$ that takes a modal base and an ordering and yields the set of top-ranked worlds as deter

---

[37]See, for example, (Kratzer, 1977, 1981, 1991, 2012).

[38]Strictly speaking, both the modal base and the ordering source are relativized to worlds, but I will generally suppress this here. Also, the term "modal base" is often used ambiguously in the literature: sometimes it is used to mean the set of propositions that are intersected to determine the modal domain (or the function from worlds to such sets), while other times it is used to mean the modal domain itself. I use it in the second way here.

mined by the ordering. We can simplify by restricting our attention to scenarios where the modal base is finite. In this case, BEST can be presented as follows:

Specification of top-ranked worlds

For any modal base $\mathcal{B}$, and ordering $\geq$:
$\text{BEST}(\mathcal{B}, \geq) = \{w' \in \mathcal{B} \mid \neg \exists w'' \in \mathcal{B} \text{ such that } w'' \geq w' \text{ and } w' \not\geq w''\}.$

The schematic entries for 'must' and 'may' are below (expressions are evaluated relative to a triple consisting of a world w, modal base $\mathcal{B}$, and ordering $\geq$):

Standard analysis of 'must'

'Must ϕ' is true relative to $\langle w, \mathcal{B}, \geq \rangle$ iff $\forall w' \in \text{BEST}(\mathcal{B}, \geq)$: ϕ is true at w'.

Standard analysis of 'may'

'May ϕ' is true relative to $\langle w, \mathcal{B}, \geq \rangle$ iff $\exists w' \in \text{BEST}(\mathcal{B}, \geq)$: ϕ is true at w'.

On this view, different modal flavors (e.g. epistemic, deliberative, deontic, bouletic, etc.) correspond to different values of the two parameters $\mathcal{B}$ and $\geq$. For deontic modals, the ordering $\geq$ is determined by the dictates of morality, or moral goodness: $w \geq w'$ iff w is at least as good as w', morally speaking. To have something a bit more concrete to work with, I'll assume that the ordering is dictated by broadly utilitarian considerations, so that $w \geq w'$ iff at least as many lives are saved in w as in w'. As for $\mathcal{B}$, I contend that its composition is key to answering **Difference**.

I maintain that claims of permission place requirements on the modal base used in their evaluation. 'S may X' can only be assessed for truth or falsity if $\mathcal{B}$ contains worlds where S Xes. In fact, not just any worlds where S Xes, but the most *counterfactually similar* worlds where S Xes. Let us call this the *counterfactual condition* on claims of permission. To make this condition more explicit, we can appeal to a similarity function $\text{SIM}(\cdot, \cdot)$ which is widely employed in theories of counterfactual conditionals (Stalnaker, 1968, Lewis, 1973). This function maps worlds (w) and propositions (p) to propositions, and intuitively takes each proposition p to the set of worlds maximally similar to w in which p is true.[39] Then the entry for permission claims is as follows:

COUNTERFACTUAL PERMISSION

'S may X' is defined relative to $\langle w, \mathcal{B}, \geq \rangle$ only if $\text{SIM}(w, S\ Xes) \subseteq \mathcal{B}$.

If defined, 'S may X' is true relative to $\langle w, \mathcal{B}, \geq \rangle$ iff $\exists w' \in \text{BEST}(\mathcal{B}, \geq)$: S Xes at w'.

---

[39]This function is usually taken to obey the following constraints: (i) *Success*: $\text{SIM}(w, p) \subseteq p$; (ii) *Strong Centering*: $\text{SIM}(w, p) = \{w\}$, if $w \in p$; and (iii) *Uniformity*: If $\text{SIM}(w, p) \subseteq q$ and $\text{SIM}(w, q) \subseteq p$, then $\text{SIM}(w, p) = \text{SIM}(w, q)$. See Starr (2022) for a detailed discussion of these conditions, among others.

The counterfactual condition is implemented as a definedness constraint, or presupposition, of permission claims. As we will see, this is key to capturing the target phenomena.

First, we want to explain the contrast between (21a) and (21b) (repeated from above):

(21)   a. Ann is permitted to save both Bill and Charlie.
          b. Ann is permitted to blow up the cruise ship and kill everyone on board.

It is natural to assume that in most normal contexts, the modal base $\mathcal{B}$ against which these claims are evaluated does not include worlds where Ann does the impossible, e.g. save both Bill and Charlie or blow up the ship. This follows from standard characterizations of the modal base appropriate for deontic modals on which this parameter is taken to track a collection of "relevant facts" (Kratzer, 1977). At least by default, it is plausible that Ann lacking the ability to save both Bill and Charlie or blow up the ship will be contextually relevant. (I will return to the standard conception of the modal base, and its default status, in §5.4.)

Given that $\mathcal{B}$ does not contain worlds where Ann either saves both Bill and Charlie or blows up the ship, COUNTERFACTUAL PERMISSION predicts that both (21a) and (21b) should be undefined, or suffer from presupposition failure. However, it is well known that presuppositions can be *accommodated*. Roughly speaking, a presupposition is accommodated if it is not accepted by the conversational participants prior to the target utterance, but comes to be accepted 'quietly and without fuss' in the course of evaluating the utterance (von Fintel, 2008).[40] For instance, it is standardly assumed that possessive constructions, e.g. 'my sister', trigger existence presuppositions, so that 'My sister is a lawyer' presupposes that I have a sister. But it has been widely observed that this presupposition is easily accommodated: even if you have no idea that I have a sister (or even if you are quite sure that I have no siblings), I can felicitously assert things like 'I'm having lunch with my sister this afternoon'. In the course of evaluating my utterance, you come to accept that I have a sister.

I submit that the counterfactual condition on claims of permission is, in general, very easy to accommodate. In effect, this means that an assertion of (21a) relative to $\mathcal{B}$ triggers domain expansion. Thus, (21a) is not evaluated relative to $\mathcal{B}$, but rather the expanded domain $\mathcal{B}_{\text{both}}$ which is obtained by combining $\mathcal{B}$ with the set of most similar counterfactual possibilities where Ann saves both Bill and Charlie ($\mathcal{B}_{\text{both}} = \mathcal{B} \cup \text{SIM}(w, \textit{Ann saves Bill and Charlie})$). Since saving both Bill and Charlie is at least as (morally) good as anything that Ann is able to do, there are some worlds in $\text{BEST}(\mathcal{B}_{\text{both}}, \geq)$ where Ann saves both Bill and Charlie.[41] That is, (21a) is true as evaluated relative to $\mathcal{B}_{\text{both}}$. More generally, (21a) will be true, if defined, thus accounting for its felicity.

---

[40]Classic work on the phenomenon of accommodation includes Stalnaker (1970, 1973, 1974), Kartunnen (1974), Lewis (1979). von Fintel (2008) provides a helpful overview.

[41]We should not assume that saving both Bill and Charlie would be *better*—in the operative sense of "better"—than anything that Ann is able to do, since then $\text{BEST}(\mathcal{B}_{\text{both}}, \geq)$ would not contain any worlds where Ann saves exactly one person. However, claims such as 'Ann is not permitted to save exactly one person' are unacceptable, and they do not improve when uttered

Similarly, an assertion of (21b) relative to $\mathcal{B}$ triggers domain expansion to $\mathcal{B}_{\text{blow-up}}$ which is obtained by combining $\mathcal{B}$ with the set of most similar counterfactual possibilities where Ann blows up the ship ($\mathcal{B}_{\text{blow-up}} = \mathcal{B} \cup \text{SIM}(w, \textit{Ann blows up the ship})$). Since blowing up the ship is (morally) worse than anything that Ann is able to do, there are no worlds in $\text{BEST}(\mathcal{B}_{\text{blow-up}}, \geq)$ where Ann blows up the ship. That is, (21b) is false as evaluated relative to $\mathcal{B}_{\text{blow-up}}$. More generally, (21b) will be false, if defined, thus accounting for its unacceptability.

It is worth emphasizing the importance of similarity in the counterfactual condition. For instance, suppose that if Ann had another lifesaver available and saved both Bill and Charlie, then this would initiate a causal chain leading to a global catastrophe. In that case, (21a) sounds much worse. Conversely, suppose that if Ann blew up the ship, then this would initiate a causal chain leading to world peace. Then (21b) starts to sound better. This shows that permission is not just sensitive to counterfactual situations where agents perform the relevant actions, but it is sensitive to the closest, or most similar such situations.

Second, note that if the counterfactual condition was implemented as a straightforward entailment of permission claims—and not as a presupposition—then the robust infelicity of examples such as (4b) and (5b) from §3.1 would not be captured:

(4b) Ann is not permitted to save both Bill and Charlie.

(5b) You are not permitted to cure cancer.

For if the counterfactual condition was a regular entailment, then it could be targeted by negation. In that case, these examples would be true when evaluated relative to any modal base that is ability constrained, e.g. $\mathcal{B}$. By contrast, as noted in §3.2.4, presuppositions are known to project from embedded environments, in particular from under negation. Thus, on my account (4b) is also predicted to presuppose that the background modal base contains the most similar worlds where Ann saves exactly two lives, i.e. it is equivalent to $\mathcal{B}_{\text{both}}$. Assuming that this presupposition is accommodated, (4b) will be evaluated relative to $\mathcal{B}_{\text{both}}$. But this sentence is false relative to $\mathcal{B}_{\text{both}}$, accounting for its infelicity.

Third, we observe that although (24a) and (24b) are both acceptable in the *Cruise* case, their conjunction (24c) is infelicitous:

(24) a.   Ann has to save exactly one life.
      b.   Ann is permitted to save exactly two lives.
      c. # Ann has to save exactly one life, but she is permitted to save exactly two.

---

after (21a). Note that the "flattening" of intuitively morally distinct options is fairly common granted the possibility of supererogatory acts: I am permitted to donate some of my money to charity, and I am permitted to donate all of my money to charity, even though there is a clear sense in which the latter is better than the former. See Heyd (2024) for discussion.

On the face of it, this pattern is quite puzzling. One way of explaining it appeals to an independent fact about modal constructions: the parameters relative to which modals are interpreted cannot easily be shifted mid-sentence. For example, consider epistemic interpretations of 'might' and 'must':

(25) # Amy must be at work but she might be at home.

Presumably (25) is unacceptable because 'might' and 'must' are being evaluated relative to the same modal base; and relative to a single modal base both conjuncts cannot be true together.[42] Similarly, given COUNTERFACTUAL PERMISSION and the standard entry for deontic necessity modals, it is straightforward to verify that relative to a single modal base the conjuncts in (24c) cannot be true together.

A related, but more technical account of the unacceptability of (24c) can be given in terms of so-called "presupposition projection" algorithms.[43] Such algorithms calculate the presuppositions of complex sentences given the presuppositions of their parts. One algorithm is provided by the "Weak Kleene" truth-table which allows presuppositions to project through Boolean connectives. Given COUNTERFACTUAL PERMISSION, this account predicts that (24c) presupposes that the background modal base is equivalent to $\mathcal{B}_{\text{both}}$. Granted the standard entry for necessity modals, the first conjunct is false relative to $\mathcal{B}_{\text{both}}$, which means that the whole conjunction will be false, if defined.[44]

Fourth, the unacceptability of (24c) is connected to a desideratum concerning the logic of deontic concepts in light of COUNTERFACTUAL PERMISSION, in particular the status of DUALITY. In our discussion so far, I (implicitly) assumed a classical notion of validity which requires preservation of truth at a point of evaluation. But now that we have presuppositions as part of the meaning of permission claims, more sophisticated notions of consequence can be formulated. More specifically, consider the notion of *Strawson validity* (von Fintel, 1999). Essentially, an argument from a set of sentences Γ to a sentence ψ is Strawson valid just in case whenever all of the $\varphi \in \Gamma$ and ψ are defined, if all of the $\varphi \in \Gamma$ are true, then ψ must be true as well. A bit more explicitly:[45]

---

[42]The point generalizes across context-sensitive expressions: there is independent evidence that intersentential context shifts are often difficult to achieve. For instance, consider the behaviour of quantifiers, "weather predicates" such as 'raining', and gradable adjectives such as 'tall':

(26) a. # Everyone is asleep but someone is awake.
      b. # It's raining, but it's not raining.
      c. # Mark is tall but he is short.

All of these expressions have been argued to be context-sensitive, so that there exist contexts such that each conjunct is evaluated as true relative to that context (Stanley and Szabó, 2000, Kennedy, 2007). Yet each conjunction is unacceptable.

[43]The literature on presupposition projection is enormous. Important work in this area includes Karttunen (1973), Heim (1983), and more recently Schlenker (2009).

[44]I appeal to the Weak Kleene projection algorithm here mainly because it is fairly straightforward, but it is not essential to our arguments. Indeed, examples such as (24c) as well as those below can be shown to be problematic on all existing projection algorithms that I am aware of (given COUNTERFACTUAL PERMISSION).

[45]Strawson validity has been used to account for the behaviour of a range of natural lan-

Strawson Validity:

$\Gamma \underset{\overline{S}}{\models} \psi$ iff there is no $\langle w, \mathcal{B}, > \rangle$ such that (i) every $\varphi \in \Gamma$ and ψ are all defined at $\langle w, \mathcal{B}, > \rangle$; (ii) every $\varphi \in \Gamma$ is true at $\langle w, \mathcal{B}, > \rangle$; and (iii) ψ is false at $\langle w, \mathcal{B}, > \rangle$.

It is straightforward to verify that DUALITY is Strawson valid (that is: $\underset{\overline{S}}{\models} OX \leftrightarrow \neg M \neg X$), given COUNTERFACTUAL PERMISSION and the standard entry for obligation from above.[46] This ensures that examples such as (7a) and (7b) from §3.2.1 will be false, if defined, accounting for their infelicity:

(7)   a.   # Ann has to save exactly one person but she is permitted to not save exactly one.

    b.   # Ann is neither permitted to save exactly two people nor does she have to not save exactly two.

Fifth, I have proposed that the counterfactual condition is part of the meaning of permission claims. Should we also place this condition on claims of obligation? That is, should we adopt COUNTERFACTUAL OBLIGATION:

COUNTERFACTUAL OBLIGATION

'S is obligated to X' is defined relative to $\langle w, \mathcal{B}, \geq \rangle$ only if $\text{SIM}(w, S \; Xes) \subseteq \mathcal{B}$.

If defined, 'S is obligated to X' is true relative to $\langle w, \mathcal{B}, \geq \rangle$ iff $\forall w' \in \text{BEST}(\mathcal{B}, \geq)$: $S \; X$es at w'.

?

An important data point here involves claims such as the following:

(27)   Ann doesn't have to save exactly two people, but she does have to save exactly one.

Observe that (27) is perfectly felicitous in the *Cruise* case. However, it is predicted to be unacceptable given COUNTERFACTUAL OBLIGATION. For existing presupposition projection algorithms—e.g. Weak Kleene—predict that this sentence presupposes that the background modal base is equivalent to $\mathcal{B}_{\text{both}}$. This means that the second conjunct will be evaluated relative to $\mathcal{B}_{\text{both}}$. But again, this conjunct will be false relative to $\mathcal{B}_{\text{both}}$, which entails that the whole conjunction should be false. (Note that examples such as (4b) and (5b) provide independent evidence that the counterfactual condition projects from negation, and thus that (27) presupposes that the modal base is $\mathcal{B}_{\text{both}}$.) This suggests

---

guage constructions, for example negative polarity items (von Fintel, 1999, Condoravdi, 2010), conditionals (von Fintel, 2001, Cariani and Goldstein, 2018, Mandelkern, 2021), temporal adverbials (Csirmaz, 2008), and bouletics (Blumberg, 2022).

[46]More generally, one can show that the set of Strawson validities for our deontic language corresponds exactly to the set of theorems of a normal modal logic for the operators O and M. It is also straightforward to verify that the rules of inference which correspond to DUALITY are Strawson valid, i.e. $OX \underset{\overline{S}}{\dashv}\underset{\overline{S}}{\vdash} \neg M \neg X$.

an interesting asymmetry between permission and obligation with respect to counterfactuality: only permission, and not obligation, has the relevant counterfactual aspect (also see §5.3 for further discussion).

Finally, it is worth reflecting on the relationship between permission and ability given COUNTERFACTUAL PERMISSION. For proponents of CONTROL (the claim that ability carries a substantial control constraint—see §3.2), accommodating the counterfactual condition involves accommodating performance not ability. This is because the counterfactual condition only requires us to include in the modal base the most counterfactually similar worlds where the subject *performs* the relevant action; but given CONTROL these needn't be the most counterfactually similar worlds where the subject is *able* to perform the action (since at such worlds they may not possess the requisite control).[47] This means that, at least for those sympathetic to CONTROL, the relationship between permission and ability will be fairly indirect, and COUNTERFACTUAL PERMISSION should not be glossed along the lines of "permission implies counterfactual ability".

5.3 Comparisons

At this point, it will be helpful to compare and contrast our treatment of permission claims with some other phenomena that, at least at first blush, appear to be related to it.

First, consider an observation involving modal expressions from Lewis (1979, 354-355). Lewis notes that the background domain of possibilities used to evaluate natural language modals can shift as the conversation evolves. In particular, domain expansion is fairly easy to induce by making possibilities conversationally salient. Here is one of Lewis's cases:

> The commonsensical epistemologist says: "I *know* the cat is in the carton—there he is before my eyes—I just *can't* be wrong about that!" The sceptic replies: "You might be the victim of a deceiving demon". Thereby he brings into consideration possibilities hitherto ignored, else what he says would be false. The boundary [of the modal base] shifts outward so that what he says is true. Once the boundary is shifted, the commonsensical epistemologist must concede defeat. And yet he was not in any way wrong when he laid claim to infallible knowledge. What he said was true with respect to the [modal base] as it then was.

This raises the question: is the phenomenon exhibited by the examples in (21)-(22) just an instance of the more general phenomenon of domain expansion displayed in Lewis's case?

---

[47]There is some evidence that this is the right result. Suppose that you are an inexperienced darts player and you have been given the opportunity to throw a dart at a dart board. Suppose also that hitting the bullseye has significant positive moral value. Then 'You are permitted to hit the bullseye' is easily heard as true. But intuitively, worlds where you become an expert darts player and come to possess control over the flight of the dart aren't relevant for assessing this claim. Instead, the only situations that appear to be relevant are ones where your skill at darts remains at its actual level, but you luckily hit the bullseye.

Crucially, there are many ways to make possibilities conversationally salient in a way that triggers the sort of expansion that Lewis is drawing attention to in the quoted passage. For instance, instead of employing a possibility modal, the sceptic could have uttered any one of the following:

(28)   a. If you were the victim of a deceiving demon, everything would look exactly as it does now.
       b. Victims of deceiving demons cannot tell the difference between appearance and reality.
       c. But holographic cats don't go in cartons.[48]

These would also have undermined the epistemologist's commonsense knowledge claim—the epistemologist could not happily make their assertion on the heels of any of these examples. It follows that if the sort of modal base expansion triggered by examples such as those in (21)-(22) was of a piece with that exhibited in Lewis's case, then we would expect that it could also be induced by other sorts of constructions. For instance, presumably the conditionals below raise to salience the possibility that Ann saves both Bill and Charlie, and the possibility that you cure cancer, respectively:

(29)   a. If Ann was able to save both Bill and Charlie, then they would forever be in her debt.
       b. If you were able to cure cancer, then that would be the best thing you could do.

Thus, we would expect that everyday obligation claims such as 'You have to visit your grandmother this afternoon' and 'Ann has to save exactly one life' would be undermined by the conditionals in (29). In fact, this is not what we find: everyday obligation claims remain robustly assertable in the face of these conditionals. Indeed, contrast (30a) with (30b)-(30c):

(30)   a.   I don't know that I have hands, since if I was the victim of a deceiving demon, everything would look exactly as it does now.
       b.   # You are not required to visit your grandmother this afternoon, since if you were able to cure cancer instead, that would be the best thing you could do.
       c.   # Ann doesn't have to save exactly one life, since if she was able to save both Bill and Charlie, then that would be the best thing she could do.

Although (30a) is felicitous, (30b)-(30c) sound bizarre. Merely raising possibilities to salience isn't enough to trigger the sort of domain expansion associated with the counterfactual condition. Thus, I conclude that the phenomenon of domain expansion exhibited by Lewis's case is distinct from the phenomenon

---

[48]This example is inspired by Moss (2012), who points out that the phenomenon of domain expansion via salience is quite general, and can be triggered by a wide variety of constructions.

captured in COUNTERFACTUAL PERMISSION. (This is not to say that the counterfactual condition is necessarily a *sui generis* semantic effect—see §5.4 for further discussion.)

Second, recall the distinction between "weak" necessity modals such as 'ought', and "strong" necessity modals such as 'must' from §2. In recent work, Silk (2021) argues that this distinction should be spelled out in counterfactual terms. One source of motivation for this is the judgment that forms such as '(Must ϕ) $\wedge\neg\phi$' are infelicitous, but forms such as '(Ought ϕ) $\wedge\neg\phi$' are not:[49]

(31)   a.   Bill ought to come tomorrow, but he won't.
       b.   # Bill must come tomorrow, but he won't (Palmer, 1990, 123).

(32)   a.   You ought to help your mother, but you won't.
       b.   # You must help your mother, but you won't (Silk, 2021, 225).

To capture these contrasts, Silk maintains that for each world w in the *context set*—the information representing what is common knowledge in the conversational context—the modal base for strong necessity modals at w must be included in the context set (226). (31b) is unacceptable because the speaker is trying to place incompatible constraints on the context set: the first conjunct implies that there will always be worlds in the context set where Bill comes tomorrow, but the second conjunct implies that there are no such worlds. By contrast, where w is a world in the context set, the modal base for weak necessity modals at w needn't be included in the context set. Silk claims that the modal base for weak necessity modals can even be disjoint from the context set. This explains why (31a) is acceptable. In order to capture this behaviour of weak necessity modals, Silk builds a counterfactual element into their meaning (219-222).[50]

The foregoing exposition of Silk's account is fairly sketchy, but we don't need to get into further details to see that the counterfactual feature he proposes for weak necessity modals is distinct from the counterfactual condition. For in the literature on the weak/strong distinction, permission modals are taken to be the duals of *strong* necessity modals, not weak necessity modals.[51] This can be motivated by considering an example from §2:

(33)   [You are on a flight and the seatbelt light has just been turned off.]
       a.   You ought to keep your seat-belt fastened, but you don't have to.
       b.   You are permitted to unfasten your seat-belt.

As noted earlier, (33a) is acceptable, illustrating the distinction between weak and strong necessity modals. But (33b) is also acceptable, which is incompatible

---

[49]This observation is discussed by Palmer (1990), Ninan (2005), Willer (2014), Charlow (2015) among others. Further motivation for the connection between 'ought' and counterfactuality is provided by the observation that in many languages, the meaning of English 'ought' is expressed by combining a strong necessity modal (e.g. the equivalent of 'must' or 'have to') with counterfactual morphology (von Fintel and Iatridou, 2008, 2023).

[50]Silk sketches several ways of incorporating a counterfactual aspect into the meaning of 'ought', but doesn't endorse any one in particular.

[51]See, for example, von Fintel and Iatridou (2008, 2023) and the discussion of Silk (2021) below.

with permission being the dual of weak necessity. This means that even if we accept that weak necessity modals have a counterfactual aspect, this will not automatically carry over to permission claims.[52] Indeed, this is reflected in Silk's entry for permission modals: as with strong necessity modals, this entry also implies that for each world w in the context set, the modal base for permission modals at w must be included in the context set (219). This predicts that claims such as (21a) ('Ann is permitted to save both Bill and Charlie') and (22a) ('You are permitted to cure cancer') cannot be true. So, even if we chose to spell out the weak/strong distinction along the lines that Silk suggests, this wouldn't provide an adequate answer to **Permission**, or yield a suitable model of the data considered in this paper.

Finally, a popular analysis of conditionals is the so-called "restrictor theory" (Kratzer, 1981, 1986). On this account, the primary function of the antecedent of a conditional is to restrict modal bases of modals which appear in the consequent of the conditional. For instance, in (34) the proposition expressed by the antecedent is taken to restrict the modal base for the epistemic modal in the consequent:

(34) If Bill is at home, then Ann must be with him.

The resulting meaning is roughly equivalent to the claim that in all the possibilities consistent with our knowledge where Bill is at home, Ann is with him.[53]

However, the restrictor analysis faces a challenge when it comes to conditionals which embed deontic modals. It predicts that forms such as 'If p, then ought p' should be trivially true:

(35) If Bill evaded his taxes, then he ought to evade his taxes.

After all, if every world in the restricted modal base for 'ought' in the consequent is one where Bill evades his taxes, then it follows that all of the most highly ranked worlds in this restricted modal base will be ones where Bill evades his taxes. But (35) is not trivially true; indeed, it will be false in many contexts.[54]

In response, some have argued that deontic modals carry a *diversity presupposition*: 'S ought X' presupposes that the background modal base $\mathcal{B}$ is *diverse* with respect to the proposition S Xes, i.e. there are worlds in $\mathcal{B}$ where S Xes, and there are worlds in $\mathcal{B}$ where S does not X (Frank, 1997, Condoravdi, 2001). Given a diversity presupposition, examples such as (35) are not trivially true. For even if this presupposition is accommodated, it is plausible that worlds where Bill pays his taxes will be more highly ranked than worlds where he evades them. It is arguable that a diversity presupposition on permission modals can capture

---

[52]It is also worth noting that Silk only allows the modal base for weak necessity modals to reach outside of the context set; but this doesn't imply that the modal base can include worlds where subjects perform actions which are actually unperformable, e.g. curing cancer.

[53]It is natural to wonder how conditionals which feature no overt modals in their consequents should be interpreted on the restrictor theory. In this case, proponents of the analysis maintain that the consequent is embedded under a covert epistemic 'must' which gets restricted by the antecedent. See Rothschild (2011) for discussion.

[54]For discussion, see Frank (1997), Zvolenszky (2002, 2006), Kratzer (2012), Willer (2014), Stojnić (2024).

the contrasts in (21)-(22).[55] So, should we understand the counterfactual condition on permission as a consequence of the more general diversity presupposition on deontics?

There are at least two reasons why not. For one thing, positing a diversity condition on deontics is problematic even when we restrict ourselves to the phenomenon that is supposed to motivate it, namely the interaction between modals and conditionals.[56] For as Zvolenszky (2002, 351-352) points out, some conditionals of the form 'If p, then ought p' are acceptable. For instance, suppose that the Dalai Lama is an exceptionally even-tempered person, and only gets angry if he has good reason to be angry. Then (36) is intuitively true:

(36)   If the Dalai Lama is angry, then he ought to be angry.

However, we can suppose that (the closest) situations where the Dalai Lama is not angry are better than (the closest) situations where he is angry. Thus, (36) would not be true if deontics carried a diversity presupposition.

Second, a general diversity presupposition suffers from similar problems to those raised above for COUNTERFACTUAL OBLIGATION: it predicts that examples such as (27) ('Ann doesn't have to save exactly two people, but she does have to save exactly one') should be unacceptable. For given a diversity presupposition, standard projection algorithms predict that this sentence presupposes that the background modal base includes (the closest) worlds where Ann saves exactly two people, and includes (the closest) worlds where she does not save exactly two people. But the second conjunct will be false relative to this modal base. And note that it is no good to respond by saying that worlds where Ann saves both Bill and Charlie are ranked below those where she saves exactly one of them, since this would predict that permission claims such as (4a) ('Ann is permitted to save both Bill and Charlie') won't be true, contrary to observation.[57]

---

[55] Although diversity isn't canonically framed in terms of similarity or closeness, it is straightforward to implement the condition this way. For instance, see Willer (2014).

[56] The restrictor theory of conditionals should not be taken for granted. In particular, there is an important tradition in philosophy which treats conditionals as devices for referring to particular worlds (Stalnaker, 1968, 1980, Schlenker, 2004, Schulz, 2014, 2017, Bacon, 2015, Mandelkern, 2018, Santorio, 2022, Schultheis, forthcoming). It has been argued that this account is better able to capture various phenomena involving conditionals than the restrictor theory—see Mandelkern (2018) for discussion.

[57] Both Willer (2014) and Stojnić (2024) develop responses to the 'If p, then ought p' problem on which the modal base relevant for the assessment of 'ought' claims includes closest worlds where the prejacent holds and closest worlds where the prejacent fails to hold, but they do not capture this condition as a presupposition. I lack the space for a detailed discussion of these accounts, but I will allow myself some brief remarks. For one thing, Willer appears to maintain that only weak necessity modals carry this condition (13). Thus, like Silk's theory considered above, on Willer's account permission claims are not predicted to have a counterfactual aspect to their meaning. Putting this to one side, given the way that Willer implements the condition, the resulting logic for obligation deviates significantly from that of a normal modal logic. In particular, given Willer's analysis it is straightforward to construct counterexamples to plausible principles such as WEAKENING (see §4). Finally, Willer handles examples such as (36) by maintaining that they involve uncertainty over which worlds are normatively ideal, and develops an account on which factual information can resolve this uncertainty (17-19). But given DUALITY, this account predicts that whenever one is uncertain as to whether one is obligated to X, one is permitted to X, which appears to be controversial.

As for Stojnić, the way she implements the condition also makes the logic for obligation

28

5.4 Interpretational defaults

I'll close this section by returning to the question **Difference** in light of COUN-
TERFACTUAL PERMISSION:

> **Difference**: Where X and Y are actions which form an inability-
> permission pair for S, what accounts for the normative difference
> between X and Y, i.e. why is S is permitted to X, but not permitted
> to Y?

As we have seen, the answer provided by COUNTERFACTUAL PERMISSION is
relativized to the background domain of possibilities $\mathcal{B}$. However, some might be
unsatisfied with this, on the grounds that it is insufficiently predictive. Without
constraints on $\mathcal{B}$, it might seem like almost any two distinct actions could form
an inability-permission pair.

We can go some way towards allaying this concern by considering the stan-
dard restrictions imposed on the modal base $\mathcal{B}$ for the interpretation of deontic
modals. A popular view is that $\mathcal{B}$ is a "circumstantial" modal base. What this
means is that $\mathcal{B}$ at a world w is determined by a set of "relevant facts" at w:
every world in $\mathcal{B}$ is one where each of these facts holds (Kratzer, 1991, 646). As
suggested in §5.2, it is plausible that facts pertaining to an agent's abilities at
w will often be relevant. After all, there is a reason why many theorists have
been attracted to the view that ability constrains the space of actions eligible
for deontic status, and endorsed principles such as OIC. We can think of ability
facts as something like a *default constraint* on $\mathcal{B}$: by default, speakers and hearer
assume that $\mathcal{B}$ is ability constrained.[58] This provides us with a more robust an-
swer to **Difference**: in many cases, where X and Y form an inability-permission
pair for S, S is permitted to X because Xing is at least as good as anything
that S is able to do, and S is not permitted to Y because Ying is worse than
anything that S is able to do.

Interestingly, the idea that ability forms something like an interpretational de-
fault on permission claims has analogues in other areas of philosophical seman-
tics, namely recent work on bouletic attitudes. A common view in philosophy
and linguistics is that what an agent desires is intimately connected to what
they believe, or more precisely: what they believe to be possible. This has been
justified on both intuitive and theoretical grounds.[59] For instance, suppose I am
at my local sandwich shop. Even if it is made salient that lobster is my favorite
food, I can perfectly well utter (37):

(37)   I want the ham and cheese.

---

very weak, and one can construct counterexamples to WEAKENING on her account (1717-
1718). Moreover, Stojnić maintains (i) that the modal base prior to counterfactual expansion
is always ability constrained, and (ii) that a permission claim 'S may X' presupposes that there
are worlds in the prior modal base where S Xes (1726-1727). This predicts that examples such
as (4a) ('Ann is permitted to save both Bill and Charlie') cannot be true.

[58]See (Sinnott-Armstrong, 1984) for a similar claim.

[59]See Heim (1992), von Fintel (1999), Levinson (2003) among many others.

Theorists have concluded that what a subject desires is bounded by their beliefs. In terms of the framework for modality introduced earlier, this is captured by maintaining (i) that 'want' is a necessity modal; and (ii) the modal base for desire ascriptions is the subject's "belief set": the set of worlds compatible with everything the subject believes. Then (37) is predicted to be true: the modal base for 'want' can only access doxastic possibilities, and the top-ranked doxastic possibilities are ones where I get the ham and cheese. Although I prefer lobster to a ham and cheese sandwich, I don't believe it's possible that I will get lobster.

However, it has been observed that such belief constraints do not sit well with the deployment of desire ascriptions in general. In some cases it seems perfectly fine to ascribe agents desires for things which the agent believes—or even knows—they will not get:

(38) a. I want this weekend to last forever (but of course I know it will be over in a few hours) (Heim, 1992, 199).
      b. Al wants to be promoted (but believes he won't be) [(Grano and Phillips-Brown, 2022) inspired by (Portner and Rubinstein, 2012)].
      c. I want lobster.

Indeed, note that (38c) can easily be heard as acceptable, even when I know there is no hope of getting lobster.

These data points set up a puzzle: why does desire seem to be intimately connected to belief—as evidenced by (37)—on the one hand, but also disconnected from it—as evidenced by, for example, (38c)—on the other? A recent answer has been given by Grano and Phillips-Brown (2022). On this approach, when p is compatible with S's beliefs, then 'S wants p' is evaluated relative to S's belief set $\mathcal{B}$. Given that I believe that it is possible that I get the ham and cheese, this explains why (37) is true. However, when p is incompatible with S's beliefs, domain expansion is triggered and 'S wants p' is evaluated relative to the set $\mathcal{B}^+$, which is obtained by combining $\mathcal{B}$ with the closest possibilities (to S's belief worlds) where p holds.[60] Since the closest possibilities where I get lobster are better (by my lights) than any of the possibilities compatible with my beliefs, (38c) is predicted to be true. One way of conceiving of Grano and Phillips-Brown's account is that the agent's beliefs form an interpretational default for the evaluation of desire ascriptions. Although the modal base for desire ascriptions isn't always the agent's belief set, this set has a default status.[61]

6 Conclusion

I focused on the question **Permission**: if S is unable to X, is S permitted to X? I considered two simple answers—always 'no' and always 'yes'—but showed

---

[60]This is a rather rough characterisation of Grano and Phillips-Brown's view. See their paper for the details.

[61]The connection between COUNTERFACTUAL PERMISSION and Grano and Phillips-Brown's entry for desire reports is suggestive of a deeper relationship between deontic modals and bouletics. For further reflections on this relationship, see Jerzak (2019), Blumberg and Hawthorne (2022).

that, granted plausible background assumptions, they have significant implications for the theory of obligation, on the one hand, and the theory of ability, on the other. Thus, I answer '(only) sometimes'. That is, there are actions X and Y such that (i) S is unable to perform either, and (ii) S is permitted to X, but S is not permitted to Y. This raises the question **Difference**: what accounts for the normative difference between the actions X and Y such that the former is permissible but the latter is not? I suggested that the answer involves counterfactual situations where the subject performs the relevant action, and developed an analysis on which permission claims have a counterfactual component to their meaning.

One of my aims in this paper has been to present a positive theory of permission; and the approach I put forward appears to be on the right track. But regardless of whether this project is ultimately successful, my broader goal has been to draw out the various consequences of answers to **Permission**. Overall, I conclude that **Permission** is an interesting and challenging question. Hopefully our discussion brings more attention to it, and opens up the space for other, possibly more compelling responses to be formulated.

References

Vuko Andrić. Objective consequentialism and the rationales of ? ?ought? implies ?can? ? *Ratio*, 30(1):72–87, 2017. doi: 10.1111/rati.12108.

Andrew Bacon. Stalnaker's Thesis in Context. *Review of Symbolic Logic*, 8(1):131–163, 2015. doi: 10.1017/s1755020314000318.

Kyle Blumberg. On preferring. *Linguistics and Philosophy*, 45(6):1315–1344, 2022. doi: 10.1007/s10988-022-09351-4. URL https://doi.org/10.1007/s10988-022-09351-4.

Kyle Blumberg and John Hawthorne. Desire. *Philosophers' Imprint*, 22(n/a), 2022. doi: 10.3998/phimp.2116.

Kyle Blumberg and John Hawthorne. Inheritance: Professor Procrastinate and the Logic of Obligation. *Philosophy and Phenomenological Research*, 106(1):84–106, 2023. doi: 10.1111/phpr.12846.

David Boylan. Does success entail ability? *Noûs*, 56(3):570–601, 2022. doi: https://doi.org/10.1111/nous.12370. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12370.

David Boylan. Putting 'ought's together. *Semantics and Pragmatics*, 16, 2023. doi: https://doi.org/10.3765/sp.16.5.

John Broome. *Rationality Through Reasoning*. Malden, MA: Wiley-Blackwell, 2013.

Frederick E. Brouwer. A difficulty with 'ought implies can'. *Southern Journal of Philosophy*, 7(1):45–50, 1969. doi: 10.1111/j.2041-6962.1969.tb02040.x.

James Brown. Moral theory and the ought–can principle. *Mind*, 86(342):206–223, 1977. doi: 10.1093/mind/lxxxvi.342.206.

Fabrizio Cariani. 'Ought' and resolution semantics. *Noûs*, 47(3):534–558, 2013. doi: 10.1111/j.1468-0068.2011.00839.x.

Fabrizio Cariani. Consequence and contrast in deontic semantics. *Journal of Philosophy*, 113(8):396–416, 2016. doi: 10.5840/jphil2016113826.

Fabrizio Cariani and Simon Goldstein. Conditional Heresies. *Philosophy and Phenomenological Research*, (2):251–282, 2018. doi: 10.1111/phpr.12565.

Fabrizio Cariani, Magdalena Kaufmann, and Stefan Kaufmann. Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy*, 36(3):225–259, 2013. doi: 10.1007/s10988-013-9134-4.

Nate Charlow. Prospects for an expressivist theory of meaning. *Philosophers' Imprint*, 15:1–43, 2015.

Gennaro Chierchia and Sally Mcconnell-Ginet. *Meaning and Grammar - 2nd Edition: An Introduction to Semantics*. MIT Press, Cambridge, Mass, second edition edition, March 2000. ISBN 978-0-262-53164-1.

Cleo Condoravdi. Temporal interpretation of modals - modals for the present and for the past. *The Construction of Meaning*, 10 2001.

Cleo Condoravdi. Npi licensing in temporal clauses. *Natural Language & Linguistic Theory*, 28(4):877–910, 2010. doi: 10.1007/s11049-010-9115-z. URL https://doi.org/10.1007/s11049-010-9115-z.

Aniko Csirmaz. *Aspect, Negation and Quantifiers*, pages 225–253. Springer Netherlands, Dordrecht, 2008. ISBN 978-1-4020-4755-8. doi: 10.1007/978-1-4020-4755-8_10. URL https://doi.org/10.1007/978-1-4020-4755-8_10.

Nadine Elzein. Determinism, ?ought? implies ?can? and moral obligation. *Dialectica*, 74(1): 35–62, 2020. doi: 10.48106/dial.v74.i1.03.

John Martin Fischer. ?ought-implies-can?, causal determinism and moral responsibility. *Analysis*, 63(3):244–250, 2003. doi: 10.1111/1467-8284.00430.

Bas C. Van Fraassen. Values and the heart's command. *Journal of Philosophy*, 70(1):5–19, 1973. doi: 10.2307/2024762.

Anette Frank. *On context dependence in modal constructions*. PhD thesis, Universit at Stuttgart, 1997.

Harry G. Frankfurt. Alternate possibilities and moral responsibility. *Journal of Philosophy*, 66(23):829–839, 1969. doi: 10.2307/2023833.

Melissa Fusco. Agential free choice. *Journal of Philosophical Logic*, 50(1):57–87, 2020. doi: 10.1007/s10992-020-09561-w.

Peter A Graham. 'Ought' and ability. *The Philosophical Review*, 120:337–382, 2011.

Peter A. Graham. *Subjective versus Objective Moral Wrongness*. Cambridge: Cambridge University Press, 2021.

Thomas Grano and Milo Phillips-Brown. (counter)factual want ascriptions and conditional belief. *Journal of Philosophy*, 119(12):641–672, 2022. doi: 10.5840/jphil20221191241.

James Griffin. The human good and the ambitions of consequentialism. *Social Philosophy and Policy*, 9(2):118, 1992. doi: 10.1017/s0265052500001436.

Ishtiyaque Haji. *Deontic Morality and Control*. Cambridge University Press, 2002.

John Hawthorne. The Epistemic Use of 'Ought'. In Lee Walters and John Hawthorne, editors, *Conditionals, Paradox, and Probability: Themes from the Philosophy of Dorothy Edgington*, page 0. Oxford University Press, February 2021. ISBN 978-0-19-871273-2. doi: 10.1093/oso/9780198712732.003.0010.

Brian Hedden. Options and the subjective ought. *Philosophical Studies*, 158(2):343–360, 2012. doi: 10.1007/s11098-012-9880-0.

Brian Hedden. *Reasons Without Persons: Rationality, Identity, and Time*. Oxford University Press UK, Oxford, United Kingdom, 2015.

Irene Heim. On the Projection Problem for Presuppositions. In P. Portner and B. H. Partee, editors, *Formal Semantics - the Essential Readings*, pages 249–260. Blackwell, 1983.

Irene Heim. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3):183–221, 1992.

Lawrence L. Heintz. Excuses and "ought" implies "can". *Canadian Journal of Philosophy*, 5 (3):449–462, 1975.

David Heyd. Supererogation. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2024 edition, 2024.

Frances Howard-Snyder. It's the Thought That Counts. *Utilitas*, 17(3):265–281, 2005. doi: 10.1017/s0953820805001640.

Nick Hughes. Guidance, obligations and ability: A close look at the action guidance argument for ought-implies-can. *Utilitas*, 30(1):73–85, 2018. doi: 10.1017/s095382081700005x.

Sabine Iatridou and Hedde Zeijlstra. Negation, polarity, and deontic modals. *Linguistic Inquiry*, 44(4):529–568, 2013. ISSN 00243892, 15309150.

Frank Jackson. Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101:461–482, 1991.

Richard C. Jeffrey. *The Logic of Decision*. University of Chicago Press, Chicago, IL, 1983. ISBN 978-0-226-39582-1.

Ethan Jerzak. Two ways to want? *Journal of Philosophy*, 116(2):65–98, 2019. doi: 10.5840/jphil201911624.

Andrew J. I. Jones and Ingmar Pörn. Ought and Must. *Synthese*, 66(1):89–93, 1986. doi: 10.1007/bf00413581.

James M. Joyce. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press, 1999.

Lauri Karttunen. Presuppositions of compound sentences. *Linguistic Inquiry*, 4(2):169–193, 1973. ISSN 00243892, 15309150.

Lauri Kartunnen. Presupposition and linguistic context. *Theoretical Linguistics*, 34(1):181–, 1974.

Christopher Kennedy. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1):1–45, 2007. doi: 10.1007/s10988-006-9008-0.

Anthony Kenny. *Will, Freedom, and Power*. Blackwell, New York, 1975.

Anthony Kenny. *Human Abilities and Dynamic Modalities*, pages 209–232. Springer Netherlands, Dordrecht, 1976. ISBN 978-94-010-1823-4. doi: 10.1007/978-94-010-1823-4_11. URL https://doi.org/10.1007/978-94-010-1823-4_11.

Sophie Kikkert. Ability?s two dimensions of robustness. *Proceedings of the Aristotelian Society*, 122(3):348–357, 2022. doi: 10.1093/arisoc/aoac002.

Alex King. 'ought implies can': Not so pragmatic after all. *Philosophy and Phenomenological Research*, 95(3):637–661, 2017. doi: 10.1111/phpr.12447.

Alex King. *What We Ought and What We Can*. New York, NY, USA: Routledge, 2019.

Justis Koon. Options must be external. *Philosophical Studies*, 177(5):1175–1189, May 2020. ISSN 1573-0883. doi: 10.1007/s11098-019-01240-0.

Angelika Kratzer. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1: 337–355, 1977.

Angelika Kratzer. The notional category of modality. In H. J. Eikmeyer and H. Rieser, editors, *Words, Worlds, and Contexts: New Approaches in World Semantics*. Berlin: de Gruyter, 1981.

Angelika Kratzer. Conditionals. *Chicago Linguistics Society*, 22(2):1–15, 1986.

Angelika Kratzer. Modality. In Arnim von Stechow and D Wunderlich, editors, *Semantik/Semantics: An International Handbook of Contemporary Research*. Berlin: de Gruyter, 1991.

Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, 2012.

E. J. Lemmon. Moral dilemmas. *Philosophical Review*, 71(2):139–158, 1962. doi: 10.2307/2182983.

Dmitry Levinson. Probabilistic model-theoretic semantics for 'want'. *Semantics and Linguistic Theory*, 13(0):222–239, 2003. ISSN 2163-5951. doi: 10.3765/salt.v13i0.2888.

David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359, 1979. doi: 10.1007/bf00258436.

David K. Lewis. *Counterfactuals*. Blackwell, 1973.

Robert Lockie. *Free Will and Epistemology: A Defence of the Transcendental Argument for Freedom*. Bloomsbury Academic, London, UK, 2018.

Matthew Mandelkern. Talking about worlds. *Philosophical Perspectives*, 32(1):298–325, 2018. doi: 10.1111/phpe.12112.

Matthew Mandelkern. If p, then p! *Journal of Philosophy*, 118(12):645–679, 2021. doi: 10.5840/jphil20211181242.

Matthew Mandelkern. Chance, ability, and control. *Philosophy and Phenomenological Research*, forthcoming.

Matthew Mandelkern, Ginger Schultheis, and David Boylan. Agentive modals. *Philosophical Review*, 126(3):301–343, 2017. doi: 10.1215/00318108-3878483.

Elinor Mason. Consequentialism and the "ought implies can" principle. *American Philosophical Quarterly*, 40(4):319–331, 2003.

Elinor Mason. *Ways to Be Blameworthy: Rightness, Wrongness, and Responsibility*. Oxford, United Kingdom: Oxford University Press, 2019.

Paul McNamara. Must I do what I ought (or will the least I can do do)? In Mark Brown and Jose Carmo, editors, *Deontic Logic, Agency and Normative Systems*, pages 154–173. Berlin: Springer-Verlag, 1996.

G. E. Moore. The nature of moral philosophy. In *Philosophical Papers*. Routledge and Kegan Paul, 1922.

Sarah Moss. On the pragmatics of counterfactuals. *Noûs*, 46(3):561–586, 2012. doi: 10.1111/j.1468-0068.2010.00798.x.

Daniel Muñoz and Jack Spencer. Knowledge of objective 'oughts': Monotonicity and the new miners puzzle. *Philosophy and Phenomenological Research*, 103(1):77–91, 2021. doi: 10.1111/phpr.12702.

Dilip Ninan. Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, editors, *New Work on Modality, MIT Working Papers in Linguistics*. 2005.

F. R. Palmer. *Modality and the English Modals*. New York: Longman, 2nd edition, 1990.

Derek Parfit. *Reasons and Persons*. Oxford University Press, 1984.

Derek Parfit. *On What Matters: Two-volume Set*. Oxford University Press, 2011.

John L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press, 1995.

Douglas W. Portmore. *Opting for the Best: Oughts and Options*. Oxford University Press, New York, NY, USA, 2019.

Paul Portner and Aynat Rubinstein. Mood and contextual enrichment. In Anca Chereches, editor, *Proceedings of SALT 22*, 2012.

H. A. Prichard. Duty and ignorance of fact. *Philosophy (London, England)*, 8(30):226–228, 1932. doi: 10.1017/S0031819100062653.

Nicholas Rescher. *Ethical Idealism*. University of California Press, 1987.

William David Ross. *Foundations of Ethics: The Gifford Lectures Delivered in the University of Aberdeen, 1935–6*. Oxford University Press, 1939.

Daniel Rothschild. Do indicative conditionals express propositions? *Noûs*, 47(1):49–68, 2011. doi: 10.1111/j.1468-0068.2010.00825.x.

Paul Saka. Ought does not imply can. *American Philosophical Quarterly*, 37(2):93–105, 2000.

Paolo Santorio. Path semantics for indicative conditionals. *Mind*, 131(521):59–98, 2022. doi: 10.1093/mind/fzaa101.

Steve F. Sapontzis. ??ought? does imply ?can?? *Southern Journal of Philosophy*, 29(3): 382–393, 1991. doi: 10.1111/j.2041-6962.1991.tb00598.x.

Philippe Schlenker. Conditionals as definite descriptions. *Research on Language and Computation*, 2(3):417–462, 2004. doi: 10.1007/s11168-004-0908-2. URL https://doi.org/10.1007/s11168-004-0908-2.

Philippe Schlenker. Local contexts. *Semantics and Pragmatics*, 2(3):1–78, July 2009. doi: 10.3765/sp.2.3.

Mark Schroeder. Ought, agents, and actions. *Philosophical Review*, 120(1):1–41, 2011. doi: 10.1215/00318108-2010-017.

Ginger Schultheis. Counterfactual probability. *Journal of Philosophy*, forthcoming.

Moritz Schulz. Counterfactuals and Arbitrariness. *Mind*, 123(492):1021–1055, 2014. doi: 10.1093/mind/fzu144.

Moritz Schulz. *Counterfactuals and Probability*. Oxford, United Kingdom: Oxford University Press, 2017.

Wolfgang Schwarz. Ability and possibility. *Philosophers' Imprint*, 20, 2020.

Wolfgang Schwarz. Objects of Choice. *Mind*, 130(517):165–197, January 2021. ISSN 0026-4423. doi: 10.1093/mind/fzaa032.

Alex Silk. Weak and strong necessity modals: On linguistic means of expressing "a primitive concept ought". In Billy Dunaway and David Plunkett, editors, *Meaning, Decision, and Norms: Themes From the Work of Allan Gibbard*, pages 203–245. Maize Books, 2021.

Walter Sinnott-Armstrong. ?ought' conversationally implies ?can'. *Philosophical Review*, 93 (2):249–261, 1984. doi: 10.2307/2184585.

Walter Sinnott-Armstrong. *Moral Dilemmas*. Blackwell, New York, NY, USA, 1988.

Aaron Sloman. 'ought' and 'better'. *Mind*, 79(315):385–394, 1970. doi: 10.1093/mind/lxxix.315.385.

Holly M. Smith. Subjective rightness. *Social Philosophy and Policy*, 27(2):64–110, 2010. doi: 10.1017/s0265052509990161.

Jesse Spafford. Does initial appropriation create new obligations? *Journal of Ethics and Social Philosophy*, 17(2), 2020. doi: 10.26556/jesp.v18i1.952.

Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457, 1973. doi: 10.1007/BF00262951. URL https://doi.org/10.1007/BF00262951.

Robert Stalnaker. Pragmatic Presuppositions. In Robert Stalnaker, editor, *Context and Content*, pages 47–62. Oxford University Press, 1974.

Robert C. Stalnaker. A theory of conditionals. In Nicholas Rescher, editor, *Studies in Logical Theory (American Philosophical Quarterly Monographs 2)*, pages 98–112. Oxford: Blackwell, 1968.

Robert C. Stalnaker. Pragmatics. *Synthese*, 22(1):272–289, 1970. doi: 10.1007/BF00413603. URL https://doi.org/10.1007/BF00413603.

Robert C. Stalnaker. A defense of conditional excluded middle. In William Harper, Robert C. Stalnaker, and Glenn Pearce, editors, *Ifs*, pages 87–104. Reidel, 1980.

Jason Stanley and Zoltan Gendler Szabó. On quantifier domain restriction. *Mind and Language*, 15(2-3):219–261, 2000.

W. Starr. Counterfactuals. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2022 edition, 2022.

Una Stojnić. The inferential constraint and ?if Φ, ought Φ? problem. *Philosophical Studies*, 181(6), 2024. doi: 10.1007/s11098-024-02127-5.

Peter F. Strawson. On Referring. *Mind*, 59(235):320–344, 1950. doi: 10.1093/mind/lix.235.320.

Bart Streumer. Does 'ought' conversationally implicate 'can'? *European Journal of Philosophy*, 11(2):219–228, 2003. doi: 10.1111/1468-0378.00184.

Bart Streumer. Reasons and impossibility. *Philosophical Studies*, 136(3):351–384, 2007. doi: 10.1007/s11098-005-4282-1.

Bart Streumer. Reasons, impossibility and efficient steps: Reply to heuer. *Philosophical Studies*, 151(1):79–86, 2010. doi: 10.1007/s11098-009-9422-6\ensuremath{\}}.

Roger Trigg. Moral conflict. *Mind*, 80(317):41–55, 1971. doi: 10.1093/mind/lxxx.317.41.

Kai von Fintel. NPI licensing, strawson entailment, and context dependency. *Journal of Semantics*, 16(2):97–148, 1999. doi: 10.1093/jos/16.2.97.

Kai von Fintel. Counterfactuals in a Dynamic Context. In *Ken Hale: A Life in Language*. The MIT Press, 04 2001. ISBN 9780262316125. doi: 10.7551/mitpress/4056.003.0006. URL https://doi.org/10.7551/mitpress/4056.003.0006.

Kai von Fintel. What is presupposition accommodation, again? *Philosophical Perspectives*, 22(1):137–170, 2008. doi: https://doi.org/10.1111/j.1520-8583.2008.00144.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1520-8583.2008.00144.x.

Kai von Fintel. The best we can (expect to) get? challenges to the classic semantics for deontic modals. Chicago, IL, 2012. Central APA. URL https://web.mit.edu/fintel/fintel-2012-apa-ought.pdf.

Kai von Fintel and Sabine Iatridou. How to Say Ought in Foreign: The Composition of Weak Necessity Modals. In Jacqueline Guéron and Jacqueline Lecarme, editors, *Time and Modality*, Studies in Natural Language and Linguistic Theory, pages 115–141. Springer Netherlands, Dordrecht, 2008. ISBN 978-1-4020-8354-9. doi: 10.1007/978-1-4020-8354-9_6.

Kai von Fintel and Sabine Iatridou. Prolegomena to a theory of X-marking, March 2023.

Peter B. M. Vranas. I ought, therefore i can. *Philosophical Studies*, 136(2):167–216, 2007. doi: 10.1007/s11098-007-9071-6.

Roger Wertheimer. *The Significance of Sense*. Cornell University Press, Ithaca, 1972.

Malte Willer. Dynamic thoughts on ifs and oughts. *Philosophers' Imprint*, 14:1–30, 2014.

Malte Willer. Two puzzles about ability can. *Linguistics and Philosophy*, 44(3):551–586, 2020. doi: 10.1007/s10988-020-09296-6.

B. A. O. Williams and W. F. Atkinson. Ethical consistency. *Aristotelian Society Supplementary Volume*, 39(1):103–138, 1965. doi: 10.1093/aristoteliansupp/39.1.103.

Seth Yalcin. Modalities of Normality. In Nate Charlow and Matthew Chrisman, editors, *Deontic Modality*, page 0. Oxford University Press, June 2016. ISBN 978-0-19-871792-8. doi: 10.1093/acprof:oso/9780198717928.003.0009.

Michael J. Zimmerman. Is moral obligation objective or subjective? *Utilitas*, 18(4):329–361, 2006. doi: 10.1017/s0953820806002159.

Michael J. Zimmerman. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press, 2008.

Michael J. Zimmerman. *Ignorance and Moral Obligation*. Oxford: Oxford University Press, 2014.

Zsófia Zvolenszky. Is a possible-worlds semantics of modality possible? a problem for kratzer's semantics. *Proceedings of Semantics and Linguistic Theory (SALT)*, pages 339–358, 2002.

Zsófia Zvolenszky. A semantic constraint on the logic of modal conditionals. In Beáta Gyuris, László Kálmán, Chris Pinon, and Károly Varasdi, editors, *Proceedings of the Ninth Symposium on Logic and Language (LoLa 9)*, pages 167–177. 2006.