Deontic Modality

Suikkanen, Jussi

DOI:

[10.1093/analys/any015](https://doi.org/10.1093/analys/any015)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Suikkanen, J 2018, 'Deontic Modality', *Analysis*, vol. 78, no. 2, pp. 354-363.

<https://doi.org/10.1093/analys/any015>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Checked for eligibility: 26/02/2018

This is a pre-copyedited, author-produced version of an article accepted for publication in *Analysis* following peer review. The version of record Jussi Suikkanen; Deontic Modality, *Analysis*, Volume 78, Issue 2, 1 April 2018, Pages 354–363 is available online at:
<https://doi.org/10.1093/analys/any015>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

English language contains a number of expressions that often stand for concepts called the ‘deontic modals’. These expressions include ‘must,’ ‘should,’ ‘ought,’ ‘may,’ ‘can,’ and the like. The deontic modal concepts, which these words are often used to express, are interesting in many ways. Firstly, there are logical connections between them – for example, they can be ordered in terms of strength. That I must go home now entails that I ought to go home now, and that proposition furthermore entails that I may go home now. Yet, the inferences in the other direction from may to ought and from ought to must are not valid. The strength of the deontic modals thus varies from the strict necessity modals (‘must’) to weak necessity (‘ought’) and possibility modals (‘may’).

Deontic modals also come in a number of different ‘flavours’. If I say that I must go home now, the intuitive idea is that I am using ‘must’ in a different sense depending on whether I express the thought that overall I have most reason to go or the thought that morality, prudence, law, or some other normative standard requires me to do so. These same words can also be used to express different non-deontic modalities. For example, when I say that the train ought to be at the station, I am using ‘ought’ in an epistemic sense that is related to how things would normally be relative to what information I have.²

This brief introduction of deontic modals makes it clear that these complex concepts are both interesting and important. Because of this, their meaning has recently been intensively debated in both linguistics and philosophy (especially in metaethics and philosophy of language). There has also been a significant recent development in these debates. Both in linguistics and in philosophy of language (and to a lesser extent even in metaethics) the project has been to develop further the familiar formal semantic models that are commonly constructed by relying on the possible world machinery. The hope has been that, by applying the resulting models to deontic modals, the meaning of these expressions – understood as their contribution to the truth conditions of the sentences in which they appear – can be stated precisely. *Deontic Modality* makes an invaluable contribution by collecting together the very best recent work done to pursue this project in both linguistics and philosophy. It contains an introduction and 15 previously unpublished state of the art articles.

In this critical notice, I will first introduce, in a rough informal outline, the standard formal semantic model used to capture the truth-conditional meaning of deontic modals (§1). §1 also briefly explains some of the reasons why that model has come to be seen as unsatisfactory. Most of the articles in *Deontic Modality* then, in one way or another, react to the perceived problems of the standard model. §2 thus explains how many of the articles of this collection try to either develop the standard model further or formulate alternatives for it.

§3 will be the main original critical contribution. It will consider the threat that the more sophisticated versions of the competing semantic models discussed in §2 turn out to be extensionally equivalent. I will consider the principles that should govern our choice of

¹ *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman. Oxford University Press, 2016, 434 pp.

² In addition, there are also, for example, dispositional and preferential modal uses of these words.

formal models in this situation and what lessons we can draw as a consequence concerning the purpose and status of the formal models. Finally, in §4, I will briefly say few things about the remaining articles and the collection.

1. The standard ordering semantics

The task at hand is to develop a formal model that would be able to capture the meaning of the words that are used to express deontic modal concepts in terms of the contribution they make to the truth conditions of the sentences in which they appear. This section first informally sketches the standard model developed by Angelika Kratzer (1977, 1991, and 2012) on the basis of David Lewis's (1973) well-known account of counterfactuals. This account is known as the Lewis-Kratzer Semantics or the Standard Ordering Semantics.³

According to the standard model, in terms of their logical form, deontic modals are propositional operators: other propositions (called 'prejacent') can be embedded within their scope. The semantic role of these operators is, on this view, then to quantify over a certain set of possible worlds. The strict necessity modal 'must' functions as a universal quantifier (Must (p) is true just when the prejacent p is true in all the possible worlds in the relevant set), whereas the possibility modal 'may' functions as an existential quantifier (May (p) is true just when the prejacent is true in at least one of those worlds).

The second part of the theory is a two-stage account of how the relevant set of worlds over which the modals quantify is determined. Firstly, we need a set of relevant alternatives called the 'modal base'. It is to be understood as a set of possible worlds (or, more precisely, as a function from a world of evaluation to a premise-set of propositions that corresponds to a set of worlds in which those propositions are true – the modal background). Secondly, we need a ranking of the worlds in the modal base according to a standard. This ranking is provided by an ordering source (more precisely, a function from a world of evaluation to an evaluative ranking of the worlds in the modal background). The idea then is that 'must' and 'may' quantify over all the highest ranked worlds in the modal base.

The third part of the standard model is the idea that the context of conversation provides both the relevant alternatives (the modal base) and the ranking of those alternatives (the ordering source). One interesting feature of Kratzer's original account is that this happens in different ways in the case of deontic and epistemic modals. With the deontic modals, the modal base is 'circumstantial': it consists of all the worlds that are similar to the world of evaluation in the objective ways that are contextually salient in a conversation. In contrast, with the epistemic modals, the modal base is 'epistemic': it consists of all the worlds that are consistent with a body of information that has certain properties salient in the conversational context (such as being what the speaker knows, what she and the audience know, or the like).

The conversational context also provides the ordering source. The idea is that, with different conversational backgrounds, different normative standards are salient. A salient standard can be, for example, what the law requires, what is in the agent's interests, or even a moral ideal depending on what deontic 'flavour' of the modal is used in the context. The worlds in the modal base are then ranked in terms of how well they meet the relevant normative standard. In contrast, with epistemic modals, the contextually salient ordering sources are usually

³ For helpful presentations of the standard model in this collection, see Charlow (ch. 2, §2), Bronfman & Dowell (ch. 4, §1), Finlay (ch. 6, §1), and Yalcin (ch. 8, §4).

meant to be ‘stereotypical’: the worlds in the epistemic modal base are ranked in terms of to what degree they conform to situations that are the normal course of events relative to the conversational background.

The standard model thus provides an elegant way to model the truth conditions of different types of modal talk and thought in different contexts. It attempts to capture precisely both what is common to all different flavours of deontic modality and also what distinguishes different types of uses in different contexts from one another. There is, however, also a growing body of evidence against the standard model. Here I will merely raise four significant problems discussed in this collection.

Firstly, the standard model does not seem to have the resources required for capturing the meaning of deontic weak necessity modals such as ‘ought’ and ‘should’.⁴ If ‘Daniel must dance’ is true when he dances in all of the relevant worlds and ‘Daniel may dance’ when he dances in some of those worlds, when is it true that Daniel should dance? The standard account itself does not answer this question. Furthermore, whatever the answer is, it should also explain how ‘should’ is logically weaker than ‘must’ but stronger than ‘may’.

Secondly, because in the standard model the modal base of deontic modals is fixed ‘circumstantially’, it is not clear that the model can explain how the truth conditions of the assertions in which deontic modals are used are in some cases *information-sensitive*.⁵ This problem is usually illustrated with the so-called miners case in which ten miners are trapped in one of two shafts, A or B, and about to be drowned by floodwaters.⁶ You do not know whether the miners are in A or in B. If you block the wrong shaft all the miners die, whereas if you block neither just one miner drowns and the rest are saved.

Intuitively, here we ought to block neither shaft. Yet, the truth of this claim seems to depend on our ignorance – our state of information and what is missing from it. However, with the deontic modals, the standard model picks out the modal base ‘circumstantially,’ in a way that does not depend on what we know. And, in all the worlds that are in the relevant ways similar to the world of evaluation, the miners are in shaft A (or in B depending on the world of evaluation) in which case the worlds in which we block shaft A (or B) are bound to be ranked highest by the ordering source. This entails that the standard model predicts that it is false that we should block neither shaft.

Thirdly, it is also objected that the standard model encodes a controversial decision rule into the meaning of deontic modals.⁷ In situations in which the outcome of our actions is uncertain, this maximax rule tells us to choose the action that has the best potential outcome (whatever the likelihood of that outcome is). The standard model encodes this decision rule into the semantics because, according to it, Must (p) is true when p is the case in the worlds that are ranked highest in the modal base according to the ordering source. However, the maximax rule is a controversial rule – in fact, most people reject it. Yet, the standard model tells us that we should understand the truth conditions of everyone’s deontic utterances to be based on the maximax rule.

⁴ For discussions, see, e.g., Charlow (ch. 2, §3.1), and Finlay (ch. 6, §4).

⁵ For discussions, see, e.g., Bronfman & Dowell (ch. 4, §2.1) and Finlay (ch. 6, §5).

⁶ Kolodny & MacFarlane (2010).

⁷ See, e.g., Charlow (ch. 2, §4.1) and Lassiter (ch. 3, §2).

Finally, an additional difficulty for the standard model follows from the thought that at least some deontic modals come in degrees.⁸ I ought to take my vitamins but this ought is presumably weaker than the one at play in the thought that I ought not to kill other people. Given that according to the standard model deontic modals merely quantify over the relevant sets of worlds, it is not clear how they could capture this type of gradability.

2. Reactions

The previous section outlined the standard model and it also introduced some of its problems. A lot more could clearly be said about both. The previous overview, however, hopefully provided a sufficient background for the articles in *Deontic Modality*. This is because most of the articles in it either develop the standard model further or formulate alternatives to it so as to avoid the previous objections to the standard model.

2.1. Sophisticated ordering semantics

Many of the articles in this collection respond to the previous objections by developing the standard model further. In chapter 1, Fabrizio Cariani argues that several of the previous objections can be responded to by (i) replacing the standard model's modal bases with fine-grained states (set of worlds and probability function pairs) and (ii) letting the ordering source rank alternatives (options available for an agent in a choice-situation) in terms of their qualities – importantly including the likelihoods of their different outcomes. Because of these amendments, Cariani's model allows that probabilistic considerations influence how the contextually salient ordering source ranks the relevant alternatives.

In a somewhat resembling way, Ralph Wedgwood (ch. 5) argues that the standard model needs to be amended by stipulating that the worlds in the contextually determined modal base must be ordered expectationally. That is, on his view, there must be a probability function that assigns the metaphysically possible worlds a probability. The contextually salient value function then orders mutually exclusive but jointly exhaustive actions in terms of their expected utility (the value function orders epistemic possibilities that partition the space of the metaphysical possibilities in the modal base). According to Wedgwood, the deontic modals can then be understood in the standard way as propositional operators that quantify over the highest ranked alternatives in the modal base.

The first advantage of these amendments to the standard model is that they entail the appropriate kind of information-sensitivity. This is because the body of information that is contextually salient can affect the relevant probability function that in part fixes the expected value of the alternatives ranked by the alternatives. On Cariani's view, furthermore, a contextually salient ordering source can be sensitive to probabilistic features of the alternatives (rank low risky ones, for example, just because they are risky). For these reasons, in the miners case for example, it can turn out that blocking neither shaft will be the highest ranked alternative relative to our body of information. In addition, Cariani's version of the standard model has the further advantage that it no longer encodes any decision theoretic rule because different contextually salient ordering sources react to uncertainty in different ways (whereas Wedgwood's semantics seems to encode the expected utility maximization rule).

⁸ See, e.g., Lassiter (ch. 3, §1), Finlay (ch. 6, §4), and Portner & Rubinstein (ch. 9).

Aaron Bronfman and J.L. Dowell (ch. 4) defend the standard model almost in its original form against the objection according to which this model fails to fit our information-sensitivity intuitions in the miners case. According to them, we can use deontic modals at least in three different senses: objective, subjective, and advisability. Objective readings are ones in which the modal base consists of the worlds that are objectively alike the world of evaluation, whereas with the subjective readings contain the worlds that are compatible with what the agent knows. Finally, the advisability readings of deontic modals are ones in which certain salient facts are added to the agent's body of information to produce an updated modal space. This can be done with conditionals such as 'If the miners are in shaft A, we ought to block A' which adds a fact about the miners whereabouts to the relevant body of information for the purposes of practical reasoning. Bronfman and Dowell then aim to show that, with these three readings, we can provide appropriate truth conditions for what we would say about different versions of the miners case (where either new information will or won't be available before actions are needed) and we can do so without assuming that the contextual background changes in the middle of a single piece of reasoning. Their goal is thus to show that the standard model has more resources to deal with the information-sensitivity objection than commonly assumed.

Paul Portner and Aynat Rubinstein (ch. 9) argue that we should understand the difference between ought and must with the model of gradable adjectives (big) and their extreme counterparts (huge). In order to capture this idea in the framework of ordering semantics, Portner and Rubinstein suggest that we should add a second ordering source to the model and also give the model a scalar structure. The basic idea is that there is a primary ordering that corresponds to what is non-negotiable in a given context. The truth of must claims then is based on what is the case in the worlds ranked highest by the primary ordering source (i.e., what is true at the extreme end of the necessity scale). The secondary source (based on what might be negotiable given other relevant factors) then operates further on the worlds already ranked according to the primary ordering source to create a new ranking. This provides us with the non-extreme part of the necessity scale: ought claims (and their degrees) then indicate what is the case in the worlds highly ranked worlds in this part of the scale. With these amendments, we have a new model that enable the orderings semantics to make sense of both the logical connections between different deontic modals and their gradability.

Finally, Matthew Chrisman (ch. 13) argues that the standard model is unable to distinguish claims about what ought to be the case (John is kissed by Mary) from claims about what an agent ought to do (Mary: you ought to kiss John). For this reason, he recommends that we should allow that also prescriptions can be embedded to the deontic modal operators. The idea then is that we think of the modal base consisting of world-norm pairs and take a given deontic modal to which a prescription is embedded to be true when the relevant prescription is justified by the contextually salient norms.

2.2. *Decision theoretic models*

Decision theoretic semantic accounts reject the standard model's basic tenet according to which deontic modals quantify over possible worlds. Daniel Lassiter (ch. 3) defends a Bayesian version of this type of an account. On his view, we first assign a value to each possible world in the set of all worlds. We then assume that the context provides a probability measure that satisfies the standard constraints that apply to probability measures. With these

two assumptions in place, we can define expected value for each of our actions in the standard way. The expected value of an option is the average value of the option's potential outcomes weighed by their probabilities. As a consequence, different options in a choice-situation are located on the expected value scale.

We can then think of the deontic modals very much like gradable adjectives. The basic idea is that Ought (p), for example, is true if p has the greatest expected value. This simple account already can explain information-sensitivity of deontic modals, their gradability, and also the non-monotonicity of deontic reasoning (the simple idea that additional premises can change an inference from a good one to bad). This is both because additional information can change the expected value of our options by changing the relevant probabilities and because expected value clearly comes in degrees.

2.3. End-relational dyadic semantics

Stephen Finlay (ch. 6) tries to avoid the problems of the standard model by making it radically simpler by getting rid of the second, ordering source parameter of the model. According to his dyadic model, the truth conditions of the utterances containing deontic modals are always 'end-relative'. Whenever a deontic modal is used, there is always some salient end that is either part of the conversational background (law, morality, prudence, and the like) or explicitly expressed in the antecedent of a conditional ("If you want to go to Harlem, you ought to take the A train"). We then first use the relevant end to update the contextually salient modal base – we remove all the worlds from the modal base in which the end does not obtain. Must (p) then expresses the idea that p is the case in all the remaining worlds (p is thus necessary for the relevant end), Ought (p) that p is the case in most of the remaining worlds (p makes the end more likely), and May (p) that p is the case in some of them.

Because Finlay's semantics is relatively simple, he lets the pragmatics do a lot of work. This is the way in which the conversational context and the norms governing conversation more generally can fix the modal bases and the salient ends to which the modals are relative in a way that provides plausible truth conditions for our utterances. Here, depending on the context, which worlds belong to the modal base can, for example, depend on what the speakers know and so adding and subtracting information can make a difference to whether in the most worlds in which the end obtains it is true that p (and so whether it is true that Ought (p)). As a consequence, this view does entail information-sensitivity. Finlay's account also has surprisingly strong resources to deal with many of the other objections to the standard model. For example, it can explain the degrees of ought and the existence of conflicting obligations in terms of how likely the obtaining of the prejacent makes the achieving of different salient ends and the importance of those ends.

3. Extensional equivalence and theory choice

The previous section hopefully gives some idea of the ways in which the articles of *Deontic Modality* try to either develop the ordering semantics further or formulate alternatives to it. Let me then begin from an observation of how the debate concerning these competing formal semantic models seems to have progressed here. We began from a relatively simple formal model. Then, several problem cases, including the miners scenario, were introduced. In these cases, the basic model appears to have unintuitive consequences: it assigns truth conditions to

utterances containing deontic modals such that they make intuitively true assertions false, and intuitively false assertions true.

There is an interesting background methodological assumption at play here. In this debate, it seems like it is generally assumed that we should interpret speakers charitably, which is why the correct semantics for deontic modals should be ‘conservative’. It should make all intuitively true deontic utterances and only them true, including especially the relevant utterances in the problem cases. It is precisely on the basis of this assumption that the defenders of the different formal models construct their objections to each other’s models. Each side tries to show how the other views fail to lead to intuitive truth conditions in the problem cases whereas their own views can do so.

What will happen next in the debate? It is likely that the debate’s current stage has not yet fully played out. New fascinating problem cases will be discovered and these cases will require new amendments to the different types of models introduced above. Thus, for some time, we will get (i) new definitions and understandings of the different elements of the old accounts, (ii) additional parameters added to them, and (iii) novel ways of using the pragmatics of deontic discourse to explain how the models can produce right predictions. However, at some point in the future, this stage of the debate will have to come to an end. This is when we will have different formal models (ordering, decision theoretic, dyadic, etc.) that provide equally intuitive isomorphic truth conditions for different sentences containing deontic modals at each context of use. I see no reason why there will not be highly sophisticated versions of these models that turn out to be both extensionally equivalent and equally well able to fit the empirical data concerning which deontic statements we would intuitively think to be true.

What should we do then? How should we choose between the extensionally equivalent models?⁹ There are two current methodological principles that will no longer be able to do work. First, we will no longer be able to rely on the value of unity. This theoretic virtue has so far been understood in terms of how wide range of cases the models are able to deal with, but now we are assuming that there is a version of each type of a model that perfectly fits all the empirical data concerning cases. Likewise, we will not be able to choose between the models by considering whether they are able to provide the right kind of an inferential role for deontic modals (see Chrisman (ch. 13)). Because these models will entail the truth of exactly the same deontic utterances and thoughts in each context, presumably they will vindicate exactly the same inferential patterns.

Here Stephen Finlay (195) makes a helpful suggestion. He suggests that we should also compare the extensionally equivalent formal models in terms of their *simplicity*. It is, of course, a sound methodological principle that, other things being equal, we should prefer simpler theories. We can ask, however: why should we do so here? Let me consider two alternatives.

First, one might think that, as competent speakers, why have a skill of how use deontic modals correctly. We have the know-how required for using these terms in communication

⁹ One alternative not discussed below would be to think that extensionally equivalent formal models are in fact notational variants of the same semantic theory (for an analogy, see Dreier (2011)). In this case, the debate over the formal models would converge from a substantial debate to a verbal dispute.

and speaking truthfully in different contexts. One understanding of the formal semantic models is that their role is to make explicit in propositional form the implicit principles that govern competent speakers use of these terms.¹⁰ The question of theory choice would then be an empirical question of which one of the formal semantic models would in some physical sense be ‘encoded’ in the speakers’ brains and their information processing. Yet, in this case, it just does not seem that we should give much weight to the structural simplicity of the formal semantic models. A much better idea would be to do empirical cognitive science to understand how the neural network processes information when speakers use deontic modals in reasoning.¹¹

The second alternative is to reject the idea that the formal semantic models could in some sense capture the implicit principles that govern our use of the deontic modals. The idea here would be that deontic utterances and thoughts have truth conditions and the different co-extensive formal semantic models can be used to represent these equally well, but there is no way in which these models would be responsible for generating the competent speakers use of these terms. What the speakers have is a pure skill – a way of going about – that isn’t based on having any of these models doing work in the brain and thus there is nothing there within the speakers to which the models could correspond. Yet, even in this situation, we could use the models instrumentally for making predictions what the competent speakers are likely to say in different contexts. If the models were understood in this way, then it would definitely make sense to value their simplicity and the resulting ease of use.

This alternative has interesting consequences too. First, which type of a formal model is easy to use for predicting behaviour presumably differs from person to person. As evidenced by logic classes, different people just find different formal frameworks easier to grasp and manipulate. In the present context, there would then not be a further fact of the matter which of the formal models would be correct and so we could let the defenders of different formal models to get on with what they are doing with their own models. Secondly, if the only purpose for which the modal systems are needed is to make systematic predictions of which deontic sentences people accept in different contexts, then it is not clear to me that *formal* systems are needed for this purpose. After all, we are not alien anthropologists but rather speakers who share the same semantic competence.

4. *Few final words*

In this critical notice, I have focused on the articles in the collection that apply the truth conditional semantic approach to deontic modals. Several articles in the collection, however, take a different route. Two of these alternative approaches are especially worth mentioning here. Firstly, several authors attempt to develop dynamic semantic accounts for capturing the meaning of deontic modals.¹² According to them, we should try to understand the meaning of these concepts at least in part in terms of how using them in conversation updates the common ground of both beliefs and other attitudes that the conversationalists share. Secondly, several authors also pursue the expressivist project in which we attempt to capture the meaning of deontic modals in terms of the practical, desire-like attitudes they conventionally

¹⁰ See Finlay (195). Wedgwood (163), in contrast, explicitly states this is not what his model is aiming to do.

¹¹ There is, however, some reason to think, on scientific grounds, that semantic models for expressions that are not explicitly indexical or demonstrative should be ‘minimal’ – free from any contextual contributions (see Borg 2004, ch. 2).

¹² See Hellie (ch. 10), Willer (ch. 11), and Starr (ch. 12).

express.¹³ Finally, perhaps most interestingly, some authors focus on combining these two approaches. They try to model formally how the role deontic modals is to update the intentions and preferences that are shared in the conversational context – how the meaning of deontic modals consists of the way in which accepting sentences that contain them transforms certain practical elements of our state of mind.¹⁴ William Starr’s (ch. 12) attempt to build a formal semantic model for deontic modals in this way, for example, is one of the most promising attempts to solve the Frege-Geach problem for the expressivists I have seen.

Finally, I should emphasise how *Deontic Modality* is a rich and very rewarding collection for anyone interesting in deontic modals. It is also a very difficult book, at least for those of us who are not experts of complicated formal systems and their technical details. There are two important ways in which the collection could have been made more accessible. Firstly, it would have been helpful to see more authors to provide informal translations of their formal definitions. Secondly, the collection would have needed an historical overview of the recent debates on the topic. This too would have greatly helped the reader. At the moment, many authors criticise earlier versions of the models without explaining bothering to explain them. But, even if reading this collection occasionally feels like climbing a mountain, it definitely is worthwhile given how much one learns.

References

- Borg, E. 2004. *Minimal semantics*. Oxford: Oxford University Press.
- Dreier, J. 2011. In defense of consequentializing. *Oxford Studies in Normative Ethics* 1: 97–119.
- Kolodny, N. and MacFarlane, J. 2010. Ifs and oughts. *Journal of Philosophy* 107: 115–143.
- Kratzer, A. 1977. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy* 1: 337–355.
- Kratzer, A. 1991. Modality. In *Semantics: an international handbook of contemporary research*, eds. Groenendijk, J. et al, Berlin: de Gruyter; 639–651.
- Kratzer, A. 2012. *Modals and conditionals*. Oxford: Oxford University Press.
- Lewis, D. 1973. *Counterfactuals*. Malden: Blackwell.

¹³ See Charlow (ch. 2, §6), Yalcin (ch. 8, §9), Hellie (ch. 10), and Starr (ch.12).

¹⁴ See Hellie (ch. 10) and Starr (ch. 12).