ARTICLE

Error, consistency and triviality

Christine Tiefensee | **Gregory Wheeler**

Frankfurt School of Finance & Management

Correspondence
Christine Tiefensee, Frankfurt School of Finance & Management.
Email: c.tiefensee@fs.de

Abstract

In this paper, we present a new semantic challenge to the moral error theory. Its first component calls upon moral error theorists to deliver a deontic semantics that is consistent with the error-theoretic denial of moral truths by returning the truth-value *false* to all moral deontic sentences. We call this the 'consistency challenge' to the moral error theory. Its second component demands that error theorists explain in which way moral deontic assertions can be seen to differ in meaning despite necessarily sharing the same intension. We call this the 'triviality challenge' to the moral error theory. Error theorists can either meet the consistency challenge or the triviality challenge, we argue, but are hard pressed to meet both.

KEYWORDS
classical modal logic, deontic semantics, meaning, metaethics, moral error theory, neighborhood semantics

1 | INTRODUCTION

According to moral error theorists, all of our moral assertions are untrue because there is nothing that could make them true: Since there are no such things as objective moral properties or categorical reasons, moral error theorists explain, no action is ever morally obligatory, forbidden, or permissible.[1]

In this paper, we shall not concern ourselves with the plausibility of error-theoretic arguments against the existence of moral properties or categorical reasons. Rather, our focus will be firmly fixed on error theorists' ability to provide a convincing account of the meaning of deontic moral

sentences.[2] More precisely, we will challenge moral error theorists to deliver a non-trivial semantics of deontic moral sentences that renders all of these sentences false while simultaneously accounting for their different semantic content. Two observations explain why.

Firstly, given their rejection of moral truths, it is clear that error theorists must be able to present a deontic semantics which never returns the truth-value *true* to any deontic moral sentence. It is equally clear that for this purpose, error theorists cannot employ mainstream deontic logic and semantics, as for instance developed by Kratzer (1977, 2012) and Chellas (1980).[3] To elaborate, the family of standard deontic logics and their associated *minimal* (Chellas, 1980) or *neighborhood* (Pacuit, 2017) models, which include Kratzer semantics, Kripke semantics, and most other mainstream modal semantics as special cases,[4] all encode the assumption that *permissibility* and *obligation* are logically dual concepts: namely, if ¬P¬φ (¬O¬φ) is true, then the positive dual Oφ (Pφ) is true, too. They leave no room, therefore, for the error-theoretic thesis that no moral assertion is ever true. This conclusion is particularly conspicuous with regard to mainstream deontic logics that accept the dual schema. Importantly, though, it is not limited to semantics which regard O and P as duals.[5] Rather, the point applies to any deontic semantics that treats *permissibility* and *obligation* as distinct but logically related notions. For, so far as a semantics offers distinct but logically related satisfiability conditions for *permissibility* and *obligation*, conditions for universally falsifying one operator will fail to falsify the other, and thus return the truth-value *true* to some deontic moral sentence. Consequently, given their position's incompatibility with this wide array of mainstream and non-mainstream deontic semantics, error theorists owe us an alternative account for the meaning of deontic moral terms which secures consistency with the error theory by returning the truth-value *false*[6] to all deontic moral assertions. We call this the 'consistency challenge' to the modal moral error theory.

However, returning the truth-value *false* to all deontic moral assertions straightforwardly leads to our second observation. It is well-known that intensional modal semantics assign the same meaning to any two sentences that share the same truth-value in all circumstances of evaluation (Speaks, 2018). Consequently, by requiring intensional deontic semantics to assign the truth-value *false* to all deontic moral sentences in all possible worlds, the error theory entails that all deontic moral sentences have the same meaning. Yet, ascribing the same meaning to assertions such as 'Parents ought to care for their children' and 'It is impermissible for the prime minister to mislead the public' is just as implausible as holding them to mean the same as '2 + 2 = 5' and 'Siblings have different parents'. In order to avoid this result, error theorists thus need to explain how to account for the different meanings of deontic moral expressions despite their being necessarily intensionally equivalent. We call this the 'triviality challenge' to the modal moral error theory.[7]

Hence, combine the consistency and the triviality observations, and you arrive at our semantic challenge to the error theory: In order to provide a convincing account of deontic moral assertions, error theorists must not only present a deontic semantics which is compatible with their error-theoretic denial of moral truths, but also explain in which way moral assertions can be seen to differ in meaning despite necessarily sharing the same intension.

You might think that meeting the semantic challenge should be easy enough. After all, the moral error theory appears to be a coherent and conceptually possible metaethical position. Accordingly, there must surely be *some* deontic semantics that succeeds in capturing the error-theoretic thesis. Similarly, if intensional semantics provides insufficient resources to account for fine-grained meaning, adopting a hyperintensional semantics, which allows for intensionally equivalent formulas to have distinct semantic content, should certainly do the trick. However, as we will show in this paper, meeting our semantic challenge to the modal moral error theory is far from easy.

Our argument to this effect is two-pronged. Based on a general modal semantics for a basic modal language (Section 2), we will first develop a novel way of meeting the consistency challenge on behalf of moral error theorists (Section 3). This proposes that error theorists should analyze moral assertions in terms of a typed language and introduces a new operator—the 'XBox'—for the moral modal fragments of that language.[8] Semantically, the XBox works exactly as error theorists say, namely by mapping all moral assertions to the truth-value *false*. Moreover, it does so without having to reject the dual schema or deny that *obligation* and *permissibility* are logically related. Rather, our XBox solution succeeds in securing consistency for the error theory while leaving mainstream deontic semantics entirely untouched. With this new and very attractive solution to the consistency problem in hand, we then turn to the triviality challenge and consider how to generalize our framework to support hyperintensional modalities in order to avoid the XBox behaving as a triviality operator within intensional semantics (Section 4). However, this will show that employing hyperintensional modals only comes at the price of no longer being able to meet the consistency challenge: More precisely, what allows error theorists to address the consistency challenge is precisely what also prevents them from meeting the triviality challenge and *vice versa*. We conclude, therefore, that far from having an easy way out of our semantic challenge, modal moral error theorists are in trouble: They can either meet the consistency challenge or the triviality challenge, but are hard pressed to meet both (Section 5).

In what follows, we will focus exclusively on *modal moral* error theory. As such, we will consider deontic operators, such as 'ought' and 'permissible', which we will furthermore interpret in moral terms. However, as will become apparent, nothing in our argument hinges on interpreting the basic modal operators as the moral deontic operators 'ought' and 'permissible'. Hence, our arguments can be adapted to any other monadic modal error theory, should another arise. Moreover, to simplify our argument we will treat 'ought' and 'permissible' as logically dual modal operators. Yet, our semantic challenge requires only a much weaker assumption, namely that there is *some* logical relation between the semantic satisfiability conditions for 'ought' statements and 'permissible' statements. Finally, we will not examine the meaning of evaluative moral expressions, such as 'good', 'bad' or 'desirable'. Accordingly, in order to arrive at a comprehensive assessment of moral error theories' semantic credentials, our study would need to be supplemented with an independent inquiry into their ability to provide a convincing semantics of evaluative terms. Still, given that deontic sentences are generally taken to play a central role in our moral practices, and bearing in mind that moral error theorists themselves frequently consider deontic moral assertions when developing their positions, our semantic challenge is a key tool in putting the error theory's semantic qualities to the test.

2 | MODAL LOGIC PRELIMINARIES

Let us start by laying the foundations for our argument and introduce a general modal semantics for a basic modal language. Here, we choose a semantics that imposes very minimal constraints, and is thus able to accommodate practically all standard deontic modal semantics as special cases. We develop our argument in this general framework to make clear that, even though there is a variety of standard semantic accounts available, these do not offer error theorists a way out of our semantic challenge. Rather, our challenge gathers momentum no matter which of these semantics error theorists choose. For example, we briefly discuss in Appendix A how to cash out in our model Kripke's relational semantics and Kratzer's semantics for graded modals.

We begin by constructing a basic propositional modal language from a finite set of atomic propositions, $\mathsf{At} = \{p, q, r, ...\}$. This set At contains the non-logical propositional symbols of our language and will be the signature of our logic. Intuitively, if $p, q \in \mathsf{At}$, then p and q are each formulas of our propositional logic $\mathscr{L}$, with signature At, and so too are $\neg p$ and $p \wedge q$. If we add a single modal operator, O, then we have the necessary ingredients for a standard propositional deontic modal logic. Formally, the well-formed formulas φ of the basic modal language are given by the grammar

$$\varphi := p \mid \bot \mid \neg \phi \mid \phi \wedge \psi \mid \mathsf{O}\phi$$

where p ranges over the set of atomic propositions At and ϕ and ψ are arbitrary well-formed formulas constructed recursively by these rules. The modal language $\mathscr{L}(\mathsf{At})$ is the smallest set generated by this grammar.

Additional logical connectives are introduced for convenience, such as using $\phi \rightarrow \psi$ for $\neg(\phi \wedge \neg\psi)$ and using $\phi \vee \psi$ for $\neg(\neg\phi \wedge \neg\psi)$. The propositional constant $\bot$, called *falsum*, is useful for denoting falsehood and is defined by $(p \wedge \neg p)$, for any $p \in \mathsf{At}$. By convention, the permissibility operator P is defined in terms of the obligation operator, $\mathsf{P}\phi := \neg\mathsf{O}\neg\phi$, and O is used in deontic logic in place of the modal box operator $\Box$. Hereafter we will freely refer to O as a box operator and P as a diamond operator.

Next we turn to the semantics for our modal language $\mathscr{L}(\mathsf{At})$. Our semantic challenge to the modal moral error theory is a general one, which motivates our choice of neighborhood semantics. One advantage of a general semantic framework is that most if not all mainstream deontic semantics—for propositional languages, at least—can be reformulated as a specific kind of neighborhood model. This allows us to put technical details that are not relevant for our general argument safely into the background, although they otherwise are important distinguishing features of different semantics. If Kratzer and Kripke each have their own neighborhoods, which they do, then differentiating details of Kratzer and Kripke need not concern us here. As an aside, even when such technical details are important, having a common semantic framework can help to make those comparisons clear.

Let W be a non-empty set of states or 'worlds', as some prefer to call them. A *neighborhood function* $\mathscr{N}(\cdot)$ is a map from a non-empty set of states W to a collection of sets of states, $\mathscr{N}(w) : W \mapsto \wp(\wp(W))$, where $\wp(W)$ is the power set of W.[9] A *neighborhood model*, $\mathbb{M} = (W, \mathscr{N}, V)$, is a tuple, where W is a non-empty set of states, $\mathscr{N}$ is a neighborhood function over W, and V is a valuation function which assigns a set of states to each atomic proposition, $\mathsf{At} \mapsto \wp(W)$.

A formula $\varphi \in \mathscr{L}(\mathsf{At})$ is *true* in a state w of a neighborhood model $\mathbb{M}$, written '$\mathbb{M}, w \vDash \varphi$', and is defined by induction on the structure of φ:

1. $\mathbb{M}, w \vDash p$ iff $w \in V(p)$
2. $\mathbb{M}, w \vDash \neg\varphi$ iff $\mathbb{M}, w \nvDash \varphi$
3. $\mathbb{M}, w \vDash \varphi_1 \wedge \varphi_2$ iff $\mathbb{M}, w \vDash \varphi_1$ and $\mathbb{M}, w \vDash \varphi_2$
4. $\mathbb{M}, w \vDash \mathsf{O}\varphi$ iff $[\![\varphi]\!]_{\mathbb{M}} \in N(w)$
5. $\mathbb{M}, w \vDash \mathsf{P}\varphi$ iff $W - [\![\varphi]\!]_{\mathbb{M}} \notin N(w)$

where $[\![\varphi]\!]_{\mathbb{M}}$ is the *truth set* of formula $\varphi \in \mathscr{L}(\mathsf{At})$ in model $\mathbb{M}$, that is the set of states in which φ is true, $\{w : \mathbb{M}, w \vDash \varphi\}$. The truth conditions for $\mathsf{O}\varphi$ and $\mathsf{P}\varphi$ are duals (Chellas, 1980, Def. 7.2), meaning that in our language $\mathbb{M}, w^* \vDash \mathsf{P}\varphi \leftrightarrow \neg\mathsf{O}\neg\varphi$, for all w^* in W, and any formula φ in $\mathscr{L}(\mathsf{At})$. We nevertheless include the satisfiability conditions for both O and P for convenience.

While we think there is a lot to recommend switching over to neighborhood semantics, neither our XBox-based solution to the consistency challenge nor the triviality challenge to the modal moral error theory depends on it. Each can be configured to run natively on a specific semantics in the class of neighborhood models.

3 | HOW TO MEET THE CONSISTENCY CHALLENGE: THE XBOX

With this general framework of modal semantics in hand, let us turn to the consistency challenge, which requires error theorists to deliver a deontic semantics that is consistent with their denial of moral truths by returning the truth-value *false* to all deontic moral sentences. Here, then, is how we propose that error theorists should meet this challenge.[10]

Remember that the principal reason for the error theory's incompatibility with standard deontic semantics lies in the latter's commitment to the dual schema. To get around this problem, the heart of our proposal is for error theorists to adopt a *two-sorted* modal language. Our proposal is simple: add to the basic modal language $\mathscr{L}(\mathsf{At})$ a new monadic modal operator, $\boxtimes$, called the XBox. The XBox is charged with the duty to make all moral modal propositions false. Although picking out moral modal propositions for special semantic treatment is a metalinguistic operation, introducing a two-sorted language allows the XBox to perform this particular task in the object language. What differentiates the XBox from ordinary box modalities like O is that the grammatical rule for forming XBox formulas is restricted to formulas of $\mathscr{L}(\mathsf{At})$. Informally, this means that our new modal language with the XBox simply consists of all the formulas of the basic modal language, $\mathscr{L}(\mathsf{At})$, plus, for each formula φ in the basic modal language $\mathscr{L}(\mathsf{At})$, a formula of the form $\boxtimes\varphi$. Technically, this means that the new language with the XBox is sorted into two types, and the formation rules of the language are restricted by type. Thus, the well-formed formulas χ of this new modal language are given by the sorted grammar

$$\varphi : p \mid \bot \mid \neg\phi \mid \phi \vee \psi \mid \mathsf{O}\phi$$
$$\chi : \varphi \mid \boxtimes\varphi$$

where $p \in \mathsf{At}$ and the variables ϕ and ψ are as before, φ denotes the formulas of our basic modal language $\mathscr{L}(\mathsf{At})$, and χ is (effectively) any formula of the basic modal language alone or prepended by a single XBox, $\boxtimes$. Call the smallest language generated by this grammar $\mathscr{XL}(\mathsf{At})$.

The austere rules for constructing XBox formulas resolve the consistency problem with the dual schema because there is no dual of the XBox. Although $\boxtimes\varphi$ is a well-formed formula schema, $\neg\boxtimes\neg\varphi$ is not, nor is $\boxtimes\chi$. Thus, there are no negated XBox formulas, no nested XBoxes inside other modal operators, and no nested XBoxes inside other XBoxes. XBoxed formulas are not conjoined to any other formula, either, although this condition could be relaxed. The XBox has a single syntactic role, namely to prepend any well-formed formula from our original deontic modal language. Note, though, that although these construction rules preclude a dual of the XBox, they leave the dual schema of standard deontic semantics untouched. As such, introducing the XBox allows us to declare that all moral deontic sentences are false by evaluating both $\boxtimes\mathsf{P}\varphi$ and $\boxtimes\mathsf{O}\varphi$ as false without having to deny, as Streumer (2017, 124–128) and Olson (2014, 11–15) do, that $\mathsf{P}\varphi \leftrightarrow \neg\mathsf{O}\neg\varphi$ and that negated deontic moral sentences such as $\neg\mathsf{O}\varphi$ and $\neg\mathsf{P}\varphi$ are moral. Our proposal thus renders the maneuvers suggested by Streumer and Olson unnecessary.

Intuitively, the satisfiability conditions for the XBox are very simple as well. Prepending an XBox to a formula φ equips error theorists with 'X-ray vision' to see inside the formula φ and

turn all and only the offending moral modal subformulas within φ false whilst leaving the truth-values of all other formulas unaltered. This basic approach allows for some customization. For example, some error theorists target all normative discourse, sweeping up moral discourse in the lot, whereas others single out moral discourse specifically and leave other normative language alone.[11] The basic semantics we present here will focus on how to single out moral modal discourse for special treatment.

While our intuitive description of the semantics sounds simple enough, there is a complication. Although we refer to modal and non-modal *subformulas*, this distinction is not reflected directly by the grammar. Instead, each well-formed formula of the language is built from applying a finite sequence of our grammatical rules, which can place modal subformulas under negations, within clauses, or nested under other modalities. The upshot is that a simple rule that replaced each modal subformula with *falsum* might flip to *verum* in the end if they appear under a wider-scope negation. Fortunately, there is a workaround.

We say that a modal formula is a *negated normal form* (NNF) if it uses only the logical connectives $\wedge$, $\vee$, and $\neg$, and negations only apply to atomic propositions. Additionally, a modal NNF may have either O or P operators. Any modal formula can be translated into a logically equivalent modal NNF by the following equivalences:

1. $\phi \rightarrow \psi \equiv \neg\phi \vee \psi$
2. $\neg(\phi \wedge \psi) \equiv \neg\phi \vee \neg\psi$
3. $\neg(\phi \vee \psi) \equiv \neg\phi \wedge \neg\psi$

4. $\neg\neg\phi \equiv \phi$
5. $\neg O\phi \equiv P\neg\phi$
6. $\neg P\phi \equiv O\neg\phi$

Intuitively, translating a modal formula into a modal negated normal form is achieved by applying the above rules to 'push' any wide-scope negations that appear in the original formula through until they attach to atomic propositions. Alternatively, since every well-formed formula is built up from atomic propositions by applying some finite number of rules of the grammar, the construction of every formula may be represented by its parse or ancestral *tree* (Enderton, 2001). The distinguishing characteristic of a modal NNF parse tree is that negations only appear, if they appear at all, directly above the leaves of the tree, which consist of atomic propositions.

For any φ of $\mathcal{XL}(At)$, suppose $\tilde{\varphi}$ is a NNF of φ and $\tilde{\varphi}^*$ denotes the formula obtained from the formula $\tilde{\varphi}$ by replacing every modal subformula of either the form $O\phi$ or $P\phi$, wherever each occurs in $\tilde{\varphi}$, by the propositional constant $\bot$. With this bit of machinery, we introduce the satisfiability condition for XBox in terms of this *modality-free NNF*, $\tilde{\varphi}^*$.

7. $\mathbb{M}, w \vDash \boxtimes\varphi$ if and only if $\tilde{\varphi}^*$ is a modality-free NNF of φ and $\mathbb{M}, w \vDash \tilde{\varphi}^*$

What the XBox satisfiability condition says is that the XBox of φ is true at state w just in case a negated normal form ($\tilde{\varphi}$) of φ whose modal subformulas have all been replaced by falsum ($\tilde{\varphi}^*$) is true. So, while syntactically the XBox is a modal operator, semantically it 'flattens' any modal formula φ it prefixes to a locally satisfiable propositional formula of the form $\tilde{\varphi}^*$. In other words, what is XBoxed at w is evaluated at w: no other element of the modal base W is involved. Even so, since condition (7) is applicable to every w in W, the XBox faithfully represents the error theorist's contention that moral claims are false at all possible worlds.[12]

The XBox satisfiability condition (7) and the six conditions we presented in Section 2 constitute the general neighborhood semantics for interpreting formulas in the augmented language $\mathcal{XL}(At)$.

Example 1. Consider the sentence *'If Boris has colleagues then Boris is not permitted to disparage them'*, symbolized by φ:

$$c \to \neg \mathsf{P}d$$

and whose negated normal form is $\tilde{\varphi} := \neg c \vee \mathsf{O}\neg d$ and modal-free negated normal form is $\tilde{\varphi}^* := \neg c \vee \bot$. The table below displays each formula along with its corresponding syntactic parse tree.

Then, $\boxtimes(c \to \neg \mathsf{P}d)$ is true at w in $\mathbb{M}$ if and only if c is false at w. Thus, the XBoxed assertion that *if Boris has colleagues then Boris is not permitted to disparage them* is true just in case Boris does not have colleagues.

To sum up, then, the XBox is a device for rendering all modal moral formulas within its scope false, and reducing the truth-conditions of a modal formula to the truth-conditions of its non-modal, propositional subformulas. If there are no such non-modal subformulas, then the formula is false. Since we interpret the XBox as a device to model the thesis that all of our moral modal assertions are false, the error theorist will seek to XBox all of our moral claims, whereas their opponents will aim to never do so. Each, nevertheless, should be interested in the use of the Xbox. For, since the XBox is built atop a range of deontic logics interpreted by some or another neighborhood model, our semantics also serves the more general purpose of allowing one to compare the logical consequences and semantic commitments of moral modal formulas expressed in that language with an XBox and without. Hence, whilst our proposal might not be the only way in which the consistency challenge can be met, we believe that it provides an elegant solution which is not only extremely general, but also does not require revision of mainstream deontic semantics.

4 | THE TRIVIALITY CHALLENGE: LOST DIFFERENCES IN MEANING

With the XBox, error theorists now have a consistent framework with which to formulate their view. Pairing the XBox modal language $\mathcal{XL}(\mathrm{At})$ with neighborhood semantics provides an extremely general and flexible platform which can accommodate practically all mainstream semantics for deontic propositional logics.[13] Moreover, the consistency challenge to the error theory has been met without having to give up the dual schema of mainstream deontic semantics.

As we have explained above, though, meeting the consistency challenge does not imply that error theorists are out of the woods just yet. Remember that intensional modal semantics assign

the same meaning to any two sentences that are intensionally equivalent. Hence, as things stand, the XBox is not only a straightforward and minimal implementation of error-theoretical commitments, but also a triviality trap: By assigning the truth-value *false* to all deontic moral sentences, XBox semantics leads to the extremely implausible result that all deontic moral sentences—such as 'It is impermissible for politicians not to have a Plan B', 'Fred ought not to have lied', or 'It is permissible to eat meat'—mean exactly the same. Put differently, then, error theorists still face the triviality challenge, and thus the task of explaining in which way different moral deontic assertions can be seen to differ in meaning despite being intensionally equivalent.

The most obvious way to tackle this challenge is to consider leaving intensional semantics behind by connecting the XBox to a *hyperintensional* semantics. Unlike intensional semantics, hyperintensional semantics are designed to account for different meanings of sentences that have the same intension. Just as was the case with regard to standard deontic semantics, there is also a considerable variety of hyperintensional semantics on offer. Historically, these are often heavily involved in the philosophical debate about the metaphysics of propositions, and thus tie their formal proposals to philosophical positions about which properties of intensional semantics ought to be weakened or abandoned altogether. For instance, some approaches to hyperintensional contexts propose extending the range of states that are semantically admissible to include impossible worlds alongside possible worlds.[14] Others propose to count how a proposition is syntactically represented as part of the semantic meaning of a formula, either to represent reasons for asserting the formula or to account for the resources required to verify it.[15] Another approach takes the semantic content of a formula to be something other than a set of states in which that formula is true, either by focusing on Russellian propositions or some other structured objects.[16]

However, there is a problem. For, despite their generality, neighborhood models and their associated modal logics do not support hyperintensional modalities. In fact, the only non-negotiable modal inference rule that is valid with respect to every neighborhood model is a rule stating that logically equivalent formulas are equally obligatory.[17] Put differently, while the flexibility of neighborhood semantics might lull us into believing that anything goes for modal operators, intensionality is non-negotiable.

Fortunately, though, Igor Sedlár has recently provided us with a way around this problem by extending neighborhood semantics in a manner that supports hyperintensional modals (Sedlár, 2019). More precisely, he introduces a generalization of neighborhood models and proves completeness results for a large family of hyperintensional monadic modal logics supported by his models. His framework features a set of semantic *contents* that are assigned to formulas and associated with states, with neighborhood models appearing as a special case when those semantic states are identical to the power set of the set of worlds used in our specification of neighborhood models above. The details of Sedlár's framework can be set aside here. Instead, what makes his framework particularly attractive for our purposes is its great generality, in that his family of hyperintensional modal logics is constructed without a need to pick a side in the philosophical debate about the metaphysics of propositions. Instead, his framework is neutral towards the different hyperintensional semantics currently on offer, in that he shows how each position's commitments can be faithfully expressed within his formalism. As such, Sedlár's framework for hyperintensional modal logics is similar in spirit to our proposal for the XBox: each is very general in scope, if not the most general account of its kind currently available.

Sedlár thus presents a very appealing approach to hyperintensional modals in response to the triviality challenge. Unfortunately for error theorists, though, they cannot employ it. The reason why is very simple: The price of using Sedlár's very general hyperintensional semantics in order to work out an account of moral modalities that renders them false but non-equivalent in meaning

is to accept that any hyperintensional model for an XBox will, necessarily, also be a model for an 'XDiamond' defined as ¬ ⊠ ¬φ. In fact, the logic of *all* hyperintensional monadic modal logics is simply classical propositional logic augmented by the dual schema (Sedlár, 2019, Theorem 1). So, if the XBox were to be upgraded to a monadic hyperintensional modal in order to meet the triviality challenge, then there would also be a hyperintensional XDiamond. In a nutshell, Sedlár's monadic hyperintensional modal logic would require that *both* ⊠φ and ¬ ⊠ ¬φ are well-formed formulas.

However, recall that in order to meet the consistency challenge and resolve the error theory's incompatibility with standard deontic semantics, the grammatical rules for constructing XBox formulas explicitly *prohibit* a dual of the Xbox: that is, ¬ ⊠ ¬φ is *not* a well-formed formula of $\mathscr{XL}$(At). Accordingly, what allows error theorists to address the consistency challenge—namely, the assumption that the XBox is without a dual—is precisely what prevents them from meeting the triviality challenge through a hyperintensional modal semantics.

5 | A SEMANTIC DILEMMA FOR THE MORAL ERROR THEORY

Let us recap. We confronted the error theory with the semantic challenge to deliver a deontic semantics that renders all deontic moral sentences false—what we have called the 'consistency challenge'—while accounting for the fact that deontic moral assertions differ in meaning—what we have called the 'triviality challenge'. We demonstrated how to meet the consistency challenge by introducing the XBox operator, which achieves compatibility between the error theory and deontic semantics in very simple and general terms.

At the same time, based on Sedlár's very general framework for hyperintenstional modalities, we showed that meeting the consistency challenge within mainstream modal semantics precludes meeting the triviality challenge and *vice versa*: The XBox prohibits a dual, while a hyperintensional XBox would go through only if a dual in the form of an XDiamond were admitted. As a result, our semantic challenge turns into a semantic dilemma for the error theory: Either, the moral error theory can be consistently formulated, but entails the meaning-equivalence of all deontic moral assertions. Or it can account for the different meanings of deontic moral assertions, but is now straightforwardly inconsistent. Either way, the modal moral error theory is in trouble.

To be clear, in this paper we have not provided an impossibility result. That is, we have not shown that it is *impossible* for modal moral error theorists to rise to our semantic challenge by finding a way out of this dilemma. However, let us close by stressing just how difficult a task this turns out to be.

For instance, as was pointed out earlier, it would be a mistake to think that the semantic dilemma is driven by the dual schema, such that abandoning this schema by appealing to a tradition in deontic logic that rejects treating O and P as duals will stop our challenge in its tracks. Firstly, remember that our XBox-based solution to the consistency challenge does not even require giving up the dual schema in mainstream deontic semantics. Rather, the XBox succeeds in securing consistency for the moral error theory while accepting that O and P are modal duals. Accordingly, it is not the dual schema of mainstream deontic semantics that generates the semantic dilemma, but rather the observation that any hyperintensional operator requires a dual.

Secondly, even if moral error theorists sought refuge in non-standard deontic semantics, it is not a given that these semantics are compatible with the error theory. For instance, one such approach in this tradition attempts to reduce deontic modals to standard, alethic modal logic by taking Oφ to express the conditional $\Box(G \rightarrow \varphi)$, which says that you are obligated to φ just in case it is necessarily the case that every state that is good (G) is one in which φ is satisfied, and

defining $P\varphi$ as $\Box(\varphi \to G)$, which says that it is necessary that either φ is not satisfied in a state or that state is good (van Benthem, 1979). Yet, while this reductionist strategy may appear to offer an escape route from our consistency challenge, it is not open to the error theorist. For, note that the formula $\Box(G \to \varphi)$ is false at a state w just in case every neighborhood $N(w)$ is such that G is true but φ is false, which would commit the error theorist to claiming that there are morally good states—a claim she is committed as an error theorist to reject.

Thirdly, while there are other proposals in this tradition that an error theorist might want to consider, let us stress once more that the aim of all of these non-standard deontic logics is to specify *how* O and P are logically related to one another, and not to deny *that* they are logically related. As long as the satisfiability conditions for *permissibility* and *obligation* remain logically related, though, it will always be the case that the falsity of some deontic sentence will entail the truth of another. As a result, our consistency challenge simply resurfaces. Consequently, although we have appealed to the dual schema when setting up the consistency challenge, our point does not depend on it. Instead, our semantic challenge should be viewed as a template that can be adapted to a wide range of deontic semantics, including those that reject that *obligation* and *permissibility* are modal duals. Hence, denying that O and P are modal duals is neither necessary nor sufficient to dissolve the semantic dilemma.

Finally, could moral error theorists seek to react to our challenge by going one step further, namely by denying that the satisfiability conditions for O and P are logically related at all?[18] Note that merely issuing such a denial is not sufficient for meeting our semantic challenge. Rather, error theorists need to present a plausible deontic semantics that severs all logical relations between O and P, successfully secures consistency for the error theory and avoids the triviality challenge. Since we know of no deontic semantics that meets these objectives and treats modal operators as logically independent, we would be curious to see one.

Alternatively, could moral error theorists respond to the generality and comprehensiveness of our semantic challenge not by seeking different interpretations of O and P, but by denying that *obligation* and *permissibility* should be treated as modal operators in the first place? That is, could they introduce a *bifurcation* in the treatment of notions such as *necessity* and *possibility* on the one hand and *obligation* and *permissibility* on the other, such that the former are taken to function as modal operators, whereas the latter are not?

Firstly, note that this maneuver would not come with a built-in guarantee to succeed as a response to our semantic challenge. Rather, moral error theorists would still have to demonstrate how such a non-modal semantics of *obligation* and *permissibility* is both compatible with the error theory and able to account for the different semantic content of moral assertions. Secondly, treating *ought* and *permissible* as discontinuous with modals such as *must* and *might* would amount to a radical departure from the established approach to deontic notions that is shared across logic, linguistics, and metaethics alike. Accordingly, if error theorists felt forced to adopt such a bifurcatory strategy in response to our semantic challenge, this would be a very significant result in itself. So far, moral error theorists have assumed, at least implicitly, that they can employ the same semantic account as other representationalists. Now, though, it would be apparent that unlike other representationalists, error theorists cannot adopt a unified approach to *obligation*, *permissibility* and other modals, but instead require a non-standard, non-modal semantics of moral deontic expressions tailor-made for the error theory. Such a non-modal approach would, moreover, come with its own severe problems, ranging from questions in compositional semantics, to the interpretation and unifying explanation of *ought* and *must* together with their different flavors, to the plausibility of potential definitions of *ought* and *permissible* in terms of reasons or value, say.[19] Hence, besides

having to deal with the semantic challenge presented here, opting for bifurcation would burden error theorists with the additional demand of having to address the numerous objections faced by non-modal approaches to *obligation* and *permissibility*.

Finally, assume that moral error theorists chose to react to our semantic challenge not by trying to meet it head-on, but by undermining it. More precisely, imagine that they attempted to turn tables by declaring that if deontic semantics is not compatible with the moral error theory, so much the worse for deontic semantics. After all, the moral error theory is clearly conceivable and conceptually possible. Our semantic challenge, in turn, seeks to rule out this theory by semantic fiat, and thus tries to pull a 'metaphysical rabbit' out of a 'semantic hat' (Loeb, 2008) by deriving metaphysical conclusions from semantic premises alone. Yet, any such endeavor is clearly hopeless. Hence, our semantic challenge collapses, or so the error theorist now argues.

It is correct that if the modal moral error theory could not meet our challenge and thus had to be rejected on semantic grounds, it would follow conceptually that at least one moral claim is true. And indeed, this might be a surprising result. Still, we should not dismiss it out of hand.

Firstly, note that this result would not necessarily entail robustly metaphysical implications about moral reality. For, truth-conditional deontic semantics is arguably also compatible with positions such as metasemantic expressivism and inferentialism, and thus views which reject the existence of robustly metaphysical moral facts and properties whilst accepting the existence of moral truths.[20] Hence, even if our semantic arguments led to the result that it is conceptually true that not all moral claims are false, this would not conjure into existence any metaphysical 'rabbits' in the form of metaphysically robust moral reasons and properties. Secondly, note that despite being surprising, this result might not be as troubling as it might seem at first sight. For instance, from our treatment of the *necessity* and *possibility* operators, $\Box$ and $\Diamond$, it also follows that it is conceptually true that some modal claim is true. Yet, we do not appear to be overly concerned about this finding. Moreover, instead of having to accept these results as brute and unexplainable semantic facts, it might well be the case that at least certain accounts of normative and non-normative modals will be able to explain why these conceptual truths hold, say by appealing to the specific function that modal terms play within our language and practices.

Finally, and most importantly, note that our semantic challenge to the moral error theory does *not* seek to end metaethical debate by semantic fiat. Rather, it should be understood as exactly that: a *challenge* to back up the claim that the error-theoretic position is both consistent and semantically plausible. After all, the fact that a position seems to be coherent does not entail that it *is* coherent. Accordingly, in order to meet our challenge, it is not sufficient merely to *declare* that the moral error theory is consistent and conceptually possible. Rather, error theorists must *demonstrate* their position to be so by presenting a semantics which is both compatible with their position and able to account for the different semantic contents of deontic sentences. As such, the onus of developing a semantics for their theory that can escape our semantic dilemma is squarely on error theorists. Unless they have discharged this task, there is no reason to trust that the modal moral error theory can provide a coherent and semantically plausible account of deontic moral language.

ENDNOTES

[1] There are different versions of the error theory. Some reject the existence of moral properties on conceptual grounds (Loeb, 2008; Streumer, 2017), others deny their existence on the basis of metaphysical considerations such as metaphysical queerness (Joyce, 2001; Mackie, 1977; Olson, 2014). Some error theorists take moral claims to conceptually entail that moral properties or categorical reasons exist (Joyce, 2001; Mackie, 1977; Olson, 2014), others understand moral language's commitment to moral facts in terms of semantic or pragmatic presupposition (Kalf, 2018). None of these differences matter for our purposes. The only distinction that is of some relevance to our arguments is that between what we can call the 'necessary' and the 'contingent' error theory (for a related discussion, see Olson, 2014, p.12, fn.17). We can understand the former as holding that moral claims are necessarily untrue, and thus untrue in all possible worlds. The latter, in turn, can be understood as submitting that whilst moral claims are untrue in the actual world and relevantly similar worlds, they might be true in others (Brown, 2013; Kalf, 2015). In what follows, we develop our semantic challenge in relation to the necessary moral error theory. However, in Section 3, footnote 12, we explain in which way a variant of our challenge also applies to the contingent moral error theory. Accordingly, adopting the error theory's contingent version does not provide an obvious way out of our semantic challenge.

[2] Moral error theorists, who are representationalists about moral language, have recently made limited forays into semantics by propounding specific theses about the absence of conceptual entailment relations between deontic terms such as 'wrong' and 'permissible' in an attempt to secure consistency for their position (Olson, 2014, 11–15; Pigden, 2007, 450–454; Streumer, 2017, 124–128). At the same time, they have not as yet provided a general semantics of deontic notions that would systematically specify how the meaning of these terms is to be understood.

[3] For a related argument, see also (Tiefensee, 2020).

[4] Appendix A includes some brief remarks about simulating Kripke's relational semantics and Kratzer's semantics for graded modalities within neighborhood models.

[5] Some error theorists, such as Streumer (2017) and Olson (2014), reject the dual schema by denying that $O\varphi$ ($P\varphi$) conceptually follows from $\neg P\neg\varphi$ ($\neg O\neg\varphi$). Similarly, there is a tradition in deontic logic that denies the dual schema (see, for example, van Benthem (1979)). As we explain here, though, just as the consistency challenge does not depend on acceptance of the dual schema, our response to this challenge does not require its rejection.

[6] It might be argued that we should not adopt the error theory's standard formulation, according to which all moral assertions are *false*, but its non-standard formulation, according to which moral assertions are *neither* true *nor* false (Kalf, 2018). Yet, since mainstream deontic semantics is based on classical logic and thus does not cater for truth-value gaps, it should be obvious that the non-standard formulation of the error theory is no more compatible with mainstream deontic semantics than its standard formulation. Simulating the departure from bivalence within some multi-modal system in order to capture the non-standard formulation of the error theory, in turn, encounters serious limitations (Wheeler & Alberti, 2011). Accordingly, adopting the error theory's non-standard formulation in response to our semantic challenge is not an avenue that we would advise error theorists to pursue. As such, we will follow the standard formulation here by understanding the error theory as holding all moral assertions to be false. For different formulations of the error theory, see Olson (2014, 11–15).

[7] It might be suggested that this problem also applies to all non-error-theoretic representationalists who hold that certain true (false) moral claims are necessarily true (false), such that error theorists could simply latch on to whichever solution these other representationalists may provide to this problem. However, non-error-theoretic representationalists and error theorists do not face the same semantic challenge, nor can they offer the same solution to their respective challenges, exactly because the former but not the latter hold that there are necessarily true as well as necessarily false moral claims. As such, our triviality challenge for the error theory is separate from observations about the familiar incapacity of intensional semantics to differentiate, semantically, among propositions that are necessarily true and those that are necessarily false, respectively. For, whilst this incapacity can arguably be addressed by non-error-theoretic representationalists through turning to hyperintensional semantics, we shall argue that for principled reasons, the error theorist cannot avail herself of *any* existing hyperintensional semantics in her response to the triviality challenge exactly because of her error-theoretic commitment to the thesis that *all* moral assertions are necessarily false. We thank an anonymous referee for *Noûs* for asking us to render this more explicit.

8   It might be hoped that rather than having to introduce new operators within our semantics, a little tweak to mainstream deontic semantics will suffice for error theorists to meet the consistency challenge. Why this hope is frustrated is explained in (Tiefensee, 2020).

9   Neighborhood models, proposed by Scott (1970) as a generalization of Kripke's relational semantics, have been applied to deontic logic and the logic of conditionals (Chellas, 1980), games (Parikh, 1985), alternating time (Alur et al., 1992), epistemic modals (Kyburg & Teng, 2001), and coalitions (Pauly, 2002), among others. For a contemporary introduction to neighborhood semantics, see Pacuit (2017).

10 One may view our proposal as a positive general framework for error theorists which imposes minimal constraints on theory construction. That we later use this generality to widen the scope of our triviality argument against moral error theories is a separate issue.

11 For a general normative error theory, see Streumer (2017). For error theories restricted to moral discourse, see Mackie (1977), Olson (2014), and Joyce (2001).

12 In footnote 1, we distinguished between necessary and contingent versions of the moral error theory and have so far focused on the former. Why a variant of our semantic challenge also applies to the contingent error theory becomes apparent once we ask which deontic semantics contingent error theorists could employ to deliver their desired result, namely that all sentences of the from $O\varphi$ and $P\varphi$ are false at the 'actual' and relevantly similar worlds but true at others. Mainstream deontic semantics clearly do not fit this bill. Nor does our XBox. For instance, in order to allow moral claims to be 'contingently' true in some states but false in others, one might modify (7) by restricting the applicability of (7) to a proper subset of the modal base, W. However, this simple proposal makes the satisfiability conditions of a formula depend on the state in which that formula is evaluated, which is not how modal logic works. For, although the interpretation of modal formulas proceeds in terms of relationships among states in W, the syntax of modal logic does not reference states. Hence, there is no way for a modal language to refer to states to determine whether or not a modal formula is to be XBoxed. Next, the standard method to address this inability to denote states of a model is *hybrid modal logic* (Areces & ten Cate, 2007), which adds a set of atomic names to a modal language that denote states and a new modal operator that applies to those names of states. However, hybrid modal logics are proper extensions of modal logics, so they validate the dual schema. More importantly still, hybrid modal languages are constructed to express whether or not a formula is evaluated at a particular state, which at minimum requires contingent error theorist to adopt a language expanded to include negated XBox formulas, thus exposing the contingent error theorist to the consistency problem anew. Hence, contingent error theorists run into problems with regard to consistency in combination with contingency *for the very same reason* that explains why necessary error theorists run into trouble with regard to consistency in combination with triviality: *Necessary* error theorists' attempt to address the *triviality* challenge presupposes precisely what meeting the consistency challenge precludes, or so we argue. We can now see that *contingent* error theorists' attempt to secure *contingency* for their position also presupposes precisely what meeting the consistency challenge precludes. Both times, the assumption being presupposed is that the XBox allows for negation. In addition, we can foresee that certain ways to secure contingency—such as the attempt to restrict the modal base to which the XBox applies to a specific set of states—will re-generate the triviality challenge even for the contingent error theory. Still, let us stress that while we think that these problems are serious, we are not claiming to have shown here that it is *impossible* to develop a deontic semantics that simultaneously (a) caters for contingency, (b) returns the truth-value *false* to all deontic moral claims in the actual and relevantly similar worlds, and (c) does so in a way that accounts for moral claims' different meanings. However, what these brief considerations do show is that adopting a contingent error theory does not provide an obvious way out of our challenge. Rather, if error theorists chose to turn to the contingent error theory in response to our semantic challenge—which would be a very significant result in itself, given that doing so would rule out prominent error-theoretic positions and arguments—the onus would be firmly on them to spell out a deontic semantics that successfully meets (a)–(c). We thank a referee for this journal for pressing us on this point.

13 Indeed, it arguably provides *the* most general platform currently on offer. At the very least, neighborhood semantics cover the two most popular intensional semantics for deontic modals.

14 Examples include Hintikka (1975), Barwise and Perry (1983) and Jago (2015).

15 Examples include Fagin and Halpern (1988), Artemov (2008) and Wheeler (2018).

16 See for instance Lewis (1970), Soames (1987) and King (2007).

17 This rule is called RE and says that if $\phi \leftrightarrow \psi$ is a theorem, then $O\phi \leftrightarrow O\psi$ is a theorem, when O is a box modal. See Chellas (1980).

[18] For this strategy, compare Pidgen's (2007), Olson's (2014), and Streumer's (2017) considerations on conceptual entailment relations.

[19] See Chrisman (2016), chapter 2 for an assessment of non-modal interpretations of *obligation* and *permissibility* and a strong defence of the modal approach to deontic notions.

[20] For truth-conditional semantics' combination with metasemantic expressivism and inferentialism, see (Chrisman, 2016; Köhler, 2018; Ridge, 2014; Schroeter & Schroeter, 2018; Silk, 2013).

REFERENCES

Alur, R., Henzinger, T. A., & Kupferman, O. (1992). Alternating-time temporal logic. *Compositionality: The significant difference*, number 1536 in LNCS (pp. 23–60). Springer.

Areces, C., & ten Cate, B. (2007). Hybrid logics. In P. Blackburn, F. Wolter, & J. van Benthem (Eds.), *Handbook of Modal Logic* (pp. 821–868). Elsevier.

Arló-Costa, H., & Taysom, W. (2005). Contextual modals. In A. Dey, B. Kokinov, D. Leake, & R. Turner (Eds.), *CONTEXT 2005: Modeling using context*, volume 3554 of *Lecture Notes in Computer Science* (pp. 15–28). Berlin: Springer.

Artemov, S. (2008). The logic of justification. *The Review of Symbolic Logic*, *1*(4), 477–513.

Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: MIT Press.

Brown, P. (2013). The possibility of morality. *Philosophical Studies*, *163*, 627–636.

Chellas, B. (1980). *Modal logic*. Cambridge: Cambridge University Press.

Chrisman, M. (2016). *The meaning of 'ought': Beyond descriptivism and expressivism in metaethics*. New York: Oxford University Press.

Enderton, H. B. (2001). *A Mathematical introduction to logic* (2nd ed.). San Diego, CA: Harcourt Academic Press.

Fagin, R., & Halpern, J. Y. (1988). Belief, awareness, and limited reasoning. *Artificial Intelligence*, *34*, 39–76.

Goble, L. (2005). A logic for deontic dilemmas. *Journal of Applied Logic*, *3*(3–4), 461–483.

Hansen, H. H. (2003). Monotonic modal logics. Research report PP-2003-24, University of Amsterdam, Institute for Logic, Language and Computation.

Hintikka, J. (1975). Impossible worlds vindicated. *Journal of Philosophical Logic*, *4*(4), 475–484.

Jago, M. (2015). Impossible worlds. *Noûs*, *49*(4), 713–728.

Joyce, R. (2001). *The Myth of morality*. Cambridge: Cambridge University Press.

Kalf, W. (2015). Are moral properties impossible? *Philosophical Studies*, *172*, 1869–1887.

Kalf, W. (2018). *Moral error theory*. Cham, Switzerland: Palgrave McMillan.

King, J. C. (2007). *The nature and structure of content*. Oxford: Oxford University Press.

Köhler, S. (2018). Expressivism, meaning, and all that. *Canadian Journal of Philosophy*, *48*, 337–356.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, *1*(3), 337–356.

Kratzer, A. (1991). Modality. In H. Steger, & H. E. Wiegand (Eds.), *Semantics: An international handbook of contemporary research* (pp. 639–650). Berlin: Walter de Gruyter.

Kratzer, A. (2012). *Modals and conditionals* (2nd ed.). Oxford: Oxford University Press.

Kripke, S. (1963). Semantical analysis of modal logic i. normal propositional calculi. *Zeitschrift fur mathematische Logik und Grundlagen der Mathematik*, *9*(5–6), 67–96.

Kyburg, Jr, H. E., & Teng, C. M. (2001). *Uncertain inference*. Cambridge: Cambridge University Press.

Lewis, D. (1970). General semantics. *Synthese*, *22*(1), 18–67.

Lewis, D. (1981). Ordering semantics and premise semantics for conditionals. *Journal of Philosophical Logic*, *10*(2), 217–234.

Loeb, D. (2008). Moral incoherentism: How to pull a metaphysical rabbit out of a semantic hat. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 355–386). Cambridge, MA: MIT Press.

Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. New York: Penguin Books.

Olson, J. (2014). *Moral error theory: History, critique, defence*. Oxford: Oxford University Press.

Pacuit, E. (2017). *Neighborhood semantics for modal logic*. Cham, Switzerland: Springer.

Parikh, R. (1985). The logic of games and its applications. *Topics in the theory of computation*, number 24 in Annals of Discrete Mathematics. Elsevier Science.

Pauly, M. (2002). A modal logic for coalitional power in games. *Journal of Logic and Computation*, *12*(1), 149–166.

Pigden, C. R. (2007). Nihilism, Nietzsche and the doppelganger problem. *Ethical Theory and Moral Practice*, *10*, 441–456.

Ridge, M. (2014). *Impassioned belief*. Oxford: Oxford University Press.

Schroeter, L., & Schroeter, F. (2018). Metasemantics and metaethics. In T. McPherson, & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 519–535). New York: Routledge.

Scott, D. (1970). Advice on modal logic. In K. Lambert (Ed.), *Philosophical problems in logic* (pp. 143–173). Dordrecht: Reidel.

Sedlár, I. (2019). Hyperintensional logics for everyone. *Synthese*.

Silk, A. (2013). Truth conditions and the meanings of ethical terms. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 8, pp. 195–222). Oxford: Oxford University Press.

Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, *15*, 47–87.

Speaks, J. (2018). Theories of meaning. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2018 ed.). Metaphysics Research Lab, Stanford University.

Streumer, B. (2017). *Unbelievable errors*. Oxford: Oxford University Press.

Tiefensee, C. (2020). 'Ought' and error. *Journal of Philosophy*, *117*(2), 96–114.

van Benthem, J. (1979). Minimal deontic logics (abstract). *Bulletin of the Section of Logic*, *8*(1), 115–121.

Wheeler, G. (2018). Bounded rationality. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2018 ed.). Metaphysics Research Lab, Stanford University.

Wheeler, G., & Alberti, M. (2011). NO revision and NO contraction. *Minds and Machines*, *21*(3), 411–430.

APPENDIX A: FINDING NEIGHBORHOODS FOR KRIPKE AND KRATZER

Kripke's relational semantics for normal modal logic (Kripke, 1963) remains the best known modal semantics. A *standard Kripke model* for the basic modal language $\mathscr{L}(\text{At})$ is a triple, $\mathscr{M} = (W, R, V)$, where W is a non-empty set of states and V is a valuation function which assigns a set of states to each atomic proposition, $\text{At} \mapsto \wp(W)$, and R is a binary relation on W (i.e., $R \subseteq W \times W$).

Briefly, a binary relation on W can be understood as a particular kind of neighborhood function, namely one that satisfies the properties of 'R-necessity' (Pacuit, 2017, Definition 2.8). In addition to Pacuit's discussion of how to simulate Kripke models within a neighborhood model, we recommend Hansen's study of monotonic modal logics and the associated class of *supplemented* neighborhood models (Hansen, 2003). Monotonic modal logics reside in a Goldilocks zone between fully classical modal logics, which are the most general but weakest monadic modal logics, and the more restrictive normal modal logics and their associated Kripke relational models. Monotone modal logics have nice logical properties, but are more flexible than Kripke-style normal modal logics. Unlike normal modal logic, for instance, monotone modal logics allow for moral dilemmas (Goble, 2005): it is not a theorem of monotone modal logics that a conjunction of obligations imply an obligation to all the conjuncts. Also, their semantics, the class of supplemented neighborhood models, $\mathbb{M}^s$, afford very elegant and intuitive satisfiability conditions for box and diamond modals, respectively:

5^s. $\mathbb{M}^s, w \vDash O\varphi$ iff $(\exists X \in \mathcal{N}(w), \forall v \in X) : \mathbb{M}^s, v \vDash \varphi$

6^s. $\mathbb{M}^s, w \vDash P\varphi$ iff $(\forall X \in \mathcal{N}(w), \exists v \in X) : \mathbb{M}^s, v \vDash \varphi$

Informally, condition (5^s) says that Oφ is satisfied just in case there is some set of states X in the neighborhood such that every state in X satisfies φ, whereas (6^s) says that Pφ is satisfied just in case every set X in the neighborhood is one that has at least one state where φ is satisfied.

§

Kratzer's semantics presents a few complications, but only because there are some clarifications that need to be made before embarking on locating the right neighborhood.

Kratzer has proposed a two-dimensional semantics to account for the role that a *conversational background* plays in determining the meaning of modal particles in natural language (Kratzer, 2012). A conversational background specifies a context, which is represented in the semantics by a state-dependent set of propositions determined by a function from every state $w \in W$ to a subset of the power set of W. Kratzer's semantics features two such state-dependent functions. The first picks out the relevant states to use to evaluate a modal formula, much as a standard Kripke model does. The other function, which Kratzer calls an *ordering source*, imposes an ordering on the states selected by the first function. Whether this too can be captured within a standard Kripke model depends on some details we will return to in a moment. The reason for imposing an ordering on states is to allow the semantics to handle *graded modalities*, thereby providing satisfiability conditions for comparative modal claims such as 'ϕ is better than ψ' or 'ϕ is more likely than ψ'. Call this scheme a *two-context* model.

There are a variety of ways to flesh out the formal semantics of a two-context model, depending on the properties you impose on each function. Unfortunately, there is some disagreement over what precisely Kratzer's two-context model is committed to. One source of confusion is due to David Lewis's misuse of the term 'partial order' (reflexive, antisymmetric, transitive) to refer to a preorder (reflexive, transitive) (Lewis, 1981, p. 220), which Kratzer regrettably follows (Kratzer, 1991, p. 644). When an ordering source is partially ordered, and there is a minimal state in the ordering source, then there are equivalent augmented neighborhood models (Arló-Costa & Taysom, 2005, Theorems 2 & 3). So, when the ordering is partial and has a bottom element, there is a Kripke model for this class of two-context models, too. If there is no minimal state in a partially ordered ordering source, then Kratzer semantics is equivalent to supplemented models closed under finite intersection and containing the unit (i.e., a *filter*) (Arló-Costa & Taysom, 2005, Theorems 4).

A second source of confusion is due to Kratzer's apparent focus on classes of neighborhoods that satisfy the properties of a filter. Although this class of models is slightly more general than Kripke models can capture, arguably they are not general enough. For when the domain of states W is finite, there are counterpart Kripke models. Moreover, a consequence of Kratzer's definition is that the range of the first function—which selects the set of states to be ordered—is a non-empty set. Yet, many of the natural language examples used to motivate her semantics do not have this feature. Thus, although the class of supplemented neighborhood models standing behind Kratzer's semantics is more general than standard Kripke models, focusing on models that satisfy the properties of filters is not general enough, for those models are ill-suited to representing deontic modals due to them all being closed under finite intersection and validating schema (**C**), O$\phi \wedge$ O$\psi \rightarrow$ O($\phi \wedge \psi$). In other words, within this class of models it is a theorem that a conjunction of obligations implies an obligation to all the conjuncts.

Rather than explore corresponding neighborhood models for variants of Kratzer's two-context semantics, one can instead begin to separate the role that ordering plays in the semantics from the role that a selection function plays for determining which states are to be used in the satisfiability

conditions. (Arló-Costa & Taysom, 2005), for example, introduces machinery from conditional logics and belief revision to add more sophisticated selection operations to model the semantics of conditional, comparative modal sentences. The upshot is that neighborhood semantics can capture the core features of Kratzer's two-context semantics while allowing us to go well beyond Kratzer's arbitrarily restrictive construction.