

## Nouwen’s Puzzle and a Scalar Semantics for Obligations, Needs, and Desires\*

Daniel Lassiter  
*Stanford University*

**Abstract** Nouwen (2010a,b) presents a puzzle involving the interaction of degree expressions and modals: a class of apparently unremarkable sentences expressing minimum and maximum requirements, obligations, desires, etc. are predicted to be trivially false or undefined, or receive otherwise incorrect truth-conditions. I suggest that the puzzle can be resolved if we treat the affected modals not as universal quantifiers over possible worlds but instead as scalar expressions which map propositions to points on a scale of expected utility. Independent arguments indicate that these modals are scalar, non-monotonic, and information-sensitive — facts which are highly problematic for quantificational theories, but predicted immediately by the proposed scalar semantics. With no additional modification, this semantics also predicts the correct truth-conditions for Nouwen’s examples, modulo some subtleties involving epistemic interpretations of minimum and maximum operators. These conclusions provide additional evidence in support of previous work arguing that the semantics of obligation and desire should be built around scales rather than quantification over possible worlds.

**Keywords:** Minimum and maximum requirements, gradability, modality, deontic modals, desire verbs, scalar modals

### 1 Nouwen’s Puzzle

Young Sammy wants to be tall when he grows up, but not too tall. Here are some reasonable desires he might have:

- (1) a. The minimum height that Sammy wants to be when he grows up is 6’.
- b. The maximum height that Sammy wants to be when he grows up is 6’8”.

---

\* I received helpful criticism on various ideas found in this paper from Luka Crnic, Rick Nouwen, Pranav Anand, Paul Portner, Angelika Kratzer, and audiences at NYU, the Institute of Philosophy (School of Advanced Study, University of London), *SALT 21*, and Georgetown. §§2-3 are closely related to parts of ch.5-6 of my dissertation (Lassiter 2011). Many thanks to my committee Chris Barker, Anna Szabolcsi, Seth Yalcin, Philippe Schlenker, and Chris Kennedy for valuable feedback on this and related material.

These are fairly unremarkable desires, but — as [Nouwen \(2010a,b\)](#) notes — standard assumptions about the semantics of degree expressions and modals would lead us to expect that Sammy could not consistently hold these desires, or else that they would mean something very different from what they in fact do mean.

The purpose of this paper is to discern which of the standard assumptions needs to be modified, and how. I’ll argue that the problem can to a large extent be traced to the assumption that modals like *should* and *want* denote universal quantifiers over some set of possible worlds, and argue for an alternative which gets the truth-conditions of sentences like (1) right without making any radical changes to the semantics of degree expressions. As it turns out, this alternative has a good deal of independent motivation, some of which will be discussed in §2 below.

### 1.1 The Problem

[Nouwen \(2010a,b\)](#) points out that, on standard assumptions about the meanings of modals and degree expressions, (2a) ought to mean something like (2b) (where  $\min(D) = \iota d \in D[\forall d' \in D : d \leq d']$ ).

- (2) a. The minimum height that Sammy wants to be when he grows up is 6'.  
 b.  $\min(\lambda d. \forall w \in \mathbf{Acc} : \text{Sammy is at least } d\text{-tall}) = 6'$

If numerals have an “at least” semantics (e.g. [von Stechow 1984](#); [Heim 2001](#); [Hackl 2001](#)), no one could have a desire like this: (2b) is trivially false because, in every world in which Sammy is 6' tall, he is also 5' tall, 4' tall, 3' tall, etc. No matter what his dreams may be, the right side of “=” should be zero, or else (2b) should be undefined (if we think that zero is not on the height scale). Of course, what (2a) actually says is quite sensible: he wants to be at least 6' tall, and hopefully taller.

One alternative is an “exactly” semantics for numerals. Recently a number of authors — even some staunch neo-Griceans — have concluded that numerals are ambiguous between an “at least” and an “exactly” meaning ([Horn 1992, 2006](#); [Geurts 2006](#); [Nouwen 2010a](#)). On this approach, (2a) can also be interpreted as:

- (3)  $\min(\lambda d. \forall w \in \mathbf{Acc} : \text{Sammy is exactly } d\text{-tall}) = 6'$

(3) does no better than (2b) at capturing the truth-conditions of (2a), though. There is no minimum degree such that Sammy is exactly that tall in every world compatible with his desires, unless he is the same height in every such world; so (3) should presuppose that Sammy has a single precise desired height, and entail that this height is exactly 6'. This is not what the sentence means either: (2a) leaves open the possibility that Sammy’s desires would be satisfied if he were 6'2''.

The puzzle is not particular to *want* or to *minimum*; it arises with various expressions of desire and obligation and a range of degree minimizers. For example:

- (4) a. The smallest amount of butter that you should use in a pie is 150 grams.  
 b. The minimum number of dishes that you must bring to the potluck is 2.  
 c. The minimum height that is required to ride on this ride is 4 feet.  
 d. The least number of children that I want to have is 3.

If *should*, *want*, *required*, and *must* are universal quantifiers over accessible worlds, all of the sentences in (4) should be trivially false or presupposition failures. On the *at least* interpretation, this is because any world in which you use 150 grams of butter is also a world in which you use 50 grams, 10 grams, 1 gram, etc., and these numbers are smaller than 150; likewise having three children entails having two and having one, and being four feet tall entails being three feet tall, two feet tall, and one foot tall. On the *exactly* interpretation, this is because the amount in question is wrongly predicted to be an upper bound. Not only is an upper-bounding inference not present, but these sentences clearly suggest the opposite, namely that it would be a good idea to use more than 150 grams of butter (that is the least that will allow you to still make a decent pie), or that the speaker wants to have *at least* three children.

The semantics also gives us the wrong truth-conditions for sentences expressing maximum desires such as (1b), repeated as (5a). This sentence has the intuitive interpretation that Sammy wants to be *no taller than 6'8"*; but the predicted interpretation is either that he wants to be *at least 6'8"* (5b) or *exactly 6'8"* (5c).

- (5) a. The maximum height that Sammy wants to be when he grows up is 6'8".  
 b.  $\mathbf{max}(\lambda d. \forall w \in \mathbf{Acc} : \text{Sammy is at least } d\text{-tall}) = 6'8''$   
 c.  $\mathbf{max}(\lambda d. \forall w \in \mathbf{Acc} : \text{Sammy is exactly } d\text{-tall}) = 6'8''$

These are clearly not the right meanings for these sentences. The puzzle also arises for a variety of maximality operators and modals, for instance:

- (6) a. The maximum amount of paint that you need to use to paint this room is three cans (but you might be able to do it with less).  
 b. The greatest portion of your diet that should be composed of high-fat foods is  $\frac{1}{5}$ .

The clear cases of the Puzzle involve deontic, bouletic, and teleological uses of modals such as *should*, *ought*, *want*, *need*, *must*, *require*, and *have to* which are usually translated as universal quantifiers over worlds. I will abbreviate this class of expressions as  $\mathcal{D}$ -MODALS, and will occasionally use  $\mathcal{D}$  as a placeholder for the denotations of these expressions in formulae.<sup>1</sup>

<sup>1</sup> It is not totally clear whether Nouwen's Puzzle also affects epistemic modals. Suppose I have been assigned to provide surveillance for a building. I have no idea who was inside when I arrived, but I have watched seven people enter the building and no one has left. The question is whether (20)

## 1.2 Plan of Attack

There are three obvious ways that we could try to deal with Nouwen's Puzzle.

- A) Find a new semantics for degree expressions;
- B) Find a new semantics for the affected modals;
- C) Do both (A) and (B).

In this paper I will mostly pursue option (B), arguing that we can make a good deal of progress on this problem by incorporating an independently motivated alternative to standard semantics on which modals like *want* and *should* are scalar and — crucially — non-monotonic. This is possible even while maintaining a fairly simple analysis of degree expressions.

In addition, there is an empirical argument for thinking that the modal semantics is the primary culprit: similar examples with universally quantified DPs clearly do not produce similar readings. That is, (7a) does not have an interpretation analogous to (7b) with either *minimum* or *maximum*.

- (7) a. The minimum/maximum amount that everyone worked was 40 hours.  
# “The person who worked the least/most worked 40 hours.”
- b. The minimum/maximum amount that Bill wants to work is 40 hours.  
✓ “In the world in which he works the least/most, Bill works 40 hours.”

If we could somehow modify the degree semantics so that modals interpreted as quantifiers over worlds produce the desired readings for (7b), we would need some sort of mechanism preventing similar readings from arising in examples like (7a). I don't have any idea what such a mechanism would look like, or how it could be motivated. However, if the affected modals are not universal quantifiers at all (as I will argue), the fact that (7a) and (7b) differ in this respect is not troubling.<sup>2</sup>

---

describes this situation well:

- (1) ?? The minimum number of people that must be in the building is seven.

Nouwen (2010b) suggests that it does not, although I have found speakers who find this description marginally acceptable, as well as a handful of naturally-occurring examples. My suspicion is that, to the extent that this description is acceptable, it is not a genuine case of Nouwen's Puzzle but due to a separate phenomenon whereby minimum and maximum operators can receive an epistemic interpretation; see §4.2 below.

<sup>2</sup> This is, incidentally, why I will not consider in detail the effects of switching to an interval-based semantics for degrees here (as in Schwarzschild & Wilkinson 2002). This approach has various motivations, and it might be able to avoid some of the problems for the interaction of  $\mathcal{D}$ -modals with the simple degree semantics noted above. To the extent that it does, though, it would seem to predict that universally quantified DPs should do the same (and epistemic modals as well). Interval semantics may well provide part of the full solution, though.

I want to be clear, however, that I am not actually arguing that the semantics for degrees that I'm using here is correct. In fact, this semantics is probably in need of refinement, and there is no shortage of proposals available. However, I don't think it will be possible to deal with Nouwen's Puzzle by modifying the degree semantics without also adopting a modal semantics along the lines I'll propose here. To keep the paper to a manageable length, then, I will settle for showing that the simplistic assumptions about degree semantics that I am making are not the root of the problem: first and foremost, we need a new semantics for  $\mathcal{D}$ -modals.

## 2 Inferences, Information, & Scales: Problems for Quantificational Theories

This section presents three types of data which are problematic for standard accounts of the semantics of  $\mathcal{D}$ -modals which treat them as quantifiers over possible worlds.

- Quantificational theories predict that  $\mathcal{D}$ -modals and universal quantifiers should license the same inferences. Several examples call this prediction into question, and in particular, the assumption that  $\mathcal{D}$ -modals are upward monotonic. I'll suggest that they are in fact **non-monotonic**.
- $\mathcal{D}$ -modals are sensitive to **fine-grained probabilistic information** in a way that is difficult to capture in quantificational theories.
- Most  $\mathcal{D}$ -modals participate in degree modification and comparison structures, a fact which strongly suggest that their semantics is — like other expressions which have these properties — built on **scales**.

These data motivate the scalar semantics to be proposed in the next section, which gives a straightforward answer to the three problems.

### 2.1 $\mathcal{D}$ -Modals are Not Universal Quantifiers

Many classic “paradoxes of deontic logic” involve divergences between the inference patterns which are intuitively valid with universal quantifiers and those which are intuitively valid with strong deontic modals. Here I will mention two inference patterns which would be valid if  $\mathcal{D}$ -modals were universal quantifiers, but admit of counter-examples; Lassiter (2011) discusses several more. In the next section I'll show that the counter-examples are predicted by the scalar alternative I'll propose.

#### 2.1.1 Professor Procrastinate

As a consequence of the upward monotonicity of the universal quantifier in its nuclear scope, the inference from  $\forall x(\phi \wedge \psi)$  to  $\forall x(\phi)$  is valid. However, Jackson

(1985); Jackson & Pargetter (1986) argue that the inference from *ought*( $\phi \wedge \psi$ ) to *ought*( $\phi$ ) admits of counter-examples, which suggests again that *ought* is not a universal quantifier. Jackson & Pargetter (1986) describe the case as follows:

Prof. Procrastinate is invited to review a book on which he is the only fully qualified specialist on the planet. Procrastinate's notable character flaw, however, is his inability to bring projects to completion. In particular, if Procrastinate accepts to review the book, it is extremely likely that he will not end up writing the review. In the eyes of the editor, and of the whole scientific community, this is the worst possible outcome. If Procrastinate declines, someone else will write the review — someone less qualified, but more reliable.

According to Jackson & Pargetter, (8a) is true in this scenario, but (8b) is false:

- (8) a. Professor Procrastinate ought to accept and write the review.  
b. Professor Procrastinate ought to accept the review.

Intuitively, the reason for this is that *ought* in (8b) takes into account not only the good worlds in which Prof. Procrastinate accepts and writes, but also what is much more likely if he accepts, that he will not write. Since this is the worst possible outcome, the likelihood of its occurrence somehow outweighs the fact that the worlds in which he accepts and writes are optimal. This example suggests that the relevant operators simply are not universal quantifiers. (It is not too difficult to construct similar cases with *want*, cf. Lassiter 2011.)

### 2.1.2 Chicken

Jackson (1985) argues that the inference from  $\mathcal{D}(\phi) \wedge \mathcal{D}(\psi)$  to  $\mathcal{D}(\phi \wedge \psi)$  admits counter-examples, even though its universally-quantified counterpart is valid.

Attila and Genghis are driving their chariots towards each other. If neither swerves, there will be a collision; if both swerve, there will be a worse collision (in a different place, of course); but if one swerves and the other does not, there will be no collision. Moreover if one swerves, the other will not because neither wants a collision. Unfortunately, it is also true to an even greater extent that neither wants to be 'chicken'; as a result what actually happens is that neither swerves and there is a collision. (Jackson 1985: 189)

As Jackson points out, all of the following are intuitively true in this scenario:

- (9) a. Attila ought to swerve.

- b. Genghis ought to swerve.
- c. It's not the case that Atilla and Genghis ought to both swerve.

Again, it is possible to construct similar examples using *want*.

If *ought* and *want* are universal quantifiers ranging over a set of worlds (however determined), it should not be possible for the judgments in (9) to be even coherent, much less intuitively correct. This is a serious problem, since the inference from  $\forall x(P(x)) \wedge \forall x(Q(x))$  to  $\forall x(P(x) \wedge Q(x))$  is of course valid.

## 2.2 *D*-Modals are Sensitive to Fine-Grained Probabilistic Information

Viewed from a certain perspective, all three of the puzzles in the last section are instances of a broader class of puzzles for quantificational semantics of desire verbs and deontics. The problem, in a nutshell, is that obligation and desire interact with graded belief in a more fine-grained way than quantificational theories can capture.

To see the connection, consider: how can *Prof. Procrastinate ought to accept the review* be false even though the best worlds are ones in which he does accept the review (and writes it too)? Because, assuming that he accepts the review, there is a high **probability** that he will not write it, and we will find ourselves in one of the worst possible situations instead of one of the best. How can it be that Atilla ought to swerve and Genghis ought to swerve, even though it would be disastrous if both of them did? Because, as part of the story, we are told that it is highly **unlikely** that either of them will in fact do what they ought. The lesson here is that, at least with respect to *ought* and *want*, probability matters. If so, we cannot treat *D*-modals as quantifiers whose domain is restricted to the best (epistemically) possible worlds. Rather, we need some kind of mechanism which tells us how to weigh good and bad outcomes against each other, taking probability into account in some way.

In fact this is what Goble (1996) and Levinson (2003) conclude regarding *ought* and *want* respectively. Goble's story, simplified considerably, goes like this. Suppose that a doctor must choose which medicine to give to a critically ill patient, *A* or *B*. *A* might produce a total cure, but will more likely kill the patient. Meanwhile *B* will save the patient's life, but will leave him slightly debilitated.

What should the doctor do? Intuition suggests that the doctor ought to choose *B*, since *A* is very risky. However, standard quantificational semantics for *ought* unhesitatingly recommends choosing *A*, because all of the **best** epistemically accessible worlds in this scenario are worlds in which the doctor gives medicine *A*. This is evidently the wrong recommendation. Again, the problem relies on the fact that, in quantificational semantics, the fact that all of the best worlds are *A*-worlds is enough to render true *The doctor ought to give A*, regardless of their improbability.

Levinson (2003) gives an example which makes the same point involving *want*.

Consider the following four worlds, representing the possible outcomes I must consider in making insurance-buying decisions as a homeowner:

- (10)  $w_1$ : I do not buy insurance and my home burns down  
 $w_2$ : I do not buy insurance and my home does not burn down  
 $w_3$ : I buy insurance and my home does not burn down  
 $w_4$ : I buy insurance and my home burns down

It seems clear that, as a homeowner, if my house burns down, I would prefer to have fire insurance:  $w_4 > w_1$ . I also do not like to spend money pointlessly, and so, assuming my home does not burn down, I prefer a state in which I do not buy insurance:  $w_2 > w_3$ . Finally, I prefer a state in which my home does not burn down to a state in which my home burns down, no matter what:  $w_2, w_3 > w_1, w_4$ . The only consistent preference order meeting these constraints is:

- (11)  $w_2 > w_3 > w_4 > w_1$

Assuming that  $w_1$ - $w_4$  are epistemic possibilities, we should be able to conclude that

- (12) I want not to buy insurance,

since all of the top-ranked worlds in (11) are worlds in which I do not buy insurance. But this is not an appropriate inference. Even though I would prefer not to buy insurance if I **knew** that my house would not burn down, I may still want to do so because I am uncertain whether it will. In particular, if I think that there is a decent chance that my house will burn down at some point, I may want to buy insurance in full knowledge that all of the worlds in which I do so are globally suboptimal.

These examples indicate that in considering what we ought to do, or what we want to do, we cannot only look at the best worlds. Instead, whether or not I want to or ought to buy insurance will depend on my judgment about how likely it is that my house will burn down, as well as factors such as the cost of insurance vs. the value of the house. Simply put, non-optimal worlds matter, and probability matters.

### 2.3 $\mathcal{D}$ -Modals are (Mostly) Gradable

On standard assumptions, gradable expressions are functions or relations between individuals, VP meanings, propositions, etc. and an ordered set of points called *degrees* which form part of a *scale*. The semantic mechanisms underlying universal quantification are quite different, and do not involve degrees or scales. In particular, if something is semantically a universal quantifier, we do not expect it to display the core empirical characteristics of gradability — participation in degree modification and comparison structures. If it does, we have some explaining to do.

Most of the  $\mathcal{D}$ -modals show a fair bit of gradable behavior. The main verbs *want*, *need*, and *require* are clearly gradable, displaying roughly the same behavior in this



respect as ordinary gradable verbs such as *like*. (These are discussed along with sources and further examples in Lassiter 2011: §5.)

- (13) a. The members of a literary group are required to have a blazer, more than they are required to have ever actually read a book.  
 b. Thus, you are very much required to have a good credit record to prove yourself as a reliable client to the insurance providers.
- (14) a. [M]any library officials want more to intimidate than to really change an institutional culture that has squelched feedback.  
 b. I am an American and I want very much to travel to Cuba.
- (15) a. But really emails need to be timely more than they need to be amazing.  
 The auxiliary *should* and the quasi-auxiliary *ought* also show limited gradability.
- (16) a. [O]nce the damage is done, Constance ought to help George ... more than she ought to help anyone else similarly situated. (Driver 1997: 853)  
 b. A war between Britain and the U.S. ought very much to be deprecated.
- (17) a. I don't think he [UFC fighter Phil Davis] should be compared to Rosholt as much as he should be to Houston Alexander.<sup>3</sup>

There are two notable exceptions to the observation that *D*-modals are frequently gradable: *must* and *have to*. For instance,

- (18) a. ?\* You must/have to leave more than John.  
 b. ?\* Sam must/has to very much be at home.

We might conclude that these items are quantifiers while the other *D*-modals are scalar. However, there is an alternative. As discussed in Lassiter 2011: §§1,6, the fact that an expression is not gradable does not show conclusively that it does not have a scalar semantics. Rather, what it shows is that the expression in question does not interact with threshold-manipulating operators as gradable expressions do; this does not tell us what the theoretical explanation for this fact is. We still have to ask on a case-by-case basis whether there are arguments for a scalar semantics.

The most straightforward reason for thinking that all *D*-modals are scalar, even if some are not gradable, is that they all take part in some or all of the other puzzles under consideration here — Ross's Puzzle, Professor Procrastinate, Chicken, information-sensitivity, and Nouwen's Puzzle. In the next section I will propose a unified scalar semantics for *D*-modals which allows us to make sense of these puzzles. A general solution of this type is probably not achievable within theories which treat these expressions — including the non-gradable ones — as quantifiers over worlds; the best we can hope for is piecemeal accounts for the various objections.

<sup>3</sup> <http://www.bloodyelbow.com/2009/12/17/1205767/ufc-signs-hot-prospect-phil-davis>

### 3 Scalar Semantics for $\mathcal{D}$ -Modals

The solution that I will advocate to the three problems just noted is a generalization of proposals due to Goble (1996) and Levinson (2003). On this account,  $\mathcal{D}$ -modals are scalar expressions which place conditions on the **probability-weighted average preference** of the worlds in their propositional argument. This is formally identical to the well-known construct of **expected utility** ( $\mathbb{E}(\cdot)$ ), though it is important to keep in mind that the semantics does not have any commitments with respect to questions about decision-making behavior. Rather, expected utility is simply a method for relating preferences over worlds to preferences over propositions, albeit one that — I will argue — has a number of empirically desirable characteristics.

#### 3.1 Preference and Information

Preferences are often represented formally in logic and semantics as binary orders over worlds. I will assume something slightly stronger, that preferences over worlds are given by an **interval order**. Essentially, an interval order contains all of the information in a binary order, including comparative facts about whether world  $w$  is preferred to world  $w'$ ; but it is richer in that it also contains information about *how much* better  $w$  is than  $w'$ . Without going into the details of the qualitative characteristics of interval orders, I will simply note here that every interval order can be given an equivalent representation by a class of “admissible” measure functions  $\mu : W \rightarrow \mathbb{R}$  which take worlds to real numbers. This means that we can use qualitative (order-theoretic) representations and representations using numerical measurements interchangeably, as long as we are careful about how we relate the two; see Krantz, Luce, Suppes & Tversky (1971); Lassiter (2011): §§2,6 for formal details and motivation for using interval orders to represent preferences. In the case of preferences, the measure functions are called *utility functions* and will be represented by  $\mathbb{U}$ .

The kinds of preferences that express using natural language are typically not preferences over individual worlds, but over sets of worlds (propositions). To give a semantics for these items we need to define an order over propositions in terms of the order over worlds just outlined. A simple and *prima facie* attractive method is:

- (19) **Possibilism:**  $\phi$  is preferred to (better than, more desirable than, ...)  $\psi$  if and only if the highest-ranked accessible worlds in  $\phi$  are preferred to (better than, more desirable than, ...) the highest-ranked accessible worlds in  $\psi$ .<sup>4</sup>

Possibilism is assumed explicitly or implicitly in a great deal of work in logic, semantics, and meta-ethics (see van Fraassen 1973; Jackson 1985 for critical discus-

<sup>4</sup> Or, if there are no maximal  $\phi$ -worlds, if there is a  $\phi$ -world that is preferred to all  $\psi$ -worlds. The term “Possibilism” is due to Goble (1996).

World	Treatment	Outcome	Utility	$prob(w \text{Treatment})$
$w_1$	Give A	Full recovery	+100	.05
$w_2$	Give A	Death	-100	.95
$w_3$	Give B	Partial recovery	+20	1

**Table 1** Utilities and probabilities in a toy model for the medicine problem.

sion). Kratzer (1981, 1991), whose theory is widely assumed in formal semantics, builds this principle into her definition of “comparative possibility”.

As Jackson (1985); Goble (1996) point out, Possibilism encounters serious difficulty in situations of the type that were discussed above: frequently we are not just interested in what happens in the *best* worlds in a proposition, but also in what happens in non-optimal worlds, and particular in whether we are *likely* to find ourselves in good worlds or bad if the proposition is realized. In cases such as the insurance and medicine puzzles discussed in §2.2, Possibilism clearly gives us the wrong result: it can be desirable to buy insurance even though the best worlds in which we do not buy insurance are better than the best worlds in which we do; and it can be morally better to choose a safe but imperfect medicine even though the best worlds in which we choose a risky alternative (and it happens to work) are better than the best worlds in which we choose the safe one.

Goble’s account of the medicine puzzle starts with the idea that *ought* is not a quantifier over worlds, but an operator which tests whether the **expected utility**  $\mathbb{E}(\phi)$  of its argument  $\phi$  is sufficiently high. Expected utility is defined in (20):

$$(20) \quad \mathbb{E}(\phi) = \sum_{w \in \phi} U(w) \times prob(\{w\}|\phi)$$

On this account, the preference order over propositions ranks the propositions  $\phi \subseteq \mathcal{P}(W)$  by calculating a weighted average of the desirabilities of the worlds  $w \in \phi$ , where the weight of  $w$  is given by the conditional probability that  $w$  will be actual if  $\phi$  comes about.<sup>5</sup> In effect, this means that worlds in  $\phi$  which are very unlikely to be actual if  $\phi$  occurs are given proportionally little weight in the calculation of  $\phi$ ’s position in the preference order over propositions. As a result, a proposition  $\psi$  which contains moderately desirable worlds which are fairly likely (conditional on  $\psi$ ) can outrank a proposition  $\chi$  with worlds that are very good but (conditionally) very unlikely if the latter also has undesirable worlds which are conditionally likely.

Goble suggests that *ought* has the denotation in (21):

<sup>5</sup> I assume that the semantics has access to probabilistic information; see Yalcin 2010; Lassiter 2010, Lassiter 2011: §§3-4 for detailed arguments involving grammatical and inferential characteristics of epistemic modals. Note also that the definition in (20) is not appropriate if  $|\phi|$  is infinite. This is not a serious limitation, but a way to simplify the mathematics a bit and clarify intuitions about what EU is.

World	Insurance?	Fire?	Utility	Prob(w Insurance?)
$w_1$	No	No	+20	.95
$w_2$	Yes	No	+10	.95
$w_3$	No	Yes	-300	.05
$w_4$	Yes	Yes	-10	.05

**Table 2** Utilities and probabilities in a toy model for the insurance problem.

(21) *ought*( $\phi$ ) is true if and only if  $\mathbb{E}(\phi) > \mathbb{E}(\neg\phi)$ .

To see how this accounts for the medicine puzzle, suppose that the probabilities and utilities of the various possible outcomes of the choice of medicine are as in Table 1. (Remember that the numerical values aren’t important, just the relative differences in magnitude.) The expected utility of giving *A* in this model is  $\mathbb{E}(\textit{give } A) = +100 \times .05 + -100 \times .95 = -90$ , while the expected utility of giving *B* is  $\mathbb{E}(\textit{give } B) = +20 \times 1 = 20$ . Since  $\mathbb{E}(\textit{give } B) > \mathbb{E}(\textit{give } A)$ , we conclude by the definition in (21) that the doctor ought to give *B*. This is intuitively the correct answer, and an improvement on theories which predict that the doctor ought to give *A* simply because the best worlds in *A* — no matter how improbable — are better than the best worlds in *B*.

Levinson’s (2003) solution to the insurance puzzle has essentially the same form. In the model sketched in Table 2,  $\mathbb{E}(\textit{do not buy insurance}) = +10 \times .95 + -300 \times .05 = -5.5$ , which is less than the expected utility of buying insurance,  $+5 \times .95 + -10 \times .05 = +4.25$ . If we suppose with Levinson that *x wants*  $\phi$  is true iff  $\mathbb{E}(\phi) > \mathbb{E}(\neg\phi)$  according to *x*’s personal probability and utility functions, then *x wants to buy insurance* will come out true in a model of this form. This holds despite the fact that the worlds with highest utility are ones in which *x* does *not* buy insurance.

These are not, to be sure, the final word on the meanings of *ought* and *want* (cf. Goble 1996; Lassiter 2011 for refinements). However, these examples illustrate the fact that an expected utility semantics for *ought* and *want* explains why probabilistic information matters in inferences involving these modals, and that this semantics is capable of returning intuitively correct predictions about specific cases.

### 3.2 Inferences and Scalarity

If *want* and *ought* place conditions on expected utility, it is reasonable to hypothesize that the other  $\mathcal{D}$ -modals do the same. On this proposal, the semantic contribution of all  $\mathcal{D}$ -modals is to map their propositional argument  $\phi$  to its expected utility  $\mathbb{E}(\phi)$  and to compare this value to a threshold  $\theta_{\mathcal{D}}$ . So, for example, *x needs*  $\phi$  is true if and only if  $\mathbb{E}(\phi)$  exceeds  $\theta_{\textit{need}}$ , a threshold whose value is determined by the lexical semantics of *need* in interaction with the context of conversation. Skeletally:

(22)  $\mathcal{D}(\phi)$  is true iff  $\mathbb{E}(\phi) \geq \theta_{\mathcal{D}}$ .

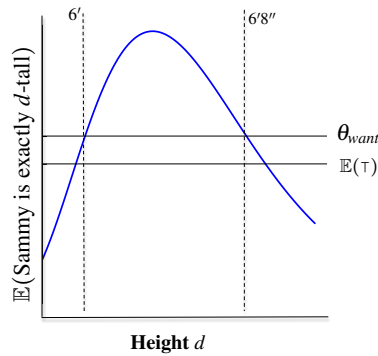
We have to acknowledge at least three parameters of variation among the  $\mathcal{D}$ -modals:

- How the preference order is established (e.g., personal preference for *want* and *need*; moral preference for *ought* and *should*; etc.);
- What body of probabilistic information is relevant;
- How the threshold  $\theta_{\mathcal{D}}$  is located relative to other items on the scale, which will differ depending on the modal strength (e.g., *may* vs. *should* vs. *must*).

For most of the rest of this paper it won't be necessary to say much about these parameters of variation. For the most part, the puzzles that we are considering can be resolved simply by adopting the skeletal proposal in (22), however the other details of the lexical semantics and context are filled in (though see fn. 6 below).

An immediate advantage of this proposal is that the three puzzles discussed in §2.1 — involving inferences that are predicted to be valid on the standard semantics, but admit intuitive counter-examples — are resolved by the proposal in (22). The key fact is that expected utility is **non-monotonic**, essentially because different conditional probability measures are used as weights in the calculation of  $\mathbb{E}(\phi)$  and  $\mathbb{E}(\phi \vee \psi)$  — the first using  $prob(\cdot|\phi)$ , and the latter using  $prob(\cdot|\phi \vee \psi)$ .

Ross' Puzzle does not arise in this semantics because  $\mathbb{E}(\phi) \geq \theta_{\mathcal{D}}$  does not imply  $\mathbb{E}(\phi \vee \psi) \geq \theta_{\mathcal{D}}$ , regardless of what  $\theta_{\mathcal{D}}$  happens to be. In contrast to the standard semantics — where the goodness of a disjunction is always equal to the goodness of the better of the disjuncts —  $\mathbb{E}(\phi \vee \psi)$  can be *less* than  $\mathbb{E}(\phi)$ , and will be whenever  $\mathbb{E}(\psi) < \mathbb{E}(\phi)$  and  $prob(\psi|\phi \vee \psi) > 0$ . The expected utility-based semantics also makes intuitively correct predictions about the other inferential puzzles we saw above. Briefly,  $\mathbb{E}(\phi \wedge \psi) \geq \theta_{ought}$  does not imply  $\mathbb{E}(\phi) \geq \theta_{ought}$ . Rather,  $\phi$  can be worse than  $\phi \wedge \psi$  when there are bad worlds which satisfy  $\phi \wedge \neg\psi$  and are likely to happen if  $\phi$  occurs. This is precisely what is going on in the Professor Procrastinate puzzle: even though worlds in which he accepts and writes the review ( $\phi \wedge \psi$ ) are very good, the fact that he probably will not write the review if he does accept it means that the undesirable ( $\phi \wedge \neg\psi$ )-worlds receive greater weight than the desirable ( $\phi \wedge \psi$ )-worlds in calculating  $\mathbb{E}(\phi)$ . Similarly, in Chicken we know that, probably, neither party will swerve; so the expected utility of Genghis' swerving is high, since the most likely scenario if he does is that Attila will not, leading to a good outcome. The same holds with the roles reversed, but it does not follow that we are in a good situation if both swerve; rather, this guarantees the worst possible outcome, with the result that  $\mathbb{E}(\phi) \geq \theta_{ought} \wedge \mathbb{E}(\psi) \geq \theta_{ought}$  is true but  $\mathbb{E}(\phi \wedge \psi) \geq \theta_{ought}$  is false. See Lassiter 2011 for a more leisurely discussion of these cases and several others.



**Table 3** An expected utility function which validates (24a) and (24b).

The explanation of gradable behavior among the  $\mathcal{D}$ -modals is likewise straightforward in this semantics: expected utility is a scale, and we can simply port over a standard semantics for gradability and comparison. The net result is that we have truth-conditions like these:

- (23) a. *Emails need to be timely more than they need to be amazing* is true iff  $\mathbb{E}(\text{Emails are timely}) > \mathbb{E}(\text{Emails are amazing})$ .  
 b. *You ought very much to leave* is true iff  $\mathbb{E}(\text{you leave})$  is much greater than  $\theta_{\text{ought}}$ .

Gradability of  $\mathcal{D}$ -modals, which is mysterious within quantificational semantics for these expressions, receives a simple and clean account in the scalar theory.

## 4 Nouwen’s Puzzle Revisited

### 4.1 Basic Account

The scalar semantics gives us an immediate solution to the Puzzle in the case of the mid-scalar items *want*, *ought*, and *should*. I’ll illustrate mostly with *want*.

- (24) a. The minimum height that Sammy wants to be is 6’.  
 b. The maximum height that Sammy wants to be is 6’8”.

The semantics sketched in the last section assigns these the following truth-conditions:

- (25) a.  $\llbracket (24a) \rrbracket = 1$  iff  $\min(\lambda d[\mathbb{E}(\text{Sammy is } d\text{-tall})] \geq \theta_{\text{want}}) = 6'$   
 b.  $\llbracket (24b) \rrbracket = 1$  iff  $\max(\lambda d[\mathbb{E}(\text{Sammy is } d\text{-tall})] \geq \theta_{\text{want}}) = 6'8''$

Whether numerical expressions receive an “at least” or an “exactly” interpretation, these truth-conditions are not trivial. The correct interpretation is generated most straightforwardly if we adopt an “exactly” interpretation, though. To see what

the world would have to be like for both of the sentences in (24) to be true, consider Table 3. (25a) and (25b) require that the expected utility of Sammy’s being exactly  $d$ -tall reaches  $\theta_{want}$  for the first time at  $d = 6'$ , that it meets or exceeds  $\theta_{want}$  at  $d = 6'8''$ , and is less than  $\theta_{want}$  for every  $d > 6'8''$ . Not only are these truth-conditions satisfiable, they appear to be precisely what the sentences in (24) mean: the least height that is desirable by Sammy’s lights is  $6'$ , and the greatest is  $6'8''$ .

Interpreted according to the scalar semantics argued for here, examples (24) do not receive trivial interpretations comparable to those generated by the standard semantics for  $\mathcal{D}$ -modals. The contrast with quantificational theories is due essentially to the fact that we are no longer universally quantifying some set of good possibilities. Instead the semantics instructs us to examine the relevant proposition to ensure that it is sufficiently desirable, without excluding the possibility that incompatible propositions are also desirable. Not surprisingly, this fact generates many different predictions; for example, as this example makes clear, we predict that it is possible for Sammy to have numerous logically incompatible desires. See Lassiter 2011: §§5-6 for arguments that this is a welcome prediction, and consideration of various other logical differences between the standard semantics and the scalar alternative.<sup>6</sup>

## 4.2 Two Kinds of Maximal Requirements

While maximum requirement sentences with modals such as *should* and *ought* operate like the example just given, similar sentences with logically stronger  $\mathcal{D}$ -modals have strikingly different interpretations.<sup>7</sup> Compare the following:

- (26) a. The earliest Sam should/ought to/wants to arrive at the party is 8PM.  
       “Sam should not arrive earlier than 8PM”  
       b. The earliest Sam must/is required to/has to arrive at the party is 8PM.  
       “Sam isn’t required to arrive earlier than 8PM”

<sup>6</sup> This solution has to be modified slightly for the logically stronger  $\mathcal{D}$ -modals *must*, *required*, and *have to*, which — on the analysis of Lassiter (2011) — require both that their propositional argument  $\phi$  be very good and that all ways of realizing its negation be fairly bad. With an *exactly* meaning for numerals and degrees, we end up with truth-conditions that are too strong for sentences like *The minimum number of dishes that you must bring to the potluck is 2*: instead of a minimum requirement the predicted meaning is a precise requirement. Space limitations prevent me from going the details here, but it is possible to capture the truth-conditions of these examples by using an “at least” meaning for degrees and numerals in these cases (which, recall, we’re assuming is freely available in addition to the “exactly” meaning, due to free enrichment by a **max** operator or something with the same effect). The net result is an entailment that the expected utility function drops off very sharply below the minimum requirement. This captures the meaning of these sentences as well as the difference between them and similar cases with weaker modals such as *ought*.

<sup>7</sup> I owe Rick Nouwen (p.c.) for the observation that sentences like (26b) have this reading.

(26a) is what we would expect from the above discussion, but (26b) needs explaining: it seems to suggest that arriving after 8PM is acceptable and that arriving before 8PM is not, almost like a *minimum* requirement.

Note, however, that this contrast matches up perfectly with an independently motivated difference in semantic strength: the modals in (26b) are logically stronger than the ones in (26a) (Sloman 1970; Horn 1972, 1989; von Stechow & Iatridou 2008). While a detailed account of (26) goes beyond the scope of this paper, I think it can be accounted for. First, the reading that the previous discussion leads us to expect for (26b) does appear when we set up a favorable context, for example:

(27) My birthday party will be in a small venue. The maximum number of people that must be there is 40.

(27) clearly indicates that no more than 40 may attend. (The expected reading of (26b) is also available in a favorable context, especially with stress on *EARLIEST*.)

Second, the prominent reading of (26b) and examples like it generally indicate speaker uncertainty about the precise requirement. This is clear in (28):

(28) The fastest Sam has to drive to get home on time is 70mph. But he might be able to get home if he drives a bit slower, I'm not sure.

In other words, sentences like (26b) and (28) are not interpreted as minimum requirements *per se*. Rather, they indicate that there is some uncertain (minimum) requirement *d*, and the value given is the maximum value of *d* that is epistemically possible by the speaker's lights: "I'm certain that Sam doesn't have to drive faster than 70mph to get home on time". This may seem surprising, but it makes sense in light of the fact that even unmodalized sentences with *minimum* and *maximum* frequently have epistemic interpretations (Nouwen 2010a). For example:

(29) The maximum amount of money that Sam has is \$50.

(29) does not mean that Sam has *exactly* \$50 as we might expect, but rather that the speaker is unsure how much he has, but is sure that it does not exceed \$50. Whatever the correct account of sentences like (29) is, it should carry over to cases like (26b) and (28). The remaining puzzle, of course, is whether examples like (26a) have an epistemic interpretation, and if so why not:

(30) The most butter you ought to use in the pie is 200 grams. ? But maybe the recipe calls for less, I don't remember.

My judgments are not totally clear, but the reading that the continuation invites does not strike me as too implausible. To the extent that it is available, we can conclude that a "maximum possibility" interpretation is generally available for maximum requirement sentences, giving rise to the misleading appearance that they sometimes express minimum requirements. I do not, however, know why the epistemic reading is less prominent in (26a) than with the stronger modals in (26b).



## 5 Conclusions

Nouwen's Puzzle is a serious problem for widely accepted theories of modality and degree semantics: more or less standard assumptions about these two domains lead to spectacularly wrong truth-conditions for a class of sentences in which they interact. This means that standard assumptions about one or both of these domains must be revised. I suggested that modal semantics should be our main target.

Modals are usually treated as restricted quantifiers over possible worlds, but a number of puzzles about missing inferences, interactions with fine-grained information, and grammatical gradability suggest that the right semantics for the modals which are affected by Nouwen's Puzzle should be non-monotonic and scalar. I argued for a semantics for deontic, bouletic, and teleological modals which treats them as operators which check whether the *expected utility* of their proposition argument meets or exceeds a threshold value — just as gradable adjectives do. This proposal resolves the problems noted for quantificational theories and gives us a simple solution to Nouwen's Puzzle. The fact that minimum and maximum requirements — which are deeply puzzling for quantificational theories — receive an immediate account within a theory of modal semantics which was proposed for independent reasons is a strong argument in favor of the new account.

## References

- Driver, Julia. 1997. The Ethics of Intervention. *Philosophy and Phenomenological Research* 57(4). 851–870.
- von Fintel, Kai & Sabine Iatridou. 2008. How to say *ought* in foreign: The composition of weak necessity modals. In Jacqueline Guéron (ed.), *Time and Modality*, 115–141. Springer.
- van Fraassen, Bas C. 1973. Values and the heart's command. *Journal of Philosophy* 70(1). 5–19.
- Geurts, Bart. 2006. Take 'five': the meaning and use of a number word. In Vogeleer & Tasmowski (eds.), *Non-Definiteness and Plurality*, 311–329. John Benjamins.
- Goble, Lou. 1996. Utilitarian deontic logic. *Philosophical Studies* 82(3). 317–357.
- Hackl, Martin. 2001. *Comparative quantifiers*: MIT dissertation.
- Heim, Irene. 2001. Degree operators and scope. In Fery & Sternefeld (eds.), *Audiatur vox sapientiae: A festschrift for Arnim von Stechow*, 214–239. Akademie Verlag.
- Horn, Laurence. 1972. *On the Semantic Properties of Logical Operators in English*: UCLA dissertation.
- Horn, Laurence. 1989. *A Natural History of Negation*. University of Chicago Press.
- Horn, Laurence. 1992. The said and the unsaid. In Barker & Dowty (eds.), *Semantics and Linguistic Theory (SALT) 2*, 163–192. Ithaca, NY: CLC Publications.

- Horn, Laurence. 2006. The border wars: A neo-Gricean perspective. In von Heusinger & Turner (eds.), *Where semantics meets pragmatics*, 21–48. Emerald.
- Jackson, Frank. 1985. On the semantics and logic of obligation. *Mind* 94(374). 177–195.
- Jackson, Frank & Robert Pargetter. 1986. Oughts, options, and actualism. *The Philosophical Review* 95(2). 233–255.
- Krantz, David H., R. Duncan Luce, Patrick Suppes & Amos Tversky. 1971. *Foundations of Measurement*. Academic Press.
- Kratzer, Angelika. 1981. The notional category of modality. In Eikmeyer & Rieser (eds.), *Words, worlds, and contexts*, 38–74. de Gruyter.
- Kratzer, Angelika. 1991. Modality. In von Stechow & Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, de Gruyter.
- Lassiter, Daniel. 2010. Gradable epistemic modals, probability, and scale structure. In Li & Lutz (eds.), *Semantics and Linguistic Theory (SALT) 20*, 197–215. Ithaca, NY: CLC Publications.
- Lassiter, Daniel. 2011. *Measurement and Modality: The Scalar Basis of Modal Semantics*: New York University dissertation.
- Levinson, Dmitry. 2003. Probabilistic model-theoretic semantics for *want*. *Semantics and Linguistic Theory* 13 222–239.
- Nouwen, Rick. 2010a. Two kinds of modified numerals. *Semantics and Pragmatics* 3(3). 1–41.
- Nouwen, Rick. 2010b. Two puzzles about requirements. In Aloni, Bastiaanse, de Jager & Schulz (eds.), *Logic, Language and Meaning: 17<sup>th</sup> Amsterdam Colloquium*, 345–354. Springer.
- Schwarzschild, Roger & Karina Wilkinson. 2002. Quantifiers in comparatives: A semantics of degree based on intervals. *Natural Language Semantics* 10(1). 1–41.
- Slooman, Aaron. 1970. 'Ought' and 'Better'. *Mind* 75(315). 385–394.
- von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3(1). 1–77.
- Yalcin, Seth. 2010. Probability Operators. *Philosophy Compass* 5(11). 916–937.

Daniel Lassiter  
Department of Psychology  
Stanford University  
420 Jordan Hall, 450 Serra Mall  
Stanford, CA 94305, USA  
[danlassiter@stanford.edu](mailto:danlassiter@stanford.edu)