Saving monotonic modals with ranked ordering sources[1]

Dee Ann REISINGER — *Johns Hopkins University*

Abstract. In this paper, I argue that Jackson (1985)'s Professor Procrastinate (PrP) problem is not inconsistent with a quantificational semantics for deontic modals. In particular, the apparent inability to infer *ought*(ϕ) from *ought*$(\phi \wedge \psi)$ is because the modals in the two sentences are interpreted with respect to two distinct ordering sources, and these contexts differ in the relative ranking that they assign to different priorities. I show how formalisms for modeling contextual priority rankings—particularly the ordered merging operation (Katz et al., 2012) and ranked ordering sources (Reisinger, 2016)—account for the PrP problem and argue that ranked ordering sources better account for priority-sensitive modals in embedded contexts.

Keywords: modality, ordering semantics, monotonicity.

1. Introduction

The main data point I consider in this paper is the Professor Procrastinate (PrP) problem, raised by Jackson (1985) to argue against the distributivity of *ought* over conjunction in deontic logic. The problem takes place in the following scenario.[2]

> Professor Procrastinate is the leading researcher in her field, and accordingly she has been invited to review a book on that topic. A review written by Prof. Procrastinate will be significantly more valuable to the scholarly community than one written by anyone else. Unfortunately, she has a habit of putting off her work, and if she accepts the invitation, she is extremely unlikely to finish the review on time. If she declines, then the editor will find someone else to write the review, and this would be a better outcome than no review at all.

In this scenario, Jackson judges (1) to be true and (2) to be false.

(1) Prof. Procrastinate ought to accept and write the review.

(2) Prof. Procrastinate ought to accept.

On this basis, he argues that the inference from *ought*$(\phi \wedge \psi)$ to *ought*(ϕ) is invalid, where in this case ϕ is *Prof. Procrastinate accepts* and ψ is *Prof. Procrastinate writes the review*. Lassiter (2011) moves this problem from deontic logic to natural language semantics as part of his argument that deontic modals are not upward monotonic. This is a critical argument against a standard quantificational semantics for modals, which predicts that modals are upward monotonic, but the argument only goes through if the modals in (1) and (2) quantify over the same set of possible worlds, which in Kratzer (1991)'s semantics corresponds to interpretation

[1]I want to thank Kyle Rawlins, Aaron Steven White, Justin Bledin, Paul Smolensky, and Daniel Lassiter for their suggestions, feedback, and encouragement at various stages of this project.

[2]This version of the PrP scenario is based more closely on Lassiter (2011)'s presentation, which, along with Jackson and Pargetter (1986), simplifies the original description of the PrP problem from Jackson (1985).

with respect to the same conversational backgrounds. I will argue in this paper that these sentences are interpreted under two different *ordering sources*, and that these ordering sources encode different relative rankings of priorities. Hence, the PrP problem is consistent with a quantificational modal semantics that can encode this contextual priority ranking difference.

I begin in Section 2 by reviewing the PrP problem in more detail, and in Section 3 I outline how the problem challenges standard quantificational accounts of modality. In Section 4, I present a solution to the PrP problem in which (1) and (2) are interpreted under ordering sources encoding different priority rankings. This solution motivates comparing different formalisms for encoding priority rankings, and hence I show in Section 5 that an embedded modification of the PrP problem supports Reisinger (2016)'s *ranked ordering source* account over the *ordered merging operation* from Katz et al. (2012). Finally, in Section 6, I consider additional data with implications for how contextual priorities are encoded and suggest applications of priority ranking formalisms to other questions in modal semantics.

2. The Professor Procrastinate problem

First, I take a closer look at the readings of the PrP problem sentences (1) and (2) in this scenario. Recall that the three possible outcomes in the PrP scenario, in decreasing order of desirability, are

- that Prof. Procrastinate accepts the invite and, against all odds, actually writes the review;
- that she declines the invite, and someone else writes the review; and
- that she accepts the invite and fails to write the review.

Furthermore, if she accepts the invite, it is much more likely that she will not write the review than that she will. In this scenario, Jackson (1985) judges (1) to be true. This intuition is straightforward, as accepting the invitation and writing the review leads to the most desirable outcome. To understand why he judges (2) to be false, we need to remember that by far the most likely consequence of accepting the invite is that no review gets written—the *least* desirable outcome—even though accepting is also a prerequisite for the *most* desirable outcome.

This is not the only reading for (2), however. One could also judge the sentence to be true if one also holds Prof. Procrastinate responsible for overcoming her procrastination habit. This contrast is clarified by the continuations in (3)

(3) a. Prof. Procrastinate ought not to accept. If she does, she almost certainly won't write the review. (Jackson's intended reading)
b. Prof. Procrastinate ought to accept. And she better actually write the review if she does!

Hence, the PrP problem places at least two empirical constraints on a theory of natural language modals. First, the theory must account for both the truth of (1) and the falsity of (2); I show

in the next section that this constraint is not trivial to satisfy with a standard quantificational modal semantics. Second, it should be flexible enough to allow both readings in (3).

3. The challenge for quantificational semantics

In this section, I review why, as Lassiter (2011) points out, the PrP problem challenges standard quantificational accounts of modality and outline how that challenge can be resolved.

First, recall the doubly relative modal semantics developed by Kratzer (1981, 1991), in which a modal expression is evaluated with respect to a pair of *conversational backgrounds*—functions from possible worlds to sets of propositions—supplied by context. One conversational background is the *modal base*, commonly denoted by f, which maps a world w to a set of relevant facts or body of information in w. The other is the *ordering source*, denoted by g, which maps w to a set of norms, desires, laws, goals, or so on, depending on the type of ordering source. The denotation of a necessity[3] modal like *ought* is given in (4).

(4) $[\![\text{ought}]\!]^{w_@,f,g} = \lambda p_{\langle s,t\rangle} \,.\, \forall w \in \max_{g(w_@)}(\bigcap f(w_@)) : p(w)$

In this denotation, $\max_P(W)$ denotes the maximal worlds in W according to the preorder $\leq_P$, which ranks a world u above another world v if every proposition in P satisfied by v is also satisfied by u:

(5) $v \leq_P u$ iff $\forall p \in P : p(v) \rightarrow p(u)$

What is important for the current discussion is that the truth of $[\![\text{ought}]\!]^{w_@,f,g}(p)$ depends on whether p holds throughout a contextually determined domain of quantification that depends on the world of evaluation $w_@$ and the conversational backgrounds f and g. This is enough to show that the modal semantics in (4) is upward monotonic in its prejacent. That is, if p and q are propositions such that $p \subseteq q$, then $[\![\text{ought}]\!]^{f,g}(p)$ implies $[\![\text{ought}]\!]^{f,g}(q)$. For if all worlds in the domain of quantification are in p, then by assumption they are also in q. Because $p \wedge q \subseteq p$ for any p and q, distribution over conjunction follows as a special case of monotonicity. Thus, *assuming a fixed choice of conversational backgrounds*, this semantics cannot account for the simultaneous truth of (1) and falsity of (2), which are reproduced as (6) and (7), respectively.

(6) Prof. Procrastinate ought to accept and write the review.

(7) Prof. Procrastinate ought to accept.

One reaction to this problem is to adopt a modal semantics that does not license monotonic inferences. For example, Lassiter (2011) proposes a scalar semantics for deontic and bouletic modals in which *ought*(p) is true if, given some probability measure on propositions and a utility random variable, the expected utility conditioned on p being true exceeds a contextually provided threshold. Such a semantics is non-monotonic in either direction. In particular, the expected utility conditioned on some $p \wedge q$ may exceed a threshold even if the expected utility

[3] For now, I gloss over the distinction between weak and strong necessity modals.

conditioned on p alone does not if the worlds in $p \wedge \neg q$ have sufficiently low utility and high probability.

An alternative strategy, which I pursue in this paper, is to argue that the contextual parameters used to interpret (6) and (7) are not actually the same. If the two sentences are interpreted under different conversational backgrounds, then their modal quantification domains will also generally differ, and hence there is no reason to expect any particular logical relation between them to hold.

Nevertheless, if there is a difference in contextual parameters used to interpret (6) and (7), it is not obvious what that difference is. If the modal bases differ, then there should be a difference in the relevant facts or body of information used to interpret the sentences. If the ordering sources differ, then some criteria by which an outcome is judged to be desirable, morally good, or legally permissible must differ. Thus, the burden falls on the proponent of this solution to explicitly describe this contextual difference; I do this in the next section.

4. Solving the problem with priority ranking

I now propose a solution to the PrP problem that relies on what I will call *priority rankings* on ordering sources. First, I lay out a simple model of the PrP scenario and establish some notation for the rest of the paper. Next, I motivate the concept of priority rankings and show how it solves the PrP problem by blocking the inference from (6) to (7) with their intended readings. I then discuss a couple of ways to implement it formally in a quantificational modal semantics.

4.1. A model of the PrP scenario

Let *Accept* be the proposition that Prof. Procrastinate accepts the invitation to write the review, let *ProfWrite* be the proposition that she writes the review, and let *OtherWrite* be the proposition that someone else writes the review. Additionally, let *Focused* be the proposition that Prof. Procrastinate is in the right state of mind, is sufficiently organized, and so on such that she would be able to complete the review if she accepted the invitation.

These propositions are sufficient to characterize the possible outcomes in the PrP scenario as well as their causes. For example, we can capture that someone else will write the review exactly if Prof. Procrastinate rejects the invite ($\mathit{OtherWrite} \leftrightarrow \neg \mathit{Accept}$), that the outcome of accepting the invitation depends on whether she procrastinates or not ($\mathit{Accept} \rightarrow (\mathit{ProfWrite} \leftrightarrow \mathit{Focused})$), and so on. These background facts that are relevant to interpreting the modals in (6) and (7) are captured by the modal base f, which should contain these propositions when evaluated at the actual world.

These propositions are also sufficient to define certain orderings on possible worlds that are relevant to the PrP scenario. One such ordering, which I denote $\leq_{DES}$, captures the relative goodness or desirability of worlds. In this case, the worlds in which Prof. Procrastinate writes

the review are ranked above those in which someone else writes the review, which in turn are ranked above those in which no one writes the review. Note that these outcomes exhaust all possibilities given the modal base.

(8) $\neg \mathit{ProfWrite} \wedge \neg \mathit{OtherWrite} <_{DES} \mathit{OtherWrite} <_{DES} \mathit{ProfWrite}$

Another, simpler ordering $\leq_{LH}$ captures the relative likelihood of events. In particular, given Prof. Procrastinate's poor work habits, it is more likely that she is not focused than that she is:

(9) $\mathit{Focused} <_{LH} \neg \mathit{Focused}$

Both of these orderings will need to be derived from appropriate ordering sources, but I leave these details to Section 4.3.

Note also that the valuation of every proposition in this scenario is completely characterized by the valuations of *Accept* and *Focused*. That is, knowing whether each of these two propositions is true in a world is enough to know the truth value of every relevant proposition in that world. Thus, there are four relevant equivalence classes of possible worlds, which I will denote AF, $\neg AF$, $A\neg F$, and $\neg A\neg F$. This model is summarized in Table 1.

World	*ProfWrite*	*OtherWrite*	*DES* rank	*LH* rank
AF	✓	✗	3	1
$\neg AF$	✗	✓	2	1
$A\neg F$	✗	✗	1	2
$\neg A\neg F$	✗	✓	2	2

Table 1: The equivalence classes of possible worlds in the PrP scenario. The first three columns give the valuations of all relevant propositions in the model, and the last two give the rankings of the equivalence classes according to the orders $\leq_{DES}$ and $\leq_{LH}$ where higher ranks are more desirable or likely, respectively.

4.2. The concept of priority ranking

One way to characterize the intuitive challenge of the PrP scenario is that it puts two competing sets of priorities or constraints into conflict. In this case, the constraint that outcomes be as desirable as possible conflicts with the constraint that outcomes be realistic or attainable. If this conflict is resolved in different ways, then the outcomes that best satisfy the competing constraints will also change. In particular, if a modal's domain of quantification depends on how such a conflict is resolved, then its truth conditions will also depend on that resolution.

This intuition motivates the concept of *priority ranking* in the interpretation of modals. At a high level, if two potentially conflicting priorities are used to compute the set of maximal possible worlds in a modal's domain of quantification, then we defer to the higher-ranked priority when the two priorities conflict. More formally, since the Kratzer's modal semantics makes use

of preorders over possible worlds, suppose that we have two such preorders $\leq_A$ and $\leq_B$. Then using the order $\leq_{A*B}$ defined in (10) corresponds to ranking $\leq_A$ over $\leq_B$.[4]

(10) $u \leq_{A*B} v$ iff $u <_A v$ or ($u \cong_A v$ and $u \leq_B v$)

That is, $\leq_{A*B}$ orders worlds by first consulting $\leq_A$ and then only falling back to $\leq_B$ if $\leq_A$ is indifferent. To see how priority rankings help us solve the PrP problem, first consider the possible rankings of $\leq_{DES}$ and $\leq_{LH}$ from the previous section. These orders are summarized in Figure 1.

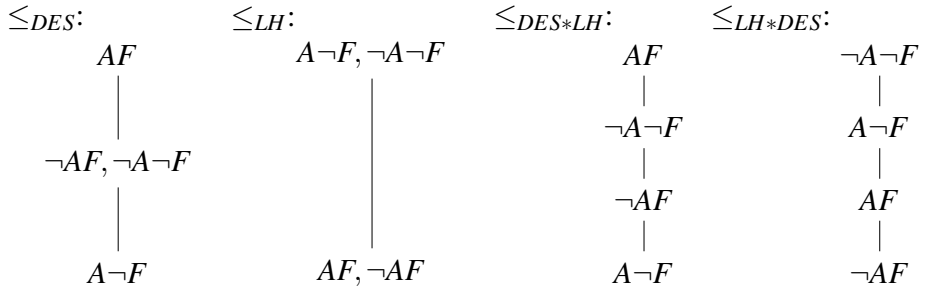

Figure 1: The orders $\leq_{DES}$ and $\leq_{LH}$ along with their possible priority rankings.

In all of the optimal worlds under $\leq_{DES*LH}$ (and in fact under $\leq_{DES}$), Prof. Procrastinate accepts the invitation and writes the review. In all of the optimal worlds under $\leq_{LH*DES}$, Prof. Procrastinate does not accept the invitation. Thus, if the modals in (11) and (12) are interpreted using the world orderings $\leq_{DES*LH}$ and $\leq_{LH*DES}$, respectively, we recover Jackson's judgments on these sentences despite using a quantificational semantics for *ought*.

(11) Prof. Procrastinate ought to accept and write the review.

(12) Prof. Procrastinate ought to accept.

Additionally, priority rankings neatly account for the different readings in (13) that were pointed out in Section 2. In particular, the first sentence in (13a) is true under the ordering $\leq_{LH*DES}$, and the one in (13b) is true under $\leq_{DES*LH}$ (or in fact under $\leq_{DES}$).

(13) a. Prof. Procrastinate ought not to accept. If she does, she almost certainly won't write the review.
b. Prof. Procrastinate ought to accept. And she better actually write the review if she does!

Not only does a priority ranking account have the formal degrees of freedom to account for both of these readings, but it does so in a way that reflects the intuitive difference between them. In particular, the reasoning that would cause someone to utter (13a) places high weight on the likelihood that desirable outcomes actually occur ($\leq_{LH*DES}$), and the rationale for (13b) must

[4]Here I use $u \cong_A v$ as an abbreviation for $u \leq_A v$ and $v \leq_A u$. In the special case where $\leq_A$ is a *partial* order, $u \cong_A v$ reduces to $u = v$.

be willing to overlook ($\leq_{DES*LH}$) or even disregard ($\leq_{DES}$) the low probability of a desirable outcome. Thus, a priority-ranking-based account can meet both of the empirical challenges raised by the PrP problem.

4.3. Implementing priority rankings in Kratzer's semantics

In the preceding discussion, I have shown how a Kratzerian modal semantics that has access to the priority ranking preorders $\leq_{LH*DES}$ and $\leq_{DES*LH}$ can account for the PrP problem. Nevertheless, in Kratzer's semantics, modals are not directly parametrized by a preorder over possible worlds. Rather, a world preorder is derived from an ordering source as described in Section 3. Thus, in this section, I review formalisms from Katz et al. (2012) and Reisinger (2016) that derive priority ranking preorders from conversational backgrounds and show how they model the PrP scenario.

4.3.1. Ordered merging

The first of these formalisms is the *ordered merging operation* from Katz et al. (2012), the first proposed formalism for implementing priority rankings over ordering sources in a Kratzerian semantics. The operation, denoted by $*$, is a binary operation on sets of propositions with the following property: If P and Q are sets of propositions that induce preorders $\leq_P$ and $\leq_Q$, then $P*Q$ is a set of propositions that induces the preorder $\leq_{P*Q}$ that ranks $\leq_P$ over $\leq_Q$. For completeness, the definition of $*$ is in (14), though the preceding property of $*$ is enough for the current discussion.

(14) $$P*Q = \\ P \cup \{q \wedge \bigwedge P \mid q \in Q\} \cup \{q \vee \bigvee P \mid q \in Q\} \cup \{(q \wedge P_n) \vee P_{n+1} \mid 0 < n < |P|, q \in Q\} \\ \text{where } P_k = \bigvee_{\substack{R \subseteq P \\ |R|=k}} \bigwedge R$$

The operation is extended from sets of propositions to ordering sources in the natural way: $(g_1 * g_2)(w) = g_1(w) * g_2(w)$.

To see how this operation applies the the PrP problem, I define two ordering sources, g_{DES} and g_{LH}, that induce the orders $\leq_{DES}$ and $\leq_{LH}$ when evaluated at $w_@$.

(15) $$g_{DES}(w_@) = \{\mathit{ProfWrite}, \neg \mathit{ProfWrite} \rightarrow \mathit{OtherWrite}\} \\ g_{LH}(w_@) = \{\neg F\}$$

After some work, we can compute the ordered merge of these sets of propositions:[5]

(16) $$g_{LH}(w_@) * g_{DES}(w_@) = \{\neg F, P \wedge \neg F, P \vee \neg F, (\neg P \rightarrow O) \wedge \neg F\}$$

Indeed, the order that $g_{LH} * g_{DES}$ induces is exactly $\leq_{LH*DES}$ from the previous section. Similarly, $\leq_{DES*LH}$ is induced by $g_{DES} * g_{LH}$, which is computed below.

[5] *ProfWrite* is abbreviated P, and *OtherWrite* is abbreviated O

(17) $g_{DES}(w_@) * g_{LH}(w_@) = \{P, \neg P \rightarrow O, \neg F \wedge P, \neg F \vee P \vee O, P \vee (\neg F \wedge O)\}$

Thus, the ordered merging operation can account for each of the PrP problem readings by using the ordering source $g_{LH} * g_{DES}$ or $g_{DES} * g_{LH}$ as appropriate.

4.3.2. Ranked ordering sources

An alternative formalism that captures priority rankings is Reisinger (2016)'s *ranked ordering sources*. In contrast to the ordered merging operation previously described, which involves no modification to Kratzer's semantics, ranked ordering sources require minor changes to the way ordering sources induce preorders on possible worlds.

In particular, ranked ordering sources are functions from possible worlds to *partially ordered* sets (posets) of propositions. Equivalently, they can be thought of as pairs $(g, \prec)$, where g is a conventional (unordered) conversational background, and $\prec$ is a function that maps w to a partial order on the set of propositions $g(w)$. For any two propositions $p, q \in g(w)$, the statement $p \prec_w q$ can be read as "q has higher priority than p." When discussing ranked ordering sources, the term *priority order* will refer to the ordering $\prec$.

Intuitively, propositions ranked higher by the priority order have greater weight in determining the preorder on possible worlds. This intuition is formalized in (18).[6]

(18) $u \leq_{P,\prec} v$ iff $\forall p \in P : (p(u) \wedge \neg p(v)) \rightarrow (\exists q \in P : q(v) \wedge \neg q(u) \wedge p \preceq q)$

That is, for each ordering source proposition p that would rank u above v, there is another proposition q that ranks v above u and that has higher priority than p.

With these definitions, it is simple to specify a model of the PrP scenario. Let $g_{LH}(w_@)$ be defined as before (with a trivial priority order), and let $g_{DES}(w_@) = \{\mathit{ProfWrite}, \mathit{OtherWrite}\}$, where $\mathit{OtherWrite} \prec_{w_@} \mathit{ProfWrite}$; note that this captures the fact that *ProfWrite* is a better outcome than *OtherWrite* using the partial order structure of g_{DES}, whereas the previous formalism had to do so using a material conditional. As before, $g_{LH}(w_@)$ and $g_{DES}(w_@)$ induce the preorders $\leq_{LH}$ and $\leq_{DES}$, respectively, via the order defined in (18).

To derive priority rankings on ranked ordering sources, Reisinger (2016) introduces the priority join operation $\sqcup$, which can be thought of as an extension of the ordered merge operation $*$ to partially ordered sets. Given two ranked ordering sources $(g, \prec)$ and $(h, \prec')$, their priority join $(g, \prec) \sqcup (h, \prec')$ is $(g \cup h, \prec^{\sqcup})$, where $\prec^{\sqcup}$ is defined in (19) and where the $\min_w$ and $\max_w$ operators compute which of a pair of ordering sources has lower or higher priority ranking in a given possible world w.[7]

[6] In the special case where $\prec$ is a total order, then computing the maximal worlds under $\leq_{P,\prec}$ is equivalent to optimization in the sense of Optimality Theory (Prince and Smolensky, 2008), where the output candidates are possible worlds, and the constraints are the propositions in P ordered by $\prec$.

[7] This formulation of the priority join differs slightly from the one in Reisinger (2016), where the join takes a third contextual parameter determining the priority ranking at each possible world. In the current formulation, these priority rankings are assumed to be specified by each possible world.

(19) $p \prec_w^{\sqcup} q$ iff $p,q \in g_{max}(w)$ and $p \prec_w^{max} q$ or
$q \in g_{max}(w)$ and $p \in g_{min}(w) - g_{max}(w)$ or
$p,q \in g_{min}(w) - g_{max}(w)$ and $p \prec_w^{min} q$

where $(g_{min}, \prec^{min}) = \min_w((g,\prec),(h,\prec'))$ and
$(g_{max}, \prec^{max}) = \max_w((g,\prec),(h,\prec'))$

That is, there are two ways for p to have lower priority than q. The first and third clauses say that if p has lower priority than q in either of the operand ordering sources, then it also has lower priority in the join. The second clause says that p has lower priority than q if p and q are members of the lower and higher priority ranked operands, respectively. Intuitively, the join $(g,\prec) \sqcup (h,\prec')$ "glues" one of the partial orders on top of the other, depending on which ordering source has higher priority in a given world.

To see how this applies to the PrP problem, let $g = g_{LH} \sqcup g_{DES}$. If g_{LH} has higher priority than g_{DES} in $w_@$—that is, if $\max_{w_@}(g_{LH}, g_{DES}) = g_{LH}$—then $g(w_@)$ has the following priority order: *OtherWrite* $\prec_w$ *ProfWrite* $\prec_w \neg F$. The reader can check that this ranked ordering source induces the preorder $\leq_{LH*DES}$. Alternatively, in a world w where g_{DES} has higher priority, the priority order is $\neg F \prec_{w_@}$ *OtherWrite* $\prec_{w_@}$ *ProfWrite*, and in this case $g(w)$ induces the preorder $\leq_{DES*LH}$. In both cases, the priority relation within g_{DES} is left intact. This join is represented schematically in Figure 2.

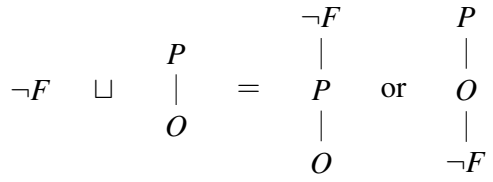

Figure 2: Possible values of the priority join $(g_{LH} \sqcup g_{DES})(w_@)$, depending on the priority ranking of g_{LH} and g_{DES} at $w_@$. In parentheses above each poset of propositions is the world preorder that it induces.

Thus, ranked ordering sources and the priority join operation, like the ordered merging operation in the previous section, are a way to implement a priority ranking solution to the PrP problem in Kratzer's modal semantics. In the next section, I discuss some differences between these two formalisms.

5. Priority rankings in embedded contexts

Many of the differences between the ordered merging and ranked ordering source implementations of priority rankings are purely formal. For example, the ordered merge operation implements priority rankings without any modification to Kratzer's ordering semantics, whereas ranked ordering sources capture priority rankings with intuitively transparent formal represen-

tations; compare the priority joins in Figure 2 with the merged ordering sources in (16) and (17).

Nevertheless, there is at least one difference that has empirical consequences. In an ordered merging account, the priority ranking is fixed by the choice of ordering source, which in turn is presumably determined by the context of utterance. For example, $g_{LH} * g_{DES}$ and $g_{DES} * g_{LH}$ are two different ordering sources that encode two different priority rankings relating g_{LH} and g_{DES}. In contrast, the priority join $g_{LH} \sqcup g_{DES}$ allows the priority ranking to vary with the world of evaluation. This difference is significant in sentences like (20), where a priority-sensitive modal (*ought*) is in an embedded context that shifts the world of evaluation (the complement of *thinks*).

(20) Prof. Procrastinate ought not to accept, but Kat thinks that she ought to.

In particular, (20) has a reading under which Kat and the speaker agree on the facts behind the PrP scenario, including the relative likelihood and utility of outcomes, but differ in the kind of moral reasoning used to determine what Prof. Procrastinate should do. In other words, (20) can entail that the speaker and Kat disagree on how to prioritize the outcome feasibility and desirability.

The ordered merging account in Section 4.3.1 does not predict that (20) has this reading. I assume here that *think* has the standard neo-Hintikkan denotation in (21), where $Dox_x(w_@)$ is the set of worlds doxastically accessible from $w_@$ by entity x.

(21) $[\![\text{think}]\!]^{f,g,w_@} = \lambda p_{\langle s,t\rangle} \,.\, \lambda x_e \,.\, \forall w \in Dox_x(w_@) : p(w)$

We can model the fact that Kat and the speaker agree on the relevant facts of the PrP scenario with the conditions in (22), that the modal base, likelihood ordering source, and desirability ordering source all have the same values in Kat's doxastically accessible worlds as in the world of evaluation.

(22) $\forall w \in Dox_{Kat}(w_@)$:
a. $f(w) = f(w_@)$
b. $g_{LH}(w) = g_{LH}(w_@)$
c. $g_{DES}(w) = g_{DES}(w_@)$

From this it follows that $(g_{LH} * g_{DES})(w)$ and $(g_{DES} * g_{LH})(w)$ agree with their values at $w_@$ for all w in $Dox_{Kat}(w_@)$, and hence the domains of quantification for both instances of *ought* in (20) are the same. Thus, this account predicts that (20) is a contradiction regardless of which priority ranking is used to interpret it, since the first conjunct requires all worlds in the domain of quantification to be $\neg A$ worlds while the second requires them to be A worlds.

In contrast, a ranked ordering source account makes no such prediction. In particular, suppose that g_{LH} has higher priority than g_{DES} in $w_@$ and that the reverse is true for each w in $Dox_{Kat}(w_@)$. Then $g_{LH} \sqcup g_{DES}$ induces the preorder $\leq_{LH*DES}$ (with maximal worlds $\neg A \neg F$) at

$w_@$ and $\leq_{DES*LH}$ (with maximal worlds AF) within $Dox_{Kat}(w_@)$, and hence both conjuncts of (20) are predicted to be true under the priority join.

Thus, on the basis of sentences like (20), I argue that we need accounts of priority rankings that, like ranked ordering sources, allow relative priority to vary with the world of evaluation.

6. Further implications and future directions

In this section, I begin by pointing out additional empirical and theoretical consequences of the priority ranking solutions to the PrP problem, particularly in comparison to non-monotonic alternatives, by focusing on the possible readings of *ought*(*Accept*), which has been taken by Jackson (1985) and others in the literature to be false. I then conclude by drawing connections between priority rankings and a variety of related topics in the modality literature for further investigation.

6.1. Revisiting Jackson's judgment

First, I claim that Lassiter (2011)'s scalar semantics for deontic modals alluded to in Section 3 accounts for the possible readings of (23) very differently from either of the quantificational priority ranking analyses.

(23) Prof. Procrastinate ought to accept.

A simplified version of Lassiter's semantics for *ought* is given in (24), where I use $\mathbb{E}(p)$ to denote the expected utility conditioned on a proposition p given some probability and utility functions on possible worlds and θ is a contextually determined threshold.

(24) $[\![\text{ought}]\!]^\theta = \lambda p_{\langle s,t\rangle} \,.\, \mathbb{E}(p) \geq \theta$

Even in this form, there are clear differences between the scalar and priority ranking accounts for (23). First, since the PrP scenario is reasonably modeled by assuming at least that $\mathbb{E}(\textit{Accept} \wedge \textit{ProfWrite}) > \mathbb{E}(\textit{Accept})$, there are three different regimes that θ can fall in. If $\theta \leq \mathbb{E}(\textit{Accept})$, then both *ought*(*Accept*) and *ought*(*Accept* $\wedge$ *ProfWrite*) are true, and if $\theta > \mathbb{E}(\textit{Accept} \wedge \textit{ProfWrite})$, both *ought* statements are false. Finally, if θ is between these two expected utility values, then *ought*(*Accept* $\wedge$ *ProfWrite*) is true while *ought*(*Accept*) is false, exactly as judged by Jackson (1985).

Thus, the only way for someone to judge (23) as true is that the standard for things one ought to do, θ, be sufficiently low, assuming that person correctly understands the PrP scenario. In constrast, under a priority ranking account, (23) can be judged true by someone who uses a different (perhaps normatively irrational) mode of moral reasoning that discounts the role of outcome plausibility.

Lassiter also posits additional conditions on the threshold θ in order to, among other goals, differentiate between strong necessity, weak necessity, and possibility modals in a scalar framework. For weak necessity modals like *ought*, θ must be significantly greater than the expected utility of some set of relevant alternatives to the prejacent. In the case of (23), the relevant alternative is $\neg Accept$. But since $\mathbb{E}(Accept) < \mathbb{E}(\neg Accept)$ in the PrP scenario, this implies that any θ low enough to make (23) true will also violate this constraint. Thus, this stronger version of Lassiter's analysis predicts that (23) *cannot* be judged true by a competent speaker of English who correctly understands the PrP scenario.

These distinct analyses of (23) are difficult to distinguish empirically as it is difficult to ensure that an informant has a correct understanding of the scenario. Nevertheless, I suspect that an account that requires (23) to be false in the PrP scenario *for semantic reasons* too strongly encodes rational norms of decision making under uncertainly into the meaning of natural language modals.

Additionally, the ordered merging and ranked ordering source accounts tell slightly different stories about the conditions under which (23) can be judged true. Using the ordered merging operation, there are two possible merged ordering sources that encode two possible priority rankings, $g_{LH} * g_{DES}$ and $g_{DES} * g_{LH}$, and (23) is false under the former and true under the latter. Using the priority join operation, however, there is only a single ordering source, $g_{LH} \sqcup g_{DES}$, that encodes a priority ranking relation between likelihood and desirability, and whether the sentence is true or not depends on which priority ranking holds in the world of evaluation (although note that the sentence is true under g_{DES} alone). This difference is subtle, but it bears on disagreement data like the dialogue in (25), where I assume A and B are both speakers who agree on the likelihood and desirability of outcomes in the PrP problem.

(25) A: Prof. Procrastinate ought to accept the invitation, right?
B: No! She really ought not to. She wouldn't write the review if she did.

If the difference between A's and B's position is just a matter of which ordering source is used to interpret the modal, as it would be using the ordered merging account, then this disagreement is odd, as A and B are asserting two mutually compatible propositions. On the other hand, under a priority join analysis, A and B disagree about whether $ought(Accept)$ is true under the ordering source $g_{LH} \sqcup g_{DES}$. That is, since the priority ranking between g_{LH} and g_{DES} is a matter of fact in a given possible world, A and B can debate which kind of world they are actually in. Thus, to the extent that a dialogue like (25) can occur without some kind of misunderstanding between the participants, I argue that the priority join is a better model of this kind of disagreement.

6.2. Priority rankings in other domains

Although in this paper and in Reisinger (2016), I have primarily motivated priority rankings through the study of deontic modals under conditions of uncertainty, there are a variety of other topics in modal semantics where formalisms similar to priority rankings have been applied or where they are likely to be appropriate.

One such topic, which was one of the motivations for the ordered merging operation in Katz et al. (2012) as well as one of Lassiter (2011)'s main empirical targets, is the study of gradable modal expressions, such as modal expressions appearing with degree modifiers or in comparative constructions. In a ranked ordering source framework, it is tempting to paraphrase a sentence like (26a) with an expression like (26b), where $\prec$ is the priority order for a deontic ordering source.

(26) a. It is more illegal to commit murder than to jaywalk.
b. λw_s . $\neg$*Jaywalk* $\prec_w$ $\neg$*Murder*

Thus, it may be desirable for a modal degree semantics that makes use of ordering sources, such as the one proposed by Portner and Rubinstein (2016), to be consistent with such priority order paraphrases (or vice versa).

Additionally, similar formalisms in which multiple ordering sources contribute to a modal domain of quantification with different degrees of priority have been proposed to model other phenomena. For example, von Fintel and Iatridou (2008) explain why weak necessity modals are logically weaker than strong necessity modals in terms of an additional ordering source, called a *secondary* ordering source, that is used to further winnow down a weak necessity modal's domain of quantification after the primary ordering source has already been applied. This is equivalent to optimizing under a priority ranked pair of ordering sources in which the primary ordering source is ranked above the secondary. Furthermore, Rubinstein (2012) shows that the difference between primary and secondary ordering sources corresponds to a difference in discourse commitments by conversational participants. In particular, primary ordering sources are those that are collectively committed to by all participants, whereas secondary ordering sources may only be committed to by the speaker. Thus, if the primary/secondary distinction can be modeled using priority rankings, it is natural to wonder whether similar discourse constraints apply to priority rankings more generally.

7. Conclusion

In this paper, I have addressed a potent objection to a quantificational semantics for natural language modals raised by the Professor Procrastinate problem. Although this problem has been proposed as evidence against any upward monotonic modal semantics, I have shown that adding a *priority ranking* as an additional contextual degree of freedom circumvents this argument. I have shown how this solution can be implemented using two related formalisms, the *ordered merging operation* of Katz et al. (2012) and Reisinger (2016)'s *ranked ordering sources*, and argued that data involving modals embedded under attitude verbs better accord with the latter. Finally, I have argued that priority rankings are not merely a stopgap measure to save quantificational modals from this one objection, as they have additional empirical and theoretical implications that differ from non-monotonic alternatives and may be applicable to several other areas of research in modal semantics.

References

Jackson, F. (1985). On the semantics and logic of obligation. *Mind 94*(374), 177–195.

Jackson, F. and R. Pargetter (1986). Oughts, options, and actualism. *The Philosophical Review 95*(2), 233–255.

Katz, G., P. Portner, and A. Rubinstein (2012). Ordering combination for modal comparison. In *Proceedings of SALT 22*, pp. 488–507.

Kratzer, A. (1981). The notional category of modality. In H. J. Eikmeyer and H. Rieser (Eds.), *Words, Worlds, and Contexts*, pp. 38–74. De Gryuter.

Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research*, pp. 639–650. De Gryuter.

Lassiter, D. (2011). *Measurement and Modality: The Scalar Basis of Modal Semantics*. Ph. D. thesis, New York University.

Portner, P. and A. Rubinstein (2016). Extreme and non-extreme deontic modals. In N. Charlow and M. Chrisman (Eds.), *Deontic Modality*. Oxford University Press.

Prince, A. and P. Smolensky (2008). *Optimality Theory: Constraint Interaction in Generative Grammar*. John Wiley & Sons.

Reisinger, D. (2016). Ranked ordering sources and embedded modality. In *Proceedings of the Linguistics Society of America*, Volume 1.

Rubinstein, A. (2012). *Roots of Modality*. Ph. D. thesis, University of Massachusetts, Amherst.

von Fintel, K. and S. Iatridou (2008). How to say *ought* in foreign: The composition of weak necessity modals. In J. Guéron and J. Lecarme (Eds.), *Time and Modality*, pp. 115–141. Springer.