Grading Modal Judgment

Nate Charlow

Contents

Grading Modal Judgment

Nate Charlow

1 Introduction

Agents can apparently bear graded attitudes (e.g., intermediate or high credence) towards epistemic modalities. Sentences expressing such graded attitudes are commonplace; consider the following triad (slightly modified from Moss 2015: 4):

(1)   It is probably the case that Bob might be hired.

(2)   It is probably the case that Bob will be hired.

(3)   Bob might be hired.

Moss remarks that "our judgments suggest that [(1)] is weaker than either [(2)] or [(3)]. Believing [(2)] is intuitively sufficient reason to bet at even odds that we will hire Bob, whereas merely believing [(1)] is not" (Moss 2015: 4). Meanwhile, while asserting (3) seems to represent the speaker as believing that Bob might be hired, (1) plainly does not.

It will be important to observe that this basic data point can be established in a variety of ways. Our judgments about (1)–(3) establish that there are constructions in natural language that apparently express graded attitudes towards epistemic modalities. In addition to evidence from natural language, graded attitudes towards epistemic modalities appear to be required by platitudes about the conversational role of assertions of epistemic modalities.

Willer (2013), for instance, observes that assertions of epistemic modalities are understood as non-trivial proposals to add information to (i.e., address a question within) a discourse. Assertion of a sentence like (3) addresses a question about whether Bob might be hired. But the notion of a question about whether Bob might be hired *presupposes* the possibility of a graded attitude (i.e., a degree of confidence greater than 0 and less than 1) toward a sentence like (3): such an attitude typically forms at least part of the cognitive basis for posing such a question.

This paper offers a model of graded modal judgment. It begins (§2) by showing why the phenomenon is so theoretically vexing: given plausible constraints on the logic of epistemic modality, it is actually *impossible* to model graded attitudes toward modal claims as judgments/ascriptions of *probability* to modalized propositions [OMITTED]. In response to this problem, this paper considers two alternative models, on which modal operators are *non-proposition-forming* operators:

- §3: Moss (2015), in which graded attitudes toward modal claims are represented as judgments/ascriptions of probability to a "proxy" proposition, belief in which would underwrite belief in the modal claim.
- §4: A model on which graded attitudes toward modal claims are represented as judgments/ascriptions of credence *to a (non-propositional) modal representation* (rather than a proxy proposition).

The second model is shown to be both semantically and mathematically tractable—a feature which does not distinguish it from Moss (2015). In §5, however, I will

argue that the second model is easily integrated into our ordinary understanding of the functional role of graded attitudes toward modal claims (in both cognition and normative epistemology)—something that, I argue, represents a positive contrast with the account of Moss (2015).

2 No Uncertainty?

All agree that the base semantic clause for the epistemic possibility modal $\diamond$ (and its dual operator $\square$) is information-sensitive—i.e., invokes reference to an information state—and that, relative to a "base" information state—for present purposes, a function from worlds of evaluation into sets of possible worlds—epistemic possibility modals quantify existentially over possibilities compatible with that state. Relative to a choice of information state σ and a choice of index of evaluation w, the appropriate semantic clause for $\diamond$ is as follows:

$$[\![\diamond\phi]\!]^{\sigma,w} = \mathsf{T} \Leftrightarrow \exists v \in \sigma_w : [\![\phi]\!]^{\sigma,v} = \mathsf{T}$$

On this understanding, a sentence of the form $\diamond\phi$ can express a possible worlds proposition, namely:

$$[\![\diamond\phi]\!]^{\sigma} = \{w : \exists v \in \sigma_w : [\![\phi]\!]^{\sigma,v} = \mathsf{T}\}$$

Such a proposition is the sort of thing to which a probability function can assign a probability, and is the sort of thing to which agents can have graded attitudes (e.g., being 10% confident in this proposition).

On the other hand, there is apparently strong evidence that sentences of the form $\diamond\phi$ *cannot* generally express possible worlds propositions with these sorts of characteristics. Assume that σ is governed by constrains of Reflexivity and Euclideanness.[1]

Reflexivity: $w \in \sigma_w$ $\qquad\qquad$ $\forall w, \sigma : [\![\square\phi \supset \phi]\!]^{\sigma,w} = \mathsf{T}$
Euclideanness: $v \in \sigma_w \Rightarrow \sigma_w \subseteq \sigma_v$ $\qquad$ $\forall w, \sigma : [\![\diamond\phi \supset \square\diamond\phi]\!]^{\sigma,w} = \mathsf{T}$

These constraints imply that information states are epistemically transparent:

Transparency: $v \in \sigma_w \Rightarrow \sigma_w = \sigma_v$
$\qquad\qquad\qquad$ $\forall w, v, \sigma : v \in \sigma_w \Rightarrow [\![\diamond\phi]\!]^{\sigma,w} = [\![\diamond\phi]\!]^{\sigma,v}$

Given Transparency, epistemic modalities are "rigid" relative to a choice of σ and w: if ϕ is a sentence of the form $\diamond\psi$ or $\square\psi$ and $[\![\phi]\!]^{\sigma,w} = \mathsf{T}$, then, for any $v \in \sigma_w$, $[\![\phi]\!]^{\sigma,v} = \mathsf{T}$. Hence, whenever $[\![\phi]\!]^{\sigma,w} = \mathsf{T}$:

$$\sigma_w \subseteq [\![\phi]\!]^{\sigma}$$

---

1 On the empirical case for Euclideanness, see Appendix B, and [OMITTED]. These are standard assumptions in both the semantics (e.g., von Fintel & Gillies 2010; Gillies 2010; von Fintel & Gillies 2011) and logic (e.g., Holliday & Icard III 2010) of epistemic modalities. The phenomena of interest in this paper will also arise for modalities of belief (axiomatized by KD45, rather than S5).

Now let us consider a probabilistically coherent agent[2] A for whom p is possible (alternatively, necessary), but who has a graded attitude—confidence $\in (0,1)$—in $\Diamond p$ (alternatively, $\Box p$). It is not possible to represent this sort of agent in the present semantic framework. Let σ_w^A represent A's information at w. If p is possible (alternatively, necessary) for A at w, then $[\![\Diamond p]\!]^{\sigma^A,w} = \mathsf{T}$ (alternatively, $[\![\Box p]\!]^{\sigma^A,w} = \mathsf{T}$). But then, in view of Transparency, $[\![\Box\Diamond p]\!]^{\sigma^A,w} = \mathsf{T}$ (alternatively, $[\![\Box\Box p]\!]^{\sigma^A,w} = \mathsf{T}$). The difficulty is this: if A is probabilistically coherent, and A's information entails $\Diamond p$ (alternatively, $\Box p$), then A must assign $\Diamond p$ (alternatively, $\Box p$) probability 1. It follows that A's confidence in $\Diamond p$ or $\Box p$ must be extremal (0 or 1).

Proof. Consider any probabilistically coherent agent A; let σ_w^A be A's information at w and Pr_w^A be A's probability measure at w. Either $\exists v \in \sigma_w : [\![p]\!]^{\sigma,v} = \mathsf{T}$ or $\forall v \in \sigma_w : [\![p]\!]^{\sigma,v} = \mathsf{F}$. If $\exists v \in \sigma_w : [\![p]\!]^{\sigma,v} = \mathsf{T}$, then $[\![\Diamond p]\!]^{\sigma,w} = \mathsf{T}$, in which case $\sigma_w \subseteq [\![\Diamond p]\!]^{\sigma}$. Since A is probabilistically coherent, $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 1$. If $\forall v \in \sigma_w : [\![p]\!]^{\sigma,v} = \mathsf{F}$, then $[\![\Diamond p]\!]^{\sigma,w} = \mathsf{F}$, in which case $\sigma_w \cap [\![\Diamond p]\!]^{\sigma} = \emptyset$. Since A is probabilistically coherent, $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 0$. Thus, either $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 1$ or $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 0$. $\qquad\square$

It is worth underlining: this is **not** an artifact of the use of possibilia to model information. So long as the class of models for an epistemically modal language is required to satisfy object language analogues of Reflexivity ($\Box\phi \supset \phi$) and Euclideanness ($\Diamond\phi \supset \Box\Diamond\phi$), the logic of epistemic modality will be constrained by the following entailments:

$$\Box\phi \dashv\vdash \Box\Box\phi$$

$$\Diamond\phi \dashv\vdash \Box\Diamond\phi$$

Let I_w^A designate A's information at w; we will *not* assume that I_w^A is a set of possible worlds. Now either $[\![\Diamond p]\!]^{I^A,w} = \mathsf{T}$ (if I_w^A is compatible with p) or $[\![\Diamond p]\!]^{I^A,w} = \mathsf{F}$ (otherwise). If $[\![\Diamond p]\!]^{I^A,w} = \mathsf{T}$, then $[\![\Box\Diamond p]\!]^{I^A,w} = \mathsf{T}$, in which case I_w^A is incompatible with $\neg\Diamond p$ (i.e., I_w^A entails $\Diamond p$). Since A is probabilistically coherent, $Pr_w^A(\Diamond p) = 1$. If, on the other hand, $[\![\Diamond p]\!]^{I^A,w} = \mathsf{F}$, then $[\![\Box\neg\Diamond p]\!]^{I^A,w} = \mathsf{T}$, in which case I_w^A is incompatible with $\Diamond p$ (i.e., I_w^A entails $\neg\Diamond p$). Since A is probabilistically coherent, $Pr_w^A(\Diamond p) = 0$. Thus, again, either $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 1$ or $Pr_w^A([\![\Diamond p]\!]^{\sigma}) = 0$.

The following commonplaces are, then, simply incompatible:

- Models for epistemic modality are Reflexive and Euclidean.
- Probabilistically coherent agents can assign non-extremal probability to epistemic modalities.

We seem to confront a choice: between a revisionary (prima facie) logic of epistemic modality, and a revisionary (prima facie) understanding of the attitudes it is possible to coherently bear toward sentences expressing subjective uncertainty.

Really, though, this is no choice at all. Even if it epistemic modalities somehow failed to logically entail their epistemic necessitations, it would remain a conceptual truth about probabilistically coherent agents that propositions describing global

---

2 A probabilistically coherent agent is one whose degrees of belief in a σ-algebra of W are representable with a probability function.

features of their information receive probability 1 or 0. Relative to any model for a probabilistically coherent agent, then, epistemic modalities would receive extremal probabilities. Which is just to say that probabilistically coherent agents cannot assign (non-extremal) probabilities to epistemic modalities. If actual agents can have graded attitudes toward epistemic modalities, therefore, we cannot represent these graded attitudes as probabilities.

Nor, apparently, can we represent graded attitudes toward epistemic modalities using *sets* of probability measures. According to the "Bayesian" proposal for representing such attitudes (Yalcin 2012; Rothschild 2012), "Where an agent assigns a determinate probability to a proposition, every measure in their credal set [i.e., the set of probability measures compatible with their information] assigns that probability to it. A probabilistic claim is true of a credal set just in case it is true on every probability measure in the set" (Rothschild 2012: 110). The difficulty is that, given the arguments of this section, a set of probability measures S is constrained so that, if ϕ is epistemically modal:

$$\forall Pr \in S : Pr(\phi) = 0 \text{ or } Pr(\phi) = 1$$

Attitudes of intermediate confidence (e.g., confidence n) toward a sentence ϕ are represented, according to the Bayesian proposal, with sets of probability measures, all of which assign probability $\geq n$ to ϕ. No probability measure assigns intermediate confidence to ϕ if ϕ is epistemically modal. And so, given the Bayesian proposal, no set of probability measures can represent attitudes of intermediate confidence toward epistemic modalities.

3   Distinguishing the Objects of Gradability and Probability?

Moss (2015) describes an elegant system that can account for embeddings like (1), without *actually* assigning probabilities to the semantic value of the epistemic modal. Moss holds that a sentence of the form $\triangle\phi$ (read: it is probable that ϕ) expresses a constraint on probability measures, namely, the constraint that one's probability measure fall within the following set:

$$[\![\triangle_1\phi]\!]^c = [\![\triangle_1]\!]^c([\![\phi]\!]^c) = \{m : m(\bigcup\{p \in g_c(1) : m|_p \in [\![\phi]\!]^c\}) > .5\}$$

On Moss' semantics, modals are interpreted relative to contextually salient partitions; numerical indices (like $_1$) are mapped to contextually salient partitions by a contextual variable assignment g_c. Thus, $\triangle\phi$ expresses a constraint on probability measures that m satisfies iff m assigns this proposition a value exceeding .5:

$$\bigcup\{p \in g_c(1) : m|_p \in [\![\phi]\!]^c\}$$

The object that ultimately receives a probability is the disjunction of those propositions in the salient partition that confirm ϕ—i.e., the disjunction of those propositions p such that, if m were conditionalized on p, m would satisfy the constraint expressed by ϕ. Thus, if ϕ is epistemically modal, $\triangle\phi$ is not represented as actually ascribing a probability to (or expressing a constraint regarding one's assignment of probability to) the semantic value of the probability operator's sentential complement. $\triangle\phi$ semantically rules out probability measures that do not regard as likely the disjunction

4

of the propositions conditionalization on which is sufficient for believing ϕ. It does **not** rule out probability measures that fail to regard ϕ as likely; indeed, in many cases (e.g., if ϕ is epistemically modal), Moss would deny that ϕ is the sort of thing that can strictly receive a probability at all (since it does not, according to Moss' semantics, express a possible worlds proposition).

Moss' system smoothly accounts for constructions *in natural language* that express graded attitudes towards epistemic modalities. Spotting ourselves the requisite compositional bells and whistles, the system can be extended to account for attitude-ascriptions ascribing such attitudes:

(4)   Alice thinks it is probably the case that Bob might be hired.

A sentence like (4) will say, roughly, that Alice is representable as satisfying the constraint expressed, on Moss' semantics, by (1). More roughly still, (4) attributes to Alice the attitude of thinking it probable that at least one of the propositions p such that belief in p is sufficient for thinking Bob might be hired, is true.

But does this sort of story provide an account of the nature of *graded attitudes* towards epistemic modalities? Yes, but only if the cognitive structure of such attitudes is assumed to parallel the semantic structure of sentences ascribing such attitudes—if, that is to say, such attitudes are represented, not as ascriptions of probability-like values to epistemically modal representations, but rather as ascriptions of probabilities to "proxy representations" (i.e., possible worlds propositions) that cognitively underwrite epistemically modal quasi-representations.

Objection: this is a substantive thesis about the structure of such cognition, one in apparent tension with these commonplaces: (i) because of their functional role in cognition, epistemically modal thoughts constitute (at least a kind of) representations (see esp. Schroeder 2011) (ii) a graded attitude (e.g., being 50% confident that ϕ) is a natural cognitive kind, with a unified type of cognitive realizer—namely, the representation that ϕ being mapped to a middle point on a bounded scale whose endpoints represent outright acceptance and outright rejection. Moss relies on a different understanding of the relevant cognitive states—one that is revisionary with respect to this familiar cognitive model. And that is a cost.[3]

Partial Reply: in Moss' semantics, sentences expressing graded attitudes (e.g., the attitude of thinking ϕ probable) do form a semantic natural kind: they all express a constraint that one satisfies iff *one assigns a proxy representation a high probability* (recall Moss' proposal for $[\![\triangle_1\phi]\!]^c$ above). When ϕ is non-modal, the proxy representation is the possible worlds proposition expressed by ϕ; when ϕ is modal, the proxy representation is the disjunction of possible worlds propositions conditionalization on which is sufficient for believing ϕ. This reveals the sense in which graded attitudes are natural cognitive kinds: in every case, thinking ϕ probable may be analyzed as thinking p probable, for some possible worlds proposition p such that conditionalizing on p is sufficient for believing ϕ.

Objection: it remains true that the only "basic" gradability in Moss' model is probabilistic gradability—gradability of the sort of that attaches to the elements of a σ-algebra over W, assessed as more or less probable, relative to a probability measure defined over that algebra. Gradability is *not* a uniform feature of the

---

3 For a similar style of argument, see Schroeder (2011).

representations semantically encoded in sentences of natural language; it is only a feature of *worldly* representations. This is a cost, prima facie, since it seems the theorist needs some way of talking, within the metalanguage, about an agent's confidence in the representation semantically encoded in a sentence like (3). The (at least, this) theorist would like to be able say that (1) expresses, and (4) ascribes, high relative confidence in the representation semantically encoded in (3).

Here is a related dimension of this problem. The theorist would also like to account for why graded attitudes towards representations of any type are apparently subject to a *single set of rational norms*—why, for instance, it is a rational mistake to believe either of (5) or (6):[4]

(5)   It is probably the case that Bob might be hired, and it's probably the case that he can't be hired.

(6)   The Eagles will probably win, and they probably won't win.

A natural explanation is that an agent's degree of confidence in representation R summed with her degree of confidence in $\neg R$ should not exceed 1. On the face of things, this sort of explanation apparently requires (i) countenancing degrees of confidence as properties of representations of any type, (ii) describing (and motivating) rational norms that govern the suite of graded attitudes agents can bear toward representations of any type.

Similarly, the theorist would like to account for how graded attitudes towards representations of any type *constrain rational action*. Consider an agent who thinks it is 50% likely that it is 25% likely that p. Such an agent should be willing to accept a bet whose payoff, if it is 25% likely that p, is more valuable to her than the payoff, if it is not 25% likely that p. It would be elegant if we could explain this normative fact by appeal to a generalization of the calculus of Expected Value. But such a generalization apparently requires (i) countenancing degrees of confidence as properties of representations of any type, (ii) describing (and motivating) rational norms that govern the relations between the suite of graded attitudes agents can bear toward representations of any type, and action.

4   Credences in Representations

Our challenge is that the semantic content of an epistemic modal ϕ, relative to an agent A's information, is apparently unfit for being the object of graded attitudes (e.g., A being 10% confident that ϕ), according to standard techniques for modeling those attitudes. Moss' view avoids the challenge, by denying that that epistemically modal representations, as such, are the object of graded attitudes. I have argued that Moss' view is revisionary, prima facie, with respect to the folk understanding of graded attitudes.[5] Let us consider how an alternative might be formulated.

---

4 This sort of problem receives discussion in Schroeder (2011); Staffel (forthcoming).

5 I don't, however, want to assume that there could be no sort of theoretical equivalence between the sorts of explanations that are statable within Moss' theory, and the sorts of explanations that will be statable within the theory pursued here. Indeed, one thing that is distinctive of Moss' view is its ability to replicate explanations ordinarily thought to presuppose a propositional (and, more generally, representational) notion of semantic content (here see esp. Moss 2013).

4.1 Introducing Credences

The most obvious "fix", if it could be pulled off, would be to introduce a new quantity, call it *credence*, for probability operators of natural language to uniformly express. To utter a sentence like (1) is simply to express high credence in the representation semantically encoded in (3).

$$\llbracket \triangle\phi \rrbracket = \llbracket \triangle \rrbracket(\llbracket \phi \rrbracket) = \{Cr : Cr(\llbracket \phi \rrbracket) > .5\}$$

Here is how I prefer to conceptualize this idea. Some credences are *probabilities*: subjective estimates of objective chance of the truth of a worldly representation (alternatively, subjective estimates of actual-worldly truth value). Some credences are *not probabilities* (when a subject's credence cannot be understood as their estimate of objective chance of the truth of a worldly representation, or as a subjective estimate of actual-worldly truth value). We can nevertheless define Cr so that it *behaves like* a probability function, whether its argument is a worldly or non-worldly representation.[6] First, assume that a *representation* is a set of objects of arbitrary semantic type—a set of alternative possibilities. Only some sets of alternative possibilities (i.e., sets of possible worlds) represent sets of candidates for actuality; other sets of alternative possibilities represent sets of candidates for attitudes like endorsement, selection, or adoption.[7]

Definition 1. *A set of representations $\{\mathcal{R}_1, ..., \mathcal{R}_n\}$ **partitions** $\mathcal{R}$ iff, for all $1 \leq i \neq j \leq n$:*

$$\mathcal{R}_i \cap \mathcal{R}_j = \emptyset$$
$$\bigcup_{i=1}^{n} \mathcal{R}_i = \mathcal{R}$$

Definition 2. *An **alternative set** for $\mathcal{R}$ is any set $\mathfrak{R}$ that partitions $\mathcal{R}$.*

Definition 3. *If $\mathfrak{R}$ is an alternative set for $\mathcal{R}$, $\mathfrak{R}$'s σ-**closure** Σ is the closure of $\mathfrak{R}$ under $\cap$, $\cup$, and $\bar{\cdot}$.*

Consider any representation $\mathcal{R}$, alternative set $\mathfrak{R} = \{\mathcal{R}_1, ..., \mathcal{R}_n\}$ for $\mathcal{R}$, and $\mathfrak{R}$'s σ-closure Σ.

Definition 4. *A **credence function** based on $\mathcal{R}$, $\mathfrak{R}$, and Σ is a function $Cr : \Sigma \mapsto$*

---

6 Though the details are very different, this general perspective draws inspiration from Bradley's "Multi-Dimensional" approach towards the probabilities of indicative conditionals (Bradley 2012), as well as remarks in Staffel (forthcoming) discussing how an Expressivist might model the descriptive and normative characteristics of gradable attitudes towards non-factual semantic contents. Jonathan Weisberg (pc) alerts me to an earlier approach to higher-order probability (Hild 1998) that is similar in both spirit and certain modeling choices to the one developed here.

7 This is only a rough first pass at stating a functional psychological role for the representation of a set of alternatives of arbitrary semantic type. It will be elaborated in §5.

[0, 1] *such that:*

$$Cr(\mathcal{R}) = 1$$

$$Cr(\bigcup_{i=1}^{n} S_i) = \sum_{i=1}^{n} Cr(S_i) \quad (i \neq j \Rightarrow S_i \cap S_j = \emptyset)$$

We will allow that, possibly, $n = \infty$, in which case Cr is constrained by Normalization $[Cr(\mathcal{R}) = 1]$ and Countable Additivity $[Cr(\bigcup_{i=1}^{\infty} S_i) = \sum_{i=1}^{\infty} Cr(S_i)]$. Ordinarily, however, $n \in \mathbb{N}$, in which case Cr is constrained by Normalization $[Cr(\mathcal{R}) = 1]$ and Finite Additivity $[Cr(\bigcup_{i=1}^{n} S_i) = \sum_{i=1}^{n} Cr(S_i)]$.

Definition 5. *Given a credence function Cr based on $\mathcal{R}$, $\mathfrak{R}$, Σ, and $\mathcal{T} \in \Sigma$ the* **conditionalization** *of Cr on $\mathcal{T}$ is a function $Cr|_{\mathcal{T}}(\cdot) : \Sigma \mapsto [0, 1]$ such that:*

i. $Cr|_{\mathcal{T}}$ is a credence function based on:

$$\langle \mathcal{R} \cap \mathcal{T}, \{\mathcal{R}' \cap \mathcal{T} : \mathcal{R}' \in \mathfrak{R}\}, \{\mathcal{R}' \cap \mathcal{T} : \mathcal{R}' \in \Sigma\} \rangle$$

ii. If $\mathcal{R} \cap \mathcal{T} \subseteq \mathcal{U}$, then $Cr|_{\mathcal{T}}(\mathcal{U}) = 1$
iii. If $\mathcal{R} \cap \mathcal{T} \cap \mathcal{U} = \emptyset$, then $Cr|_{\mathcal{T}}(\mathcal{U}) = 0$
iv. Otherwise, $Cr|_{\mathcal{T}}(\mathcal{U}) = \frac{Cr(\mathcal{U} \cap \mathcal{T})}{Cr(\mathcal{T})}$

Definition 6. *Given a credence function Cr based on $\mathcal{R}$, $\mathfrak{R}$, and Σ, a* **conditional credence function** *based on $\mathcal{R}$, $\mathfrak{R}$, and Σ is a two-place function $Cr(\cdot|\cdot) : \Sigma \mapsto (\Sigma \mapsto [0, 1])$ such that $Cr(S|\mathcal{T}) = Cr|_{\mathcal{T}}(S)$.*

4.2 *A Semantics of Representations*

The guiding semantic idea is that sentences of natural language *semantically encode representations*. Consider a language containing a denumerable stock of propositional atoms A, Boolean compounds of sentences, the indicative conditional $\rightarrow$, the 'probably' operator $\triangle$, and the epistemic possibility modal $\Diamond$.

$$\phi :: A \mid \neg\phi \mid \phi \wedge \psi \mid \phi \rightarrow \psi \mid \triangle\phi \mid \Diamond\phi$$

An interpretation function for this language maps sentences into representations. The obvious clauses would be as follows:

$$[\![p]\!] = \{w : w(p) = 1\} \ (p \in A)$$
$$[\![\neg\phi]\!]_\tau = \mathbb{U}_\tau - [\![\phi]\!]_\tau \ (X_\tau := X \text{ is a set of objects of semantic type } \tau)$$
$$[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$$
$$[\![\phi \rightarrow \psi]\!] = \{Cr : Cr([\![\psi]\!]|[\![\phi]\!]) = 1\}$$
$$[\![\triangle\phi]\!] = \{Cr : \sum_{S \in [\![\phi]\!]} Cr(\{S\}) > .5\}$$
$$[\![\Diamond\phi]\!] = \{S : S \cap [\![\phi]\!] \neq \emptyset\}$$

The "obvious" clauses, alas, do not quite work for our purposes. We have generalized the probability calculus to credence functions, by requiring any credence function to

8

be determined relative to a (i) "base" representation $\mathcal{R}$, (ii) an alternative set $\mathfrak{R}$ for $\mathcal{R}$ that partitions $\mathcal{R}$, (iii) $\mathfrak{R}$'s σ-closure Σ. The semantic clauses above do not reflect this relativity, and so must be revised.

We will therefore say that the representation expressed by such sentences is determined relative to a base representation $\mathcal{R}$, an alternative set $\mathfrak{R}$ for $\mathcal{R}$, and $\mathfrak{R}$'s σ-closure Σ. We will call a triple $\langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$ with these characteristics a **space**, and we will allow a context c to determine (via a contextually determined variable assignment g_c) a space (of the requisite semantic type) for each space-sensitive expression of our language.[8]

$$\llbracket \phi \rightarrow_1 \psi \rrbracket^c = \lambda Cr : Cr \text{ is based on } g_c(1) \, . \, Cr(\llbracket \psi \rrbracket^c | \llbracket \phi \rrbracket^c) = 1$$

$$\llbracket \triangle_1 \phi \rrbracket^c = \lambda Cr : Cr \text{ is based on } g_c(1) \, . \, \sum_{S \in \llbracket \phi \rrbracket^c} Cr(\{S\}) > .5$$

$$\llbracket \Diamond_1 \phi \rrbracket^c = \lambda S : S \text{ is based on } g_c(1) \, . \, S \cap \llbracket \phi \rrbracket^c \neq \emptyset$$

The representation expressed by $\phi \rightarrow_1 \psi$ relative to $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$ is a property of credence functions $[\lambda Cr.Cr(\llbracket \psi \rrbracket^c | \llbracket \phi \rrbracket^c) = 1]$ (equivalently, where types require, the set of credence functions with this property), which is *undefined* for any Cr not based on $\mathcal{R}$, $\mathfrak{R}$, and Σ. The idea is the same for $\triangle \phi$. $\Diamond \phi$ expresses a property undefined for any $S \notin \mathfrak{R}$—for any representation, that is to say, not among the alternative possibilities that are "live" from the point-of-view of $\mathfrak{R}$ (cf. Yalcin 2011). For any $S \in \mathfrak{R}$, S satisfies this property iff S is compatible with the representation expressed by ϕ.

4.3   Examples

Example 1.   Consider the case of $\triangle$ scoping over a propositional atom:

$$\llbracket \triangle_1 p \rrbracket^c = \{Cr : \sum_{w \in \llbracket p \rrbracket^c} Cr(\{w\}) > .5\}$$
$$= \{Cr : Cr(\{w : w(p) = \mathsf{T}\}) > .5\}$$

Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$. Here, the semantic types require that $\mathcal{R}$ be a set of worlds, e.g., $\{w, v, u\}$; $\mathfrak{R}$ is a partition of $\mathcal{R}$, e.g., $\{\{w, v\}, \{u\}\}$; Σ is $\mathfrak{R}$'s σ-closure. As intended, a sentence of the form $\triangle p$ expresses the property a credence function (based on $\mathcal{R}$, $\mathfrak{R}$, and Σ) has when it assigns the worldly representation encoded in p a value $> .5$.

Example 2.   Next consider an example involving $\triangle$ iterated over $\Diamond$.

$$\llbracket \triangle_1 \Diamond_2 p \rrbracket^c = \{Cr : \sum_{S \in \llbracket \Diamond_2 p \rrbracket^c} Cr(\{S\}) > .5\}$$

---

8 Though this notation (and the implementation via contextually determined variable assignments) is from Moss, note that our variable indices play a very different role. (Note, for instance, that the indices in Moss' semantics uniformly resolve to partitions of W.)

Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$.[9] Here the semantic types require that $\mathcal{R}$ be a set of sets of worlds, e.g., $\{\{w\}, \{w, v\}\}$; $\mathfrak{R}$ is a partition of $\mathcal{R}$; Σ is $\mathfrak{R}$'s σ-closure. As intended, $\triangle\Diamond p$ expresses the property a credence function has when it assigns the non-worldly representation encoded in $[\![\Diamond p]\!]^c$ a value $> .5$.

Example 3. Next consider the reverse iteration:

$$
\begin{aligned}
[\![\Diamond_1 \triangle_2 p]\!]^c &= \{\mathcal{S} : \mathcal{S} \cap [\![\triangle_2 p]\!]^c \neq \emptyset\} \\
&= \{\mathcal{S} : \mathcal{S} \cap \{Cr : \sum_{w \in [\![p]\!]^c} Cr(\{w\}) > .5\} \neq \emptyset\} \\
&= \{\mathcal{S} : \mathcal{S} \cap \{Cr : Cr(\{w : w(p) = \mathsf{T}\}) > .5\} \neq \emptyset\}
\end{aligned}
$$

Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$. Here the semantic types require that $\mathcal{R}$ be a set of credence functions; $\mathfrak{R}$ is a partition of $\mathcal{R}$; Σ is $\mathfrak{R}$'s σ-closure. As intended, $\Diamond\triangle p$ expresses the property a set of credence functions has when it contains a credence function that assigns the worldly representation encoded in p a value $> .5$.

Example 4. Finally two examples involving iterated epistemics:

$$
\begin{aligned}
[\![\Diamond_1 \Diamond_2 p]\!]^c &= \{\mathcal{S} : \mathcal{S} \cap [\![\Diamond_2 p]\!]^c \neq \emptyset\} \\
&= \{\mathcal{S} : \mathcal{S} \cap \{\mathcal{T} : \mathcal{T} \cap [\![p]\!]^c \neq \emptyset\} \neq \emptyset\}
\end{aligned}
$$

Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$. Here the semantic types require that $\mathcal{R}$ be a set of sets of worlds (i.e., a set of worldly propositions); $\mathfrak{R}$ is a partition of $\mathcal{R}$; Σ is $\mathfrak{R}$'s σ-closure. As intended, $\Diamond\Diamond p$ expresses the property a set of sets of worlds has when it contains a set of worlds that is compatible with p.

$$
\begin{aligned}
[\![\Diamond_1 \Diamond_2 \Diamond_3 p]\!]^c &= \{\mathcal{S} : \mathcal{S} \cap [\![\Diamond_2 \Diamond_3 p]\!]^c \neq \emptyset\} \\
&= \{\mathcal{S} : \mathcal{S} \cap \{\mathcal{U} : \mathcal{U} \cap \{\mathcal{T} : \mathcal{T} \cap [\![p]\!]^c \neq \emptyset\} \neq \emptyset\}\}
\end{aligned}
$$

Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$. Here the semantic types require that $\mathcal{R}$ be a set of sets of worldly propositions; $\mathfrak{R}$ is a partition of $\mathcal{R}$; Σ is $\mathfrak{R}$'s σ-closure. As intended, $\Diamond\Diamond\Diamond p$ expresses the property a set of sets of worldly propositions has when it contains a set of worldly propositions that is compatible with $\Diamond p$.

4.4 Compositionality and Polymorphic Types

The interesting operators of our language ($\rightarrow$, $\triangle$, $\Diamond$) uniformly take set-type meanings (representations) as arguments. This gives our system a veneer of composition-

---

9 I will generally suppress the role of space-sensitivity for embedded modals. Strictly speaking:

$$
[\![\triangle_1 \Diamond_2 p]\!]^c = \lambda Cr : Cr \text{ based on } g_c(1) . \sum_{\mathcal{S} \in [\![\Diamond_2 p]\!]^c} Cr(\{\mathcal{S}\}) > .5
$$

Notice: $[\![\Diamond_2 p]\!]^c$ is defined for $\mathcal{S}$ only when $\mathcal{S} \in g_c(2)$. Therefore:

$$
[\![\triangle_1 \Diamond_2 p]\!]^c = \lambda Cr : Cr \text{ based on } g_c(1) . \sum_{\mathcal{S} \text{ based on } g_c(2) \wedge \mathcal{S} \cap [\![p]\!]^c \neq \emptyset} Cr(\{\mathcal{S}\}) > .5
$$

This starts to strain the eye, so I will generally leave it to the reader to fill in such formal details.

ality, but, for now, only the veneer. Set-type meanings are, strictly speaking, *not* typically the semantic values of these operators' complements; the semantic values of the sentences of our language, in fact, comprise a manifold of *functional* types.[10] Here is an illustration: $\Diamond$ can semantically combine with a worldly representation $[\![p]\!] :: \langle s, t \rangle$,[11] an epistemically modal representation $[\![\Diamond p]\!] :: \langle\langle s, t \rangle, t \rangle$, an epistemically modal representation with epistemically modal content $[\![\Diamond\Diamond p]\!] :: \langle\langle\langle s, t \rangle, t \rangle, t \rangle$, and so on, *ad infinitum*.

The interesting operators of our language, therefore, have the (perhaps surprising) property of being *unselective* as to the semantic type of their complements, so long as that semantic type is isomorphic to a set (i.e., so long as that semantic type is of the form $\langle \tau, t \rangle$, for some type τ). This means that they will have a recursive (*polymorphic*[12]) semantic type τ^*:

$$\tau^* ::= \langle \alpha, \langle \alpha, t \rangle \rangle \qquad\qquad \alpha ::= \langle s, t \rangle \mid \langle \gamma, t \rangle \mid \langle \alpha, t \rangle$$
$$\gamma ::= \langle \alpha, v_{[0,1]} \rangle$$

The easiest way to think about polymorphic types is this. An expression like $\Diamond$ has a semantic type, in two guises: qua *expression-type* (in which case its type is polymorphic) and qua *expression-token* (in which case its type, as tokened on an occasion of use, is a type drawn from the polymorphic type hierarchy). The semantic type of, e.g., $\Diamond$, as tokened on an occasion of use will "depend" (very loosely speaking[13]) on the semantic type of its complement (but will always be drawn from the hierarchy of types introduced here).[14]

5   Two Aspects of Our Mental Lives

The last section showed that the notion of credences *in* epistemically modal and probabilistic representations, constrained by the probability axioms, is both mathematically and semantically tractable. But—and I intend this question seriously—does it *make sense*? We have introduced a semantic hierarchy of representations with no upper bound on the complexity of the semantic type of a representation. Is this cognitively realistic? (Here, I will argue: yes.) We have assumed that objects at any

---

10 I here assume that semantic composition is always via Function-Argument Application.

11 Notation: s is the type of worlds, t is the type of truth values. A function of type $\langle \tau, \tau' \rangle$ is a function from objects of type τ into objects of type τ'.

12 For another application of polymorphic types, see Charlow (forthcoming).

13 This is no violation of compositionality: the semantic type of $\Diamond$, as tokened on an occasion of use, is not semantically *determined by*, or selected *in virtue of*, the semantic type of its complement. It is simply to say that, if $\Diamond$ occurs in a semantically well-formed expression, its semantic type must be drawn from the hierarchy of types defined above, and must be of the right type to compose, by Function-Argument Application, with the semantic value of its sister.

14 In order to be fully compositional, our system also requires a understanding of *semantic coordination* (with, e.g., $\wedge$). We currently understand $\wedge$ as expressing $\cap$, but there are two reasons this will not work. First, the semantic values of $\wedge$'s arguments are functions, not sets. (This is trivial to fix, and I will continue to talk as if the difference between a characteristic function and a set is no difference at all.) Second, the semantic values of $\wedge$'s arguments are frequently sentences of different semantic type. This is less trivial to fix: we will require a generalized understanding of conjunction that allows it to coordinate constituents of different semantic type, as in Partee & Rooth (1983). To keep the main discussion maximally simple, I will ignore this sort of complication here (though I will address it briefly in Appendix A). For further discussion, see [OMITTED].

level of the type hierarchy can receive credences (where credences are constrained by assumptions of Normalization and Countable Additivity). Is this normatively plausible—do the standard justifications for Normalization and Countable Additivity apply, when credences do not take worldly representations as their arguments? (Here, I will argue: yes, or close enough.)

5.1 Cognitive

Moss offers an argument against a proposed extension of the Bayesian proposal pursued in Rothschild (2012); Yalcin (2012) to graded modal judgments:

> [I]t is hard to imagine a reason for ruling that embeddings of epistemic vocabulary beyond a certain level of complexity are are semantically uninterpretable. In the absence of such a reason, our theory should deliver semantic values for embeddings of arbitrary complexity. Hence in order to repair the [Bayesian] proposal, we would have to model subjects as having not just sets of sets of measures as mental states, but sets of sets of sets of measures, and so on. It is difficult to independently motivate such an arcane model of our mental life. (Moss 2015: 30)

While our proposal isn't quite Bayesian in the sense of Rothschild (2012); Yalcin (2012), Moss' critique clearly applies.

But the charge that this is an "arcane model" does not, I'll argue, really bite. To illustrate, recall that, on the view defended here:

$$\llbracket \Diamond_1 \Diamond_2 p \rrbracket^c = \{\mathcal{S} : \mathcal{S} \cap \{\mathcal{T} : \mathcal{T} \cap \llbracket p \rrbracket^c \neq \emptyset\} \neq \emptyset\}$$

On our view, $\Diamond\Diamond p$ expresses the property a set of sets of worlds (i.e., a set of worldly propositions) has when it contains a set of worlds that is compatible with p. To think or call such a sentence *probable* is to express a property of credences in sets of sets of sets of worlds (i.e., sets of sets of worldly propositions)—namely, the property of assigning a credence $> .5$ to $\llbracket \Diamond_1 \Diamond_2 p \rrbracket^c$.

Agents can treat *any set of objects* as a set of alternatives **for cognitive purpose P**. Agents can represent sets of possible worlds for the purpose of representing different abstract alternatives (individual possibilia) for accurately representing the world. Agents can represent sets of sets of possible worlds (i.e., sets of propositions) in order to represent different alternatives (propositions)—*not* for the purpose of accurately representing the world, instead for the purpose of representing alternative ways of representing the world (e.g., alternatives that treat *p* as possible *versus* those that treat *p* as impossible). Agents can represent sets of sets of sets of possible worlds (i.e., sets of sets of propositions) in order to represent different alternatives (sets of propositions)—*not* for the purpose of accurately representing the world, *nor* for the purpose of representing alternative ways of representing the world, instead for the purpose of representing alternative ways of representing alternative ways of representing the world.

Representations, as we understand them, have an iterative, or recursive, structure (but I do not think this should be viewed as surprising). But the cognitive state of *representing R for purpose P* is not arcane: it is the attitude of representing the various alternatives of $\mathcal{R}$ as candidates for fulfilling P. We have understood

the attitude of representation more expansively than is traditional[15]—in particular, we have relativized representations to cognitive purposes, and have declined to assume that the functional role of representation is uniformly about representing individual possibilia as candidates for actuality. Generalizing a familiar notion need not, however, render it arcane. Indeed, given this generalized understanding of representation, representing a set of alternatives as candidates for fulfilling P describes a sort of familiar cognitive activity in which agents might plausibly engage.

5.2 Normative

Why represent agents as having credences in non-worldly representations? There were several motivations, one theoretical, two normative (§3). The theoretical motivation: to vindicate metalinguistic ascriptions of confidence *in* epistemically modal representations. The normative motivations: to describe and justify rational norms governing confidence in epistemically modal representations; and to describe and justify rational norms governing the relationship between confidence in epistemically modal representations, and action. Our present account satisfies the theoretical motivation. What about the normative?

Proposal: the rational norms governing the relationship between credence and action are a generalization of the theory of Expected Value. The Expected Value of any action can be represented in a standard form. Consider a decision problem Π. (C_i is a relevant contingency, $\mathbf{A}_j$ an available action, $Val(\mathbf{A}_j|C_i)$ a numerical representation of the value associated with performing $\mathbf{A}_j$ if C_i.[16])

Π	$Cr(C_1|\mathbf{A}_j)$	...	$Cr(C_n|\mathbf{A}_j)$
$\mathbf{A}_1$	$Val(\mathbf{A}_1|C_1)$	...	$Val(\mathbf{A}_1|C_n)$
...	...	...	...
$\mathbf{A}_m$	$Val(\mathbf{A}_m|C_1)$	...	$Val(\mathbf{A}_m|C_n)$

Next, consider any space $\langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$.

Definition 7. Π *is **based on** $\langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$ iff (i) $C_i \in \mathfrak{R}$ (1 ≤ i ≤ n), (ii) Cr is based on $\mathcal{R}$, $\mathfrak{R}$, and Σ.*

Consider a decision problem Π with the above characteristics, and which is based on $\langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$.

Definition 8. *The **expected value** of $\mathbf{A}$ in Π is a weighted sum of credences multiplied by values:*

$$\sum_{i=1}^{n} Cr(C_i|\mathbf{A})Val(\mathbf{A}|C_i)$$

---

15 The view I am attempting to generalize here is the one developed in Stalnaker (1984).

16 How precisely to formalize a conditional value function is a matter of controversy in the literature—one I will bracket here. I will here assume Paul Weirich's informal notion of conditional value, according to which it is degree of desirability, under indicative supposition (Weirich 1980). For relevant further discussion, see Joyce (1999: Ch. 4).

Why should a rational agent maximize expected value, thus defined? More specifically—since I want to bracket controversies about how best to formulate a mathematical theory of rational action—why should an agent who wants to maximize expected value compute expected value *using a credence function* (the properties of which are constrained by Definition 4)?

There are two main ways of answering this question in the literature. First, Dutch Book Arguments, on which, roughly, agents who have incoherent credences are irrational because subject to sure losses (for an overview, see Hájek 2009). Second, Accuracy Arguments, on which, roughly, agents who want to maximize expected *epistemic* value (roughly, the proximity of one's credences to the truth), but who have incoherent credences, are irrational because coherent credences are always more proximal to the truth (originating with Joyce 1998). Let us see about the prospects of extending these answers to the present account.

Matters are, not surprisingly, less than straightforward with Accuracy Arguments. *Accuracy* is fundamentally a worldly notion: a representation is said to be accurate when it is satisfied ("true") as evaluated against a possibility taken to represent *actuality* (a point also noted by Staffel forthcoming). Accuracy Arguments purport to show that subjective estimates of objective chance that violate the axioms of probability are rationally defective, since, for any such estimate, there is another way of estimating chances that (i) satisfies the axioms of probability and (ii) is guaranteed to be overall more accurate in *w*, for any possible world *w* to which the agent assigns some credence (see esp. the accuracy theorem of Joyce 1998). In order to adapt Accuracy Arguments to the framework proposed here, we would require a non-worldly proxy for the notion of Accuracy (as well as a non-worldly proxy for the notion of actuality). The prospects here strike me as very dim—particularly given the conceptualization of our theory suggested in §4 [OMITTED].

Happily, matters are clearer with Dutch Book Arguments, which appear to generalize straightforwardly to this application. Nothing in the mathematics of the "Dutch Book Theorem" appears to require that decision-theoretic contingencies are worldly propositions (here see Hájek 2009).[17]

As illustration, here is a Dutch Book for an agent who commits herself to (5)—or more generally a sentence of the form $\triangle \diamond \phi \wedge \triangle \neg \diamond \phi$. Relative to a salient partition $\Re$ of salient representation $\mathcal{R}$, no matter which element of this partition is "selected" (whether it be for the purpose of representing the world accurately, or for the purpose of representing alternative ways of representing the world, or, indeed, for the sake of whatever cognitive purpose), the agent will regard as fair a

---

17 Staffel (forthcoming) remarks that, in an Expressivistic system (alike in some, but not all, respects to the one proposed here), "wins and losses can't be determined by checking *what the world is actually like*" (if the relevant contingencies are not worldly propositions that can be "checked" for truth against the actual world). But if, as seems correct, the conditional value *Val*(**A**|*C*) is like the conditional probability *Cr*(*A*|*C*)—in that both track degrees of desire or belief, under the *indicative supposition* that *C*—there is no immediate need for worldly matters to "determine" wins and losses. The degree to which an agent who indicatively supposes *C* desires to perform **A** will determine *Val*(**A**|*C*)—nothing worldly required, so long as the degree to which an agent desires to perform **A** can depend on a non-worldly representation. (It is easy to imagine this sort of dependence: if Bob might be hired for a job requiring business professional attire, Bob will prefer keeping a business suit to donating it. If the tap water might not be potable, I prefer bottled water to tap. Etc.) There may yet be a need for worldly matters to determine wins and losses, for a theorist who wants to use the notion of conditional desirability *to run a Dutch Book argument*. More on this just below.

series of bets that, taken together, guarantee a "loss". Assume, just for the sake of illustration, that $\triangle$ expresses a credence $> .6$ in its complement representation. Then each of the following bets—which, taken together, guarantee the agent a loss, no matter which representation in $\{[\![\Diamond\phi]\!]^c, [\![\neg\Diamond\phi]\!]^c\}$ is adopted for the relevant cognitive purpose—will appear fair to an agent who commits herself to $\triangle\Diamond\phi \wedge \triangle\neg\Diamond\phi$:

	$[\![\Diamond\phi]\!]^c$	$[\![\neg\Diamond\phi]\!]^c$
Bet 1	40	-60
Bet 2	-60	40

This is, I submit, a defect of rationality. An agent imposes a partition on $\mathcal{R}$, thereby generating an alternative set $\mathfrak{R}$ for $\mathcal{R}$, for the sake of representing alternatives whose adoption is relevant for cognitive purpose P. In this case, however the agent opts to fulfill P, she is in position to see she will be subject to a sure loss. The agent must revise her credences, then, if she does not wish to be subject to a sure loss if she selects an alternative from $\mathfrak{R}$ to fulfill P. The constraints on generalized credence functions we have introduced are, therefore, motivated by pragmatic considerations (although in the case of probability measures over worldly propositions, they can still be motivated by considerations of accuracy).

Staffel (forthcoming) develops both Accuracy-style and Dutch Book-style arguments for coherent credences in non-worldly representations (while also registering doubts that such arguments actually meet the theoretical needs that prompt them). In Staffel's Expressivistic Dutch Book—which is in certain respects similar to the one advanced here—an "underconfident" agent (e.g., one who assigns both $\Diamond p$ and its negation $\neg\Diamond p$ credence .4)...

> can avoid a sure loss by not becoming opinionated. The fact that the underconfident agent *would lose money if she became opinionated* does not point to any obvious rational defect. There are many things I might do that would put me at a great disadvantage in particular circumstances. But *if I have no reason to think I'll find myself in those circumstances, then I have little or no reason to avoid those actions*. (Staffel forthcoming: PAGE)

This difficulty certainly does threaten Staffel's Dutch Book argument (see esp. Staffel forthcoming: PAGE). It might also seem to threaten the version I have pursued here. The "irrationality" that, I claimed, characterizes an incoherent agent is as follows: relative to an alternative set $\mathfrak{R}$ that represents the various candidate representations that fulfill purpose P, the agent is subject to a loss <u>if</u> she selects an alternative from $\mathfrak{R}$ to fulfill P. However, if she does not select an alternative from $\mathfrak{R}$ to fulfill P—if, in the case of (5), she does not settle on any particular way of resolving the relevant question (i.e., whether Bob might be hired)—the negative conditional desirability (of, e.g., Bets 1 and 2 conditional on $\Diamond p$) is never "actualized". The "loss" here is of a purely hypothetical character: if the agent does this or that, she will lose; if she declines to do this or that, however, she will not. What is irrational about that?

In reply: we said that, in the case of (5):

- An agent entertains a set $\mathcal{R}$ of type $\langle\langle s, t\rangle, t\rangle$ (a set of $\langle s, t\rangle$-type objects).
- She partitions $\mathcal{R}$ into:
    - A cell of $\langle s, t\rangle$-type objects compatible with Bob being hired.

15

– A cell of $\langle s, t \rangle$-type objects incompatible with Bob being hired.

The purpose the agent tries to achieve in partitioning $\mathcal{R}$ as in $\mathfrak{R}$ is, we said in the prior section, to represent alternative ways of representing the world (e.g., alternatives that treat Bob's hire as possible *versus* those that treat it as impossible). Conditional on *either* way of representing the world—i.e., conditional on Bob's hire being possible, *and* conditional on Bob's hire being impossible—the agent is subject to a loss (more precisely, a deficit in desirability). The irrationality here is, I claim, manifest: the agent is trying to achieve goal g—in this case, g is the goal of figuring out how to represent Bob's hire (i.e., as possible or impossible)—but her credences are such that any way of achieving g presents her with a deficit in desirability; her credences are structurally such that *any way of achieving what she is trying to achieve* leaves her worse off. Claim: if your credences in context c are structurally such that they prevent you from doing what you're trying, in c, to do without being subject to sure "losses", your credences in c are irrational in c.

6 Conclusion

This paper began by observing that standard models of the semantics of epistemic modals render the phenomenon of graded modal judgment, whether in thought and language, unintelligible. In response, this paper developed a model of graded modal judgment, in both thought and language—one that represented graded modal judgment as a generalization of our cognitive capacity for reasoning with hypotheses about objective chance (i.e., our cognitive capacity for probabilistic reasoning). The generalization was developed as a package of interrelated semantic, cognitive, and epistemological theses:

- **Semantic:** modals compose with representations of arbitrary type. (§4.2)
- **Cognitive:** agents entertain representations of arbitrary type for specific cognitive purposes; the state of *bearing a graded attitude toward a representation of arbitrary type* is a natural cognitive kind (instances of which are, broadly, governed by the purpose for which the agent is entertaining the relevant representation). (§5.1)
- **Epistemological:** part of the functional role of credences in representations of arbitrary type (entertained for cognitive purpose P) is to determine fair "prices" for bets against ways of representing that fulfill P. Agents whose credences violate Normalization or Additivity are thus subject to Dutch Books. (§5.2)

According the model of graded modal judgment developed here, modal sentences are semantically evaluated against complex construction out of possibilia. But the sentences of our language—the interesting ones, anyway—were not semantically evaluated relative to *individual possibilia*. And so our model will exhibit the characteristic insensitivity of logics axiomatized by S5 to a choice of possible world taken to represent "indicative actuality" (as in Kaplan 1989), or to a choice of possible world taken to represent a non-actual circumstance of evaluation. I take this to be one of the main virtues of the present theory: it can accommodate many of the intuitions that motivate axiomatizing the logic of epistemic modality with S5,

without rendering the notion of graded modal judgment, whether in thought or in language, unintelligible (for a bit more detail, see Appendix B).

We will conclude on a polemical note. It is ironic to observe that "non-factual" theories—theories that do not take modalities of the relevant type to be proposition-forming operators, a description satisfied by both our theory and Moss'—offer the theorist at least two broadly workable models of the cognition, semantics, and epistemology of graded modal judgment. "Factual" accounts of these modalities, so long as they are constrained by S5—and, indeed, even a weaker logic like KD45—are able to offer none of these attractions (see Appendix B). It is probably time to move past the philosophical preoccupation with the ability of non-factual theories of operators in natural language to account for environments *embedding* these operators (whether in language, or in thought). If anyone has such problems, it seems to be the theorists who have pushed this objection, rather than their targets.

A Indicatives

A.1 Scope-Taking and Type-Raising with Indicatives

As intended, $\triangle(p \to q)$ expresses the property a credence function has when it assigns the non-worldly representation encoded in $p \to q$ a value $> .5$.

$$[\![\triangle_1(p \to_2 q)]\!]^c = \{Cr' : \sum_{Cr \in [\![p \to_2 q]\!]^c} Cr'(\{Cr\}) > .5\}$$

$$= \{Cr' : \sum_{Cr([\![q]\!]^c | [\![p]\!]^c)=1} Cr'(\{Cr\}) > .5\}$$

Handling the narrow-scope representation $p \to \triangle q$ is trickier. A first attempt:

$$[\![p \to_1 \triangle_2 q]\!]^c = \{Cr : Cr([\![\triangle_2 q]\!]^c | [\![p]\!]^c) = 1\}$$

$$= \{Cr : \frac{Cr([\![\triangle_2 q]\!]^c \cap [\![p]\!]^c)}{Cr([\![p]\!]^c)} = 1\}$$

But this attempt fails, since $Cr([\![\triangle_2 q]\!]^c \cap [\![p]\!]^c)$ is undefined in the present system, as $[\![\triangle_2 q]\!]^c$ and $[\![p]\!]^c$ are of different semantic types. Following Partee & Rooth (1983), we can address this by raising the type of $[\![p]\!]^c$:

$$\text{raise} X_{\langle \tau, t \rangle} = \lambda Y_{\langle \tau, t \rangle} . Y \subseteq X \qquad \textbf{(Raise)}$$

$$_\text{p} X_{\langle \tau, t \rangle} = \lambda \gamma . \gamma(\bigcup X) = 1 \qquad \textbf{(Probabilify)}$$

If $[\![\phi]\!]^c :: \langle \tau, t \rangle$, then $\text{raise}[\![\phi]\!]^c :: \langle \langle \tau, t \rangle, t \rangle$. That is to say, raising the type of a worldly representation $[\![p]\!]^c$ generates a set of worldly representations (equivalently, again, a characteristic function of worldly representations). In particular, it generates the set of worldly representations that involve representing $[\![p]\!]^c$ as true. Therefore, if $[\![\phi]\!]^c :: \langle \langle \tau, t \rangle, t \rangle$, then $_\text{p}[\![\phi]\!]^c :: \langle \gamma, t \rangle$.[18] Probabilifying a raised worldly represen-

---

18 γ is the type of credence functions (§4.4). So a function of type $\langle \gamma, t \rangle$ is of type $\langle \langle \langle \tau, t \rangle, v_{[0,1]} \rangle, t \rangle$. Our Probabilify rule is a generalization of the Type-Shifting rule introduced at Moss (2015: 34)—i.e., Moss' Type-Shifting rule is captured as a special case of Probabilification. It is natural to assume that whatever credence someone assigns $[\![p]\!]^c$ *determines* (or perhaps *rationally constrains*—I do not yet have a good sense of what issues are at stake here) their credence in $\text{raise}[\![p]\!]^c$: if you think of

tation ($_{\texttt{praise}}$-ing) p yields the set of credence functions that assign probability 1 to some way of representing that p.

Type-raising in hand, we have the following:

$$[\![p \rightarrow_1 \triangle_2 q]\!]^c = \{Cr : Cr([\![\triangle_2 q]\!]^c|_{\texttt{praise}}[\![p]\!]^c) = 1\}$$

$$= \{Cr : \frac{Cr([\![\triangle_2 q]\!]^c \cap {}_{\texttt{praise}}[\![p]\!]^c)}{Cr({}_{\texttt{praise}}[\![p]\!]^c)} = 1\}$$

$$= \{Cr : \frac{Cr(\{Cr' : \sum_{\mathcal{S} \in [\![q]\!]^c} Cr'(\{\mathcal{S}\}) > .5\} \cap \{Cr' : Cr'(\bigcup{}_{\texttt{raise}}[\![p]\!]^c) = 1\})}{Cr(\{Cr' : Cr'(\bigcup{}_{\texttt{raise}}[\![p]\!]^c) = 1\})} = 1\}$$

As intended, $p \rightarrow \triangle q$ expresses the property a credence function has when the ratio of the credence it assigns the representation $[\![\triangle_2 q]\!]^c \cap {}_{\texttt{praise}}[\![p]\!]^c$ to the credence it assigns the representation ${}_{\texttt{praise}}[\![p]\!]^c$ is 1.

A.2 Domain Restriction and Triviality

Another, ultimately simpler, possibility for representing the scopal interactions of indicatives and modal operators is to treat the latter as binary (i.e., *restrictable*[19]) operators (Kratzer 1981, 1986):

$$[\![\triangle_1(\phi)(\psi)]\!]^c = \lambda Cr : Cr \text{ is based on } g_c(1) \, . \, Cr([\![\psi]\!]^c|[\![\phi]\!]^c) > .5$$

$$[\![\Diamond_1(\phi)(\psi)]\!]^c = \lambda \mathcal{S} : \mathcal{S} \text{ is based on } g_c(1) \, . \, \mathcal{S} \cap [\![\phi]\!]^c \cap [\![\psi]\!]^c \neq \emptyset$$

Kratzer (1986) also denies that the indicative conditional contributes its own quantificational force; rather, indicative conditionals are syntactic devices for making explicit the restriction argument of a restrictable quantifier.[20] There is no semantic distinction between the "wide scope" $\triangle(p \rightarrow q)$ and the "narrow scope" $p \rightarrow \triangle q$: both are represented using the restricted modal $\triangle(p)(q)$.

One motivation for adopting Kratzer's (1986) analysis of indicative conditionals is explaining the sorts of judgments of equivalence that Stalnaker's Thesis (Stalnaker 1970) attempts to unify—e.g., the judgment that (7) and (8) are equivalent. According to Stalnaker's Thesis, the probability that an indicative conditional $A \rightarrow C$ is true equals the conditional probability of C on A. Supposing that probability operators in natural language semantically express degrees of conditional probability, Stalnaker's Thesis predicts, correctly, that (7) and (8) are equivalent.

(7)   Rain is likely, given that atmospheric pressure is low.

(8)   It is likely that it will rain if atmospheric pressure is low.

In line with Stalnaker's Thesis, the Kratzerian story about probability operators under consideration here renders (7) and (8) equivalent—more precisely, is able to generate equivalent logical forms for these sentences.

Regardless of whether Stalnaker's Thesis holds in its full generality, no version of the Thesis—even massively restricted—can be accommodated without taking

---

representation p as i-likely, then you are i-certain in some way of representing p (though typically there is no particular way of representing p such that you are i-certain of *it*).

19 Restrictable quantifiers are Generalized Quantifiers, in the sense of Barwise & Cooper (1981).

20 Except when no quantifier is provided, in which case a silent restrictable quantifier—which Kratzer (1986), e.g., took to be an epistemic necessity modal—is posited in logical form.

probability operators (and, by extension, modal operators) to be *binary* operators. That is because a language with only unary probability operators provably lacks the resources to express a sufficiently wide range of conditional probabilities.[21]

It would be difficult to overstate the importance of this idea: it allows the theorist to accommodate (a perhaps appropriately restricted version of) Stalnaker's Thesis, *without* signing onto the Thesis in the form in which it is usually presented:

$$\forall Pr : Pr(A \rightarrow C) = Pr(C|A)$$

That is because the Kratzerian analysis does *not* represent $\triangle(A \rightarrow C)$ as expressing a probability judgment whose object is a conditional proposition; rather, it expresses a *restricted probability judgment*: that C is likely (as assessed against the representation expressed by A).

> [I]n saying 'there is one chance in two that if A then C', the conditional 'if A then C' does not express any self-standing proposition. A different way to cast this observation is to go in the direction of Kratzer's analysis, namely to argue that the word 'if' does not act directly as a proposition-forming operator. However, this remains compatible with the idea that if-clauses are devices of quantifier restriction. In the scope of an operator, if-clauses do make a systematic truth-conditional contribution to the whole sentence. (Égré & Cozic 2011: 22)

This would seem to be exactly what is required to avoid challenges to Stalnaker's Thesis on grounds of Triviality results in the mold of Lewis (1976) (for discussion, see Rothschild 2015; Charlow 2016).

The dialectic in this neighborhood of issues is, however, a great deal more vexed than this quick summary would suggest. Charlow (2016) shows that Triviality results in the mold of Lewis (1976) arise for restricted operators (and that such results do not depend on the understanding of logical form embodied in Stalnaker's Thesis). Indeed, as Charlow (2016) argues, obstacles of Triviality arise for any treatment of restricted quantification that takes $[\![A]\!]$, $[\![C]\!]$, and $[\![\triangle(A)(C)]\!]$ to be *elements of the same semantic algebra*—i.e., any treatment that takes $[\![A]\!]$, $[\![C]\!]$, and $[\![\triangle(A)(C)]\!]$ to be of the same base semantic type. *This is precisely the assumption that the analysis in this paper discards*. I take this to be another argument in favor of this paper's analysis: unlike any competitor account of which I am aware, it allows the theorist to accommodate the intuitions of equivalence that underlie Stalnaker's Thesis, while also avoiding the specter of Triviality.

B   (In)validating S5

If we introduce epistemic modal operators that do not raise the semantic type of their complements, we will observe that, for such operators, the axioms of S5 are

---

21 See the theorem of Hájek (1989) and (especially) Égré & Cozic (2011)'s adaption of Hájek's theorem to an inexpressibility result for a language with unary probability operators. Sketch of the proof: consider a fair three-ticket lottery, with tickets numbered ONE, TWO, and THREE. The conditional likelihood of ONE winning if THREE doesn't is 1/2. But no Boolean combination of the relevant propositions (that ONE wins, that TWO wins, that THREE wins) is such that it has probability 1/2.

validated. Let $g_c(1) = \langle \mathcal{R}, \mathfrak{R}, \Sigma \rangle$.

$$\llbracket \blacklozenge_1 \phi \rrbracket^c = \{ \mathcal{S} \in \mathcal{R} : \mathcal{R} \cap \llbracket \phi \rrbracket^c \neq \emptyset \}$$

Either $\llbracket \blacklozenge_1 \phi \rrbracket^c = \mathcal{R}$ (if $\mathcal{R} \cap \llbracket \phi \rrbracket^c \neq \emptyset$), or else $\llbracket \blacklozenge_1 \phi \rrbracket^c = \emptyset$ (if $\mathcal{R} \cap \llbracket \phi \rrbracket^c = \emptyset$). It is clear that the logic of $\blacklozenge$ and its dual $\blacksquare$ is the logic of S5 (so long as we require, as we should—see [von Fintel & Gillies (2010)](#)—that $\llbracket \blacksquare_1 \phi \rrbracket^c \subseteq \llbracket \phi \rrbracket^c$):

$$\llbracket \blacksquare_1 \phi \rrbracket^c = \llbracket \blacksquare_1 \blacksquare_1 \phi \rrbracket^c$$
$$\llbracket \blacklozenge_1 \phi \rrbracket^c = \llbracket \blacksquare_1 \blacklozenge_1 \phi \rrbracket^c$$

But, of course, neither of the following equivalences hold for the type-raising modal operators of our language:

$$\llbracket \square_1 \phi \rrbracket^c \neq \llbracket \square_2 \square_1 \phi \rrbracket^c$$
$$\llbracket \lozenge_1 \phi \rrbracket^c \neq \llbracket \square_2 \lozenge_1 \phi \rrbracket^c$$

To illustrate, let $g_c(1) = \langle \mathcal{R}_1, \mathfrak{R}_1, \Sigma_1 \rangle$, $g_c(2) = \langle \mathcal{R}_2, \mathfrak{R}_2, \Sigma_2 \rangle$. Then:

$$\llbracket \lozenge_1 p \rrbracket^c = \lambda \mathcal{S}_{\langle s,t \rangle} : \mathcal{S} \text{ is based on } g_c(1) \, . \, \mathcal{S} \cap \llbracket p \rrbracket^c \neq \emptyset$$
$$\llbracket \lozenge_2 \neg \lozenge_1 p \rrbracket^c = \lambda \mathcal{S}_{\langle \langle s,t \rangle, t \rangle} : \mathcal{S} \text{ is based on } g_c(2) \, . \, \mathcal{S} - \llbracket \lozenge_1 p \rrbracket^c \neq \emptyset$$
$$= \lambda \mathcal{S}_{\langle \langle s,t \rangle, t \rangle} \, . \, \mathcal{S} - \{ \mathcal{T} : \mathcal{T} \cap \llbracket p \rrbracket^c \neq \emptyset \} \neq \emptyset$$

$\lozenge_1 p$ expresses the property a sets of worlds has, when it contains a p-world; $\lozenge_2 \neg \lozenge_1 p$ expresses the property a set of sets of worlds has, when it contains a set of worlds that is incompatible with p.

But, if $\lozenge_1 p$ expresses a property utterly distinct from the property expressed by $\square_2 \lozenge_1 p$, why does the following sound so borderline? [OMITTED]

(9)   ??It may be raining, but maybe it can't be.

Not because it is Moore-Paradoxical (contra [Weatherson ms](#)): note that Moore Paradoxicality dissolves in unasserted environments ([Yalcin 2007](#)):

(10)   ??Suppose it may be raining, but maybe it can't be.

(11)   Suppose it is raining, but you don't know it's raining.

Explanation: sentences of the form $\lozenge_1 p \wedge \lozenge_2 \neg \lozenge_1 p$ **are**, in fact, semantically anomalous! Notice that semantically coordinating $\llbracket \lozenge_1 p \rrbracket^c$ and $\llbracket \lozenge_1 \neg \lozenge_2 p \rrbracket^c$ requires intersecting $\llbracket \lozenge_1 \neg \lozenge_2 p \rrbracket^c$ with:

$$\texttt{raise} \llbracket \lozenge_1 p \rrbracket^c = \lambda X_{\langle \langle s,t \rangle, t \rangle} \, . \, X \subseteq \llbracket \lozenge_1 p \rrbracket^c$$

$\texttt{raise} \llbracket \lozenge_1 p \rrbracket^c$ denotes the property a set of propositions $\mathfrak{F}$ has iff each q in that set is compatible with p. As noted above, $\llbracket \lozenge_2 \neg \lozenge_1 p \rrbracket^c$ denotes the property a set of propositions $\mathfrak{F}$ has iff some q in that set is incompatible with p. Obviously, no $\mathfrak{F}$ satisfies both properties:

(12) $\llbracket \lozenge_1 \phi \wedge \lozenge_2 \neg \lozenge_1 \phi \rrbracket^c = \emptyset$

And so sentences of the form $\Diamond_1 p \wedge \Diamond_2 \neg \Diamond_1 p$ are predicted, on independent grounds, to be semantically anomalous (in spite of the fact that the left conjunct expresses a property utterly distinct from that expressed by the right conjunct).

The data from natural language, therefore, *do support a version of Euclideanness*, namely, the version in (12). This represents a strong empirical edge over the "Factualist" standard. Standard "Factualist" frameworks cannot, on the face of things, explain why sentences like (9) are semantically anomalous: in such frameworks, the semantic anomalousness of (9) must be explained by a Euclideanness constraint *on epistemic accessibility* from the "actual" world w (recall §2):

$$\textbf{Euclideanness}: v \in \sigma_w \Rightarrow \sigma_w \subseteq \sigma_v \qquad \forall w, \sigma : [\![\Diamond\phi \supset \Box\Diamond\phi]\!]^{\sigma,w} = \mathsf{T}$$

This makes vivid the dilemma confronting accounts of epistemic modals in the Factualist mold:

- Accommodate Euclideanness (at the price of rendering graded modal judgment unintelligible).
- Accommodate graded modal judgment (at the price of rendering (9) semantically impeccable).

The account defended here skirts the dilemma: it accommodates both the clear semantic intuitions motivating Euclideanness, without sacrificing an intelligible model of graded modal judgment.

References

Barwise, Jon & Robin Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4: 159–219. doi:10.1007/BF00350139.

Bradley, Richard. 2012. Multidimensional possible-world semantics for conditionals. *The Philosophical Review* 121: 539–571. doi:10.1215/00318108-1630921.

Charlow, Nate. 2016. Triviality for restrictor conditionals. *Noûs* 50: 533–564. doi:10.1111/nous.12111.

Charlow, Simon. forthcoming. A modular theory of pronouns and binding. In *Proceedings of LENLS 14*.

Égré, Paul & Mikaël Cozic. 2011. If-clauses and probability operators. *Topoi* 30: 17–29. doi:10.1007/s11245-010-9087-y.

von Fintel, Kai & Anthony S. Gillies. 2010. *Must* ... stay ... strong! *Natural Language Semantics* 18: 351–383. doi:10.1007/s11050-010-9058-2.

von Fintel, Kai & Anthony S. Gillies. 2011. *Might* made right. In A. Egan & B. Weatherson (eds.) *Epistemic Modality*, 108–30. Oxford: Oxford University Press.

Gillies, Anthony S. 2010. Iffiness. *Semantics and Pragmatics* 3: 1–42. doi:10.3765/sp.3.4.

Hájek, Alan. 1989. Probabilities of conditionals—revisited. *Journal of Philosophical Logic* 18: 423–428. doi:10.1007/BF00262944.

Hájek, Alan. 2009. Dutch book arguments. In P. Anand, P. Pattanaik & C. Puppe (eds.) *Oxford Handbook of Rational and Social Choice*, 173–195. Oxford: Oxford University Press.

Hild, Matthias. 1998. Auto-epistemology and updating. *Philosophical Studies* 92: 321–361.

Holliday, Wesley H. & Thomas F. Icard III. 2010. Moorean phenomena in epistemic logic. In L. Beklemishev, V. Goranko & V. Shehtman (eds.) *Advances in Modal Logic*, vol. 8, 178–199. College Publications.

Joyce, James. 1998. A nonpragmatic vindication of probabilism. *Philosophy of Science* 65: 575–603. doi:10.1086/392661.

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

Kaplan, David. 1989. Demonstratives. In J. Almog, J. Perry & H. Wettstein (eds.) *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.

Kratzer, Angelika. 1981. The notional category of modality. In H. Eikmeyer & H. Rieser (eds.) *Words, Worlds, and Contexts*, 38–74. Berlin: De Gruyter.

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

Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. *The Philosophical Review* 85: 297–315. http://www.jstor.org/pss/2184045.

Moss, Sarah. 2013. Epistemology formalized. *The Philosophical Review* 122: 1–43. doi:10.1215/00318108-1728705.

Moss, Sarah. 2015. On the semantics and pragmatics of epistemic modals. *Semantics and Pragmatics* 8: 1–81.

Partee, Barbara H. & Mats Rooth. 1983. Generalized conjunction and type ambiguity. In R. Bäuerle, C. Schwarze & A. von Stechow (eds.) *Meaning, Use, and Interpretation of Language*, 361–383. Berlin: Walter de Gruyter.

Rothschild, Daniel. 2012. Expressing credences. *Proceedings of the Aristotelian Society* CXII, Part 1: 99–114. doi:10.1111/j.1467-9264.2012.00327.x.

Rothschild, Daniel. 2015. Conditionals and propositions in semantics. *Journal of Philosophical Logic* 44: 781–791. doi:10.1007/s10992-015-9359-5.

Schroeder, Mark. 2011. Two roles for propositions: Cause for divorce? *Noûs* 47: 409–430. doi:10.1111/j.1468-0068.2011.00833.x.

Staffel, Julia. forthcoming. Expressivism, normative uncertainty, and arguments for probabilism. In *Oxford Studies in Epistemology*.

Stalnaker, Robert. 1970. Probability and conditionals. *Philosophy of Science* 37: 64–80.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge: MIT Press.

Weatherson, Brian. ms. Moore, bradley, and indicative conditionals. http://brian.weatherson.org/mbaic.pdf.

Weirich, Paul. 1980. Conditional utility and its place in decision theory. *The Journal of Philosophy* 77: 702–715. doi:10.5840/jphil1980771110.

Willer, Malte. 2013. Dynamics of epistemic modality. *The Philosophical Review* 122: 45–92. doi:10.1215/00318108-1728714.

Yalcin, Seth. 2007. Epistemic modals. *Mind* 116: 983–1026. doi:10.1093/mind/fzm983.

Yalcin, Seth. 2011. Nonfactualism about epistemic modality. In A. Egan & B. Weatherson (eds.) *Epistemic Modality*, 295–332. Oxford: Oxford University Press.

Yalcin, Seth. 2012. Bayesian expressivism. *Proceedings of the Aristotelian Society* CXII, Part 2: 123–160. doi:10.1111/j.1467-9264.2012.00329.x.