Confidence Reports

Fabrizio Cariani, Paolo Santorio, Alexis Wellwood

Abstract

We advocate and develop a states-based semantics for both nominal and adjectival confidence reports, as in *Ann is confident/has confidence that it's raining*, and their comparatives *Ann is more confident/has more confidence that it's raining than that it's snowing*. Other examples of adjectives that can report confidence include *sure* and *certain*. Our account leverages a neodavidsonian account of adjectival comparatives in which the adjectives denote properties of states, and measure functions are introduced compositionally. We further explore the prospects of applying these tools to the semantics of probability operators. We emphasize three desirable and novel features of our semantics: (i) probability claims only exploit qualitative resources unless there is explicit compositional pressure for quantitative resources; (ii) the semantics applies to both probabilistic adjectives (e.g., *likely*) and probabilistic nouns (e.g., *probability*); (iii) the semantics can be combined with an account of belief reports that allows thinkers to have incoherent probabilistic beliefs (e.g. thinking that A & B is more likely than A) even while validating the relevant purely probabilistic claims (e.g. validating the claim that A & B is never more likely than A). Finally, we explore the interaction between confidence-reporting discourse (e.g., *I am confident that...*) and belief-reports about probabilistic discourse (e.g., *I think it's likely that...*).

1 Introduction

A confidence report is a sentence like (1a). A reasonable assumption is that such a report expresses a confidence relation between an attitude holder and a proposition. Importantly, however, such attitudes are gradable, as evidenced by the clear acceptability and interpretability of (1b). Reports with *confident* are thus part of a broader family of gradable attitude expressions, which includes verbs like *want* and adjectives like *sure*/*unsure* and *certain*/*uncertain*.

(1) a. Ann is confident that it's raining.
    b. Ann is more confident that it's raining than that it's snowing.

Despite extensive research in linguistics and philosophy on the nature of propositional attitudes, the semantics of confidence reports is relatively under-theorized. One reason for this might be the expectation that, once we have accepted a certain theory of gradable adjectives, we will have said everything of semantic significance about confidence reports. For example, the derivation of (1b) in one standard framework would involve a function that maps propositions p and attitude-holders a to a degree representing a's confidence that p is true—in other words, *confident* would lexically express a confidence measure.

We take a somewhat broader view on the distribution of confidence reports, and suggest that a simple extension of the scalar analysis to *confident* is insufficiently general; for example, it fails to extend neatly to confidence reports with nominal *confidence*, (2a)-(2b). We thus advocate and develop an alternative semantics that interprets *confident* as a property of states, measures of which can be introduced compositionally in constructions like (1b) (Wellwood 2014, Wellwood 2015). This alternative is motivated primarily by an intuitive semantic equivalence between nominal and adjectival confidence reports in the comparative, as well as independent data supporting a common Davidsonian analysis of *confident* and *confidence*.

(2) a. Ann has confidence that it's raining.
    b. Ann has more confidence that it's raining than that it's snowing.

The Davidsonian turn supports a degree of semantic flexibility that we profitably exploit in our analysis of the bare, or 'positive', occurrences of *confident* (e.g., (1a), and their nominal counterparts) and when they occur in their comparative forms. For instance, according to our proposal, the positive form involves reference to the initial ordering on states provided by *confident*/*confidence*, while the comparative form involves reference to degrees introduced by *more*. This analysis allows us to implement a novel account of the characteristic context-dependent interpretation of such adjectival occurrences that does not make use of a covert POS morpheme.

The resulting semantics is also flexible in that it allows us to model confidence ascriptions to some agents who fall short of perfect rationality. Recent work has led to the development of a variety of probabilistic frameworks for semantic theories of probability operators. Embedding our account of confidence reports within these frameworks would likely make it hard to assign confidence states to non-probabilistic agents. The state-based framework we develop, by contrast, can model confidence ascriptions to a variety of non-probabilistic agents. Moreover, it does not require that the contents of confidence states be propositions, and thus it does not require confidence ascriptions to be tied to particular kinds of contents.

Exploring this semantics can, as we show, provide a novel perspective on the semantics of adjectives like *likely* and *probable*, e.g. (3). We argue that state-of-the art proposals about the semantics and logic of probability operators can be injected within our states-based framework. Furthermore, our proposal has some unique advantages. For one thing, nominal probability claims, such as (4), are generally overlooked in discussions of probability operators, but receive a natural analysis on the states-based proposal that parallels adjectival probability claims.

(3) a. It is likely that it's raining.
    b. It is more likely that it's raining than that it's snowing.

(4) a. There's a chance that it's raining.
    b. There is more likelihood that it's raining than that it's snowing.

For another, there seem to be inferential relations between comparative confidence reports like (1b) and reports that an agent believes a certain probabilistic content. So in particular, (1b) sounds equivalent to *Ann thinks that it's more likely to rain than to snow*. We show how this equivalence can be captured, once our semantics is supplemented with an account of the interaction between attitude verbs and epistemic vocabulary. According to our proposal, both *confident/confidence* and *likely/likelihood* are sensitive to 'information structures' (in a sense that we will define), but that associated with *confident* is relativized to an attitude-holder.

Finally, we claim that the extension of our framework to capture probability operators provides an original and valuable insight into some vexed methodological questions concerning their semantics. That is, we vindicate two desiderata that seem impossible to reconcile on existing views. On the one hand, we avoid appeal to probability functions in stating the semantics of simple sentences where *likely* appears in sentences like *It's likely to rain*. On the other, we vindicate appeals to probability where they do seem needed, i.e. in an account of sentences like (5).

(5) a. It is 85% likely that it will rain.
b. It is three times more likely that it will rain than that it will snow.
c. There is a 10% likelihood of snow.

2 Confidence reports

We consider reports both with adjectival *confident* and nominal *confidence*, and their interpretation in the positive and comparative forms, with an emphasis on the latter. We show that extant thinking about gradable adjectives and mass nouns pulls us in two different directions: the literature on gradable adjectives suggests we should assign a measure function-based interpretation to *confident*, but the literature on mass nouns suggests we should assign a property-based interpretation to *confidence*. Our observations and arguments suggest that a uniform analysis should be given. We propose an analysis that assimilates the adjectival interpretation to the nominal.

An initial attempt

Evidence abounds that *confident* (and the other adjectives in our target class) are gradable. In addition to their comfortable occurrence in the comparative, (1b), *confident* combines with the full panoply of comparative forms (e.g., *as confident*, *too confident*, *confident enough*, etc.), and with modifiers like *very* and *100%*. Thus a plausible initial thought is that *confident* lexicalizes a degree semantics directly.

For example, assume that gradable adjectives like *tall* are assigned a measure function type (e.g. Kennedy 1999), i.e. the type that maps individuals to degrees.[1] In this framework, it would be straightforward to assign *confident* the interpretation in (6), where **conf** maps a proposition p to x's degree of confidence in the truth of that proposition, and g is a variable assignment.[2] Importantly, we needn't assume that degrees of confidence are probabilities, or indeed that they have any structure besides what is provided by the standard scalar framework.

[1] A prominent alternative in the style of Heim 1985 analyzes them as having a higher quantificational type that embeds a measure function.

[2] This treatment parallels recent analyses of *likely*, such as Yalcin 2010, Lassiter 2011, 2015, 2016.

(6) $[\![\text{confident}]\!]^g = \lambda p_{st} \lambda x_e.\mathbf{conf}_x(p)$
type $\langle\langle s,t\rangle,\langle e,d\rangle\rangle$

Positing such lexical entries is standardly accompanied by the assumption that a covert morpheme, usually indicated by 'POS', is present whenever the adjective occurs without overt degree morphology.[3] While the details vary, the general function of this morpheme is to introduce comparison to a contextually salient standard for the property denoted by the adjective. Given these assumptions, the LF of (1a) has the additional structure indicated in (7a), and the truth conditions in (7b). This says that Ann's confidence in the truth of the proposition expressed by *it's raining* is greater than her standard for confidence in the context.[4]

(7) a. Ann is POS confident that it's raining.
b. $[\![(1\text{a})]\!]^g = \text{true iff } \mathbf{conf}_a(\mathbf{rain}) \geq \mathbf{standard}(\mathbf{conf}_a)$

By design, adopting this framework makes simple work of the interpretation of a comparative like (1b). In the style of Kennedy (1999), adjectival *more* is interpreted as in (8), the core contribution of which is a strict greater-than relation between degrees. Given the semantic value of the *than*-clause, d, and a measure function μ (i.e., a map from the subject of the matrix clause to their degree of μ-ness), a comparative like (1b) will express that Ann's degree of confidence in the truth of proposition **rain** strictly exceeds her degree of confidence in **snow**.

(8) $[\![\text{more}]\!]^g = \lambda d_d \lambda \mu_{ed} \lambda x_e.\mu(x) > d$
type $\langle d,\langle\langle e,d\rangle,\langle e,t\rangle\rangle\rangle$

Now, so long as *confident* saturates its propositional argument prior to combination with *more*, everything will proceed as it should: the usual compositional mechanisms will then deliver (9) as the interpretation of (1b).[5] This is read, 'Ann's degree of confidence in the proposition expressed by *it's raining* is greater than her degree of confidence in the proposition expressed by *it's snowing*.'

(9) $\mathbf{conf}_a(\mathbf{rain}) > \mathbf{conf}_a(\mathbf{snow})$

As promised, such a proposal is simple and elegant. But we will show that it is insufficiently general.

[3]But see Rett (2015) for an attempt to deploy Gricean pragmatics to dispense with POS.

[4]The lexical entry for POS on a Kennedy-style account, in the present context of adjectives that take propositional arguments, would look like:

(1) $[\![\text{POS}]\!]^g = \lambda d_d \lambda g_{\langle st,ed\rangle}\ \lambda p_{st}.\lambda x_e.\ g(p)(x) > d$

[5]Actually, the LF must look like the below, ignoring the copular verb and the internal structure of the *than*-clause. We make the standard assumption that *more* is discontinuous with the *than*-clause due to obligatory extraposition (Bresnan 1973; cf. Bhatt and Pancheva 2004, Alrenga et al. 2012).

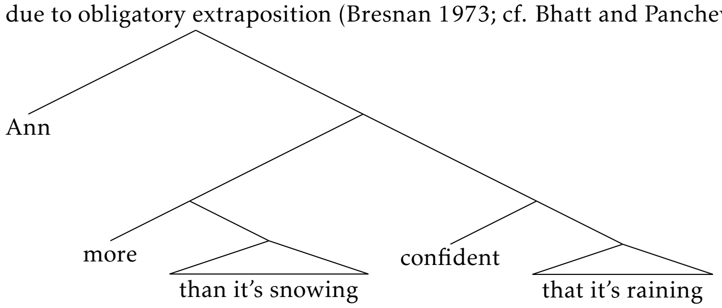

Confidence

Confidence reports do not require gradable adjectives for their expression. One can as easily express a confidence report with the mass noun *confidence*, (10), including in the comparative form, (11). It is interesting to observe that (11) is intuitively equivalent to (1b). We say more about equivalences below.

(10) Ann has confidence that it's raining.

(11) Ann has more confidence that it's raining than that it's snowing.

To say that *confidence* is a mass noun is, so far, just to say that it has a certain syntactic-semantic distribution: it appears comfortably with *much*, (12a), just like other mass nouns both concrete (*much mud*) and abstract (*much justice*). Like these nouns, it fails to appear comfortably with plural morphology, (12b), cardinal number words, (12c), or distributive quantifiers, (12d).[6]

(12) a. The men didn't express much confidence that the globe is warming.
b. ? The women expressed their confidences that the globe is warming.
c. ? The reports suggested two confidence(s) that the globe is warming.
d. ? Each confidence was high.

What of the type of *confidence*? The more often studied 'substance' mass nouns like *water* and *coffee* are typically assigned type $\langle e,t\rangle$, where it is understood that the particular e-type things that such nouns are true of are very different than the types of things that nouns like *traffic cone* and *cup* are true of. The same can be true of a contrast like that between *confidence* and (the noun) *party*: the former is true of things that hold while the latter is true of things that happen, etc.

These considerations suggest a blueprint for the semantic analysis. While we will revise some details of this blueprint in the next section in response to a variety of constraints, it will be useful to give a basic overview of the main ideas.

First, *confidence* expresses a property of states, the 'mass' subtype of the type of eventualities, v.[7]

(13) $[\![\text{confidence}]\!]^g = \lambda s_v.\textbf{confidence}(s)$
type $\langle v,t\rangle$

Then, the bare confidence report in (10) is interpreted in line with other neo-davidsonian treatments of eventuality predicates.[8] In (14), an existential claim about states involves Ann as the 'holder' or 'bearer' of the state, and the proposition (if that's what it is) expressed by **rain** is a thematic dependent which, for now,

[6] (12a) is negative because such occurrences of *much* have an NPI-like distribution; see Solt 2015.

[7] We choose the formulation 'states' for simplicity. Related possibilities for valuing the *s*s are, for example, tropes tropes (e.g., Moltmann 2009) or abstract substances (Francez and Koontz-Garboden 2015). We have little to say about the metaphysics of such entities. All we require is that, whatever they are, they show ordering relations and they have thematic participants.

[8] Assume that syntactic arguments map to conjuncts in logical form (cf. Castaneda 1967, Parsons 1990, Schein 1993, Pietroski 2005). Often, following especially Kratzer 1996, it is assumed that while the external argument—in our cases, the phrase indicating the holder of the state—corresponds to a conjunct in logical form, the internal argument fills an argument slot in the relation lexically denoted by—in this case—a noun like *confidence*.

we say plays the θ role in the state. For concreteness, $\theta(s,p)$ may be read 'p is the theme of s', or even 'p is the content of s'.[9]

(14) $\exists s_v[\textbf{holder}(s,a) \wedge \textbf{confidence}(s) \wedge \theta(s,\textbf{rain})]$

The comparative with *confidence* looks just like that with other mass nouns, modulo the distinction between entities and eventualities carried by the types e and v. In (15), a component of nominal *more* introduces the measure function (e.g. Heim 1985, Bhatt and Pancheva 2004, Bale and Barner 2009; cf. Solt 2015). Here, $g(\mu)$ stands for the value of g at index μ (one way of encoding context-sensitivity[10]), s is an eventuality to be measured, and d is provided by the *than*-clause.

(15) $[\![\text{more}_\mu]\!]^g = \lambda d_d \lambda s_v . g(\mu)(s) > d$
type $\langle d, \langle v,t\rangle\rangle$

Given these assumptions, the comparative (11) is interpreted as in (16), where δ abbreviates the *than*-clause degree. This may be read, 'there is a state of confidence s that Ann is in with respect to the proposition expressed by *it's raining*, the measure μ of which is greater than δ.' Given standard assumptions about the derivation of the meaning of the *than*-clause, the value of δ in (11) is equivalent to $max(\lambda d.\exists s_v[\textbf{holder}(s,a) \wedge \textbf{confidence}(s) \wedge \theta(s,\textbf{snow}) \wedge g(\mu)(s) \geq d])$, i.e., the maximal degree to which Ann is in a state of confidence whose theme is **rain**.

(16) $\exists s_v[\textbf{holder}(s,a) \wedge \textbf{confidence}(s) \wedge \theta(s,\textbf{rain}) \wedge g(\mu)(s) > \delta]$

The upshot is that nominal confidence reports are states-based. Comparative confidence reports in the nominal domain still involve confidence measures (i.e., $g(\mu)$ will map confidence states to degrees); yet, these measures are introduced compositionally, not by the lexical noun.

Bifurcate or assimilate?

The semantics we have given so far for *confident* and for *confidence* add up to a strange two-headed creature; call this the *bifurcation analysis*. We think there are good reasons for rejecting this analysis.

First, the bifurcation analysis obscures what appears to be an intuitive equivalence between comparative confidence reports that strand both sides of the nominal/adjectival divide. As far as we can tell, the requirements for the truth of *Ann is more confident that it's raining* and *Ann has more confidence that it's raining* are identical. On the bifurcation analysis, this equivalence is obscured: the adjectival comparative in (17a) involves the measure of a proposition, while the nominal comparative in (17b) involves the measure of a state.

(17) a. $[\![\text{Ann is more}_\delta \text{ confident that it's raining.}]\!]^g =$
$\textbf{conf}_a(\textbf{rain}) > \delta$

[9] The notion of 'content' maps between (Davidsonian) states and propositions, as we discuss more below. This contrasts with Hacquard 2006 proposal that the content of (say) a belief state is a set of propositions.

[10] Wellwood 2014 discusses empirical and theoretical issues surrounding the handling of this component of *more*'s indeterminacy.

b. $[\![\text{Ann has more}_{\mu,\delta}\text{ confidence that it's raining.}]\!]^g =$
$\exists s_v[\textbf{holder}(s,a) \wedge \textbf{confidence}(s) \wedge \theta(s,\textbf{rain}) \wedge g(\mu)(s) > \delta]$

Second, bifurcationism requires a disunified semantics for *more* and other degree morphemes that appear both with gradable adjectives and mass nouns. When *more* occurs in adjectival comparatives like (17a), it would have type $\langle d, \langle\langle e,d\rangle, \langle e,t\rangle\rangle$, (18a), while nominal *more* (at least when it occurs with properties of eventualities) would have type $\langle d, \langle v,t\rangle\rangle$, (18b). A reasonable thought is that the bifurcation analysis thus amounts to an ambiguity analysis, yet there is little cross-linguistic evidence for such an ambiguity (see Bale 2006 for related discussion).

(18) a. $[\![\text{more}]\!]^g = \lambda d_d \lambda \mu_{ed}.\lambda x_e.\mu(x) > d$
b. $[\![\text{more}_\mu]\!]^g = \lambda d_d \lambda s_v.g(\mu)(s) > d$

Nevertheless, the question of the semantic equivalence between comparative confidence reports based on the relevant adjectives and nouns remains. All else being equal, an account on which the equivalence arises due to identity is to be preferred over an account on which the equivalence is accidental. The alternative that can do this—the *assimilation analysis*—would take one of two forms: either (i) the adjectival case is assimilated to the nominal case, providing a states-based account of gradable adjectives like *confident*; or (ii) the other way around.

We pursue option (i) for two reasons. The first is that assimilating the adjective case to the nominal case allows us to leverage existing work rendering adjectival comparatives formally on a par with nominal comparatives (see Wellwood 2015 for arguments). The second reason is that a number of standard arguments for a davidsonian treatment of *confidence* replicate for *confident*, suggesting that both equally well involve something non-propositional that can be measured.

We present these data next. It should be noted, though, that they are not offered in the spirit of a refutation of option (ii); rather, we hope only to suggest that these data provide the materials for a straightforward implementation of option (i).

Reference and quantification

First, sentences 'about' John's confidence, reported using either the adjective or the noun, introduce something that can be the antecedent for anaphors like *that*, (19), and *it*, (20). That is, (19) means that John's confidence in this respect surprised Ann, and (20) means that Bill noted something both about John's confidence, and the effect that that confidence had on Ana. One's confidence can also last, (21). Importantly, such occurrences cannot be uniformly paraphrased using explicit reference to facts, propositions, or degrees.

(19) John was confident that it would snow. That surprised Ann.

(20) Bill noted John's confidence that it would snow, and that it surprised Ann.

(21) John's confidence lasted until he saw the clouds disperse.

Those things that can be surprising or noted must be of the sort that can figure into explicitly causal language. In (22a), the deadjectival nominalization of *confident* is something of which *make it snow* can predicate, however falsely. This nominalization seems equivalent to the parallel example with *confidence* in (22b).

(22a) and (22b) have no licit paraphrase involving explicit reference to propositions, facts, etc.[11] Importantly, such expressions can also be used to label effects, (23a)-(23b).

(22) a. John's being confident that it would rain made it snow.
b. John's confidence that it would rain made it snow.

(23) a. The clouds' appearance made John confident that it would rain.
b. The clouds' appearance gave John confidence that it would rain.

Another argument for the presence of an eventuality argument is given by the interaction of confidence reports with *because*-clauses.[12] On Kawamura's (2007) analysis, *p because q* says that the *q*-eventuality caused the *p*-eventuality. Our observation is that (24a) and (24b) are equally ambiguous between two readings, roughly paraphrased as in (25a) and (25b). If both the matrix and dependent clauses in (24a) and (24b) introduce eventualities, then the multiple readings can be attributed to an ambiguity in where the *because*-clause attaches.

(24) a. Ann is confident that Mary is in Paris because Gary is in Paris.
b. Ann has confidence that Mary is in Paris because Gary is in Paris.

(25) a. 'Ann is confident that: Mary is in Paris because Gary is in Paris'
b. 'Because Gary is in Paris, Ann is confident that Mary is in Paris'

These cases can be handled straightforwardly assuming that both *confident* and *confidence* introduce a property of states. For example, take (26) as the logical form of the two sentences in (19). By any assumption that gets the fact that *he* in the discourse *A man walked in. He sat down* refers to the man that walked in, the state that verifies the first logical form will value $g(t)$ in the second. This analysis can be extended to (20) by embedding forms like this under *note*.

(26) $\exists s_v[\mathbf{holder}(s,j) \wedge \mathbf{confidence}(s) \wedge \theta(s,\mathbf{snow})]$
$\exists s_v[\mathbf{experiencer}(s,a) \wedge \mathbf{surprise}(s) \wedge \mathbf{stimulus}(s,g(t))]$

Along the same lines, (27) can stand equally well as the interpretation for (22a) and (22b). Here and below, $\mathbf{cause}(x,y)$ is read, 'x is the cause of y'. (27) says, then, that there is a making of snow, and the cause of it was John's confidence that it would rain. Knowing what we know about weather patterns and the attitudes of mere mortals, (27) is unlikely to ever be true.

(27) $\exists e_v[\mathbf{making}(e) \wedge \mathbf{patient}(e,snow) \wedge$
$\mathbf{cause}(e,\iota s[\mathbf{confidence}(s) \wedge \mathbf{holder}(s,j) \wedge \theta(s,\mathbf{rain})])]$

Finally, the cases with *because*-clauses are more complex, but their treatment is nonetheless fairly straightforward on the states-based analysis of *confident* and *confidence*. When the *because*-clause attaches low, (28a), Ann is confident that Gary's

[11] The cases in (19) and (20) are less straightforward in this respect; they certainly sound atrocious when the referent of the pronoun is spelled out with *the proposition that...*, but neither sounds too bad with *the fact that...*. (See Kratzer 2012 for discussion of facts as particulars.)

[12] We learned about this feature of *because*-clauses from a homework problem designed by Chris Kennedy, who notices a similar ambiguity regarding negation.

being in Paris is the cause of Mary's being in Paris. When the *because*-clause attaches high, (28b), Gary's presence in Paris causes Ann to be confident that Mary is in Paris.

(28) a. $\exists s_v[\mathbf{confident}(s) \wedge \mathbf{holder}(s,a) \wedge \theta(s) =$
$\exists s'_v[\mathbf{in}(s',p) \wedge \mathbf{holder}(s',g) \wedge \exists s''_v[\mathbf{in}(s'',p) \wedge \mathbf{holder}(s'',m) \wedge \mathbf{cause}(s',s'')]]]$

b. $\exists s'_v[\mathbf{in}(s',p) \wedge \mathbf{holder}(s',g) \wedge$
$\exists s_v[\mathbf{confidence}(s) \wedge \mathbf{holder}(s,a) \wedge \mathbf{cause}(s',s) \wedge \theta(s) =$
$\exists s''_v[\mathbf{in}(s'',p) \wedge \mathbf{holder}(s'',m)]]]$

This last solution exploits the neodavidsonian assumption that thematic relations like **holder** and our θ are exhaustive and unique—they are in fact *functions*. On this assumption, the expressions, e.g., $\theta(s,a)$ and $\theta(s) = a$ are equivalent (see discussion and references in Williams 2015). This assumption implies that any confidence state has a unique holder (that named in the subject position of the local clause), and a unique theme (that named in the syntactic complement of the A or N introducing the states).

We retain this assumption in what follows.

3 The semantics of confidence reports

Basics: the positive form

As anticipated, we mostly build on the neodavidsonian semantics for gradable adjectives in Wellwood (2014, 2015), on which they express properties of ordered states. Novel to our implementation is a certain characterization of the interpretation of the positive and comparative forms. To foreshadow the general approach, independent of the specifics of confidence reports, we draw a basic distinction between states of height and states of tallness (i.e., the potential satisfiers of the predicate **tallness**). Every state of tallness is a state of height but not vice versa; positive occurrences of *tall* apply to height states in the 'positive region' of this ordered set, and forms like *taller* apply to any state of height.

We illustrate these details using *tall*, where the presentation will be more intuitive, and then return to our discussion of confidence reports. Following Wellwood, computing the interpretation of a simple sentence like (29) involves interpreting *tall* as a property of states, as in (30a), with the resulting logical form for (29) in (30b), according to the usual neodavidsonian assumptions.

(29) Mary is tall.

(30) a. $[\![\text{tall}]\!]^g = \lambda s_v.\mathbf{tallness}(s)$

b. $\exists s_v[\mathbf{holder}(s,m) \wedge \mathbf{tallness}(s)]$

In our proposal, the domains of gradable adjectives consist of elements in the domain of an ordering on states. Formally, this ordering is modeled as a pair $\langle D, \succsim\rangle$ of a set of states and (at least) a total pre-order on those states. For concreteness, assume that this ordering is tracked via a presupposition on the domain of the function expressed by the adjective, (31). (The subscript '**height**' flags that the relevant states are states of height—more on this momentarily.)

(31) $[\![\text{tall}]\!]^g = \lambda s_v : s \in Dom(\langle D_{\mathbf{height}}, \succsim\rangle).\mathbf{tallness}(s)$

Crucially, we make a distinction between 'height states' and 'tallness states'. So for example, the tallest individuals in a context will instantiate states of height and tallness, while the shortest individuals will instantiate height but not tallness states. Thus, the meaning of *tall* isolates which of the height states count as tallness states. Going forward, we will call the broad domain that the adjective invokes the *background ordering*, and the set of states that an adjective like *tall* is true of the *positive region* of that ordering.[13] The relationship between the background ordering and the positive region can be seen in Figure 1.

Figure 1: A pre-order on height states, some of which are tallness states.

$$\overbrace{\dots s_1 \quad s_2 \quad s_3 \quad s_4 \quad s_5 \quad s_6 \quad \underbrace{s_7 \quad s_8 \quad s_9 \quad s_{10} \quad s_{11} \quad s_{12}\dots}_{\text{positive region (states that count as } \textit{tall}\text{)}}}^{\text{height ordering}}$$

Under this approach, the well-known context-dependence of gradable adjectives is reflected in how the positive region of the background ordering is determined. As a reminder, the central datum here is that a sentence like *Ann is tall* is true just in case Ann exceeds some standard for tallness in the context. We implement this context-sensitivity in the states-based framework making use of a contextual index on the gradable adjective. To delimit the positive region of the ordering, we explicitly define the positive region in terms of a function that we label **contrast**. As we show below, this implementation is useful for capturing patterns of judgment that are particular to confidence reports, as well as for capturing the logic of those reports.

Concretely, we assume the lexical entry for *tall* in (32a), with the predicate **tallness** indexed by C. In general, the truth-conditional effect of this is given in (32b), for any gradable property, **g-ness**, associated background ordering $\gtrsim_{\textbf{g-ness}}$, and state s. The function **contrast** maps the target state s to a salient *contrast state* s'. In the case of (29), we now have the logical form in (33), which says that (29) is true just in case Mary is in a height state s ordered at least as high as the contrast state s' in that ordering.[14]

(32) a. $[\![\text{tall}]\!]^g = \lambda s_v : s \in Dom(\langle D_{\textbf{height}}, \gtrsim\rangle).\textbf{tallness}_C(s)$

b. $\textbf{g-ness}_C(s)$ is true iff $s \gtrsim_{\textbf{g-ness}} \textbf{contrast}_C(\textbf{g-ness})(s)$

(33) $\exists s_v[\textbf{holder}(s,m) \wedge \textbf{tallness}_C(s)]$

$\equiv \exists s_v[\textbf{holder}(s,m) \wedge s \gtrsim_{\textbf{height}} \textbf{contrast}_C(\textbf{tallness})(s)]$

This proposal is interesting for a number of reasons. First, our account of the positive form doesn't invoke measures, unlike those that utilize POS. Furthermore, we can already see how relativizing the **contrast** function to the gradable property can be useful: it allows different cut-off points for different adjectives, even if they

[13]Our terminology is meant to overlap with that of vagueness-based approaches like that of Klein 1980, 1982. Such approaches also posit an ordering on the domain of the adjective, though this is between individuals.

[14]We have given **tallness**, unrelativized by C, to the **contrast** with **g-ness** because, as we will see, different properties that invoke the same background ordering nonetheless have different 'cutoff' points.

plausibly share the same background ordering. Thus, we have in place a natural framework for modeling the relationships among clusters of gradable adjectives, like **cool** ∼ **warm** ∼ **hot**, or **doubtful** ∼ **unsure** ∼ **sure**. The further left in these lists, the lower the cut-off point required for a heat or credence state to count as an instance of the property.

Basics: the comparative form

Interpreting gradable adjectives as measure functions (or relations between individuals and degrees) generally leads to the problem of how to 'discharge' their degree-relativity when there is no degree morphology. In the standard scalar setting, POS is posited to do this work: it relates the subject's degree of **g-ness** to a standard of **g-ness** in the context. This meaning contribution is usually made separately from the adjective in order to avoid unwanted inferences in the comparative form. Our proposal also separates these meaning contributions, but it takes a different approach. Instead of postulating additional covert material in the positive form, we propose that comparative morphology discards the contextual information encoded in the **contrast** function—thus allowing us to access the entire background ordering.

For an instance of the kind of issue that POS and background orderings are meant to solve, consider the sentence in (34). This does not entail that Mary exceeds any standard for tallness. If the lexical meaning of *tall* invoked such standards, and if that lexical item is transparently involved in the calculation of (34), this observation would be unexpected. Our approach dispenses with standard-relativity by supposing that the comparative morpheme (*-er* or *more*) takes the gradable expression as an argument, and uses it only to extract information about the background ordering from it—there is no claim that the context-sensitive property holds of the subject.

(34) Mary is taller than Sue.

This approach also dovetails well with the observation that even entities with very little height can be compared using *taller than*. By way of illustration we add a useful piece of notation. We assume a function **background**(·) that maps states to a background ordering on which those states are located. For example, **background**(·) maps states of tallness to a background ordering of states of height, which includes also states of height that don't qualify as tall.

For ease of exposition, it's convenient to be abuse notation and allow the function **background**(·) to apply to individual states, sets of states, and characteristic functions of the latter. In all these cases, it returns the background ordering on which the relevant states lie. This allows us to use the denotations of adjectives like *tall* directly as the argument of **background**(·).

In the entry for *more* in (35), **background** is invoked in a presupposition on the state argument: the relevant state is presupposed to be part of the background ordering. Thus, *-er*/*more* combines with a degree (the denotation of the *than*-clause), and a property of states G (in the present cases) to return a property of states that are in **background**(G).

(35) $[\![\text{more}_\mu]\!]^g = \lambda d_d \lambda G_{vt} \lambda s_v : s \in \mathbf{background}(G).g(\mu)(s) > d$

In (35), $g(\mu)$ is a contextual assignment of values to a variable over measure function types, here $\langle v,d\rangle$. The selection of these values must meet one further condition which guarantees that the (strict) ordering relations on the measured domain are preserved in the corresponding degree ordering. On the assumption that **background**$(G)(s)$ is only defined if there is a background ordering, $\langle D,\succsim\rangle$, such that s is in D, the condition may be stated as in (36) (cf. Schwarzschild 2006, Nakanishi 2007, and Wellwood 2014, 2015). This says that if any two states are strictly ordered in a certain way in the background ordering, then their μ-measures are ordered in the same way in the degree ordering.

(36) $\forall s,s' \in Dom(\langle D,\succsim\rangle)$, if $s \succ s'$, then $g(\mu)(s) > g(\mu)(s')$.

For illustration, the compositional process underlying (34) is sketched in (37). Abbreviating the contribution of the *than*-clause using δ, the property expressed by the degree phrase, *taller than Sue*, is a property of states in the background ordering of $\mathbf{tallness}_C$ whose $g(\mu)$-measure is greater than δ, (37c). Combining the rest, (34) is interpreted as in (38), which says that Mary is in a height state (i.e., a state in the domain of the background ordering for $\mathbf{tallness}_C$), the $g(\mu)$ measure of which is greater than that of a corresponding state of Sue.[15]

(37) a. $[\![\text{than Sue}]\!]^g = \delta$
b. $[\![\text{er [than Sue]}]\!]^g = \lambda G_{\langle vt\rangle}.\lambda s_v : s \in \mathbf{background}(G).g(\mu)(s) > \delta$
c. $[\![\text{[tall [er than Sue]}]\!]^g = \lambda s_v : s \in \mathbf{background}(\mathbf{tallness}_C).\ g(\mu)(s) > \delta$

(38) $[\![(34)]\!]^g = \exists s_v : s \in \mathbf{background}(\mathbf{tallness}_C)\ [\mathbf{holder}(s,m) \wedge g(\mu)(s) > \delta]$

The general idea is that the function expressed by *tall* have a structured domain (of states), and the structure of that domain must be preserved in the mapping to degrees by the comparative operator. The details might look complicated, but they just extend how nominal comparatives work in simpler cases. For instance, the analysis of *more coffee* assumes that (i) *coffee* expresses a function whose domain is ordered by a part-of relation on portions of coffee, and (ii) permissible values of μ in the comparative (e.g., a volume or weight measure) respect strict part-whole relations on those portions.

Comparative confidence reports

Next, we extend these ideas to *confident* and *confidence*. Assume that the schematic comparative sentences in (39) have the same interpretation, i.e. (40). Both express an existential statement about confidence states whose holder is Ann, whose theme is (so we are supposing) the proposition p, and whose measure μ is greater than δ (the value provided by the *than*-clause, implicit in (39)). This analysis assumes that **confident** and **confidence** apply to the same subset of the same background ordering, though we revisit this assumption in the next section.

(39) a. A is more confident that p.
b. A has more confidence that p.

[15] We have indicated the restriction to the background ordering on height states as a restriction on the domain of the existential quantifier in (38).

(40) $\exists s_v : s \in \textbf{background}(\textbf{confidence}_C)\ [\textbf{holder}(s,a) \wedge \theta(s,p) \wedge g(\mu)(s) > \delta]$

The relevant background ordering ranks states by how confident a holder *a* is in the truth of the state's theme. Thus, the background ordering includes states of high as well as low confidence. Both *confident* and *confidence* single out a positive region of this ordering that includes the upper bound. The ordering is tracked, as before, via presupposition, as shown in (41). Here, the superscript $\mathbf{h}(s)$ abbreviates $\textbf{holder}(s)$ and the subscript **conf** indicates that the ordering concerns states of confidence.[16]

(41) $[\![\text{confident}]\!]^g = [\![\text{confidence}]\!]^g =$
$$\lambda s_v : s \in Dom(\langle D^{\mathbf{h}(s)}_{\mathbf{conf}}, \gtrsim\rangle).\textbf{confidence}_C(s)$$

In the end, logical forms like (40) say just that A's confidence with respect to *p* is greater than δ. This is a roundabout way of saying what the standard scalar analysis can say much more directly. However, taking the detour through states makes it explicit that it is confidence (and not the proposition that the confidence relates to) that is measured. It also allows us to capture the intuitive equivalence between the nominal and adjectival comparative forms as a matter of (propositional) identity, all the while maintaining a univocal analysis of *more*.

Even though we state the lexical semantics in terms of an ordering of states, it is occasionally more intuitive to think of confidence orderings as orderings of propositions. Since we assume that thematic relations functionally connect states with propositions, we can reserve for ourselves the ability to speak both ways—in terms of the basic orderings of states, and in terms of the orderings of the propositions that are the themes of those states.

Positive confidence reports

The machinery that we established for the interpretation of the positive form with **tallness** or for **g-ness** in general extends straightforwardly to the case of *confident* and *confidence* as well. Assume, as before, that the schematic positive sentences in (42) have the same interpretation—i.e., (43). Both express an existential statement about confidence states *s* whose holder is Ann, whose theme is the proposition *p*, and which is ordered higher than its contrast state s' according to the background ordering on confidence states.

(42) a. A is confident that *p*.
b. A has confidence that *p*.

(43) $\exists s_v[\textbf{holder}(s,a) \wedge \textbf{confidence}_C(s)]$
$$\equiv\ \exists s_v[\textbf{holder}(s,a) \wedge s \gtrsim^{\mathbf{h}(s)}_{\mathbf{conf}} \textbf{contrast}_C(\textbf{confidence})(s)]$$

The fact that this analysis predicts that (42a) and (42b) have equivalent truth-conditions might look surprising. Suppose for instance that the weatherman is 20% confident that it will rain. In such a context, there is little temptation to accept *the weatherman is confident that it will rain*. Some judge, however, that *the*

[16] Relativizing the confidence structure to a holder is one of the main differences between the interpretations we assign to *confident* (/*confidence*) and *likely*, as we show below.

weatherman has confidence that it will rain is acceptable here. Indeed, the latter sentence seems to suggest that the weatherman has *some* confidence that it will rain. Taken at face value, such judgments are in tension with our prediction that sentences with *confident* and *confidence* are equivalent.

Nonetheless, we think that there are good reasons not to take these judgments at face value. For one thing, if *having confidence* means the same as *having some confidence*, then negating a sentence of the form *A has confidence that p* should require that A has no confidence whatsoever in p. But this doesn't seem right. For example, (44) doesn't require that the weatherman assign zero confidence to the possibility of rain in order to be judged true.

(44) The weatherman does not have confidence that it will rain.

Another reason to reject any identification between *having confidence* and *having some confidence* is that it would predict *Ann has more confidence that it's raining than that it's snowing* to entail *Ann has confidence that it's raining*. This seems wrong, too: merely having more confidence in p than in q doesn't seem sufficient for having confidence in p.

Perhaps, however, one might reject the equivalence between (42a) and (42b) while denying that *A has confidence that p* is equivalent to *A has some confidence that p*. Nominal confidence reports might demand a lower (but non-zero) threshold than adjectival reports. Officially, we are agnostic on this point, but it is worth noting that our framework allows us to model it as a possibility. For instance, we might treat *confident* and *confidence* on the model of *hot* and *warm*—i.e., that there are two different properties of states, **confident** and **confidence**. They are based on the same background confidence ordering, just the **contrast** function maps them to different benchmarks on that ordering in the positive form.

The logic of confidence reports

The analysis we propose has interesting consequences for the logic of reports with *confident* and *confidence*.

On the one hand, we impose virtually no constraints on what a subject's confidence ordering looks like. This allows us to regard as true any confidence report that describes states that cannot be represented via a probability function. Here is one example: by the way probability functions are defined, the probability of a conjunction is a lower bound on the probability of a conjunct. Yet it is well-known[17] that, under determinate circumstances, subjects appear to routinely violate this constraint. As a result, it appears that, under the right circumstances, the sentences in (45) can be true together.

(45) a. John is not confident that Linda is a bankteller.
     b. John is confident that Linda is a feminist bankteller.

Our semantics is perfectly equipped to vindicate this. Similarly, nothing in our proposal dictates that a subject should be fully confident of tautologies; at the same time, probability functions assign tautologies full probability by design.

[17] The *locus classicus* for this claim is Tversky and Kahneman 1983.

On the other hand, some other predictions are hardwired in our semantics. These predictions don't track logical relations between the *contents* of confidence states. Rather, they track logical relations between *confidence states* themselves. For example, the inference from (46a) to (46b) is validated by our semantics.

(46) a. S is confident that p, S is more confident of q than of p
b. S is confident that q

Below is an (incomplete) list of logical properties of our semantics for *confident* and *confidence*.

Transitivity. (47a) and (47b) entail (47c).

(47) a. S is more confident (/has more confidence) that p than that q.
b. S is more confident (/has more confidence) that q than that r.
c. S is more confident (/has more confidence) that p than that r.

Antisymmetry. (48a) and (48b) entail (48c).

(48) a. S is at least as confident of p as of q.
b. S is at least as confident of q as of p.
c. S is equally confident of p and q.

Connectedness. (49) is a logical truth.

(49) Either S is at least as confident of p as of q, or S is at least as confident of q as of p.

Let us emphasize that the reason why we choose to encode these properties in the logic of *confident*/*confidence* is empirical. These properties seem to be encoded in the grammar of the relevant words. We pointed out that the sentences in (45) seem perfectly consistent, though of course they describe a subject whose cognitive state is not fully rational. Conversely, violations of transitivity, antisymmetry, and connectedness strike us as problematic from a grammatical point of view. For example, the discourse in (50) appears to be contradictory.

(50) Aidan is more confident that it will rain than that it will snow, and more confident that it will be windy than that it will rain. # But he's not more confident that it will be windy than that it will snow.

Any plausible semantics for POS and the comparative will vindicate the entailment in (46). Conversely, the pattern in (51) requires more controversial assumptions. If we continue to suppose that the contrast functions invoked by POS always map a proposition to its negation (a corollary of which is that one cannot be both confident that p and confident that $\neg p$), then we predict the entailment in (51): (51a) would require confidence-that-p states are ranked higher than confidence-that-not-p states in the state structure, and *more* will always map higher-ranked states to higher degrees, in line with its monotonicity condition.

(51) a. S is confident (/has confidence) that p.
b. S is more confident (/has more confidence) that p than that $\neg p$.

These seem welcome consequences of assuming that the contrast of a proposition is always is negation. Nevertheless, we prefer to remain agnostic about this assumption, since there are also substantial reasons to doubt it. Here are two.

First, the assumption might overgenerate. We did not assume that an attitude holder's confidence states are probabilistic, hence it may happen that Carlo has extremely low confidence in p and even lower confidence in $\neg p$. In this case, if we hold on to the assumption that the contrast of a proposition is always its negation, Carlo will still count as confident in p. But this seems wrong.[18]

Second, there seem to be intuitive cases where an agent is confident in a proposition, even though they are more confident of its negation.[19] Suppose that Clara, who is probabilistically coherent, believes that the Warriors have a 49% chance of winning the NBA finals this year, and that each of the other teams has at most a 3% chance of winning. Now suppose that we are having a discussion about which teams have the best shot at winning the finals. (52) seems true in this context, despite the mass of Clara's confidence favoring some team other than the Warriors.

(52) Clara is confident that the Warriors will win the NBA finals.

Conditional confidence

To conclude our discussion of confidence reports, we quickly touch on one last complication. Confidence reports interact with conditional antecedents in ways that are not entirely predicted by the system we have set up. Two specific kinds of facts stand out: it is possible for one to self-ascribe conditional confidence in A even if one is not confident of A, as in (53a). Such conditional ascriptions sound roughly equivalent to self-ascriptions of confidence in the conditional, as in (53b).

(53) a. If Lisa is in town, I am confident that she is at the lab.
b. I am confident that if Lisa is in town she is at the coffee lab.

A first stab at this kind of account is to make the background ordering sensitive to an information state i which can be operated on by conditional antecedents. Specifically, revise (41) to (54)—indexing the ordering $\gtrsim$ with i. Then, assume a semantics for conditional (in the style of Kolodny and MacFarlane 2010) in which conditional antecedents restrict i. Since (54) make the ordering $\gtrsim$ dependent on i, this predicts that one can assert (53a) even if one isn't unconditionally confident that Lisa is at the lab.

(54) $[\![\text{confident}]\!]^g = [\![\text{confidence}]\!]^g = \lambda s_v : s \in Dom(\langle D^{\mathbf{h}(s)}_{\mathbf{conf}}, \gtrsim_i\rangle).\mathbf{confident}_C(s)$

[18] We might try to fix this by letting the contrast proposition be some benchmark propositions (e.g. a proposition that the agent feels fifty-fifty about p and q). This would fix the current problem and it will make the same predictions as the original proposal when the attitude holder's confidence structure is probabilistic. But it loses the entailment (51) and its corollary, as there might be some non-probabilistic attitude holders according to which both p and $\neg p$ exceed the benchmark. Ultimately, we should expect some unintuitive predictions from any theory that attempts to model less than perfectly coherent agents. It is hard to say, when an agent's confidence structure is incoherent in the way Carlo's is, what exactly the semantics should predict.

[19] For a similar point about *likely*, cf. Yalcin 2010 and the discussion in Hawthorne et al. 2016, p. 1400.

4 Probability operators in a states-based framework

Beyond *confidence*

Confidence reports belong to a category which we roughly characterized as gradable attitude reports. We can think of such gradable attitude reports as claims about some attitude holder's being in a certain credal state. This naturally raises the question of how the work we carried out in the case of gradable attitude reports relates to the recent explosion of work on gradable epistemic modality. For instance, much recent work has tackled the semantic analysis of claims such as those in (55).[20]

(55) a. It is likely to rain.
b. It is more likely to rain than to snow.

The first thing to observe here is that there are important asymmetries between the semantics of gradable attitude reports and the semantics of probabilistic modals like *likely*. These asymmetries have already been used to draw a contrast between *believe*, on the one hand, and *might*/*must*, on the other. They can be easily replicated for *confident* and *likely*.

For example, consider the contrast in (56), where (56a) is clearly consistently acceptable while (56b) sounds like a contradiction (Yalcin 2007).

(56) a. Suppose it's raining but I believe it is not raining.
b. # Suppose it's raining but it might not be raining.

The same contrast shows up with the pairs in (57) and (58) that involve our target expressions.

(57) a. Suppose it's raining but I am confident it is not raining.
b. # Suppose it's raining but it is probably not raining.

(58) a. Suppose it's raining but I am more confident that it's snowing than that it's raining.
b. # Suppose it's raining but it's more likely that it's snowing than that it's raining.

These considerations caution us against assuming that we can just extend our lexical entry for *confident* to probability operators. Probability operators are importantly different from attitude verbs.

At the same time, there are important reasons to explore a states-based analysis of probabilistic language. For one thing, there appears to be a near equivalence between gradable confidence reports and qualitative belief ascriptions involving certain probabilistic contents. For instance, there is a reading of *confident* that makes the sentences in (59) sound roughly equivalent, and similarly for the sentences in (60).[21]

[20] For a non-exhaustive list, see Yalcin 2010; Swanson 2007; Lassiter 2011, 2015, 2016; Holliday and Icard 2013; Klecha 2014; Moss 2015, 2018; Santorio and Romoli 2017. For an approach to gradable modality that may have implications for gradable epistemic modals (even though it was not initially developed in that context) see Portner and Rubinstein 2016.

[21] We do not, of course, deny that there is another reading of (59a) on which this equivalence fails. Given this other reading, "I am confident that p" means roughly that I have faith that p will happen.

(59) a. I am confident (/have confidence) that it will rain.
b. I believe it is likely that it will rain.

(60) a. I am more confident that Masaya will teach syntax than that he will teach semantics.
b. I believe that it is more likely Masaya will teach syntax than that he will teach semantics.

We talk about 'rough equivalence' here because we encountered a variety of judgments on the matter. While some people judge them to be equivalent, others hear (59a) as somewhat stronger than (59b). This intuition is supported by the observation that (61) can be heard as consistent.

(61) I think it's likely that it will rain, though I'm not confident that it will rain.

We are going to return to this point. For now, let us observe that such worries are generally targeted to the positive form—sentences in (59). The intuitions in favor of equivalence are much stronger with the comparative sentences in (60). Those are all we need to support our point that we need to make sense of an inferential link between gradable attitude expressions and probability operators.

There are additional, more direct considerations for extending the states-based approach to probability operators. The most important one is that, as for confidence reports, related probability claims can be expressed using nominal forms, for example (62).

(62) a. There is a chance I will drink tea today.
b. There is more chance/likelihood that it will rain than that it will snow.
c. There is a good chance/probability of snow above 5000 ft.

The distributional facts that inspired the states-based approach to *confidence* replicate in this case too. Nouns like *chance*, *probability* and *likelihood* combine comfortably with *much*, but not with cardinal number words (?*two chances*, ?*two likelihoods*), or distributive quantifiers (?*each chance*, ?*each likelihood*), etc.[22]

Moreover, the array of data that suggested the presence of an event argument in confidence reports can be replicated for likelihood claims. First of all, these nominal forms can introduce causes, (63), and effects, (64). This suggests that, as for confidence reports, that the sentences in (63) and (64) involve concrete entities linked to likelihood that can enter into causal relations.

(63) a. The likelihood of snow led me to wear boots.
b. The probability of snow led me to wear boots.
c. Rain's being likely caused me to bring an umbrella.

(64) a. God banging his drum increased the likelihood of snow.
b. Warm air currents lessened the probability of snow.
c. God's banging his drum made snow likely.

[22] One potential counter-example is, *the chances of winning are low*. But this seems to be an isolated property of *chance* rather than a general feature of nominal probability operators.

Finally, again like confidence reports, likelihood reports show the type of ambiguity with *because*-clauses that we reported for *confident/confidence*. In particular, (65) is ambiguous between the two readings paraphrased in (65a) and (65b). If *likely* introduces a Davidsonian argument, then this ambiguity can similarly be seen to reduce to an ambiguity in the attachment site for the *because*-clause.

(65) It's likely (/there is some likelihood) that Mary is in Paris because Sue is in Paris.
   a. 'It's likely (/there is some likelihood) that: Mary is in Paris because Sue is in Paris'
   b. 'It's likely (/there is some likelihood) that Mary is in Paris, given that: Sue is in Paris'

Taking stock, these considerations support the project of devising a states-based semantics for probability operators. Furthermore, the apparent equivalence between gradable confidence reports and qualitative belief ascriptions involving probabilistic contents suggests development within a single, unified framework. We have to be careful not to push the connections too far, though: we must also be able to capture the difference between confidence reports and claims involving probability operators, e.g., the differences under *suppose* observed in (57)-(58).

Qualitative probability: A primer

At the heart of the framework we developed for confidence reports was a simple design principle. Degrees (e.g. representing confidence) are only recruited compositionally by degree expressions like *more*. But they are absent from the positive forms, both adjectival and nominal. In developing the semantics for confidence reports, we just assumed some background ordering of states of confidence— orderings satisfying some very minimal conditions.

We will follow the same design principle when it comes to the semantics for probability operators. First, there is an important difference between the two cases. To ensure that our probability operators have enough of their standard logic, assume that the background ordering has a distinctive structure. In particular, assume that probability operators are evaluated against the background of a qualitative probability structure, which we represent as an index coordinate.

A finite qualitative probability structure $\mathcal{Q} = \langle W, \gtrsim \rangle$ is a pair where W is a finite set of points–e.g. worlds, and $\gtrsim$ is an ordering on the elements of W by their comparative probability. (We focus on the finite case solely for reasons of simplicity of exposition.) Let $\top$ denote the set containing every world in W and $\bot$ be the empty set. With these conventions, let $\gtrsim$ be a total pre-order over the members of $\mathcal{P}(W)$ satisfying three further conditions:

1. *non-triviality*: $\top > \bot$;
2. *non-negativity*: $\forall \mathsf{A} \in \mathcal{P}(W)$, $\mathsf{A} \gtrsim \bot$;
3. *qualitative additivity*: if $(\mathsf{A} \cap \mathsf{B}) = (\mathsf{A} \cap \mathsf{C}) = \emptyset$, then $(\mathsf{B} \gtrsim \mathsf{C})$ iff $(\mathsf{A} \cup \mathsf{B}) \gtrsim (\mathsf{A} \cup \mathsf{C})$.

Here is an example of a qualitative probability structure generated by four worlds:

$$
\begin{array}{r|}
\top \\
\mathsf{A} \vee \mathsf{B}, \mathsf{A} \vee \neg\mathsf{B}, \neg\mathsf{A} \vee \mathsf{B} \\
\mathsf{A}, \mathsf{B} \\
\mathsf{A} \wedge \mathsf{B}, \neg\mathsf{A} \vee \neg\mathsf{B} \\
\neg\mathsf{A}, \neg\mathsf{B} \\
\mathsf{A} \wedge \neg\mathsf{B}, \neg\mathsf{A} \wedge \mathsf{B}, \neg\mathsf{A} \wedge \neg\mathsf{B} \\
\bot
\end{array}
$$

In the particular case of this structure, we can reverse engineer a *unique* probability function P that represents the structure in the sense that for any A, B, $\mathsf{A} \succ \mathsf{B}$ iff $P(\mathsf{A}) > P(\mathsf{B})$. Specifically, we can reason that $P(\mathsf{A} \wedge \mathsf{B}) = 1/2$, since this conjunction is equiprobable with its negation. Furthermore, since this means $P(\neg\mathsf{A} \vee \neg\mathsf{B})$ and all of $\mathsf{A} \wedge \neg\mathsf{B}, \neg\mathsf{A} \wedge \mathsf{B}, \neg\mathsf{A} \wedge \neg\mathsf{B}$ are equiprobable, we can infer that the probability of each of those atoms is 1/6. These facts are sufficient to fully identify P.

In the general case, however, qualitative probability structures carry strictly less information than individual probability functions. In the first instance, the axioms do not guarantee that for every qualitative probability structure, there is a probability measure that represents it (Kraft et al., 1959). That is, they do not guarantee that there is a probability function P' such that for every A, B, $\mathsf{A} \succsim \mathsf{B}$ iff $P'(\mathsf{A}) \geq P'(\mathsf{B})$. Furthermore, when there are probability functions that represent these orderings there generally are *multiple* such functions.

The first of these failures is of direct significance for our account. We have argued that numerical measures, though absent from the positive form, can be compositionally introduced by degree morphemes. The failure of representability teaches us that *sometimes* the measure functions that are compositionally introduced are not probability functions, but measures of a more general sort. In particular, we can say that for every qualitative probability structure $\mathcal{Q}$, there is a function m that represents $\mathcal{Q}$ in the sense above. We can choose this function so as to have additional properties, for example that $m(\top) = 1$ and $m(\bot) = 0$. We cannot, as noted, require it to be additive.[23] But we can require it to satisfy a principle corresponding to qualitative additivity (Holliday and Icard, 2013, §7), namely:

(66) if $\mu(\mathsf{A}) \geq \mu(\mathsf{B}) \Leftrightarrow \mu(\mathsf{A} - \mathsf{B}) \geq \mu(\mathsf{B} - \mathsf{A})$

However, while in general there is no guarantee that the representing function will be a probability function, the failures of representability are rather isolated. What we mean by this is that we can safely assume that in virtually all ordinary discourse, the relevant qualitative probability orderings are of the representable kind (see the related discussion in §10.1 of Holliday and Icard 2013). That is to say

[23]In the strict mathematical sense of 'measure', this means that m is not a measure–for additivity is part of what it is to be a measure. However, in the theory of gradable adjectives 'measure' is used to pick out functions that map their inputs to degrees, which m obviously does.

then that the idea that probabilistic discourse involves probability measures can stand empirically, as a generealization about virtually all discourse, even if it is not guaranteed conceptually.

States-based semantics for probability operators

For the purpose of sketching out a semantics of probability words, we assume that *likely* and *probable* are perfect synonyms, and similarly for *likelihood* and *probability*. For brevity, we focus on the adjective *likely* and on the noun *likelihood.*

At a structural level, our analysis of *likely* and related items mirrors our treatment of *confident*. Claims involving *likely* in the positive form, like (67a), are analyzed as stating the existence of a probability state whose theme is the prejacent, the proposition **snow**. Claims involving *likely* in the comparative, like (68a), are analyzed as stating the existence of such a state and claiming that its measure is greater than δ, here understood as in (69).

(67) a. It is likely that it will snow.

b. $\exists s_v[\theta(s, \mathbf{snow}) \wedge \mathbf{probability}(s) \wedge \mathbf{pos}(s)]$

(68) a. It is more likely that it will snow than that it will rain.

b. $\exists s_v[\theta(s, \mathbf{snow}) \wedge \mathbf{probability}(s) \wedge g(\mu)(s) > \delta]$

(69) $max(\lambda d.\exists s_v[\theta(s, \mathbf{rain}) \wedge \mathbf{probability}(s) \wedge g(\mu)(s) \geq d])$

At the lexical level, we assume that both *likely* and *likelihood* denote properties of states. The difference is that, unlike with *confident* and *confidence*, we assume the background structure against which we evaluate them to be probabilistic. In particular, we assume it to be a qualitative probability structure $\langle W, \succsim\rangle$, represented as an element in the index of evaluation. Our lexical entries for *likely*/*likelihood* are as in (82).

(70) $[\![\text{likely}]\!]^{\langle W,\succsim\rangle,g} = [\![\text{likelihood}]\!]^{\langle W,\succsim\rangle,g}$
$= \lambda s_v : \theta(s) \in Dom(\langle W, \succsim\rangle).\mathbf{probability}(s)$

These entries result very roughly in the following truth-conditions. (67a) says that there is a state of probability whose theme is the proposition that it snows, **snow**, such that **snow** is ranked higher by the background probability structure than the contrast proposition. If we continue operating under the assumption that the contrast proposition to p is $\neg p$, the proposal will generally mirror scalar theories that settle the probability threshold for the truth of *likely p* at 0.5. Meanwhile, (68a) is true if there is a state of probability whose theme exceeds δ, according to the relevant measure (i.e., the maximum degree to which there is a state of probability whose theme is **rain**).

Remarks on the theory

The theory we just sketched is non-committal on how the relevant qualitative probability (henceforth, QP) structure is fixed. In particular, it is compatible with a variety of positions that have been defended in the debate about epistemic modality.[24] Contextualists maintain that the value of the QP structure parameter is fixed (or, in the terminology of MacFarlane (2014), 'initialized') by the context of utterance. Relativists maintain that it is fixed by a second context, i.e. the context in which an utterance is assessed for truth or falsity. Finally, expressivists claim that the value of the QP parameter is not fixed by a context at all. Rather, sentences involving *likely* are assigned contents that directly carve a space of QP structures. Our setup is entirely neutral between these assumptions.[25]

Moving on to a different issue, the theory we just sketched steers an interesting middle course in an important recent debate on the semantics of the language of subjective uncertainty. The question that sparks this debate is whether we should invoke the resource of the mathematical theory of probability when theorizing about probabilistic language. The affirmative answer is typically (though not universally) endorsed by scalar theorists. The negative answer tends to focus on the idea that it is implausible that the language faculty might resort to something as complex as that. Instead, it is generally proposed that the semantics of probability operators be seen as emerging from qualitative comparisons between worlds (Kratzer, 2012; Holliday and Icard, 2013) These comparisons are then 'lifted' to comparisons of probability among propositions. Consequently, some scalar theorists (Lassiter, 2015, see, e.g. the discussion) take qualitative theorists to task for lacking the resources to model claims like *it is 35% likely that it will rain*. Qualitative theorists on the other hand, question whether we need the extremely powerful resources of mathematical probability to model anything except for a few isolated domains of mathematical discourse.

Our approach differs substantially from all of these: unlike the qualitative theorists, we believe that measures are involved when there are explicit degree morphemes. Concretely, the truth-conditions of *it is 35% likely that ...* may well involve a measure function. However, we agree with the qualitative theorists that qualitative probability is just fine for any contexts that do not involve degree-introducing words. Finally, we are not motivated to build our qualitative probability structure out of orderings of worlds. It is acceptable to us to start directly with an ordering of propositions (in other words, we advance a version of what Holliday and Icard (2013) call 'event-ordering semantics').

QP structures differ from confidence structures in that their are subject to further constraints. Earlier on we highlighted that we assumed little about the structure of confidence relations. Not so for probability operators. The requirement that *likely* and *likelihood* have to be interpreted relative to a QP structure imposes tighter constraints. Keeping context fixed, the claims in (71) cannot be true together (since qualitative probability satisfies a form of additivity, hence the probability of $A \& B$ works as a lower bound on the probability of A).

(71) a. It's not likely that Linda is a bankteller.
b. It's likely that Linda is a feminist bankteller.

[24] For some representative papers in this debate, see, among many: DeRose 1991 for the contextualist position; MacFarlane 2009, 2014 for the relativist position; and Yalcin 2007, 2011 for the expressivist position. See also Weatherson and Egan 2011 for a helpful introduction to the debate.

[25] This said, for concreteness it's easier to assume one of these accounts in the background, so sometimes we speak as if the value of the QP is fixed by the context of utterance.

More generally, our semantics vindicates the entire inferential gauntlet laid down by Yalcin (2010). In the Appendix to the paper we have listed all the valid and invalid entailments discussed by Yalcin (the validites are listed in the format Vn, with n a number; the invalidities are listed as In). The central claim, which we merely flash here and defend more fully in the Appendix, is that all of Yalcin's validities are valid and all of Yalcin's invalidities are invalid in our system.

Supporting these claims depends, of course, on a definition of entailment. We define entailment as preservation of truth at context-initialized indices. We complete this definition by assuming that, given a language $\mathcal{L}$, for each qualitative probability structure Q (based on $\mathcal{L}$) and admissible assignment function g (suitable for $\mathcal{L}$), $\langle Q, g \rangle$ is an eligible index—and thus within the scope of the definition of entailment.

As we anticipated in our discussion of confidence reports, we can get extra empirical coverage by relaxing our approach to the **contrast** function. In its more general form, this function inputs a state s with theme p and outputs a set of states $\{s_1, ..., s_n\}$ whose themes are the contrast propositions to p. This allows us to verify the truth of utterances like *It is likely that the Warriors will win the NBA title* even against probability structures in which this prejacent proposition has probability below 50%. We can model this sort of example by assuming that the contrasts of *the Warriors will win the NBA title* are the propositions corresponding to each individual alternative team winning.

This change does have the effect of significantly weakening the logic. Many inferences that are valid under the simpler conception of the **contrast** function become invalid once we generalize the framework in this way. The simplest such inference is the one with premise *Rain is likely* and conclusion *Rain is more likely than no rain*.

5 Vindicating the motivating equivalence

Introducing semantics for belief

Part of our motivation for a states-based analysis of probability operators was the desire to capture the near-equivalences between the reports in (72) and (73) (repeated from above). We first show how our semantics can treat the pairs in question as fully equivalent, and then introduce some hedging for the positive case.

(72) a. I am confident (/have confidence) that it will rain.
b. I believe it is likely that it will rain.

(73) a. I am more confident that Masaya will teach syntax than that he will teach semantics.
b. I believe that it is more likely Masaya will teach syntax than that he will teach semantics.

Any account of these examples must start from an account of belief states. Here we follow the general recipe we have been following and treat attitude verbs like *believe* as expressing properties of states.[26] In particular, the LF corresponding to (72b) is in (74).

[26] See Kratzer 2006 for pioneering work in this theoretical direction.

(74) $\exists s_v[\mathbf{holder}(i,s) \wedge \mathbf{belief}(s) \wedge \theta(s, \mathit{likely(rain)})]$

In addition, to validate the equivalences above, we postulate that there is a kind of semantic interaction between *believe* and *likely*. In particular, when appearing in a belief report, *likely* will not operate on a qualitative probability ordering. Rather, it will operate directly on the confidence state of the subject. This kind of interaction is familiar from recent work on epistemic modality (see e.g. Yalcin 2007, Stephenson 2007), but it needs to be adapted to our neodavidsonian framework. To this end, we make three related nonstandard assumptions.

First, we adopt a somewhat complex model of belief states. In particular, we assume that every belief state is also a confidence state. As a result, having a belief state involves (among other things) having a confidence ranking on propositions. The assumption is motivated by the idea that one cannot believe that it's raining without being confident that it's raining. In further support of the assumption, we note some natural language judgments: *Lara believes that p but she isn't confident that p* sounds contradictory. This suggests (though, of course, doesn't entail) that being in a belief state with content p involves being confident of p.

Our second assumption dovetails with our picture of belief states. On a simple view, the theme of a belief state is just a proposition. We suggest instead that the theme of a belief state is a more complex object, which we can think of as a function from pairs of a $\langle D, \succsim\rangle$-structure and a world to truth-values. Using 'S' for the semantic type of $\langle D, \succsim\rangle$-structures, we say that the type of themes of confidence structures is $\langle\langle S, s\rangle, t\rangle)$.

For illustration, let's consider a couple of concrete examples. (75b) and (76b) are the themes contributed by a simple declarative sentence, (75a), and a sentence involving *likely*, (76a). Both (75b) and (76b) contribute higher-type functions than simple propositions: specifically, they contribute functions from $\langle D, \succsim\rangle$ and worlds to truth values. But this uniformity holds merely for technical reasons. In (75b), the ordering parameter is idle and doesn't bind any variables. Conversely for (76b), where the world parameter is idle.

(75) a. It's raining.

b. $\lambda\langle D, \succsim\rangle.\ \lambda w.\ \exists s_v :\ \mathbf{rain}(s)$ at w

(76) a. It's likely that it's raining.

b. $\lambda\langle D, \succsim\rangle.\ \lambda w.\ \exists s_v :\ \theta(s) \in Dom(\langle D, \succsim\rangle).\ [\mathbf{likelihood}(s) \wedge \theta(s, \mathbf{rain}) \wedge \mathbf{contrast}_C(\mathbf{likelihood})(s) > s'\ ]]$

Nothing in (76b) settles whether the relevant structure is a QP-structure or a confidence structure. (76b) merely imposes a constraint on a $\langle D, \succsim\rangle$-structure. This connects to our third and last assumption. We assume that the theme relation for the case of attitude verbs should be understood in the following way. For a state to have a certain content as theme is for that content to be true at the $\langle D, \succsim\rangle$ structure and all the worlds compatible with the content of that the state determined by that state. To state this formally, we define an auxiliary function: **cstructure** inputs an individual and outputs the confidence structure $\langle D, \succsim\rangle$ of that individual.

(77) $\forall s : s$ is a belief state. $\theta(s, P)$ iff $\forall w'$ compatible with the beliefs of $\mathbf{holder}(s)$, $P(\mathbf{cstructure}(s), w') = 1$

Informally, (77) can be glossed as:

(78) For any belief state s: a content P is theme of s ($\theta(s,p)$) iff p is true at the confidence structure of the holder and the world of the state.

This completes our exposition of the semantics for belief. Next, we briefly discuss some of its consequences. However, the details are pretty involved and since some readers may be unfamiliar with the neodavidsonian system we have been using, in the following section we will restate the basic ideas in an event semantics framework with more familiar notation.

Weakening the equivalence for the positive case

Our semantics is designed to predict that the pairs in (72) and (73) are fully equivalent.

Some readers may take issue with the prediction for the positive case (72). A first worry is that *S believes it's likely that p* might be weaker than *S is confident that p*. *likely* seems to merely requires that the relevant state is ranked above the midpoint of the relevant ordering, but that *confident* requires a higher threshold. Suppose that Irma learns that there is a 51% chance of rain. Then (79a) might be true, while (79b) is false.

(79) a. Irma thinks it's likely that it will rain.
b. Irma is confident that it will rain.

A second worry is that *confident*, but not *likely* might be subject to pragmatic encroachment.[27] Imagine that Oleg, a week away from retirement, considers a bet of all of his retirement savings on an event that has a 80% chance of happening. His friend Laurent might say to Oleg that it is likely that (if he bets) Oleg won't lose his savings. Since Oleg trusts Laurent, (80a) seems true; but (80b) seems false.

(80) a. Oleg thinks it's likely that he would win the bet.
b. Oleg is confident that he would win the bet.

To address these worries, we may simply introduce a minor change in the semantics. We can postulate that the set of states in the extension of *likely* denotes is a superset of the set of states in the extension of *confident*. This assumption dovetails with the idea that *likely* denotes a set of states that are past the midpoint of the relevant structure, while *confident* denotes states that are close to the upper endpoint. As a result, the prediction that the sentences in (79) and (80) are equivalent is weakened to a prediction that they merely have similar meanings.

This change doesn't affect the equivalence of sentences that involve *likely* and *confident* in the comparative form. This seems a welcome prediction, as you can see by comparing the comparative forms of the sentences above.

Semantics for belief in a more traditional framework

It's helpful to see the entry for *believe* we would use in a more traditional system that combines a Hintikka-style possible worlds treatment of belief verbs with davidsonian event variables. In particular, we take as our reference point the system in Hacquard 2006. The innovation we introduce consists in letting the attitude

[27] We thank *[name omitted for blind review]* for raising this possibility.

verb control the ordering that *likely* operates on (again, following domain semantics accounts in the style of Yalcin 2007).

Assume that the interpretation function is relativized to a $\langle D, \succsim \rangle$-structure and a world. We use again 'S' for the semantic type of $\langle D, \succsim \rangle$ orderings, 'P' as a metavariable over functions from $\langle D, \succsim \rangle$-structures to propositions, and '$\textsc{dox}_x$' to denote the set of doxastic alternatives of x (i.e., x's belief worlds). The entry for *believe* is then as in (81).

(81) $[\![\text{believe}]\!]^{g,\langle D,\succeq\rangle} = \lambda s.\ \lambda P_{\langle S,st\rangle}.\ \lambda x.\ \lambda w.\ \mathbf{holder}(x,e) \wedge \mathbf{believe}(e,w) \wedge$
$\forall w' \in \textsc{content}(s)[P(\langle \textsc{dox}_x, \succeq_x\rangle)(w') = 1]$

The clausal arguments of *believe* have the same type as (75b) and (76b).[28] The denotation of a belief report involving *likely* in the complement clause is, schematically, as in (82).

(82) $[\![x \text{ believes likely } p]\!]^{\langle W,\succsim\rangle,g} = \exists s_v[\mathbf{holder}(x,s) \wedge \mathbf{belief}(s) \wedge \forall w' \in \textsc{content}(s):$
$[\exists s' :\ \theta(s') \in Dom(\langle \textsc{dox}_x, \succsim\rangle)\ [\mathbf{likelihood}(s') \wedge \theta(s',p) \wedge$
$\mathbf{contrast}_C(\mathbf{likelihood})(s')]](w') = 1$

6 Conclusion

This paper has considered the semantics of confidence reports across their nominal and adjectival uses. We posited that *confident* and *confidence* express (neodavidsonian) properties of states, type $\langle v,t\rangle$. In the bare, or positive form, an attitude-holder's confidence state with respect to a given proposition p is contrasted with their confidence with respect to (at least in some cases) the proposition $\neg p$. In the comparative form, *more* compositionally introduces a mapping from confidence states to degrees.

Next, we considered how the proposal for confidence reports might extend to better-studied reports headed by gradable adjectives like *likely* and *probable*. Essentially, we applied the (independently motivated) neodavidsonian recipe here, and got some similarly interesting results. First, we found that we could inject modern insights about the semantics of such expressions into our states-based framework. Importantly, the Davidsonian framework allowed us to do this for the nominal forms as well, without doing violence to the background assumptions about the distribution and interpretation of the relevant nouns.

In both cases, we posit a division of labor with regards to degree semantics. Expressions like *confident* and *likelihood* carry along information about the sorts of state structures they are relative to. We exploited these background orderings in our formulation of a (degree-less) $\textsc{pos}$ morpheme, but we only exploited scales (i.e., ordered sets of degrees) when dealing with the comparative operator. In this sense, our architectural proposals echo those made by Bale (2008), who posits a base ordering associated with the adjective.[29]

[28] This may be achieved in various ways. The most obvious one is to adopt a composition rule in the style of Heim and Kratzer's (1998) Intensional Functional Application.

[29] There are some important differences, though. For one thing, Bale's initial ordering is over individuals (here, the holders of the states), and the initial ordering is subsumed in a lexical bundle with a homomorphic mapping to degrees, whereas our structure-preserving map is part of the meaning of *more*.

An interesting aspect of our proposal—a consequence of our adoption of earlier views on the semantics of *more*—is that it leaves open the possibility that different measures may be exploited for the same sentence in different contexts of use—a possibility so far not open to the theorist who lexically encodes a measure function as part of the meaning of the adjective or noun. Such cases have been reported for concrete mass nouns (e.g., *more coffee* can be used to say something about volume *or* weight) and verb phrases (e.g., *run more* can be used to say something about duration or distance). Reported cases with adjectives are thinner on the ground, but Dunbar and Wellwood (2016) offer suggestive examples involving comparatives with *taller*, *redder*, and *more expensive*.

So far, we haven't found a case that makes crucial use of this flexibility in the case of *confident* or *likely*. Such a case would consist in a scenario where the confidence/probability structures are fixed, yet one and the same comparative form could be true or false depending on the choice of measure. Such a case would then represent for *confident*/*confidence* or *likely*/*likelihood* the equivalent of (83) (based on an example from Cartwright 1975), which reverses the order of the expressions introducing the measuranda, yet both can be simultaneously true.

(83) a. There is more water than sand (by volume).
b. There is more sand than water (by weight).

Appendix

Below are the validites and invalidities discussed in Yalcin (2010). One minor differences is that Yalcin presents some of these as inferences, we list everything as a single statement. We convert inferences into material conditional statements whose antecedents are the conjunctions of the premises in Yalcin's inferences and whose consequents are the conclusions of those inferences. We abbreviate *likely A* as $\triangle$A and A *is more likely than* B as A $\geq$ B.

	pattern	gloss
V1	$\triangle A \supset \sim\triangle\sim A$	not both A and $\sim$A are likely
V2	$\triangle(A \,\&\, B) \supset (\triangle A \,\&\, \triangle B)$	distribution over conjunction
V3	$\triangle A \supset \triangle(A \vee B)$	closure under disjunction intro
V4	$A \geq \bot$	A is as likely as $\bot$
V5	$\top \geq A$	$\top$ is as likely as A
V6	$\Box A \supset \triangle A$	necessities are likely
V7	$\triangle A \supset \Diamond A$	likely things are possible
V8	$(\Box(A \supset B)) \supset (\triangle A \rightarrow \triangle B)$	chancy MP
V9	$(\Box(A \supset B)) \supset (\sim\triangle B \rightarrow \sim\triangle A)$	chancy MT
V10	$(\Box(A \supset B)) \supset (B \geq A)$	if to $\geq$
V11	$(B \geq A \;\&\; \triangle A) \supset \triangle B$	Positive form transfer (v1)
V12	$(B \geq A \;\&\; A \geq \sim A) \supset (B \geq \sim B)$	Positive form transfer (v 2)
I1	$(A \geq B \,\&\, A \geq C) \supset (A \geq (B \vee C))$	Union/Disjunction Property
I2	$(A \geq \sim A) \supset (A \geq B)$	if A is likely, A is as likely as anything, v.1.
I3	$\triangle A \supset (A \geq B)$	if A is likely, A is as likely as anything, v.2

Instead of giving complete proofs, we sketch how the they can be derived from the framework in the main text. V1 depends on our claim that **contrast** maps a proposition to its negation *likely* A. If we want to allow different contrast functions, V1 will go invalid, but this is no different from what happens within Yalcin's framework.

V2-V5 all depend only on the fact that we evaluate probability claims against a qualitative probability structure and the fact that we expect that there are as many probability states as there are sets of worlds in the qualitative probability structure. So in particular if there is a probability state whose theme is the conjunction A & B then there must also be a probability state whose theme is A. The inference in V2 can only fail if it is possible to simultaneously have $A \;\&\; B \succ \neg(A \;\&\; B)$ and $(\neg A \succsim A)$; but this is ruled out by the qualitative additivity principle.

For V6-V10, we need a semantics for the modal operators $\Box$ and $\Diamond$.[30] Suppose then that:

$$[\![\Box]\!]^{\langle W, \succsim\rangle, g} = \lambda A_{\langle s,t\rangle}.\forall w \in W, A(w)$$

[30] In V8-V10 Yalcin actually uses conditional premises. We simplify our presentation here by just supposing that the relevant conditionals are strict conditionals. This choice captures the essence of Yalcin's premises, even though it does not capture the full generality of his claim.

$$[\![\Diamond]\!]^{\langle W, \succsim \rangle, g} = \lambda \mathsf{A}_{\langle s,t \rangle}.\exists w \in W, \mathsf{A}(w)$$

Under these assumptions V6-V10 are immediately seen as valid.

For example V6, says that if A is necessary, then it must be likely. But if A is necessary, then it is equivalent, modulo the qualitative probability structure, to the tautological proposition. Qualitative additivity ensures that equivalents must be ranked equally. Moreover, the three principles together entail that $\top$ is likely, so A is likely. This argument also establishes V7 (which is equivalent to $\Box \mathsf{A} \supset \neg \triangle \neg \mathsf{A}$. And indeed, we have already emphasized that when A is necessary $\mathsf{A} \approx \top$ and so cannot be outranked by its negation. Similar additivity-based arguments establish V8-V10.

Finally, the transfer principles are guaranteed by the semantics together with the claim that the measure functions that are invoked in the comparative must respect the qualitative probability ordering. That is to say, the claim that B is more likely than A is true according to our semantics relative to a qualitative probability structure Q if and only if B outranks A in Q. (Call this the *correspondence lemma*.)

The correspondence lemma is sufficient to establish the invalidity of $I1$-$I3$. Consider, for instance, the disjunction inference I1. There are qualitative probability structures that do not satisfy the analogous conditions. Let Q' be such a structure Consider a model on which Q' is the index-initial qualitative probability structure. The the correspondence lemma is going to guarantee that the antecedent of $I1$ is true while its consequent is false.

Bibliography

Peter Alrenga, Chris Kennedy, and Jason Merchant. A new standard of comparison. In Nathan Arnett and Ryan Bennett, editors, *Proceedings of the 30th West Coast Conference on Formal Linguistics*, pages 32–42, Somerville, MA, 2012. Cascadilla Proceedings Project.

Alan Bale. *The universal scale and the semantics of comparison*. PhD thesis, McGill University, Montreal, Quebec, April 2006.

Alan Bale. A universal scale of comparison. *Linguistics and Philosophy*, 31(1):1–55, 2008. doi: 10.1007/s10988-008-9028-z. URL `http://dx.doi.org/10.1007/s10988-008-9028-z`.

Alan Bale and David Barner. The interpretation of functional heads: Using comparatives to explore the mass/count distinction. *Journal of Semantics*, 26(3):217–252, 2009.

Rajesh Bhatt and Roumyana Pancheva. Late merger of degree clauses. *Linguistic Inquiry*, 35(1):1–46, 2004. doi: 10.1162/002438904322793338.

Joan Bresnan. Syntax of the comparative clause construction in English. *Linguistic Inquiry*, 4(3):275–343, 1973.

Helen Cartwright. Amounts and measures of amount. *Noûs*, 9(2):143–164, 1975. doi: 10.2307/2214598.

Hector-Neri Castaneda. Comments on Donald Davidson's 'The logical form of action sentences'. In Nicholas Rescher, editor, *The Logic of Decision and Action*, pages 104–112. Pittsburgh University Press, Pittsburgh, 1967.

Keith DeRose. Epistemic possibilities. *Philosophical Review*, 100(4):581–605, 1991.

Ewan Dunbar and Alexis Wellwood. Addressing the 'two interface' problem: The case of comparatives and superlatives. *Glossa: a journal of general linguistics*, 1 (1):5.1–29, 2016.

Itamar Francez and Andrew Koontz-Garboden. Semantic variation and the grammar of property concepts. *Language*, 91:533–563, 2015. doi: 10.1353/lan.2015.0047.

Valentine Hacquard. *Aspects of Modality*. PhD thesis, MIT, September 2006.

John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173:1393–1404, 2016.

Irene Heim. Notes on comparatives and related matters. Unpublished manuscript, University of Texas, Austin, 1985.

Irene Heim and Angelika Kratzer. *Semantics in generative grammar*. Blackwell, Malden, MA, 1998.

Wesley Holliday and Thomas Icard. Measure semantics and qualitative semantics for epistemic modals. *Proceedings of SALT 2013*, pages 513–534, 2013.

Tomoko Kawamura. *Some Interactions of Focus and Focus Sensitive Elements*. PhD thesis, Stony Brook University, 2007.

Chris Kennedy. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Garland, New York, 1999.

Peter Klecha. *Bridging the divide: scalarity and modality*. PhD thesis, University of Chicago, 2014.

Ewan Klein. A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4:1–45, 1980.

Ewan Klein. The interpretation of adjectival comparatives. *Journal of Linguistics*, 18(1):113–136, 1982.

Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, 107(3): 115–143, 2010. ISSN 0022-362X.

Charles H Kraft, John W Pratt, and Abraham Seidenberg. Intuitive probability on finite sets. *The Annals of Mathematical Statistics*, 30(2):408–419, 1959.

Angelika Kratzer. Severing the external argument from its verb. In Johan Rooryck and Laurie Zaring, editors, *Phrase Structure and the Lexicon*, pages 109–137. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1996.

Angelika Kratzer. Decomposing attitude verbs. Talk given at The Hebrew University of Jerusalem, July 2006. URL `http://semanticsarchive.net/Archive/DcwY2JkM/attitude-verbs2006.pdf`.

Angelika Kratzer. Facts: Particulars or information units? In *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press, 2012.

Daniel Lassiter. *Measurement and Modality: The Scalar Basis of Modal Semantics.* PhD thesis, NYU, 2011.

Daniel Lassiter. Epistemic comparisons, models of uncertainty, and the disjunction puzzle. *Journal of Semantics*, 32:649–684, 2015.

Daniel Lassiter. *Graded Modality*. Oxford University Press, 2016.

John MacFarlane. Epistemic modals are assessment-sensitive. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2009.

John MacFarlane. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press, 2014.

Friederike Moltmann. Degree structure as trope structure: A trope-based analysis of positive and comparative adjectives. *Linguistics and Philosophy*, 32(1):51–94, 2009.

Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, March 2015. doi: 10.3765/sp.8.5.

Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, 2018.

Kimiko Nakanishi. Measurement in the nominal and verbal domains. *Linguistics and Philosophy*, 30:235–276, 2007. doi: 10.1007/s10988-007-9016-8.

Terence Parsons. Events in the semantics of English: A study in subatomic semantics. In *Current Studies in Linguistics Series no. 19*, page 334. MIT Press, Cambridge, Massachusetts, 1990.

Paul Pietroski. *Events and semantic architecture*. Oxford University Press, Oxford, UK, 2005.

Paul Portner and Aynat Rubinstein. Extreme and non-extreme deontic modals. In Nate Charlow and Matthew Chrisman, editors, *Deontic modals*, pages 256–282. 2016.

Jessica Rett. *The Semantics of Evaluativity*. Oxford University Press, 2015.

Paolo Santorio and Jacopo Romoli. Probability and implicatures: A unified account of the scalar effects of disjunction under modals. *Semantics and Pragmatics*, 10, 2017.

Barry Schein. *Plurals and events*. MIT Press, Cambridge, Massachusetts, 1993.

Roger Schwarzschild. The role of dimensions in the syntax of noun phrases. *Syntax*, 9(1):67–110, 2006. doi: 10.1111/j.1467-9612.2006.00083.x.

Stephanie Solt. Q-adjectives and the semantics of quantity. *Journal of Semantics*, 32 (221-273), 2015.

Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525, 2007.

Eric Swanson. How not to theorize about the language of subjective uncertainty. In A Egan and B Weatherson, editors, *Epistemic modality*. Oxford University Press, Oxford, UK, 2007.

Amos Tversky and Daniel Kahneman. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4):293–315, 1983.

Brian Weatherson and Andy Egan. Epistemic modals and epistemic modality. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 1–18. Oxford: Oxford University Press, 2011.

Alexis Wellwood. *Measuring predicates*. PhD thesis, University of Maryland, College Park, August 2014.

Alexis Wellwood. On the semantics of comparison across categories. *Linguistics and Philosophy*, 38(1):67–101, 2015. doi: 10.1007/s10988-015-9165-0.

Alexander Williams. *Arguments in syntax and semantics*, volume Key Topics in Syntax. Cambridge University Press, 2015.

Seth Yalcin. Epistemic modals. *Mind*, 116(4):983–1027, 2007.

Seth Yalcin. Probability operators. *Philosophy Compass*, pages 916–937, 2010.

Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and B. Weatherson, editors, *Epistemic Modality*. Oxford University Press, 2011.